Static VWAP: A Comparative Analysis

Pragma Trading

***Abstract*—One of the most popular trading strategies is the Volume-Weighted Average Price (VWAP) trading algorithm. About 50% of institutional flow is executed via some VWAP variant. A VWAP algorithm tries to achieve an average execution price which is as close as possible to the realized VWAP in the market. Static VWAP trading algorithms use a pre-defined trading schedule that tries to approximate the volume pattern realized over the trading day. This schedule is calculated by averaging historical trading patterns. In this paper we explore the effects of various averaging methods on the expected performance of the VWAP algorithm. We conclude that updating curves can decrease the dispersion of shortfalls by about 10%. Using special curves for Fed announcement days produces a similar benefit. Using special curves for different sectors or market caps or for option expiration days has no effect.**

I. Preliminaries

A. Introduction

Over the last several years, the financial landscape has changed dramatically, moving from high-touch human traders to black-box systems where people are not involved in the actual trading process. Whether humans or machines are doing the actual trading, their performance is measured relative to some benchmark. Over the years many benchmarks have been proposed, but probably the most popular one is the volume-weighted average price, or VWAP. The VWAP represents the average price per share that was paid in the market during the life of the trade. As such it represents an ideal price that a trader would like to capture for the principal. This simple and intuitive benchmark is very appealing to many market participants. As a result, trading algorithms that try to achieve the VWAP are now responsible for about 50% of institutional trading flow [1].

There are many trading algorithms that try to achieve an average trade price close to the market-wide VWAP for the day. These algorithms are divided into two main groups: static and dynamic. Static algorithms use a pre-defined trading policy which does not change in response to market conditions. Dynamic algorithms use various real-time indicators in order to decide at what rate to trade. In essence, a dynamic strategy makes small deviations from the commands supplied by a static algorithm, based on current market conditions.

The most common static VWAP strategy is to follow the historical intra-day seasonality. While intuitively simple, the task of estimating intra-day patterns is tricky. There are two approaches: estimate the intra-day seasonality on a stock by stock basis [2], or average over many stocks [4]. The first approach requires averaging over a long history in order to achieve a stable estimate for each stock. This long period results in a mis-specification, i.e., errors, in the estimated curves due to changing market conditions. The second approach requires a much shorter history, which results in a much faster adaptation to changing market conditions. However, this faster adaptation does not come for free as it introduces an error due to inter-stock variability.

The published literature on intra-day seasonality is sparse. Moreover, even the papers that do explore this issue more often than not treat it as a nuisance that needs to be addressed in order to solve a different problem. Apart from the two articles cited above, it is worth mentioning [3], which investigated a very close question to the one we consider here, and reached conclusions similar to ours.

B. VWAP, Price Moves, and Performance Criteria

A key issue in the analysis of VWAP strategies is how to measure their performance. In practice, VWAP traders try to minimize the difference, or shortfall, between the average price of their trade and the market-wide VWAP. Considering that traders can take either side of a trade, the *average* shortfall over a large number of trades is approximately zero for any curve. So, what characterizes a good VWAP strategy? The answer is that good VWAP strategies have a low shortfall *dispersion*—ideally, the shortfall is always near zero. In this document, we compare the performance of various VWAP curves using standard deviation of VWAP shortfall as a measure of dispersion. In Section II we also mention briefly a second dispersion metric, the 95% quantile. A mathematical discussion is presented in the appendix.

In what follows we ignore two effects that have an impact on VWAP shortfall. One is the effect of commissions and fees, which are costs that depend on the size of the order but not on the strategy. Therefore we can safely ignore them for the purpose of comparing strategies.

Secondly, we assume that we are able to capture the market's VWAP over short time intervals, e.g. over a minute. In practice, our average execution price over one minute may not be exactly equal to the market's average price, because of small-scale timing issues or because we need to pay for liquidity. In this study we focus on determining the best shape for the trading curve over the entire day, and not on micro issues.

C. Scope and Main Conclusions

In this document we examine the performance of static VWAP algorithms. Recall that in a static VWAP algorithm we trade based on a pre-defined schedule, one that tries to match the intra-day seasonality as much as possible. As was mentioned, there are various ways of constructing this schedule, e.g., averaging cross-sectionally, averaging temporally, etc. In this document we examine various ways for constructing the trading schedule curve and the effect these schedules have on the performance.

We assess the performance of each algorithm by looking at both the shortfall standard deviation and the worst 5% shortfall in our sample. The latter measure is referred to as the 95% percentile. We compute shortfalls by simulating trading according to various curves, using historical price and volume data for 2008-09.

In our investigation we have found that the effect of using relatively old curves on the performance is negligible. This should not come as a surprise. The average curve is slowly changing from period to period and these changes are very small relative to the daily variability in the realized curves. Hence, it is the daily random changes in the curve that are responsible for most of the shortfalls, and not the mis-specification due to using an old curve.

One could also argue that different groups of stocks behave differently. We examine this issue and we demonstrate that the use of curves tailored to narrower groups does not improve the results in a meaningful way. For example, assume the universe of interest is composed of the largest 100 stocks. We can use a curve constructed specifically for that group. However, using a curve constructed from the largest 1000 stocks instead does not have any meaningful influence on the performance. Another possible grouping is economic sectors, i.e., use different curves for different sectors. This study shows that, for some time following the Lehman Brothers collapse, using sector-specific curves was indeed a good strategy. However, this improvement is not persistent. Considering the risks in using sector-specific curves (e.g., sensitivity to outliers), one may be better off by trading according to one universal curve.

