Optimal Execution of a VWAP Order: a Stochastic Control Approach

Christoph Frei
Dept. of Math. and Stat. Sciences
University of Alberta
Edmonton AB T6G 2G1
Canada
cfrei@ualberta.ca

*Nicholas Westray**
Deutsche Bank AG
1, Great Winchester Street
London EC2N 2DB
United Kingdom
nicholas.westray@db.com

This version: April 15, 2013
(Mathematical Finance 25 (2015), 612–639)

Abstract

We consider the optimal liquidation of a position of stock (long or short) where trading has a temporary market impact on the price. The aim is to minimize both the mean and variance of the order slippage with respect to a benchmark given by the market VWAP (volume weighted average price). In this setting, we introduce a new model for the relative volume curve which allows simultaneously for accurate data fit, economic justification and mathematical tractability. Tackling the resulting optimization problem using a stochastic control approach, we derive and solve the corresponding Hamilton-Jacobi-Bellman equation to give an explicit characterization of the optimal trading rate and liquidation trajectory.

Running title: Optimal execution of a VWAP order
Key words: Optimal trade execution, VWAP, HJB equation, gamma bridge

1 Introduction

In investment banks today algorithmic trading is rapidly becoming the preferred method for clients to acquire and liquidate positions of stock. Typically a computer based algorithm is used to buy (or sell) a position while attempting to stick to a client selected benchmark. One of the oldest and most popular of these algorithms is VWAP (volume weighted average price). The popularity of the VWAP benchmark for both brokers and clients stems from several reasons. Firstly, it is very simple to calculate, facilitating easy

*The opinions and ideas expressed in this article are those of the authors alone, and do not necessarily reflect the views of Deutsche Bank AG, its subsidiaries or affiliates. For compliance reasons, all technical questions should be addressed solely to Christoph Frei.

post-trade reporting. Secondly, it encourages the splitting of larger orders into smaller orders, reducing demand for large liquidity and hence market impact/volatility. Finally, given a time interval, it is considered a "fair" benchmark price, in the language of [4], VWAP is a price which

> "...is an unbiased estimate of prices that could be achieved by any randomly selected nonstrategic trader."

Beating market VWAP would thus be considered as a "good" execution, see [25] for further detailed discussion.

The current article proposes a stochastic control approach to tackle the question of how a broker should optimally schedule a VWAP benchmarked trade. In reality the client specifies either a buy (or sell) quantity as well as a start and end time and the broker must then acquire (liquidate) the position attempting to minimize the mean and variance of the difference of the volume weighted price achieved with the market VWAP over the order lifetime (slippage). The main motivation for the present research is attempting to improve execution efficiency. Due to the huge notional volumes being traded algorithmically, small gains arising from the application of results obtained can lead to substantial increase in profits for both brokers and their clients.

The question of optimal execution with an arrival price benchmark is well studied in the literature going back to [5], see also the seminal papers of Almgren and co-authors [1], [2] and [3]. In contrast, perhaps due to the stochastic nature of the benchmark, there is significantly less literature related to the present problem. The first work in this area is the article by Konishi [23] who derives the optimal execution trajectory for single and basket VWAP executions when the price is given by a Brownian motion. The strategy is then assessed against actual trade data from the Tokyo stock exchange. Following this article VWAP tracking has been attacked using a variety of different methods. McCulloch and Kazakov [26] view it as a quadratic hedging problem under partial information whereas Kakade et al. [21] and Białkowski et al. [6] use online learning and dynamic volume approaches and Humphery-Jenner [20] gives a VWAP trading rule which takes intraday noise into consideration. Finally Bouchard and Dang [8] formulate it as a stochastic target problem and derive a viscosity solution characterization of the value function. Note that the above articles (excluding [8]) do not take into account the market impact of a trade and none of them impose any parametric structure on the intraday volume curve.

The present article has several contributions. Firstly, we extend previous literature in this area by allowing for a linear temporary market impact model. More general models of price impact have been studied both theoretically, e.g. by Gatheral [14], and empirically, e.g. by Bouchaud et al. [9] (see Gatheral and Schied [16] for a good overview) however the linear model leads to a tractable problem. Secondly, we provide a parametric model for relative volume which fits real data well, reflecting meaningful underlying economic assumptions and simultaneously being tractable enough to perform optimization. Finally, although the optimization problem is involved due to the use of VWAP as benchmark, we are able to explicitly characterize the optimal control thus providing a closed-form solution for the optimal trading rate. This final result opens up, for the first time, a rigorous mathematical approach to the determination of commission for guaranteed VWAP trades, similar to that done for Implementation Shortfall in [2].

The paper is organized as follows. In the next section, we introduce and justify our model which uses VWAP as a benchmark in optimal trade execution. We present the main result in Section 3, Theorem 3.1, which provides the explicit solution for the optimal trading rate. Deferring the proof to Section 5, we explain in Section 3 two crucial properties of the optimal trading rate. Its sign can switch only once from negative to positive and never the other way, and the optimal trading rate can be decomposed into two parts, with one being a deterministic TWAP (time weighted average price) strategy and the other reflecting the adjustment necessary due to jumps in the relative volume curve. In Section 4, we show and discuss how well the parametric model fits to trading volume.

2 A Framework for Using a VWAP Benchmark

Here we describe the model formulation and the key assumptions. We begin with our trading strategies, without loss of generality we consider a buy program for Y shares. The situation of a sell can be considered by reversing the time.

2.1 Trading Costs

We are given a start and end time by the client which we assume (without loss of generality) to be given by $t_0 = 0$ and T, respectively. We must complete the purchase of stock by T and will be benchmarked to the market VWAP over the period $[0, T]$, for simplicity the reader may think of $T = 1$, corresponding to a day VWAP order. As is standard in the literature, we work on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, P)$ satisfying the usual hypotheses of right continuity and completeness, and we denote by $X^u(t)$ our share holdings at t where the trading strategy u is an adapted and integrable process typically referred to as trading rate. In particular our holdings evolve according to

$$\mathrm{d}X^u(t) = u(t)\,\mathrm{d}t, \qquad X^u(0) = 0, \qquad X^u(T) = Y.$$

The asset price process $\big(P(t)\big)_{0\le t\le T}$ is assumed to be an arithmetic Brownian motion

$$P(t) = P(0) + \sigma W(t),$$

where $\sigma > 0$ represents the daily volatility in dollars and $\big(W(t)\big)_{0\le t\le T}$ is a standard Brownian motion on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{0\le t\le T}, P)$. Using an arithmetic as opposed to a geometric Brownian motion in the present setup is possible since we are considering an intraday trading horizon, so that $P(t)$ is negative with an extremely small probability and there are negligible differences between the two models, see Gatheral and Schied [15].

As is well known, when we trade, we do not realize the above price $P(t)$ (which could be thought of as the midquote) but instead we pay

$$P^u(t) = P(t) + \kappa u(t),$$

where κ is the coefficient of the linear (temporary) market impact model. It represents an instantaneous linear premium on the price due to how fast we trade. This model (amongst many more complicated ones) is studied in detail in the literature, see Gatheral [14] or Bouchaud et al. [9]. The linear form is necessary as it leads to quadratic trading costs which then produces a tractable problem. However it is also important to note that given the extremely low predictive accuracy of market impact models (typically < 5% R^2), the cost of increased complexity arising from moving away from a linear model would outweigh any gains from better describing market impact. Moreover, a linear price impact is supported in the recent empirical study by Cont et al. [11]. The above will introduce some dependence on impact into the solution, which is its purpose.

Similarly to Section 1.1 of Almgren [1], our total expenditure TE^u to buy the shares Y using a control u is thus given by

$$\mathrm{TE}^u = \int_0^T P^u(t)\,\mathrm{d}X^u(t) = YP(0) + \sigma \int_0^T W(t)u(t)\,\mathrm{d}t + \kappa \int_0^T u^2(t)\,\mathrm{d}t,$$

where we applied the relation $\int_0^T u(t)\,\mathrm{d}t = X^u(T) - X^u(0) = Y$. Using $X^u(T) = Y$ and that X^u is of finite variation, the product rule yields

$$\int_0^T W(t)u(t)\,\mathrm{d}t = \int_0^T W(t)\,\mathrm{d}X^u(t) = -\int_0^T X^u(t)\,\mathrm{d}W(t) + YW(T)$$

so that

$$\mathrm{TE}^u = YP(0) - \sigma \int_0^T X^u(t)\,\mathrm{d}W(t) + \sigma YW(T) + \kappa \int_0^T u^2(t)\,\mathrm{d}t. \qquad (1)$$

2.2 The VWAP Benchmark

Given a series of prices $(P_i)_{i=1,\dots,N}$ together with volumes $(V_i)_{i=1,\dots,N}$ executed at those prices, the VWAP is defined to be

$$\mathrm{VWAP} = \frac{\sum_{i=1}^N V_i P_i}{\sum_{i=1}^N V_i}.$$

If we define $\tilde{V}_i = \sum_{j=1}^{i} V_j$ as the cumulative volume (with $\tilde{V}_0 = 0$), we have

$$\text{VWAP} = \frac{\sum_{i=1}^{N}(\tilde{V}_i - \tilde{V}_{i-1})P_i}{\tilde{V}_N} \approx \int_0^T P(t)\frac{\mathrm{d}\tilde{V}(t)}{\tilde{V}(T)}.$$

In particular, to model VWAP, we need a continuous-time process for $\frac{\tilde{V}(t)}{\tilde{V}(T)}$. This will be nondecreasing and satisfy $\frac{\tilde{V}(T)}{\tilde{V}(T)} = 1$ as well as $\frac{\tilde{V}(0)}{\tilde{V}(T)} = 0$. A natural process to use is the gamma bridge.

Definition 2.1. *1) A* gamma process $\big(L(t)\big)_{0\leq t\leq T}$ *is a process with independent and identically distributed increments such that $L(0) = 0$ and $L(t)$ is gamma distributed with mean $mt\theta$ and variance $mt\theta^2$ for some $m > 0, \theta > 0$.*
2) For a gamma process $\big(L(t)\big)_{0\leq t\leq T}$, the gamma bridge $\big(\gamma(t)\big)_{0\leq t\leq T}$ *is defined by $\gamma(t) = L(t)/L(T)$.*

There are several reasons for our choice of modelling the intraday relative volume curve by a gamma bridge. Firstly, we will see in Section 4 that our model fits well to real stock data provided the stock is sufficiently liquidly traded as well as being finite variation, like real data. Secondly, we can think of the cumulative trading volume as analogous to the accumulation of dam rain, similarly to Gani [13] considering the arrival of insurance claims as analogous to the accumulation of dam rain. The latter can be modelled by a gamma process as pointed out by Moran [27] so that similarly the relative amount will be a gamma bridge. Finally, we can prove that the intraday volume curve must be a gamma bridge if we assume that trading volume is independent and stationarily distributed through the day and the relative intraday volume is independent of the total volume. This link is based on the following theoretical result on gamma processes.

Proposition 2.2. *Let $\big(L(t)\big)_{0\leq t\leq T}$ be a Lévy process with $L(T) > 0$ a.s. and non-deterministic (i.e., $P[L(T) = c] < 1$ for all c). Then the following are equivalent:*

(i) L is a gamma process;

(ii) there exists $t \in (0, T)$ such that $L(T)$ and $L(t)/L(T)$ are independent;

(iii) for all $t \in [0, T]$, $L(T)$ and $L(t)/L(T)$ are independent.

In particular, the proposition shows that the gamma process is the only positive Lévy process whose intermediate relative values are independent of the terminal value. We immediately get from Proposition 2.2 the following application to the relative volume curve.