Finally, we consider using special curves for days that are known in advance to be special. First, we examine the optimal strategy during days in which there are scheduled Federal Reserve announcements. During these days we do observe improved performance for a specific curve. Second, we examine days in which equity options expire. Here we find that using special curves yields a negligible performance improvement.

II. Single-Name Tasks

In general, one can use a single intra-day pattern for all stocks and dates, or use tailor-made curves for specific classes of stocks and/or dates. Tailor-made curves may yield benefits but they introduce software complexity and they are necessarily "noisier", as less historical data is available.

In this section, we investigate the performance of tailor-made curves for various classes of stocks and dates. Specifically we consider combinations of:

- Using a single static curve for all stocks.
- Using a single moving-average curve that changes every day.
- Classifying stocks by market cap and using a different curve for each class.
- Classifying stocks by sector and using a different curve for each sector.
- Using special curves for scheduled Fed announcement dates.

In what follows we define an estimation universe and a simulation universe. The estimation universe is the universe of stocks we use for estimating the average curve. The simulation universe is the universe of stocks we use for testing the performance of the estimated curve. For example, we consider the case when our estimation universe is based on the 3000 stocks with the largest market cap, while we examine the performance over a simulation universe constituted by the 100 stocks with the largest market cap. In addition, we assume that all our tasks are full-day VWAP tasks, and that one can capture the VWAP price within each bar. Small deviations in the executed price relative to the average market price over time scales of minutes introduce negligible increases in the shortfall standard deviation, which we ignore.

In our study we examine the performance over the period between February 1, 2008 and March 19, 2009. We examine the performance on a two-month rolling basis. This will allow us to examine both the absolute performance of the various algorithms and the temporal behavior of the performance.

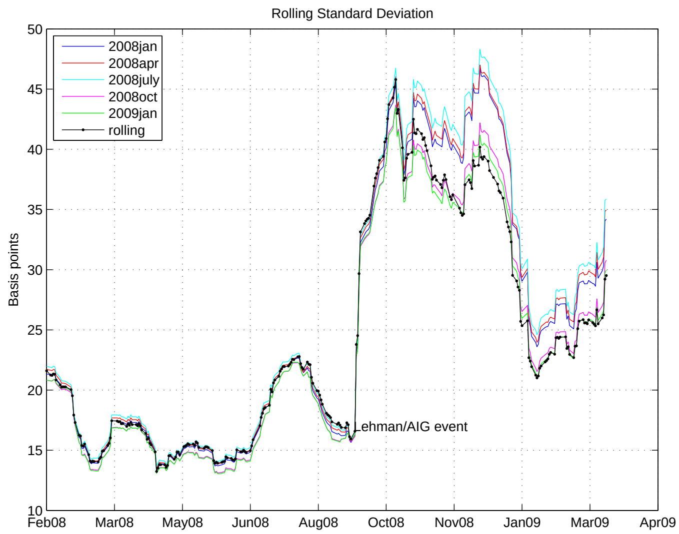

Fig. 1. Standard deviation of VWAP shortfall over a universe of 500 stocks. The standard deviation is computed over moving one-month windows. The curves labeled with a month were constructed with data from that month; the remaining curve is updated every day using the trailing one month of data.

A. Curve Drift

In this section we examine the possibility that the curve drifts over time, and the effects of temporal mismatch in the curve on the performance of the algorithm. We conduct two kind of experiments:

- Fit a curve using one month's worth of trading data, then freeze the curve and use it for all trading days.
- Fit a curve with a moving window with the last month of trading data.

Figure 1 shows the results of this experiment. The figure depicts the standard deviation as a function of time, when the algorithm uses the various curves. The standard deviation is computed over a moving two-month window. The estimation and simulation universes were the 500 stocks with the largest market cap. A snapshot with the results as of March 2009 is shown in Table I.

We note that under normal market conditions (before September 2008), there is little difference between any two curves. Second, more-recent curves are not necessarily preferable: a curve fit with January 2009 data actually does better than all the others, even in 2008. A curve fit with July 2008 data is the worst one through the summer of 2008. Third, the curve updated on a rolling basis performs close to average over stable periods and is nearly the best during volatile periods, with a difference of up to 7 basis points (15%) relative to the worst performer, and about 4 basis points (10%) relative to the average curve. This suggests that frequent curve updates may be able to improve performance over fixed curves.

TABLE I
STANDARD DEVIATION OF SHORTFALL FOR CURVES FIT ON DIFFERENT DATES, AVERAGED OVER TWO MONTHS PRIOR TO MARCH 19, 2009. EACH ROW REFERS TO A CURVE FIT ON A PARTICULAR DATE. THE STANDARD ERROR IS 1 BP.