Corollary 2.3. *Assume that the cumulative trading volume has independent and stationary increments, its terminal value is non-deterministic and strictly positive (i.e., trading happens a.s.) and the relative volume curve is independent of the total trading volume. Then the cumulative volume is a gamma process and the relative volume curve is a gamma bridge.*

The assumption that the relative volume curve is independent of the total volume is not too unrealistic; it is hard to imagine that the relative volume curve on a given day depends significantly on the total traded volume on that day. This is supported by the left panel of Figure 1 for a major US and European stock. A consequence of modelling volume with a gamma process and prices with a Brownian motion is their mutual independence; see Lemma 15.6 of Kallenberg [22]. The right panel of Figure 1 shows that the intraday correlation between volume and price changes varies a lot but is small on average. Hence it is difficult to incorporate in a model, leading to only a small potential increase in performance and thus justifying the independence assumption.

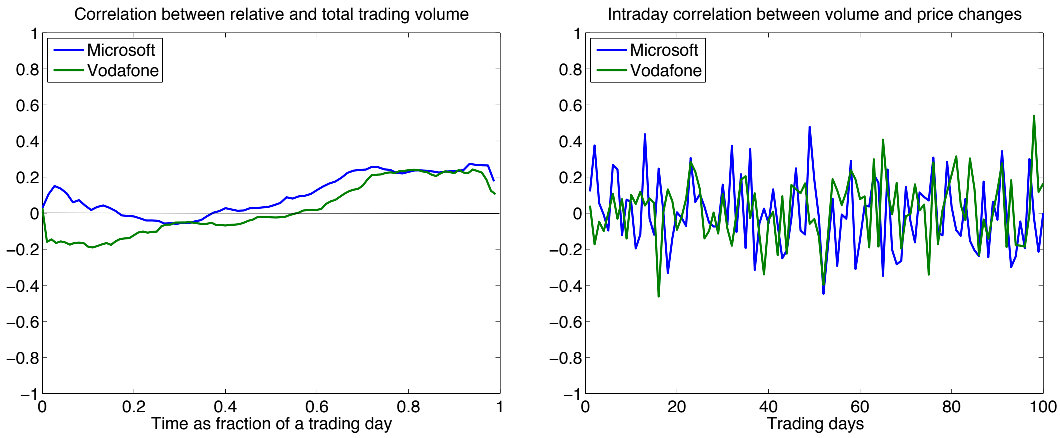

Figure 1: Correlation between relative and total trading volume (left panel) and intraday correlation between volume and price changes (right panel) for the stocks of Vodafone and Microsoft during the first 100 trading days of 2012, using 5-minute intraday data.

Looking at the historical volume data and noting its episodic nature the hypothesis that volume is independent and stationarily distributed throughout the day is idealistic. However, we can think of the gamma bridge as a zero-order approximation to the true relative volume curve, which would be exact should such an assumption hold. This then provides a tractable basis upon which one can incorporate more realistic features. As a final point, let us note that the fit to actual data is sufficiently good to conclude that the effect of deviation from such a hypothesis is (perhaps surprisingly) not too severe.

Proof of Proposition 2.2. The implication "$(i) \implies (iii)$" is shown in Proposition 3.2 of Brody et al. [10], and (iii) clearly implies (ii) so that it remains to show "$(ii) \implies (i)$". Let $t \in (0, T)$ be such that $L(T)$ and $L(t)/L(T)$ are independent. Since L is a Lévy process, it is enough to show that $L(t)$ is gamma distributed. If we set $A = L(t)$ and $B = L(T) - L(t)$, we have that A, B are strictly positive, non-deterministic and independent. By assumption, $A + B$, $A/(A + B)$ are independent, hence so are $A + B$, A/B by using the measurable mapping $x \mapsto x/(1-x)$. It follows from Lukacs' proportion-sum independence theorem [24] that both A and B have gamma distributions. $\square$

Since we will only consider a gamma bridge but not its underlying gamma process, by scaling we can (and henceforth do) set $\theta = 1$ without loss of generality. Moreover, in the later simulations, we fix $T = 1$ and use m as the model parameter. Under the assumption of gamma bridge for the relative volume curve, the expenditure (in dollars) to buy Y shares of market VWAP is given by

$$\begin{aligned} \mathrm{VWAP} &= \int_0^T P(t)\,\mathrm{d}\big(Y\gamma(t)\big) = YP(0) + \sigma Y \int_0^T W(t)\,\mathrm{d}\gamma(t) \\ &= YP(0) - \sigma Y \int_0^T \gamma(t-)\,\mathrm{d}W(t) + \sigma Y W(T), \end{aligned}$$

where we argued similarly to the derivation of (1). The VWAP benchmark should be thought of as an average market price over the lifetime of the order and it ignores the market impact because the mid quote represents a reasonable estimate of the current fair market price. Observing that we can scale by $1/Y$, we thus can (and do) assume that the client wants to buy 1 share.

One may object to the model formulation in that $L(T)$ is unknown. Whilst this is a valid objection, it misunderstands the objective of the present article. All major brokers who engage in algorithmic trading require models for the intraday relative volume curve. The primary application for such models is in the execution of VWAP and PVol (percentage of volume) orders. When evaluating the viability of such a model in a trading application, one can compare the performance when using the new model against that in the "perfect information" case, modelled here by a gamma bridge. The framework presented here allows closed-form solutions for this case and thus allows new volume curve models to be assessed in a way relevant to their use rather than just a pure goodness-of-fit test. A second benefit is in post-trade analysis, the solution presented here provides a broker independent benchmark for the execution of VWAP trades allowing the scope for relative comparison as well as the common absolute performance. We will address the issues of comparing and evaluating different VWAP strategies in future work. In summary, the focus here is not a real-time (i.e. based on an adapted estimator of relative volume) strategy for trading VWAP but rather an effort to set up a framework in which one can get closed-form and implementable solutions to the VWAP trading problem which are of use for setting upper bounds on performance and providing broker independent benchmarks. As a consequence, the filtration $(\mathcal{F}_t)_{0\le t\le T}$ we are working with is such that the gamma bridge $\big(\gamma(t)\big)_{0\le t\le T}$ and the Brownian motion $\big(W(t)\big)_{0\le t\le T}$ are adapted to it.

The reader will notice that our trading is not taken into account when calculating the relative volume curve, so that it becomes exogenous. This is clearly a simplification and is effectively equivalent to assuming that the (unnormalized) size Y is small relative to the total volume traded, i.e. that we are dealing with small orders. Since typical algorithmically traded VWAP orders are traded with an effective (when scaled to a day) size of 10–15% of the daily volume this assumption is not too restrictive and hugely simplifies the problem under consideration. As explained in Hu [19], who analyzes the difference in VWAP with and without own trading, it might even be desirable to exclude our own trading from the market VWAP calculation in order to compare our performance with a benchmark which is not affected by us. The interesting case of large VWAP trades (>50% daily volume) will be considered elsewhere.

Remark 2.4. It is important to be precise about what is meant here by "volume". In recent years the number of trading facilities with visible order books (lit venues) has increased substantially. For example in 2008 one could only trade Vodafone Group PLC (VOD) on the LSE (its primary exchange) whereas today one can trade this on BATS, CHI-X and Turquoise (amongst others). Since most modern smart order routers are able to access these MTFs (Multilateral Trading Facilities), it makes sense to consider everything in a consolidated view, so that (for example) day volume in this article should be understood as the total volume traded in a day on all (lit) venues. $\diamond$

2.3 Algorithm Performance

Now that the quantity of shares is normalized to $Y = 1$, our costs can be thought of as per share in dollars. Our first aim is to minimize the expected slippage, defined as

$$\mathrm{slip}^u = \mathrm{TE}^u - \mathrm{VWAP} = \sigma \int_0^T \big(\gamma(t-) - X^u(t)\big)\,\mathrm{d}W(t) + \kappa \int_0^T u^2(t)\,\mathrm{d}t.$$

From a broker's perspective, simply having small expected slippage, whilst good, is not the sole goal. From a post-trade perspective, it will be necessary to explain to clients why slippage was significantly far from that expected, should that occur. Our second aim is therefore to attempt to minimize the variance of the slippage. We use the approximation

$$\begin{aligned}\mathrm{Var}[\mathrm{slip}^u] &\approx E\bigg[\bigg(\sigma \int_0^T \big(\gamma(t-) - X^u(t)\big)\,\mathrm{d}W(t)\bigg)^2\bigg] \\ &= \sigma^2 E\bigg[\int_0^T \big(\gamma(t) - X^u(t)\big)^2\,\mathrm{d}t\bigg],\end{aligned}$$

where we assumed for the last equation that $\int X^u\,\mathrm{d}W$ is a square integrable martingale and used that γ is bounded and has only countably many jumps.

The main reason for this approximation is tractability, without this the difficulty would be increased significantly. The mathematical reason for the difficulty of the original problem is its time inconsistency. As explained in Section 1.2 of Björk and Murgoci [7], time inconsistence of mean-variance problems of the form $E[X] + \lambda \mathrm{Var}[X]$ is caused by the term $(E[X])^2$ in $\mathrm{Var}[X] = E[X^2] - (E[X])^2$. While standard time consistent problems are allowed to have expected values of nonlinear functions such as $E[X^2]$, the term $(E[X])^2$ is a nonlinear function of the expected value and not an expected value of a nonlinear function. However, in our particular problem, the time inconsistency is mild in the sense that the value of $\mathrm{Var}[\mathrm{slip}^u]$ is close to the variance of $\int_0^T \big(\gamma(t) - X^u(t)\big)\,\mathrm{d}W(t)$ which has zero mean, hence leading to a time consistent formulation. Therefore, such an approximation is appropriate in our situation.

Hence, we study the minimization of

$$\kappa\, E\bigg[\int_0^T u^2(t)\,\mathrm{d}t\bigg] + \lambda\,\sigma^2 E\bigg[\int_0^T \big(\gamma(t) - X^u(t)\big)^2\,\mathrm{d}t\bigg]$$

as an approximation of the mean-variance problem

$$\inf_u \big(E[\mathrm{slip}^u] + \lambda \mathrm{Var}[\mathrm{slip}^u]\big)$$

for a given mean-variance tradeoff parameter $\lambda > 0$. In further justification of this approximation we note that similar to Almgren [1] and Konishi [23] it is the case that the main driver of $\mathrm{Var}[\mathrm{slip}^u]$ is typically the volume curve γ and not the trading rate u, so that we capture the dominant term.

To get an idea of the approximation error, we used a Monte Carlo simulation to compare $\mathrm{Var}[\mathrm{slip}^u]$ and $\sigma^2 E\big[\int_0^T \big(\gamma(t) - X^u(t)\big)^2\,\mathrm{d}t\big]$. Our calculation showed that, when using the optimal strategy $\hat{u}$ for the approximated problem from Theorem 3.1, the relative error

$$\frac{\mathrm{Var}[\mathrm{slip}^{\hat{u}}] - \sigma^2 E\big[\int_0^T \big(\gamma(t) - X^{\hat{u}}(t)\big)^2\,\mathrm{d}t\big]}{\mathrm{Var}[\mathrm{slip}^{\hat{u}}]}$$

was always very small. For example, choosing $\sigma = 0.01$, $\kappa = 10^{-8}$ (see Remark 2.5), $\lambda = 1$, $m = 25$ and $T = 1$, the relative error was less than $10^{-3} = 0.1\,\%$.