	Standard deviation (bp)	95th percentile (bp)
January 2008	34	65
April 2008	35	67
July 2008	36	69
October 2008	31	58
January 2009	30	56
rolling	30	56

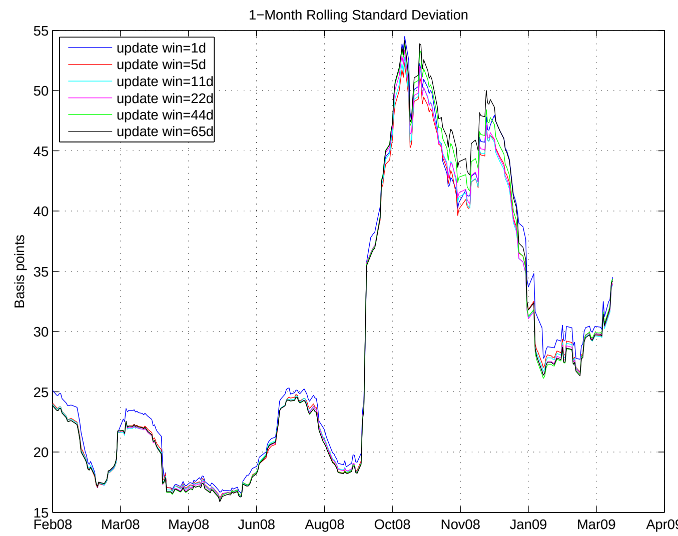

Fig. 2. Standard deviation of VWAP shortfall when the rolling pattern is updated using 1, 5, 11, 22, 44 trading days.

This conclusion is the same when other simulation or estimation universes are considered.

B. Averaging Window

The use of a moving average for the purpose of computing a VWAP curve begs the question of the optimal averaging length. Generally, shorter averages react more quickly to changes but produce noisier results.

Figure 2 shows the results of a simple experiment in which different averaging lengths were used, ranging from 1 day to 3 months. The results are not surprising. Under stable market conditions, the 1-day curve performs worst, as can be seen e.g. in the March-April 2008 period. In volatile periods, e.g. November-December 2008, the 3-month curve is worst. The differences can be up to 10%. A 1-month curve performs uniformly well.

C. Market Cap Effects

In this subsection we explore the effects on performance of dividing the stocks into different market cap

TABLE II

STANDARD DEVIATION OF SHORTFALL (IN BP) FOR VARIOUS CURVE/STOCK UNIVERSE COMBINATIONS, AVERAGED OVER TWO MONTHS PRIOR TO MARCH 19, 2009. EACH ROW REFERS TO A CURVE FROM A GIVEN ESTIMATION UNIVERSE. EACH COLUMN REFERS TO A GIVEN SIMULATION UNIVERSE. THE STANDARD ERROR IS 1 BP.

	Top 100	Top 500	Top 1000	Top 3000
Top 100	25	34	40	71
Top 500	23	30	35	68
Top 1000	23	29	34	67
Top 3000	23	30	35	67

groups. We divided the stock population into four (overlapping) classes by market cap: top 100, top 500, top 1000 and top 3000 stocks. We used each of these classes as an estimation universe, and then applied the estimated curves to all four classes as simulation universes. In other words, we compute one curve per class but test each curve against all classes. The curves were computed using a 1-month moving-average as in the rolling curve of the previous section.

We expect the four classes to behave differently: the large stocks are very liquid and have more or less predictable patterns, whereas the smaller stocks are thinly traded and are subject to wide day-to-day variations in volume. Therefore, we expect that the standard deviation of VWAP shortfall will be lowest for the top-100 class and highest for the top-3000 class, regardless of which curve is used.

On the other hand, we might expect that each curve would do best when used for stocks in its corresponding class. However, figures 3-6 show that this is not the case. Each panel refers to shortfall values for stocks in one of the four classes, and contains results for the four curves under test. Clearly, the differences between the curves are dwarfed by the time variations driven by market volatility. A close-up also reveals that in each class the top performer is *not* the curve fit to that class. In fact, the only curve with consistently worse performance is the curve computed from the largest 100 stocks (the difference is less than 10%, however.) The other three curves are practically indistinguishable. Note that in order to compute an average curve, the patterns of all stocks have been equally weighted, regardless of market cap or daily volume.

D. Sectors

Another way to divide the universe of stocks in various groups is by economic sector. Here we compute sector-specific VWAP curves by using each of the ten GICS sectors as a separate estimation and simulation universe.

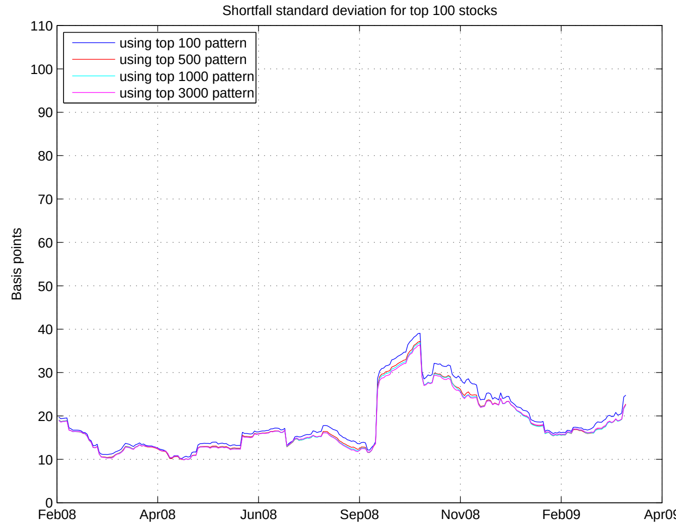

Fig. 3. Standard deviation of VWAP shortfall for the largest 100 stocks by market cap. Four curves were tested, each obtained by averaging a different number of stocks. The average is taken over a moving two-month window.

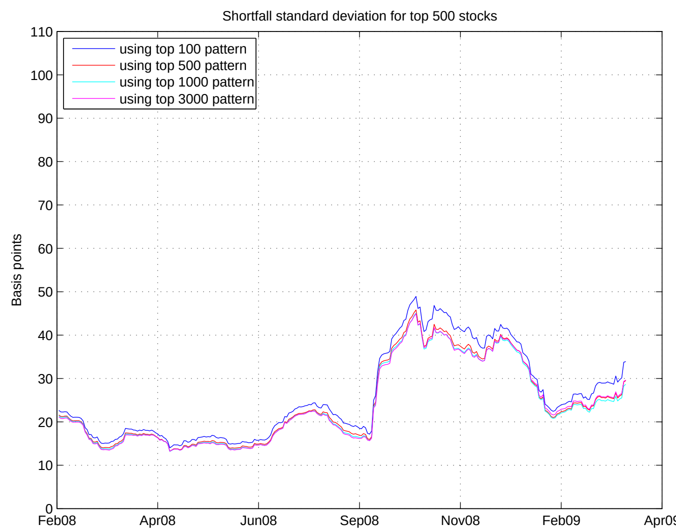

Fig. 4. Standard deviation of VWAP shortfall for the largest 500 stocks by market cap. Four curves were tested, each obtained by averaging a different number of stocks. The average is taken over a moving two-month window.

TABLE III

95TH QUANTILE OF VWAP SHORTFALLS. FOR EACH CURVE/STOCK UNIVERSE COMBINATION, THE QUANTILE IS THE VALUE s OF SHORTFALL (IN BP) SUCH THAT 95% OF HISTORICAL SHORTFALLS ARE SMALLER THAN s

	Top 100	Top 500	Top 1000	Top 3000
Top 100	50	64	81	118
Top 500	44	56	69	109
Top 1000	46	55	67	107
Top 3000	44	57	70	109

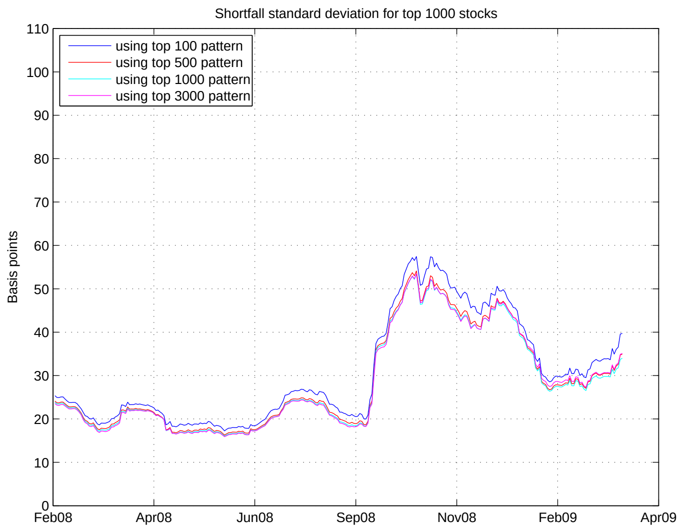

Fig. 5. Standard deviation of VWAP shortfall for the largest 1000 stocks by market cap. Four curves were tested, each obtained by averaging a different number of stocks. The average is taken over a moving two-month window.

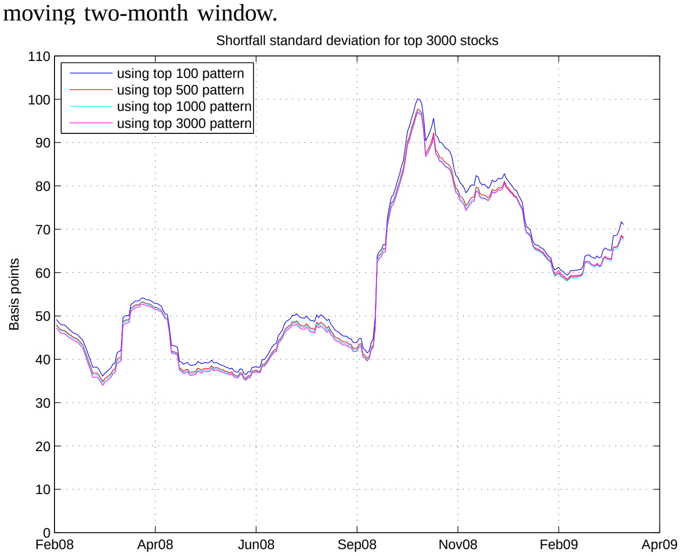

Fig. 6. Standard deviation of VWAP shortfall for the largest 3000 stocks by market cap. Four curves were tested, each obtained by averaging a different number of stocks. The average is taken over a moving two-month window.

We then estimate the overall performance of a sector-specific strategy by pooling the 10 standard deviations obtained with the 10 sector-specific VWAP curves. We compare the pooled standard deviation with the results of using a single moving-average curve for the entire population.