Remark 2.5. It is a simple exercise to calibrate a market impact model and determine κ, however as a first order approximation we simply use the trading rule of thumb, see [14] (amongst others), that trading one day's volume costs approximately one day's volatility in basis points. This is quite accurate in most cases and suffices for our purposes. In particular this implies the following relation for our linear impact model:

$$\kappa \int_0^T u^2(t)\,\mathrm{d}t \approx \sigma V$$

where V is the daily volume (shares) and σ is again the volatility in (\$). Assuming a linear execution for V (and noting $T = 1$), this reduces to $\kappa \approx \frac{\sigma}{V} \approx 10^{-8}$ (for a stock like Vodafone). $\diamond$

Dynamically formulated, the *value function* for the optimization problem is given by

$$v(t,x,\gamma) = \inf_u E\left[\kappa \int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T \big(\gamma(s) - X^u(s)\big)^2\,\mathrm{d}s\right], \qquad (2)$$

where $\big(\gamma(s)\big)_{t\le s\le T}$ is a gamma bridge with $\gamma(t) = \gamma$ and the infimum is over all adapted and integrable u with

$$\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \qquad X^u(t) = x, \qquad X^u(T) = 1.$$

In conclusion, we are considering a mean-variance formulation of minimizing slippage from VWAP.

Remark 2.6. The framework here is similar to the seminal paper of Almgren and Chriss [2]. The key extra technical difficulty in the present case comes from the switch from an arrival price benchmark ($P(0)$) to a VWAP benchmark (VWAP). This introduces extra stochastic complexity coming from the gamma bridge.

Under a gamma bridge $\big(\gamma(s)\big)_{t\le s\le T}$ with $\gamma(t) = \gamma$ we understand a process of the form $\gamma(s) = \gamma + \frac{L(s)-L(t)}{L(T)-L(t)}(1-\gamma)$ for a gamma process L so that starting with γ at t, we take the remaining part $1-\gamma$ proportional to the remaining relative portion of L. Note that $\frac{L(s)-L(t)}{L(T)-L(t)}$, $t \le s \le T$, is again a gamma bridge and independent of $\frac{L(t)}{L(T)}$ by page 673 of Émery and Yor [12].

The reader will notice that we have not considered auctions in the current model. Indeed we have formulated the VWAP tracking problem for orders executed in the so called "continuous" (non-auction) trading phase. For US stocks it would be a reasonable approximation to simply ignore auctions, since the average volumes traded there are a small percentage of total volume. In Europe where auction volumes are typically much higher, one could imagine a small modification to the above framework where a dynamic model would be used to estimate the fractions executed in the open and close auction and the above model could then be used to execute the remainder. Since the focus of this article is not the prediction of auction volumes, we assume that the historical mean values have been used and the total amount to be traded has been reduced accordingly in a preliminary step so that the results presented here apply equally to US and European stocks. $\diamond$

3 Main Result

Our main result is an explicit characterization of the value function v and the optimal control in (2).

Theorem 3.1. *The value function v is given by*

$$v(t,x,\gamma) = a(t)x^2 + b(t)\gamma x + c(t)x + d(t)\gamma^2 + f(t)\gamma + g(t)$$

for $t \in [0,T)$, $x \geq 0$ *and* $\gamma \in [0,1]$ *with the functions* $a, b, c, d, f, g : [0,T) \to \mathbb{R}$ *defined by*

$$
\begin{aligned}
a(t) &= \sqrt{\kappa\lambda\sigma^2}\,\frac{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} + \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}}{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}},\ \ b(t) = -2a(t) + \frac{2\kappa}{T-t},\ \ c(t) = -\frac{2\kappa}{T-t},\\
d(t) &= \int_t^T \Big(\lambda\sigma^2 - \frac{1}{4\kappa}b^2(s)\Big)\frac{T-s}{T-t}\,\frac{T-s+1/m}{T-t+1/m}\,\mathrm{d}s,\\
f(t) &= \frac{1}{T-t}\int_t^T \Big(b(s) + 2d(s)\frac{T-s}{T-s+1/m}\Big)\mathrm{d}s,\\
g(t) &= a(t) - (T-t)\lambda\sigma^2\\
&\quad - \int_t^T \bigg(\frac{b(s)}{2}\Big(\frac{a(s)}{\kappa}(T-s) + 1\Big) - f(s) - \frac{d(s)/m}{T-s+1/m}\bigg)\frac{1}{T-s}\,\mathrm{d}s.
\end{aligned}
$$

The optimal control $\hat{u}$ *and the corresponding share holdings* $\hat{X}$ *are given by*

$$
\begin{aligned}
\hat{u}(s) &= -\frac{v_x}{2\kappa}\big(s, \hat{X}(s), \gamma(s)\big) = -\frac{2a(s)\hat{X}(s) + b(s)\gamma(s) + c(s)}{2\kappa}, \qquad (3)\\
\hat{X}(s) &= x\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} - \frac{1}{2\kappa}\int_t^s \big(b(z)\gamma(z) + c(z)\big)\exp\Big(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\Big)\,\mathrm{d}z.
\end{aligned}
$$

We postpone the proof of Theorem 3.1 to Section 5 and instead give a non-rigorous motivation.

Suppose we replace (2) by the approximate problem,

$$
v^n(t,x,\gamma) = \inf_u E\bigg[\kappa\int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^T \big(\gamma(s) - X^u(s)\big)^2\,\mathrm{d}s + n\big(X^u(T) - 1\big)^2\bigg], \qquad (4)
$$

where the infimum is over all adapted and integrable u with

$$
\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \qquad X^u(t) = x.
$$

Observe now that X^u has no fixed terminal condition. Using that $\gamma(./m)$ is a gamma bridge on $[tm, Tm]$ with underlying mean growth rate equal to 1, it follows from Corollary 1 of Émery and Yor [12] that the infinitesimal generator of the gamma bridge γ is given by

$$
m\int_0^1 \Big(f\big(\gamma(t) + \big(1-\gamma(t)\big)z\big) - f\big(\gamma(t)\big)\Big)(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z,
$$

where f is a function on $\mathbb{R}_+$ with bounded variation on compacts. If we assume that v^n is sufficiently regular, it should satisfy

$$
\begin{aligned}
&v_t^n + \lambda\sigma^2(\gamma - x)^2 + \inf_{u\in\mathbb{R}}(v_x^n u + \kappa u^2)\\
&\quad + m\int_0^1 \big(v^n\big(t,x,\gamma + (1-\gamma)z\big) - v^n(t,x,\gamma)\big)(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z = 0,\\
&v^n(T,x,\gamma) = n(x-1)^2.
\end{aligned}
$$

Due to the quadratic structure, we propose the ansatz

$$v^n(t,x,\gamma) = a^n(t)x^2 + b^n(t)\gamma x + c^n(t)x + d^n(t)\gamma^2 + f^n(t)\gamma + g^n(t).$$

After some algebra, we derive that $a^n, \ldots, g^n$ should solve

$$\begin{aligned}
a^n_t - \frac{1}{\kappa}(a^n)^2 + \lambda\sigma^2 &= 0, & a^n(T) &= n,\\
b^n_t - \frac{1}{\kappa}a^n b^n - 2\lambda\sigma^2 - b^n\varphi_0 &= 0, & b^n(T) &= 0,\\
c^n_t + b^n\varphi_0 - \frac{1}{\kappa}a^n c^n &= 0, & c^n(T) &= -2n,\\
d^n_t - 2d^n\varphi_0 + d^n\varphi_1 - \frac{1}{4\kappa}(b^n)^2 + \lambda\sigma^2 &= 0, & d^n(T) &= 0,\\
f^n_t - \frac{1}{2\kappa}b^n c^n + 2d^n\varphi_0 - f^n\varphi_0 - 2d^n\varphi_1 &= 0, & f^n(T) &= 0,\\
g^n_t - \frac{1}{4\kappa}(c^n)^2 + f^n\varphi_0 + d^n\varphi_1 &= 0, & g^n(T) &= n,
\end{aligned}$$

where $\varphi_0(t) = \frac{1}{T-t}$ and $\varphi_1(t) = \frac{1}{T-t} - \frac{1}{T-t+1/m}$. It is intuitively sensible to expect that $v^n \uparrow v$ as well as that $a^n \to a$, $b^n \to b$ etc. However if we let $n \uparrow \infty$ we would end up with difficult singular conditions in c and g. To avoid this, observe that the ratio $\frac{a^n(T)}{c^n(T)}$ and the difference $a^n(T) - g^n(T)$ are all constant and independent of n. Due to the convergence we would expect this to hold in the limit, this is precisely what the functions a, c and g satisfy; see Lemma 5.1 below. A second justification is to see that as $n \uparrow \infty$ the terminal condition behaves like a quadratic function in x with one root at 1 which is infinite elsewhere. That is to say we have

$$\lim_{t \nearrow T}\big(a(t)x^2 + c(t)x + g(t)\big) = \lim_{t \nearrow T} v(t,x,0) = \infty$$

for all $x \neq 1$. We thus expect that g behaves like a, and c behaves like $-2a$ for $t \nearrow T$, again consistent with Lemma 5.1.

Remark 3.2. Note that our main result gives explicit formulae for both the optimal trading rate $\hat{u}$ and the optimal holdings $\hat{X}$. This desirable formula for $\hat{X}$ is due to our requirement that the holdings be absolutely continuous with respect to t. In contrast, this is not enforced in McCulloch and Kazakov [26] so that one may not directly compare their results with ours.

Observe that the optimal trading rate depends on the volume curve but not on the price process. We give some intuition as to why this is a natural consequence of using a VWAP optimization criterion together with a Brownian price process. When comparing the VWAPs of two strategies, only price movements but not the absolute level of the price process are relevant. Since in our model the price movements are independent from past prices (they are given by Brownian increments), information about past prices is not included in the optimal strategy. $\diamond$

Let us now describe in further detail the structure of the optimal control. It is intuitively clear that there should be a buy and sell region; more precisely from (3) we can see that the sign of $\hat{u}(s)$ depends on $\hat{X}(s)$. In particular, $\hat{u}(s)$ is positive if and only if $\hat{X}(s) < \frac{-b(s)\gamma(s)-c(s)}{2a(s)}$. Indeed, if we have low holdings $\hat{X}(s)$, we will make purchases to come closer to our target while for high $\hat{X}(s)$, it can be beneficial to temporarily reduce the holdings to come closer to $\gamma(s)$. This leads us to define the *frontier* ζ by

$$\zeta(t,\gamma) = \frac{-b(t)\gamma - c(t)}{2a(t)}, \quad t \in [0,T], \gamma \in [0,1]$$

so that we have

$$\begin{aligned}
\hat{u}(t) < 0 \text{ on } \hat{X}(t) > \zeta\big(t,\gamma(t)\big),\\
\hat{u}(t) = 0 \text{ on } \hat{X}(t) = \zeta\big(t,\gamma(t)\big),\\
\hat{u}(t) > 0 \text{ on } \hat{X}(t) < \zeta\big(t,\gamma(t)\big).
\end{aligned}$$

A key practical requirement is that for the original buy program (i.e. when we start with 0 shares) we should not sell. Indeed most clients would typically be unhappy if during their execution the stock holdings were not monotone increasing. In addition, in US markets such behaviour is actually prohibited by the regulators.[1] The following proposition shows that in our formulation, for any parameter values, this is indeed the case, underlining the model's applicability and compliance with this important regulatory aspect.

Corollary 3.3. *Both partial derivatives of ζ are positive on $[0,T) \times [0,1)$ and it holds that $\zeta(0,0) > 0$. Starting with $\hat{X}(0) = 0$ and $\gamma(0) = 0$, the process $\hat{u}(t)$ is nonnegative for all t.*

We postpone the proof to Subsection 5.3. The idea of the second part is that when the process $\hat{X}(t)$ is below $\zeta\big(t,\gamma(t)\big)$, it will not cross $\zeta\big(s,\gamma(s)\big)$ at a later time point due to the properties of the frontier ζ. Figure 2 illustrates this behaviour of $\hat{X}$, that it can cross ζ only from above but not from below. We can also see that for different starting points, the paths of $\hat{X}$ are on very similar trajectories after only a short time. This is due to the multiplication by $\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r}$ of x in the definition (3) of $\hat{X}$, which makes the impact of the starting value x vanish soon since a diverges to ∞.

Recalling that the gamma bridge satisfies $E[\gamma(t)] = \frac{t}{T}$, a natural question to ask is how our current optimization is related to the deterministic problem obtained by replacing $\gamma(t)$ by its mean, namely

$$\inf_u \left(\kappa \int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T \Big(\frac{s}{T} - X^u(s)\Big)^2 \mathrm{d}s \right), \tag{5}$$

[1] see FINRA directive 5310 on Best Execution and Interpositioning, http://finra.complinet.com/en/display/display_main.html?rbid=2403&element_id=10455

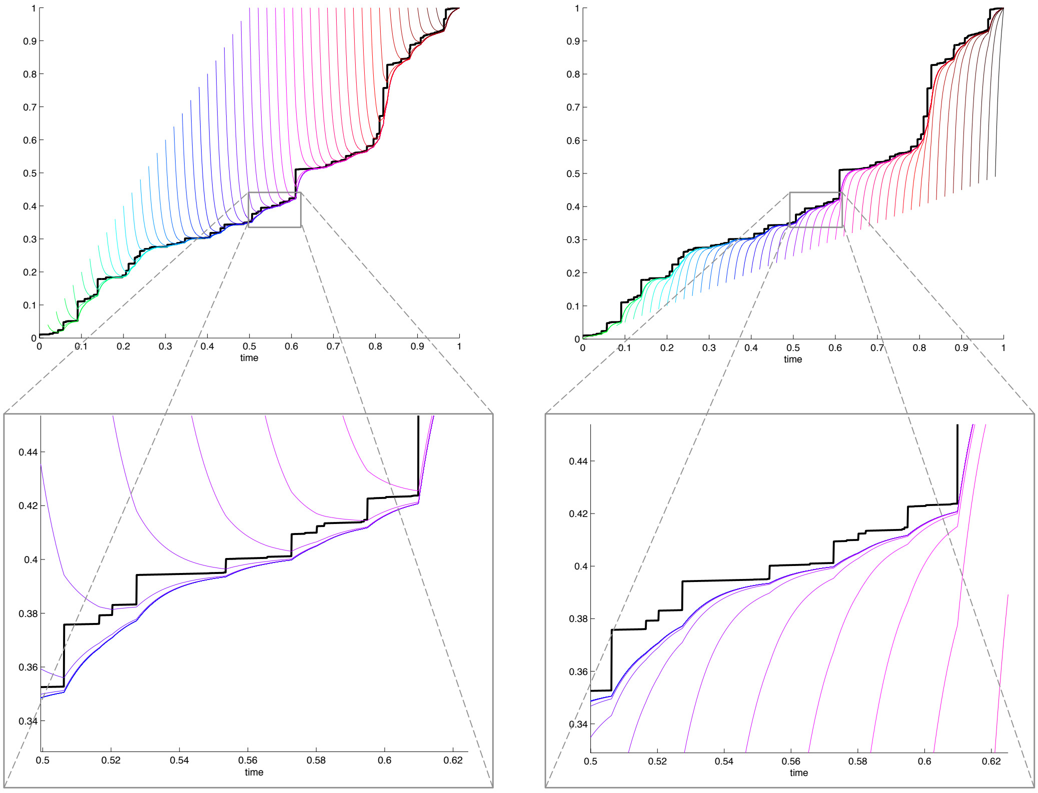

Figure 2: Simulations of $\hat{X}$: each colored curve corresponds to and starts at one initial condition $(t, \hat{X}(t))$ while the path of the gamma bridge is always the same. $\hat{X}$ can cross the frontier ζ (black curve) only from above (left panels) but not from below (right panels). The parameters are $\kappa = 10^{-8}$ (compare Remark 2.5), $\sigma = 0.01$, $\lambda = 1$, $m = 25$ and $T = 1$.

where the infimum is over all integrable u with

$$\mathrm{d}X^u(s) = u(s)\,\mathrm{d}s, \qquad X^u(t) = \frac{t}{T}, \qquad X^u(T) = 1.$$

The optimal control for the problem (5) is constant and since time is measured in calendar units in our model this corresponds to a TWAP (time weighted average price) execution. A nice consequence of attempting to relate (5) to our original problem is that we are able to show that the solution of the original problem can be decomposed into two components.

Corollary 3.4. *The optimal trading rate $\hat{u}$ and the holdings $\hat{X}$ of the optimization problem* (2) *can be decomposed as*

$$\hat{u} = \hat{u}_1 + \hat{u}_2 \quad \textit{and} \quad \hat{X} = \hat{X}_1 + \hat{X}_2, \tag{6}$$

where $\hat{u}_1(s) = 1/T$, $\hat{X}_1(s) = s/T$ for $s \geq t$ is the solution to (5). *The processes $\hat{X}_2$ and $\hat{u}_2$ are given by*

$$\hat{X}_2(s) = \left(x - \frac{t}{T}\right)\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s b(z)\left(\frac{z}{T} - \gamma(z)\right)\exp\left(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\right)\mathrm{d}z$$

and $\hat{u}_2(s) = \hat{X}_2'(s) = -\frac{1}{2\kappa}\big(2a(s)\hat{X}_2(s) + b(s)\big(\gamma(s) - \frac{s}{T}\big)\big)$.

The proof of Corollary 3.4 is contained in Subsection 5.3. This decomposition of the optimal control can be interpreted as follows: the first part corresponds to the TWAP execution. The second part shows how we need to deviate from this deterministic strategy due to the randomness and jumps of $\gamma(s)$. Such a result corresponds nicely to heuristic algorithm design where a "historical" or average amount is always executed with an adaptive correction due to volume spikes.

Using $a > 0$ and the form of $\hat{u}_2$, we see that $\hat{X}_2$ has a zero-reversion property: if $\hat{X}_2$ equals to a big positive (negative) value, $\hat{u}_2$ will become negative (positive) so that $\hat{X}_2$ will be decreasing (increasing).

We also note that for $\kappa T \gg 1$, the optimal strategy $\hat{u}$ may be approximated by $\hat{u}_1$ because $\hat{u}_2$ depends on $\frac{1}{\kappa}\big(\gamma(s) - \frac{s}{T}\big)$, whose variance vanishes as $\kappa T \to \infty$. This behaviour is inline with Section 1.2 of Almgren [1] and to be expected from the original optimization problem (2): if κ is huge, the first term dominates the second, which leads to an approximately uniform distribution of the stock purchases due to the assumption of a linear market impact. Similarly, a huge T means that $\gamma(s)$ will be close to s/T (low variance of $\gamma(s)$) and hence the second term will again have a minor impact for a linear $\hat{X}(t)$.

4 Fitting the Model to Data

To fit our gamma bridge based model to data, one could attempt to fit m, using for example least squares. We note that since the mean increase of the gamma bridge over the day is independent of m due to the fact that $E[\gamma(t)] = \frac{t}{T}$ this fit must be done on the variance or standard deviation. In reality it is known that trading volume is U-shaped (higher volume at the beginning and end of the day) which implies an approximate cubic shape for cumulative relative volume, see Figure 5. To incorporate this feature in our model, we make a deterministic time change given by a polynomial

$$t \mapsto G(t) = at^3 + bt^2 + ct + d$$

for some constants a, b, c, d. We require that G be an increasing bijection of $[0, T]$, so we have $d = 0$ and $c = 1 - aT^2 - bT$. In the next subsection, we discuss how the parameters a, b and m can be chosen to fit data. In Subsection 4.2, we explain how this time change affects our model and results.

4.1 Estimating the Model Parameters

We exemplify the estimation of the parameters on the stocks of Vodafone Group PLC (VOD) and Microsoft Corp (MSFT).[2] We also analyzed other

[2] All data is used with permission of Bloomberg L.P.

liquid European and US stocks and the results were similar. We used intraday volume observed at 5 minute frequencies, a reasonable duration for a volume prediction, although the below parameters were not sensitive to this choice. We applied a method of moments estimator computed using nonlinear regression and based on the first and second moments of the intraday volume curves for the first 60 trading days of 2012. Since for our problem, the relative and not the absolute volume is relevant, we study the fit to the relative volume. Figure 3 displays the resulting curves, which led to the estimations:

	$\hat{a}$	$\hat{b}$	$\hat{m}$
VOD	1.3538	-1.6467	45.2344
MSFT	1.0739	-1.8151	84.9270

We see in Figure 3 that we get overall a good fit for the mean and standard deviation of the intraday volume curve, given that we only have three model parameters.

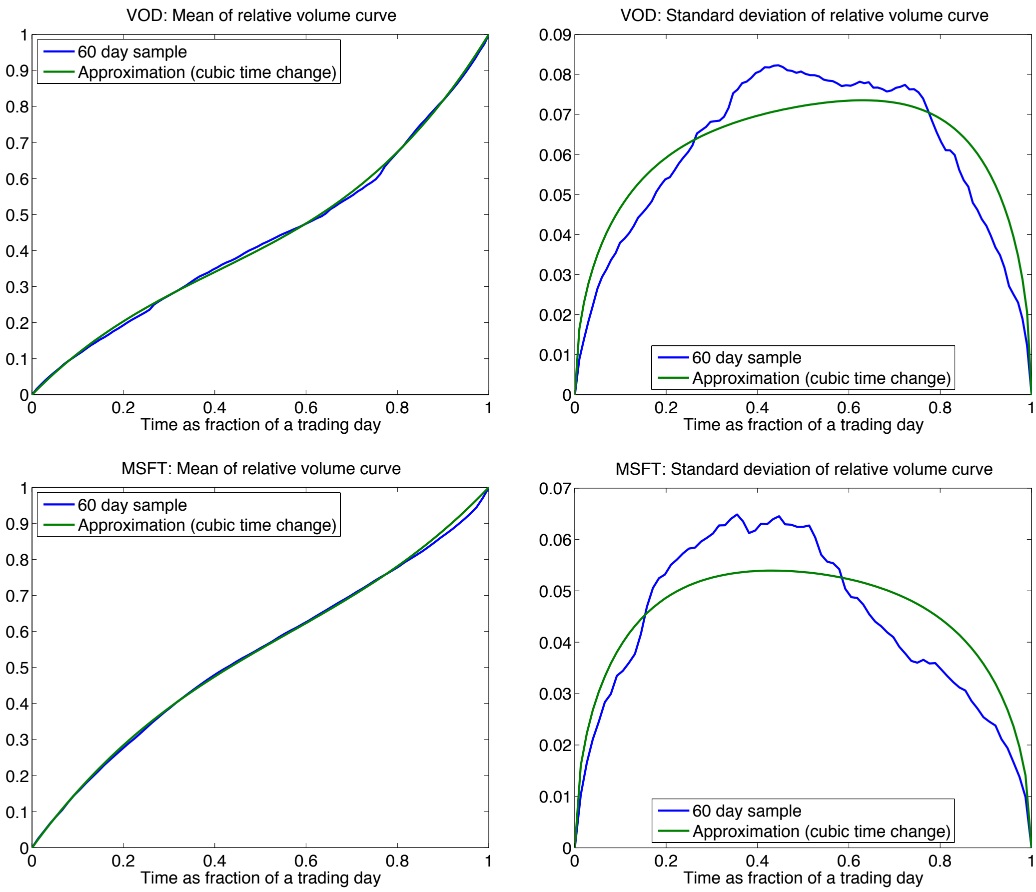

Figure 3: Example fits for the mean and standard deviation for one US and one EU stock.

To verify that the model does not exhibit any extreme seasonal dependence, we calibrated the parameters a, b and m with a 60 day rolling window for the next 30 days. The time series are shown in Figure 4. Note that the variance of the gamma bridge is proportional to $\frac{1}{mT+1}$ so that the large range of m is not as severe as might be first assumed.