In all cases we use a moving average over the last month, and we consider only stocks in the group of the largest 1000 by market cap. We omit sector 8, telecommunication services, because it contains only about 20 names and its average pattern is very noisy.

The results are depicted in Fig. 7. For normal market conditions, i.e. before September 2008, the performance of the common curve is indistinguishable from the sector-specific curves. After September 2008, when the standard deviation increases to 30-50 basis points, the sector-specific curves provide an improvement of 1-2 basis points.

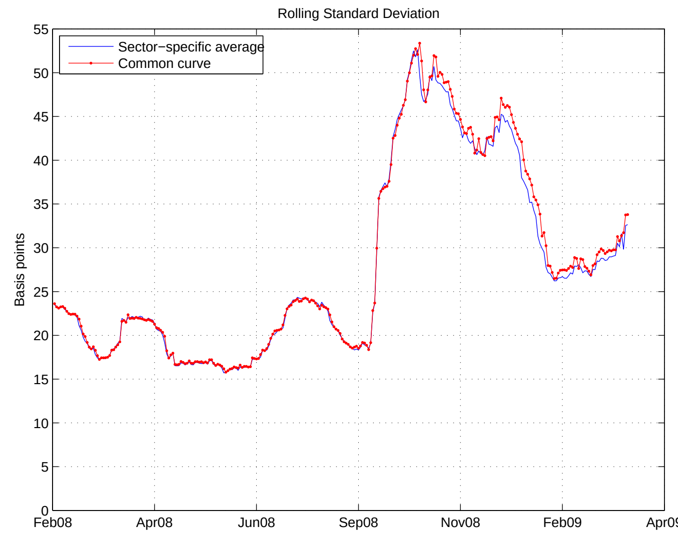

Fig. 7. Standard deviation of shortfall using sector-specific curves and a common rolling curve.

Figure 8 shows a comparison for two sectors, Energy and Financials. For each sector we show the rolling standard deviation of shortfall obtained from using a common curve and a sector-specific curve. First, it can be seen that there is wide variation between sectors, especially after September 2008. Second, the sector-specific curves can improve or worsen performance. For Financials, the sector-specific curve is never worse than the common curve, and it provides an improvement of up to 7 bps, or 8%. For Energy, the common curve is in fact better, by up to 3 bps, or 7%. Thus, in the case of Energy the positive effect of having a more-specific curve is negated by the random errors introduced by having fewer stocks to compute the average.

E. Scheduled Fed days

Days on which it is known that the Federal Reserve will make an announcement are of particular interest, since much volume is driven by the announcement, in the afternoon. Thus one might consider using a special curve for those days.

A major drawback of this idea is the relative scarcity of data: the Fed meets regularly only 8 times a year. In order to obtain more-stable estimates, we used data for

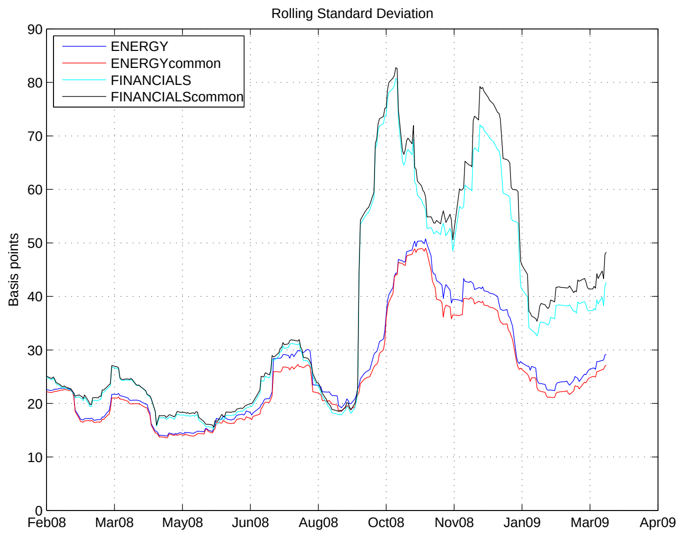

Fig. 8. Standard deviation of shortfall using sector-specific curves and a common rolling curve for two specific sectors.

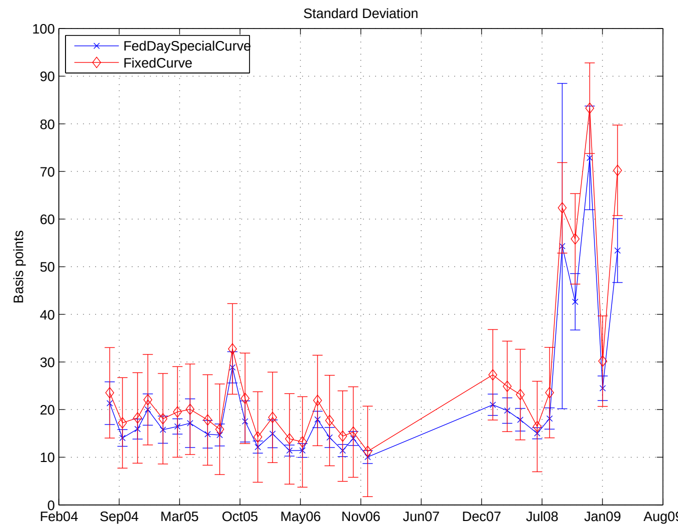

Fig. 9. Standard deviation of shortfall for scheduled Fed days, 2004-09. The vertical bars indicate the 2σ confidence interval for each value. The blue curve was obtained by averaging the pattern for Fed days only. The red curve comes from using a fixed standard pattern.