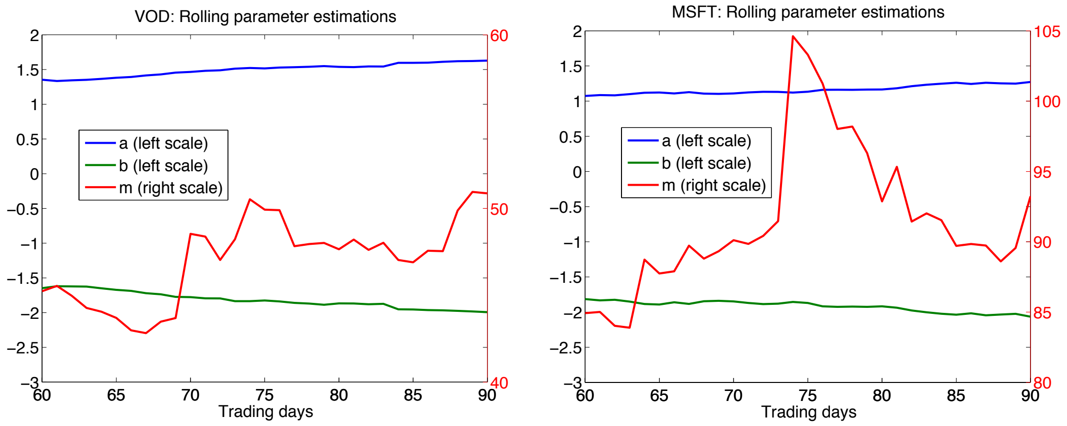

Figure 4: Rolling parameter estimations of a, b and m.

We next discuss not only the fit in the moments but the general fit. In a first qualitative analysis, we can compare the intraday volume curves of the 60 trading days with a sample of 60 trajectories based on the gamma bridge with a time change using the estimated parameters. We see in Figure 5 that the patterns of the true and sample trajectories look similar although the simulated paths have a more erratic behaviour.

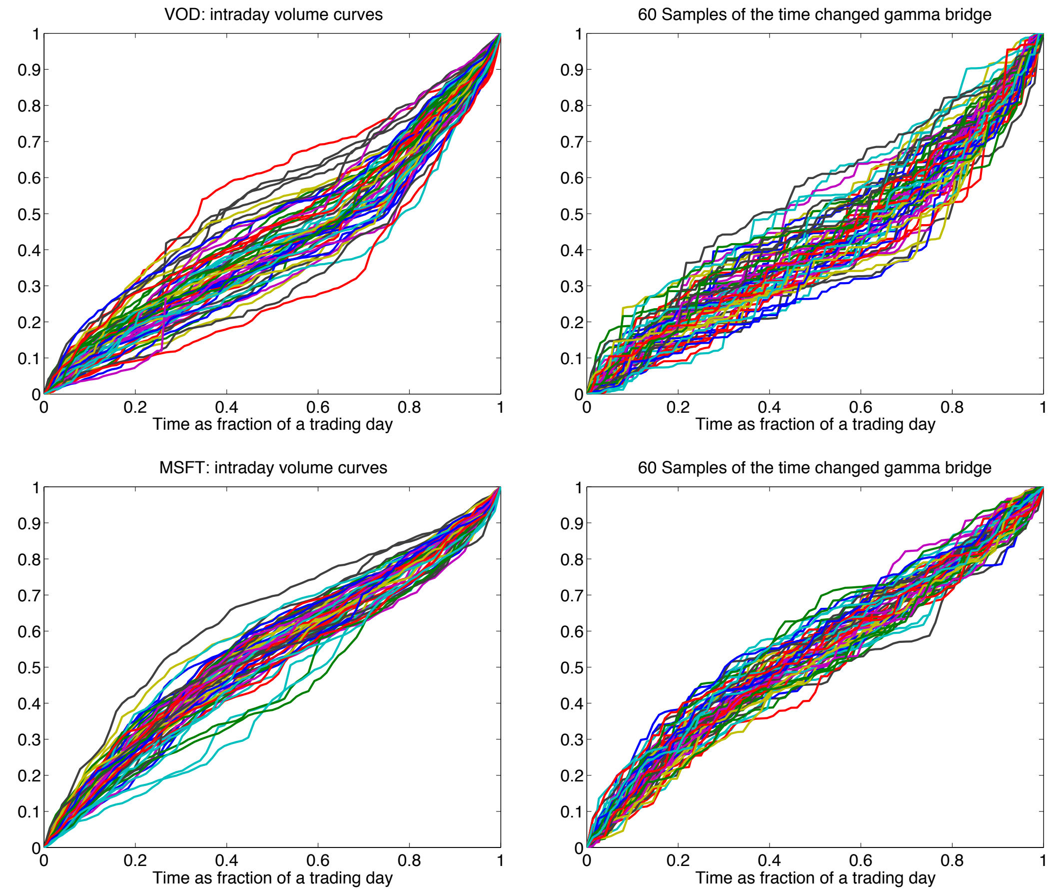

Figure 5: Comparison of the intraday volume curves with sample paths.

To make a more quantitative statement, we consider two goodness-of-fit tests. We again use the estimated parameters $\hat{a}$, $\hat{b}$ and $\hat{m}$, given in the above table and based on the first 60 trading days of 2012. We then take the following 30 days as an out-of-sample test data set. Under our assumption of a gamma bridge for the relative volume, we have, at each time t, 30 i.i.d. observations from a beta distribution with parameters

$$\alpha(t) = \hat{m}\big(\hat{a}t^3 + \hat{b}t^2 + \big(1 - \hat{a} - \hat{b}\big)t\big) \quad \text{and} \quad \beta(t) = \hat{m}T - \alpha(t).$$

The first test uses that the sample mean is approximately normally distributed with mean $\frac{\alpha(t)}{\alpha(t)+\beta(t)}$ and variance $\frac{\alpha(t)\beta(t)}{30(\alpha(t)+\beta(t))^2(\alpha+\beta+1)}$ by the central limit theorem. Performing a z-test leads to the p-values shown in Figure 6.

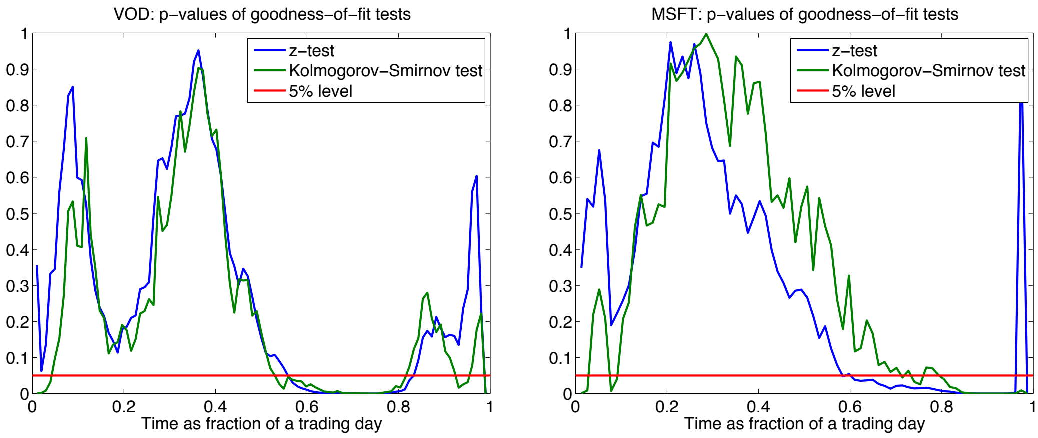

Figure 6: p-values based on a 30 days out-of-sample data set.

The second test, displayed also in Figure 6, is the well known Kolmogorov-Smirnov test, which assesses the null hypothesis that the sample is drawn from a Beta$\big(\alpha(t), \beta(t)\big)$-distribution. Observe that we now consider the whole empirical cumulative distribution and not just the sample mean. Since there is dependence across different time points, one should interpret Figure 6 in the following way: if one selects a time interval during the day at random, the corresponding y-values are the p-values of the two tests at that given time. Generally, in the morning we do not reject the null hypothesis whereas in parts of the afternoon we do. This conclusion of an imperfect fit in parts of the day is not very surprising. Indeed, relative volume is not distributed as a gamma bridge, being subject to idiosyncratic factors that the model does not (and cannot) fully capture. The model proposed here aims to capture the main stylized features whilst being tractable enough for optimization, compare with the use of Brownian motion as a model for the price despite the non-Gaussianity of returns as well as the Black-Scholes framework for option pricing. Taking these considerations into account, we conclude that the model is suitable for our purposes.

One final point related to fitting the model to data concerns the choices of the parameters κ and λ. A calculation shows that one consequence of our model formulation is that as $\kappa \downarrow 0$, since the costs vanish, we end up being able to track VWAP perfectly. In a real trading environment this is of course not possible since each execution of a market order incurs a per share cost of one half the spread (vs. mid-quote). One way to resolve this issue would be a reformulation of the impact costs to include a fixed cost of one half the spread. This naturally leads to an impulse control formulation which is beyond the scope of this paper. A more pragmatic solution observes that the key parameter driving the optimal control/trajectory is $\frac{\kappa}{\lambda\sigma^2}$. In practice one can therefore simply choose a value of λ so that this ratio is in an appropriate range for the model to be applicable. Broadly speaking, one is choosing an "effective λ" to take into account spread costs, similarly as is done in trading systems employing the Implementation Shortfall framework described in [2].

4.2 Integrating the Time Change into the Model

As we have seen in the above discussion, the gamma bridge fits well to the relative volume curve, but only after a deterministic time transformation. We next see that our main result still holds when the model undergoes such a time change. We first explain the reason in a simple example and give afterward the mathematical argument. Assume that one has calculated the optimal control and the value function in the model of Section 2. Now it is given that the model fits well except that the expected trading frequency in the second half of the day is double that of the first half. How should the optimal trading strategy be modified? A natural answer is to scale the original strategy accordingly by simply changing the speed of trading execution. In the morning, it is reduced by a third so that in the middle of the day one has the same position as one would have had after one third of the day with the original strategy. In the afternoon, the trading frequency is then correspondingly increased. In summary, the new strategy is just the time changed original strategy scaled by the derivative of the time change.

It is also sensible to expect that the value of the optimization problem will not change because the algorithm performance (minimal slippage from VWAP) should not change under a deterministic time change. Indeed we may take the time change into consideration when choosing our strategy so that there is no additional information present. Of course, these arguments use that the time change is deterministic and hence known in advance, however this is acceptable for most practical purposes.

Let us now make the above arguments rigorous by adapting the model of Section 2. The time change is assumed to be given by a differentiable deterministic function $G : [0, T] \to [0, T]$ with $G(0) = 0$, $G(T) = T$ and $G'(t) \geq C$ for some constant $C > 0$ and all $t \in [0, T]$. The relative volume curve is modelled by $\eta(t) = \gamma\big(G(t)\big)$ where $\big(\gamma(t)\big)_{0\leq t\leq T}$ is a gamma bridge whose underlying gamma process has a general parameter m. The time change means that trading frequency depends on the time of the day, e.g., higher at the beginning and closing of the day.

There is a well known link between traded volume and prices, we thus expect that the varying trading frequency will also affect the price process so that the asset price at time t equals

$$\tilde{P}(t) = P(0) + \sigma W\big(G(t)\big).$$

Observe that now we have that for two time points t_0 and t_1, the volatility of returns is proportional to

$$\sqrt{E\left[\big(\tilde{P}(t_1) - \tilde{P}(t_0)\big)^2\right]} = \sigma\sqrt{\int_{t_0}^{t_1} G'(s)\mathrm{d}s}.$$

Assuming that the difference $t_1 - t_0$ is small and fixed, this can be approximated by $\sigma\sqrt{t_1 - t_0}\sqrt{G'(t_0)}$. Since G' reflects the expected intraday trading frequency, it is typically U-shaped and hence $\sqrt{G'}$ is also U-shaped (or V-shaped). Indeed, a U-shaped G' means that its derivative G'' is negative at the beginning and increases to become positive at the end, and this property translates to $\frac{\mathrm{d}}{\mathrm{d}t}\sqrt{G'} = \frac{G''}{2\sqrt{G'}}$ as well. Therefore, we expect the instantaneous volatility to be U-shaped. Apart from being a mathematical consequence of making a time change to our model, this phenomenon is very well documented in the empirical finance literature across many different stock markets, dating back to Wood et al. [28] and Harris [17]. Also the link between the U-shaped forms of intraday volume and volatility is well known and goes back, at least, to Harris [18]. The common theoretical explanation is that the patterns of both volume and volatility are related to the flow of information, which is not constant over time. For example, in a market with asymmetrically informed participants, trading volume itself conveys information so that a U-shaped volume will lead to a U-shaped flow of information. This non-constant flow of information is captured in our model by the time change.

Similarly, the coefficient of market impact at time t is now $\tilde{\kappa}(t) = \kappa/G'(t)$ because a decrease (or increase) in G' means a slowdown (acceleration) in overall market trading frequency so that the market impact of our trades is increased (decreased).

Proceeding along the same lines as in Section 2 and additionally using $\mathrm{d}\langle W \circ G\rangle_s = \mathrm{d}G(s) = G'(s)\,\mathrm{d}s$, we see that the new value function to the optimization problem is given by

$$w(\tau, x, \gamma) = \inf_y E\bigg[\int_\tau^T \tilde{\kappa}(s)y^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_\tau^T \big(\eta(s) - X^y(s)\big)^2 G'(s)\,\mathrm{d}s\bigg],$$

where $\eta(s) = \gamma\big(G(s)\big)$ for a gamma bridge $\big(\gamma(s)\big)_{G(\tau)\le s\le T}$ with $\gamma\big(G(\tau)\big) = \gamma$ and the integrable y is such that

$$\mathrm{d}X^y(s) = y(s)\,\mathrm{d}s, \qquad X^y(\tau) = x, \qquad X^y(T) = 1$$

and y is adapted to the time-changed filtration $(\mathcal{F}_{G(s)})_{0\le s\le T}$. We associate such a y with a process u by $u(s) = \frac{y(G^{-1}(s))}{G'(G^{-1}(s))}$, which is adapted to the original filtration $(\mathcal{F}_s)_{0\le s\le T}$. We then have

$$\begin{aligned}
&E\bigg[\int_\tau^T \tilde\kappa(s) y^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_\tau^T \big(\eta(s)-X^y(s)\big)^2 G'(s)\,\mathrm{d}s\bigg]\\
&= E\bigg[\int_t^T \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^T \big(\gamma(s)-X^u(s)\big)^2\,\mathrm{d}s\bigg]
\end{aligned}$$

for $t = G(\tau)$ and

$$X^y\big(G^{-1}(s)\big) = x + \int_\tau^{G^{-1}(s)} y(r)\,\mathrm{d}r = x + \int_t^s u(r)\,\mathrm{d}r = X^u(s).$$

Since this holds for any such y, we obtain

$$w(\tau,x,\gamma) = v\big(G(\tau),x,\gamma\big),$$

where v is the original value function from (2), which is characterized in Theorem 3.1. We can also see that the optimal control of the time changed problem is given by

$$\begin{aligned}
\hat y(\tau,x,\gamma) &= \hat u\big(G(\tau),x,\gamma\big)G'(\tau)\\
&= -\frac{2a(G(\tau))x + b(G(\tau))\gamma + c(G(\tau))}{2\kappa}G'(\tau)\\
&= -\frac{2a(G(\tau))x + b(G(\tau))\gamma + c(G(\tau))}{2\tilde\kappa(t)},
\end{aligned}$$

again by using Theorem 3.1.

5 Proof of the Main Result

We split the proof of Theorem 3.1 into two parts: we first show some properties of the functions a, b, c, d, f, g and the candidate for the value function, and then we verify that this candidate indeed gives rise to the value function v. Finally, we will provide the proofs of Corollaries 3.3 and 3.4.

5.1 Properties of the Auxiliary Functions

Lemma 5.1. *The functions $a, b, c, d, f, g : [0,T) \to \mathbb{R}$ are well defined and there exist constants k_1, k_2, k_3 such that*

$$\lim_{s\nearrow T} a(s)(T-s) = \kappa, \tag{7}$$

$$|b(t)| + |d(t)| + |f(t)| \le k_1(T-t), \tag{8}$$

$$\left|\frac{c(t)}{a(t)} + 2\right| \le k_2(T-t)^2, \tag{9}$$

$$|g(t) - a(t)| \le k_3(T-t) \tag{10}$$

for all $t \in [0,T)$.

Proof. We consider $a,\dots,g$ and their limiting behaviour sequentially.

1) The functions a,b,c. The functions a,b,c are well defined for $t\in[0,T)$, and an application of L'Hôpital's rule yields (7). To study the behaviour of b, we use a series expansion

$$\begin{aligned} b(t) &= -2a(t)+\frac{2\kappa}{T-t} \\ &= -2\sqrt{\kappa\lambda\sigma^2}\,\frac{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1}{\sum_{n=1}^\infty 2^n(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/n!}+\frac{2\kappa}{T-t} \\ &= \frac{2\kappa}{T-t}\,\frac{\sum_{n=0}^\infty 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!-\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}-1}{\sum_{n=0}^\infty 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!} \\ &= 2\kappa(T-t)\frac{\sum_{n=2}^\infty 2^n(T-t)^{n-2}(\lambda\sigma^2/\kappa)^{n/2}\big(2/(n+1)!-1/n!\big)}{\sum_{n=0}^\infty 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!}. \end{aligned}$$

Using $2/(n+1)!-1/n!\le 0$ for all n, we see that

$$0>b(t)\ge -C(T-t) \tag{11}$$

for all $t<T$ and some constant C. For the function c, we obtain

$$\begin{aligned} \frac{c(t)}{a(t)}+2 &= -2\frac{\sum_{n=0}^\infty 2^{n+1}(T-t)^n(\lambda\sigma^2/\kappa)^{n/2}/(n+1)!}{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1}+2 \\ &= 2(T-t)^2\frac{\sum_{n=2}^\infty 2^n(T-t)^{n-2}(\lambda\sigma^2/\kappa)^{n/2}\big(1/n!-2/(n+1)!\big)}{\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}}+1}, \end{aligned}$$

which shows (9).

2) The function d. Since b is bounded by (11) and we have $\frac{T+1/m-s}{T+1/m-t}\le 1$ and $\frac{T-s}{T-t}\le 1$ for $s\in[t,T]$, the integrand in

$$d(t)=\int_t^T\Big(\lambda\sigma^2-\frac{1}{4\kappa}b^2(s)\Big)\,\frac{T-s}{T-t}\,\frac{T+1/m-s}{T+1/m-t}\,\mathrm{d}s$$

is bounded so that d is well defined and we can deduce $|d(t)|\le C(T-t)$ for some constant C.

3) The function f. From (11) and the boundedness of d, we derive that

$$\Big|b(s)+2d(s)\frac{T-s}{T-s+1/m}\Big|\le C(T-s)$$

for some constant C. Hence, f is well defined and

$$|f(t)|\le\frac{1}{T-t}\int_t^T\Big|b(s)+2d(s)\frac{T-s}{T-s+1/m}\Big|\mathrm{d}s\le\frac{C}{2}(T-t).$$

4) The function g. Thanks to (7), the function $s\mapsto a(s)(T-s)$ is continuous and bounded on $[0,T)$. Together with (8), this yields

$$\Big|\frac{b(s)}{2}\Big(\frac{a(s)}{\kappa}(T-s)+1\Big)-f(s)-\frac{d(s)/m}{T-s+1/m}\Big|\le C(T-s)$$

for some constant C. Therefore, the function g is well defined and

$$
\begin{aligned}
&|g(t)-a(t)| \\
&= \left| T-t+\int_t^T \left(\frac{b(s)}{2}\Big(\frac{a(s)}{\kappa}(T-s)+1\Big)-f(s)-\frac{d(s)/m}{T-s+1/m}\right)\frac{1}{T-s}\,\mathrm{d}s\right| \\
&\le k_3(T-t)
\end{aligned}
$$

for some constant k_3, which shows (10). $\square$

Lemma 5.2. *The function* $\varphi : [0,T)\times\mathbb{R}_+\times[0,1]\to\mathbb{R}$ *defined by*

$$
\varphi(t,x,\gamma)=a(t)x^2+b(t)\gamma x+c(t)x+d(t)\gamma^2+f(t)\gamma+g(t) \tag{12}
$$

satisfies

$$
\begin{aligned}
&\varphi_t+\lambda\sigma^2(\gamma-x)^2-\frac{\varphi_x^2}{4\kappa} \\
&\quad +m\int_0^1 \big(\varphi\big(t,x,\gamma+(1-\gamma)z\big)-\varphi(t,x,\gamma)\big)(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z=0.
\end{aligned} \tag{13}
$$

Proof. We can write

$$
\begin{aligned}
&m\int_0^1 \big(\varphi\big(t,x,\gamma+(1-\gamma)z\big)-\varphi(t,x,\gamma)\big)(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z \\
&=\big(b(t)+f(t)+2\gamma d(t)\big)(1-\gamma)\varphi_0(t)+d(t)(1-\gamma)^2\varphi_1(t),
\end{aligned} \tag{14}
$$

where

$$
\begin{aligned}
\varphi_0(t)&=m\int_0^1(1-z)^{Tm-tm-1}\,\mathrm{d}z=\frac{1}{T-t},\\
\varphi_1(t)&=m\int_0^1 z(1-z)^{Tm-tm-1}\,\mathrm{d}z=\frac{1}{T-t}-\frac{1}{T-t+1/m}.
\end{aligned}
$$

We also have

$$
-\frac{\varphi_x^2}{4\kappa}=\frac{1}{\kappa}(-a^2x^2-abx\gamma-acx)-\frac{1}{4\kappa}(b^2\gamma^2+2bc\gamma+c^2), \tag{15}
$$

where we write a for $a(t)$, etc. By means of a straightforward calculation, we can check that

$$
\begin{aligned}
a_t&=\frac{1}{\kappa}a^2-\lambda\sigma^2, & d_t&=2d\varphi_0-d\varphi_1+\frac{1}{4\kappa}b^2-\lambda\sigma^2,\\
b_t&=\frac{1}{\kappa}ab+2\lambda\sigma^2+b\varphi_0, & f_t&=\frac{1}{2\kappa}bc-2d\varphi_0+f\varphi_0+2d\varphi_1,\\
c_t&=\frac{1}{\kappa}ac-b\varphi_0, & g_t&=\frac{1}{4\kappa}c^2-f\varphi_0-d\varphi_1.
\end{aligned}
$$

Using this in calculating φ_t together with (14) and (15), we derive (13). $\square$

5.2 Verification

We next relate the candidate φ to the value function v by using the properties (7)–(10). We start with two auxiliary results.

Lemma 5.3. *The candidate optimal process $\hat{X}$ given by*

$$\mathrm{d}\hat{X}(t) = -\frac{1}{2\kappa}\big(2a(t)\hat{X}(t) + b(t)\gamma(t) + c(t)\big)\,\mathrm{d}t, \qquad \hat{X}(0) = x \tag{16}$$

is bounded.