Fed days between August 2004 and March 2009[1].

We simulated the effect of using a specially-designed Fed-day VWAP curve for those days, versus using a standard VWAP curve. For each of the 30 scheduled Fed days in our sample, we computed the shortfalls obtained by trading the 1000 largest stocks with those curves. The standard deviation of the computed shortfalls is shown in Fig. 9. We show the estimated standard deviation as well as 2σ error bars to indicate the degree of uncertainty in the estimation, which is particularly bad here because of the limited sample size.

For each date in the sample, the Fed-specific curve was fit using all *other* Fed dates. This prevents in-sample bias, but it also means that a slightly different curve was used for each of the dates shown. On the other hand the fixed curve is the same in all cases, and it was fit without using any days in the sample.

The error bars are so large that one curve fits within the other's error margins. That said, Fed-specific curves do better on all days in the sample. The improvement is a few bps, 16% of standard deviation on average.

Figure 10 shows the average Fed-specific pattern computed from our sample of Fed days. The afternoon volume hump is clearly visible.

F. Option Expiration Days

The third Friday of each month is the last trading day before the expiration of equity options. Therefore we may expect unusual trading activity and perhaps a special volume pattern. To test this theory, we compared

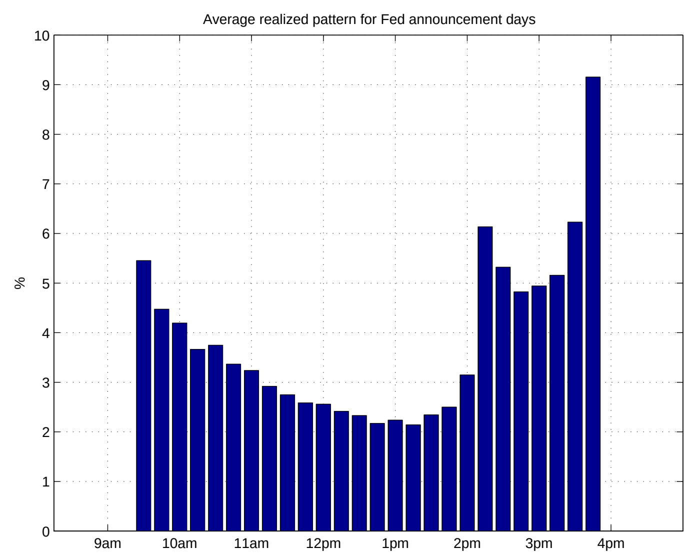

Fig. 10. Average realized volume pattern for scheduled Fed days, 2004-2009. The curve does not exhibit the typical U pattern of normal days.

the performance of a standard curve and special-purpose curves for these days.

Specifically, we considered the dates corresponding to the third Friday of each month between July, 2004 and April, 2009[2]. For each date, we built a VWAP curve by averaging volume patterns for all *other* dates, to prevent in-sample bias. Then we computed VWAP shortfalls for the top 1,000 stocks, using this special-purpose curve as well as a standard, fixed curve.

[1] excluding 2007.

[2] except March 21, 2008, which was a holiday.

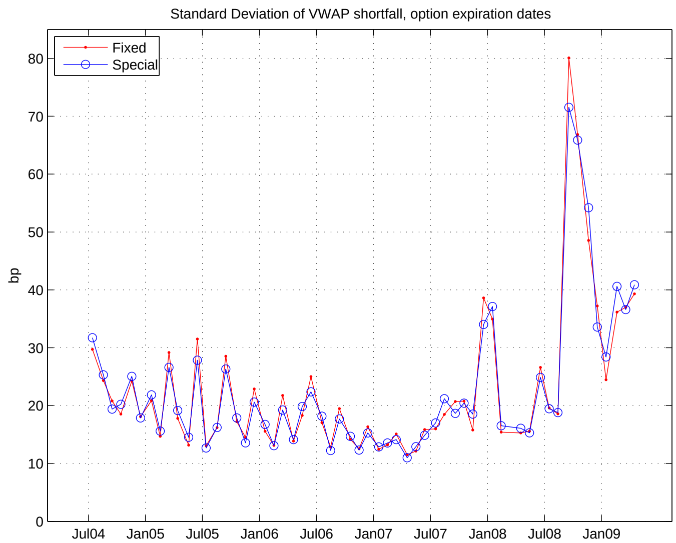

Fig. 11. Standard deviation of VWAP shortfall for option expiration days. The red points correspond to a fixed curve. The blue points were obtained with an average pattern for all other option expiration days. The standard error is 2.5 basis points.