Proof. Solving (16) for $\hat{X}$ yields

$$\hat{X}(t) = x\mathrm{e}^{-\frac{1}{\kappa}\int_0^t a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_0^t \big(-b(s)\gamma(s) - c(s)\big)\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s.$$

Recalling that b is bounded by (8), $a > 0$ and γ is a gamma bridge, we see that it is enough to show that

$$\int_0^t |c(s)|\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s \text{ is bounded uniformly in } t.$$

Because of $a > 0$ and $|c(s)| \leq a(s)\big(k_2(T-s)^2 + 2\big)$ by (9), this follows from

$$\begin{aligned}\int_0^t a(s)\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\mathrm{d}s &= \kappa\exp\left(-\frac{1}{\kappa}\int_s^t a(r)\,\mathrm{d}r\right)\Bigg|_{s=0}^{s=t} \\ &= \kappa - \kappa\exp\left(-\frac{1}{\kappa}\int_0^t a(r)\,\mathrm{d}r\right) \\ &\leq \kappa.\end{aligned}$$

□

We next establish some apriori estimates for φ.

Lemma 5.4. *There exists a constant K such that*

$$-K(T-t)(x+1) \leq \varphi(t,x,\gamma) - a(t)(x-1)^2 \leq K(T-t)(x+1)$$

for all $t \in [0,T)$, $x \geq 0$ and $\gamma \in [0,1]$. In particular, for every compact set $M \subset \mathbb{R}_+$, there exists a constant $\overline{K}$ such that

$$-\overline{K}(T-t) \leq \varphi(t,x,\gamma) - a(t)(x-1)^2 \leq \overline{K}(T-t)$$

for all $t \in [0,T)$, $x \in M$ and $\gamma \in [0,1]$.

Lemma 5.4 shows that $\varphi(t,x,\gamma) - a(t)(x-1)^2$ loses the x-dependence in the limit behaviour $t \nearrow T$. This can be interpreted as cancelling the term $n\big(X^u(T) - 1\big)^2$ in the auxiliary optimization problem (4) in the limit $n \to \infty$.

Proof. We deduce from (8)–(10) that

$$\begin{aligned}
&\left|\varphi(t,x,\gamma) - a(t)(x-1)^2\right| \\
&= \left|b(t)\gamma x + d(t)\gamma^2 + f(t)\gamma + x\big(c(t)+2a(t)\big) + g(t) - a(t)\right| \\
&\le (T-t)\big(k_1(x+1) + k_3 + k_2 a(t)(T-t)\big).
\end{aligned}$$

From (7), it follows that $a(t)(T-t)$ is bounded, which implies the first claim. For the second part, it is enough to set $\overline{K} := K\sup_{x\in M}(1+x)$. □

We are now in position to prove Theorem 3.1.

Proof of Theorem 3.1. We first show that $\varphi \le v$, where φ is defined in (12) and v is the value function. This claim is equivalent to

$$\varphi(t,x,\gamma) \le E\left[\int_t^T \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T \big(X^u(s) - \gamma(s)\big)^2\,\mathrm{d}s\right] \tag{17}$$

for all controls u such that $X^u(t) = x$ and $X^u(T) = 1$ a.s., $\gamma(t) = \gamma$, where $t \in [0,T]$, $x \ge 0$, $\gamma \in [0,1]$. Without loss of generality, we can assume that $E\big[\int_t^T \kappa u^2(s)\,\mathrm{d}s\big] < \infty$ since otherwise, the result holds trivially. Using Itô's formula, we have, for any stopping time τ_ℓ valued in $[t,T]$,

$$\begin{aligned}
&\varphi\big(\tau_\ell\wedge T, X^u(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)\big) - \varphi(t,x,\gamma) \\
&= \int_t^{\tau_\ell\wedge T} (\varphi_t + u\varphi_x)\big(s, X^u(s), \gamma(s)\big)\,\mathrm{d}s \\
&\quad + \sum_{s\in(t,\tau_\ell\wedge T]} \Big(\varphi\big(s, X^u(s), \gamma(s)\big) - \varphi\big(s, X^u(s), \gamma(s-)\big)\Big).
\end{aligned}$$

Applying Proposition 4 of Émery and Yor [12] separately to ψ^+ and ψ^-, where

$$\psi(sm,\omega,y) := \varphi\big(s, X^u(s)(\omega), y + \gamma(s-)(\omega)\big) - \varphi\big(s, X^u(s)(\omega), \gamma(s-)(\omega)\big),$$

we deduce that the process

$$\begin{aligned}
&\sum_{s\in(tm, m\tau_\ell\wedge .]} \psi\Big(s,\omega,\Delta\gamma\Big(\frac{s}{m}\Big)(\omega)\Big) \\
&- \int_{tm}^{.}\int_0^1 \psi\Big(s,\omega,\Big(1-\gamma\Big(\frac{s-}{m}\Big)(\omega)\Big)z\Big)\frac{(1-z)^{Tm-s-1}}{z}\,\mathrm{d}z\mathrm{d}s
\end{aligned}$$

is a local martingale on the interval $[tm, Tm]$ in the time-changed filtration $(\mathcal{F}_{\frac{s}{m}})_{tm\le s\le Tm}$ to which $\big(\gamma(\frac{s}{m})\big)_{tm\le s\le Tm}$ is adapted. Equivalently, the process

$$\begin{aligned}
&\sum_{s\in(t,\tau_\ell\wedge .]} \psi\big(sm,\omega,\Delta\gamma(s)(\omega)\big) \\
&- m\int_t^{.}\int_0^1 \psi\big(sm,\omega,\big(1-\gamma(s-)(\omega)\big)z\big)\frac{(1-z)^{Tm-sm-1}}{z}\,\mathrm{d}z\mathrm{d}s
\end{aligned} \tag{18}$$

is a local martingale on $[t,T]$ in the standard filtration $(\mathcal{F}_s)_{t\le s\le T}$ $\big($to which $\big(\gamma(s)\big)_{t\le s\le T}$ is adapted$\big)$. Using Lemma 5.2 together with $-\frac{\varphi_x^2}{4\kappa} \le \varphi_x u + \kappa u^2$ for all $u \in \mathbb{R}$, we derive

$$\begin{aligned}&\varphi_t + \lambda\sigma^2(\gamma - x)^2 + \varphi_x u + \kappa u^2\\ &\quad + m\int_0^1 \big(\varphi\big(t,x,\gamma+(1-\gamma)z\big) - \varphi(t,x,\gamma)\big)(1-z)^{Tm-tm-1}\frac{1}{z}\,\mathrm{d}z \ge 0.\end{aligned}$$

In particular, choosing a sequence $(\tau_\ell)_{\ell\in\mathbb{N}}$ of stopping times such that $\tau_\ell \nearrow T$ and the stopped process (18) is a true martingale, we obtain

$$\begin{aligned}\varphi(t,x,\gamma) \le{}& E\bigg[\int_t^{\tau_\ell\wedge T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{\tau_\ell\wedge T}\big(X^u(s)-\gamma(s)\big)^2\,\mathrm{d}s\bigg]\\ &+ E\big[\varphi\big(\tau_\ell\wedge T, X^u(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)\big)\big].\end{aligned} \tag{19}$$

Monotone convergence implies

$$\begin{aligned}&\lim_{\ell\to\infty} E\bigg[\int_t^{\tau_\ell\wedge T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{\tau_\ell\wedge T}\big(X^u(s)-\gamma(s)\big)^2\,\mathrm{d}s\bigg]\\ &= E\bigg[\int_t^{T} \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^{T}\big(X^u(s)-\gamma(s)\big)^2\,\mathrm{d}s\bigg].\end{aligned}$$

Applying Lemma 5.4, we have

$$\begin{aligned}&E\big[\varphi\big(\tau_\ell\wedge T, X^u(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)\big)\big]\\ &\le KE\Big[(T-\tau_\ell\wedge T)\Big(\sup_{s\in[t,T]} X^u(s)+1\Big)\Big] + E\big[a(\tau_\ell\wedge T)\big(X^u(\tau_\ell\wedge T)-1\big)^2\big].\end{aligned}$$

The first integrand converges to zero by dominated convergence, using

$$\sup_{s\in[t,T]} |X^u(s)| \le x + \int_t^T |u(s)|\,\mathrm{d}s \le x + T - t + \int_t^T u^2(s)\,\mathrm{d}s \in L^1.$$

For the second term, we observe that

$$\begin{aligned}\big(X^u(\tau_\ell\wedge T)-1\big)^2 = \big(X^u(\tau_\ell\wedge T) - X^u(T)\big)^2 &\le \bigg(\int_{\tau_\ell\wedge T}^T |u(s)|\,\mathrm{d}s\bigg)^2\\ &\le (T-\tau_\ell\wedge T)\int_{\tau_\ell\wedge T}^T u^2(s)\,\mathrm{d}s\end{aligned}$$

by Hölder's inequality. Using that $a(\tau_\ell\wedge T)(T-\tau_\ell\wedge T)$ is uniformly bounded by (7), we obtain, for some constant C,

$$E\big[a(\tau_\ell\wedge T)\big(X^u(\tau_\ell\wedge T)-1\big)^2\big] \le CE\bigg[\int_{\tau_\ell\wedge T}^T u^2(s)\,\mathrm{d}s\bigg] \to 0$$

as $m \to \infty$ by dominated convergence. Together this shows that

$$\varphi(t,x,\gamma) \le E\bigg[\int_t^T \kappa u^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T \big(X^u(s) - \gamma(s)\big)^2\,\mathrm{d}s\bigg]$$

and concludes the proof of (17). To prove the reverse inequality $\varphi \ge v$, we deduce similarly to (19) that

$$\begin{aligned}\varphi(t,x,\gamma) = E\bigg[\int_t^{\tau_\ell\wedge T} \kappa\hat{u}^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^{\tau_\ell\wedge T} \big(\hat{X}(s) - \gamma(s)\big)^2\,\mathrm{d}s\bigg] \\ + E\big[\varphi\big(\tau_\ell \wedge T, \hat{X}(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)\big)\big]\end{aligned} \tag{20}$$

for the candidate optimal control $\hat{u}$ corresponding to $\hat{X}$ given in (16) — this time we have an equality from Lemma 5.2 using $\hat{u}(s) = -\frac{\varphi_x}{2\kappa}\big(s, \hat{X}(s), \gamma(s)\big)$. By Lemma 5.3, $X^{\hat{u}}$ is bounded and hence we obtain from Lemma 5.4 that

$$\begin{aligned}&E\big[\varphi\big(\tau_\ell\wedge T, \hat{X}(\tau_\ell\wedge T), \gamma(\tau_\ell\wedge T)\big)\big] \\ &\ge E\big[a(\tau_\ell\wedge T)\big(\hat{X}(\tau_\ell\wedge T) - 1\big)^2\big] - \overline{K}E[T - \tau_\ell\wedge T].\end{aligned} \tag{21}$$

Since the first term in (21) is nonnegative and the second converges to zero, monotone convergence yields

$$\varphi(t,x,\gamma) \ge E\bigg[\int_t^T \kappa\hat{u}^2(s)\,\mathrm{d}s + \lambda\sigma^2 \int_t^T \big(\hat{X}(s) - \gamma(s)\big)^2\,\mathrm{d}s\bigg].$$

Using the admissible control $\bar{u}(s) = \frac{1-x}{T-t}$, one can see that

$$\varphi(t,x,\gamma) \le v(t,x,\gamma) \le \kappa\frac{(1-x)^2}{T-t} + M' < \infty \tag{22}$$

for some constant M'. We conclude the proof by showing $\hat{X}_T = 1$ a.s. Using (20), (21) and (22), we deduce

$$\kappa\frac{(1-x)^2}{T-t} + M' \ge E\big[a(\tau_\ell\wedge T)\big(\hat{X}(\tau_\ell\wedge T) - 1\big)^2\big] - \overline{K}E[T - \tau_\ell\wedge T].$$

Since a is increasing, it follows, for every $t_0 \in [t,T)$, that

$$\kappa\frac{(1-x)^2}{T-t} + M' \ge a(t_0)E\big[\big(\hat{X}(\tau_\ell\wedge T) - 1\big)^2 \mathbb{1}_{\{\tau_\ell \ge t_0\}}\big] - \overline{K}E[T - \tau_\ell\wedge T].$$

Using that $\hat{X}$ is bounded, dominated convergence yields

$$\kappa\frac{(1-x)^2}{T-t} + M' \ge a(t_0)E\big[\big(\hat{X}(T) - 1\big)^2\big]$$

for all $t_0 \in [t,T)$, and hence

$$\kappa\frac{(1-x)^2}{T-t} + M' \ge \lim_{t_0 \nearrow T} a(t_0)E\big[\big(\hat{X}(T) - 1\big)^2\big].$$

Since $\lim_{t_0\nearrow T} a(t_0) = \infty$, this can only hold if $E\big[\big(\hat{X}(T) - 1\big)^2\big] = 0$, which means $\hat{X}(T) = 1$ a.s. The uniqueness of the optimal control is a consequence of the linearity of X^u in u and the strict convexity of the optimization problem in the control. □

5.3 Proofs of Corollaries 3.3 and 3.4

We conclude by proving Corollaries 3.3 and 3.4.