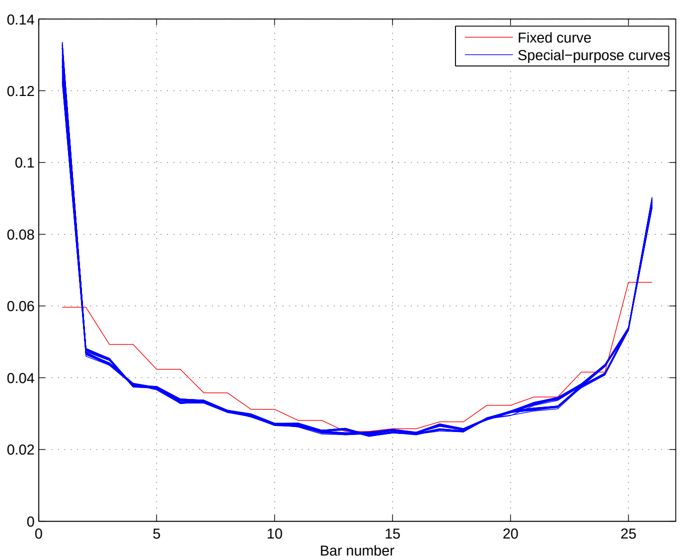

Fig. 12. Average VWAP curves from option expiration days (in blue). In red, the fixed curve used in the experiments.

The results are plotted in Fig. 11. The standard error is 2.5 basis points. It can be seen that the performance of the standard curve is very close to that of the special-purpose curves. The overall average standard deviation is 26 basis points in both cases.

Figure 12 depicts the fixed curve used (in red) and the various special curves (in blue). Clearly much higher volume than normal is traded during the first and last bar. Nevertheless, this systematic difference does not translate into lower variability.

It can be seen in Fig. 11 that the shortfall standard

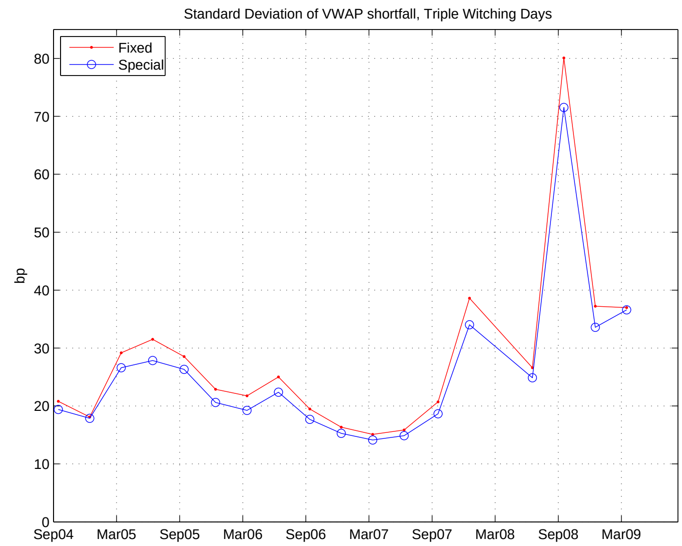

Fig. 13. Standard deviation of VWAP shortfall for triple-witching days.

deviation spikes every three months, in March, June, September and December. This is especially obvious in 2004-06. This phenomenon is known as the “Triple Witching Day” and refers to those days in which options, index futures, and options on index futures expire concurrently. Volume tends to be even higher those days.

Figure 13 shows the shortfall standard deviation for Triple Witching Days. The special-purpose curves outperform the fixed curve systematically by 2-3 basis points. The overall standard deviations are 29 and 32 basis points, respectively. This improvement is within one standard error of measurement, so it is deemed insignificant.

III. Baskets

The previous section dealt with single-name trades: all statistics were computed based on the shortfalls from trading individual names over a day.

In this section we consider multi-stock baskets. Here the shortfall for one day is computed as the difference in the weighted average price for the executed basket and the weighted average price for the same basket in the overall market. We expect some correlation, but by and large baskets should have an averaging effect, so we expect smaller standard deviations.

However, we can ask the same questions as before. In particular, is there curve drift? Do we benefit from updating the VWAP curve frequently? To answer this question we simulated VWAP shortfalls for two baskets:

- An equal-weighted, single-sided basket with the top 100 stocks by market cap.

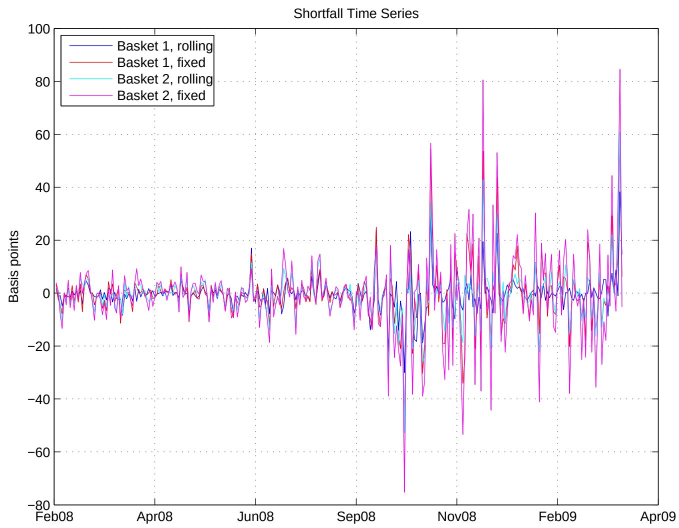

Fig. 14. Time series of shortfalls for two patterns, one rolling and one fixed, and two baskets.