Proof of Corollary 3.3. From $b(t) < 0$ for $t < T$ by (11) and $a(t) > 0$, it follows $\zeta_\gamma(t,\gamma) = \frac{-b(t)}{2a(t)} > 0$ for all $t < T$. We next write

$$\zeta(t,\gamma) = \frac{-b(t)\gamma - c(t)}{2a(t)} = \gamma + \kappa(1-\gamma)\frac{1}{(T-t)a(t)}.$$

To prove $\zeta_t(t,\gamma) > 0$, it is enough to show that $\frac{1}{(T-t)a(t)}$ is increasing or, equivalently, that $(T-t)a(t)$ is decreasing. Hence, we consider

$$\begin{aligned}
&\frac{\mathrm{d}}{\mathrm{d}t}(T-t)a(t) = -a(t) + (T-t)a_t(t)\\
&= -\sqrt{\kappa\lambda\sigma^2}\,\frac{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} + \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}}{\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}}} + (T-t)\frac{4\lambda\sigma^2\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2}\\
&= \frac{\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}\big(-\sqrt{\kappa\lambda\sigma^2}\big(\mathrm{e}^{2(T-t)\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{-2(T-t)\sqrt{\lambda\sigma^2/\kappa}}\big) + 4\lambda\sigma^2(T-t)\big)}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2}\\
&< \frac{\mathrm{e}^{2(T+t)\sqrt{\lambda\sigma^2/\kappa}}\big(-\sqrt{\kappa\lambda\sigma^2}\,4(T-t)\sqrt{\lambda\sigma^2/\kappa} + 4\lambda\sigma^2(T-t)\big)}{(\mathrm{e}^{2T\sqrt{\lambda\sigma^2/\kappa}} - \mathrm{e}^{2t\sqrt{\lambda\sigma^2/\kappa}})^2} = 0,
\end{aligned}$$

using $\mathrm{e}^x - \mathrm{e}^{-x} = 2\sinh(x) = 2\sum_{n=0}^\infty \frac{x^{2n+1}}{(2n+1)!} > 2x$ for all $x > 0$. This proves $\zeta_t(t,\gamma) > 0$. Moreover, we have

$$\zeta(0,0) = \kappa\frac{\varphi_0(0)}{2a(0)} = \sqrt{\kappa}\frac{\mathrm{e}^{2T/\sqrt{\kappa}} - 1}{2T(\mathrm{e}^{2T/\sqrt{\kappa}} + 1)} > 0.$$

For the second part of the lemma, we define a process $\xi(t) = \zeta\big(t,\gamma(t)\big) - \hat{X}(t)$, which satisfies $\xi(0) = \zeta\big(0,\gamma(0)\big) = \zeta(0,0) > 0$. We have $\xi(t) \geq 0$ a.s. for all t by the following reason. Fix t and consider the event $\xi(t) < 0$. Since $\hat{X}$ is continuous and $\zeta_\gamma > 0$ and γ has only positive jumps, there needs to exist a random $\tau < t$ such that $\xi(\tau) = 0$ and $\xi(s) \leq 0$ for all $s \in [\tau,t]$. However, we have

$$\xi(t) = \underbrace{\xi(\tau)}_{=0} + \underbrace{\zeta\big(t,\gamma(t)\big) - \zeta\big(\tau,\gamma(\tau)\big)}_{\geq 0} - \int_\tau^t \underbrace{\hat{u}(s)}_{\leq 0}\,\mathrm{d}s \geq 0$$

since $\zeta\big(s,\gamma(s)\big)$ is increasing and $\hat{u}(s) \leq 0$ if $\xi(s) \leq 0$. Therefore, $\xi(t)$ cannot become negative and $\hat{u}(t)$ will be nonnegative for all t. $\square$

Proof of Corollary 3.4. Let $\hat{X}_2$, $\hat{u}_2$ be defined as in the corollary and set $\hat{u}_1(s) = \hat{X}_1'(s)$ and

$$\hat{X}_1(s) = \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s \Big(-b(z)\frac{z}{T} - c(z)\Big)\exp\Big(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\Big)\mathrm{d}z$$

so that (6) is satisfied by construction; compare with (16). By definition, we have

$$-b(z)\frac{z}{T} - c(z) = a(z)\frac{2z}{T} - \frac{2z\kappa}{(T-z)T} + \frac{2\kappa}{T-z} = a(z)\frac{2z}{T} + \frac{2\kappa}{T}$$

so that

$$\begin{aligned}\hat{X}_1(s) &= \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{1}{2\kappa}\int_t^s \Big(a(z)\frac{2z}{T} + \frac{2\kappa}{T}\Big)\exp\Big(-\frac{1}{\kappa}\int_z^s a(r)\,\mathrm{d}r\Big)\mathrm{d}z \\ &= \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} + \frac{s}{T} - \frac{t}{T}\mathrm{e}^{-\frac{1}{\kappa}\int_t^s a(r)\,\mathrm{d}r} = \frac{s}{T}\end{aligned}$$

by integration by parts. Finally, we check that $\hat{u}_1(t) = \hat{X}_1'(t) = \frac{1}{T}$ is the optimizer to (5). To this end, we calculate

$$\begin{aligned}&\kappa\int_t^T u^2(s)\,\mathrm{d}s + \lambda\sigma^2\int_t^T \Big(\frac{s}{T} - X^u(s)\Big)^2\mathrm{d}t \geq \kappa\int_t^T u^2(s)\,\mathrm{d}s \\ &\geq \frac{\kappa}{T-t}\Big(\int_t^T u(s)\,\mathrm{d}s\Big)^2 = \frac{\kappa}{T-t}\Big(1-\frac{t}{T}\Big)^2 = \kappa\frac{T-t}{T^2}\end{aligned}$$

by Jensen's inequality using the probability measure $\frac{1}{T-t}\mathrm{d}s$. Equality holds for the choice $u = \hat{u}_1 = \frac{1}{T}$ with corresponding $X^u(s) = \hat{X}_1(s) = \frac{s}{T}$. This shows that $\hat{u}_1$ is indeed the minimizer of (5). □

Acknowledgments

We thank three anonymous referees for very useful comments, which enabled us to significantly improve the paper. We thank Mark DiBattista and Christian Hesse for helpful comments and suggestions. Christoph Frei gratefully acknowledges financial support by the Natural Sciences and Engineering Research Council of Canada.

References

[1] R. Almgren (2012): Optimal trading with stochastic liquidity and volatility, *SIAM J. Finan. Math.* **3**, 163–181

[2] R. Almgren and N. Chriss (2000): Optimal execution of portfolio transactions, *J. Risk* **3**, 5–39

[3] R. Almgren, C. Thum, E. Hauptmann and H. Li (2005): Equity market impact, *Risk* **18**, 57–62

[4] S. Berkowitz, D. Logue and E.Noser (1988): The total cost of transactions on the NYSE, *J. Finance* **43**, 97–112

[5] D. Bertsimas and A. Lo (1998): Optimal control of execution costs, *J. Finan. Markets* **1**, 1–50

[6] J. Białkowski, S. Darolles and G. Le Fol (2008): Improving VWAP strategies: a dynamic volume approach, *J. Banking Finance* **32**, 1709–1722

[7] T. Björk and A. Murgoci (2010): A general theory of Markovian time inconsistent stochastic control problems, *preprint*, available at SSRN http://ssrn.com/abstract=1694759

[8] B. Bouchard and N.-M. Dang (2013): Generalized stochastic target problems for the pricing and partial hedging under loss constraints - Application in optimal book liquidation, *Finance Stoch.* **17**, 31–72

[9] J.-P. Bouchaud, Y. Gefen, M. Potters and M. Wyart (2004): Fluctuations and response in financial markets: the subtle nature of 'random' price changes, *Quant. Finance* **4**, 176–190

[10] D. Brody, L. Hughston and A. Macrina (2008): Dam rain and cumulative gain, *Proc. R. Soc. A* **464**, 1801–1822

[11] R. Cont, A. Kukanov and S. Stoikov (2013): The price impact of order book events, *forthcoming in J. Financ. Econom.*

[12] M. Émery and M. Yor (2004): A parallel between Brownian bridges and gamma bridges, *Publ. RIMS, Kyoto Univ.* **40**, 669–688

[13] J. Gani (1957): Problems in the probability theory of storage systems, *J. R. Statist. Soc. B* **19**, 181–206

[14] J. Gatheral (2010): No-dynamic-arbitrage and market impact, *Quant. Finance* **10**, 749–759

[15] J. Gatheral and A. Schied (2011): Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework, *Int. J. Theoretical Appl. Finance* **14**, 353–368

[16] J. Gatheral and A. Schied (2013): Dynamical models for market impact and algorithms for order execution. *In:* Handbook on Systemic Risk (eds.: J.-P. Fouque and J. Langsam), Cambridge University Press, 579–602

[17] L. Harris (1986): A transaction data study of weekly and intradaily patterns in stock returns, *J. Finan. Econ.* **16**, 99–117

[18] L. Harris (1987): Transaction data test of the mixture of distribution hypothesis, *J. Finan. Quant. Anal.* **22**, 127–141

[19] G. Hu (2007): VWAP cost excluding own trades, *J. Trading* **2**, 30–34

[20] M. Humphery-Jenner (2011): Optimal VWAP trading under noisy conditions, *J. Bank. Finance* **35**, 2319–2329

[21] S. Kakade, M. Kearns, Y. Mansour and L. Ortiz (2004): Competitive algorithms for VWAP and limit order trading, Proceedings of the 5th ACM Conference on Electronic Commerce (EC), 189–198

[22] O. Kallenberg (2002): Foundations of Modern Probability, *Probability and its Applications*, Second Edition, Springer

[23] H. Konishi (2002): Optimal slice of a VWAP trade, *J. Finan. Markets* **5**, 197–221

[24] E. Lukacs (1955): A characterization of the gamma distribution, *Ann. Math. Stat.* **26**, 319–324

[25] A. Madhavan (2002): VWAP Strategies, *Trading* **1**, 32–39

[26] J. McCulloch and V. Kazakov (2012): Mean variance optimal VWAP trading, *preprint*, available at SSRN: http://ssrn.com/abstract=1803858

[27] P. Moran (1956): A probability theory of a dam with a continuous release, *Quart. J. Math. Oxf.* **7**, 130–137

[28] R. Wood, T. McInish and J. Ord (1985): An investigation of transaction data for NYSE stocks, *J. Finance* **25**, 723–739