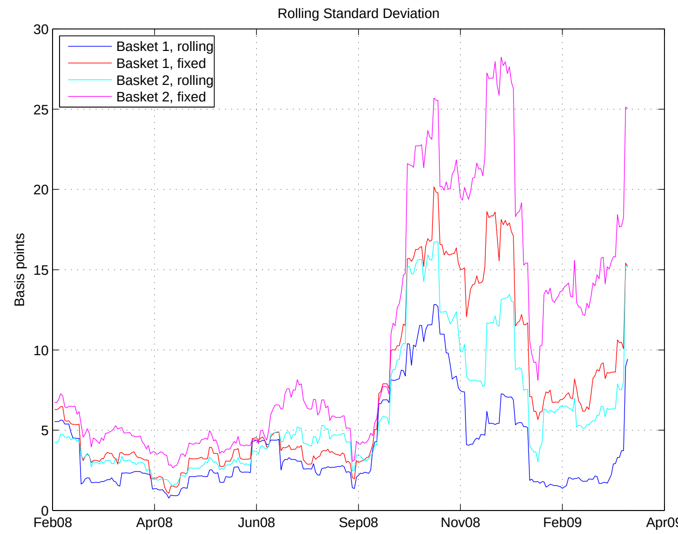

Fig. 15. Standard deviation of shortfalls for two patterns, one rolling and one fixed, and two baskets. Generally the volatility of baskets is lower than that of single stocks.

TABLE IV
STANDARD DEVIATION OF SHORTFALL (IN BASIS POINTS) FOR TWO BASKETS AND TWO CURVES, AVERAGED OVER FEB-MARCH 2009.

	Rolling Curve	Fixed Curve
Basket 1	6 ± 4	12 ± 5
Basket 2	11 ± 6	20 ± 7

- A real basket traded in the past with particularly unfavorable realized shortfall.

We simulated two VWAP strategies: a moving-average curve over the last month, and a fixed curve fitted with old data.

Figure 14 shows time series for the shortfalls. Since now we have only one shortfall per day, it is possible to visualize how the shortfall changes day by day (in the previous section, we had 100-3000 shortfall values per day). It is immediately clear that shortfall magnitudes have increased dramatically since September 2008. It is also apparent that the day-to-day variations are much larger than the difference between using one or another VWAP curve.

Figure 15 depicts the standard deviation of shortfall computed over a rolling two-month window. It can be seen that the difference between strategies is negligible up to September 2008. Then one observes differences: the moving-average curve outperforms the static curve for both baskets. The difference is not big, however, as shown in Table IV. The table shows the estimated standard deviation plus/minus two standard errors.

IV. Summary and Conclusions

In this study we examined the performance of various static VWAP strategies. The performance of a VWAP trade is measured by the shortfall between the executed average price and the market-wide realized VWAP. We showed that the average shortfall over many trades is zero, and the relevant performance metric is the shortfall dispersion. In this study we measured the dispersion as the shortfalls' standard deviation.

We estimated the various algorithms' performance by simulation using historical 2008-09 volumes and prices, and compared VWAP curves constructed using several combinations of estimation and simulation stock universes. In particular we examined the following cases:

- Fixed curves versus moving-average curves.
- Curves constructed by stock market cap.
- Curves constructed by stock sector.
- Curves constructed specifically for Fed announcement days.
- Curves constructed specifically for option expiration days.

We computed shortfalls both for single-stock orders and for baskets.

In general, customizing curves for shorter time spans or smaller stock universes introduces a fundamental trade-off. Curves for specific dates/stocks *decrease* dispersion because they apply to a more-homogeneous group. But since less data is available, there is less opportunity for averaging out any idiosyncratic errors and this may *increase* dispersion.

From the simulation results we can conclude that

- *Updating the curves on a rolling basis provides a benefit, both for single stocks and baskets.* The rolling curve provides average performance at worst, and at best it can beat the average curve by up to 4 basis points in standard deviation, or about 10%. A window of about 1 month (22 trading days) should be used for averaging.

- *Slicing stocks by market cap provides no benefit* beyond a minimum size necessary to ensure stable estimation. An estimation universe of 500 or more stocks provides best performance.
- *Slicing stocks by sector provides negligible benefit*, in the order of 1-2 basis points on average, or 5%, at best. Specific curves may in fact be detrimental for some sectors.
- *Special-purpose curves for Fed dates may reduce standard deviation* by up to 16%, but error bars are large. If a special curve is used, care must be taken to include as many Fed dates as possible.
- *Special-purpose curves for option expiration days provide no benefit*, for general option expiration days. For triple-witching days there is a visible but negligible improvement.

REFERENCES

[1] "Marching up the learning curve: The second buy-side algorithmic trading survey," Bank of America, 12 February 2007.

[2] R. Engle and J. Russel, "Autoregressive Conditional Duration: A New Model for Irregularly Spaced Transaction Data," Econometrica 66, 1127-1162, 1998.

[3] D. Hudson, "VWAP and Volume Profiles," Journal of Trading, 1(1) 38-42, 2006.

[4] J. McCulloch, "Relative Volume as a Doubly Stochastic Binomial Point Process," Research Paper 146, University of Technology Sidney, October 2005.

APPENDIX

Below we give a mathematical argument of why standard deviation of shortfall is the most relevant metric to quantify the performance of a VWAP strategy.

First, we note that the performance of a VWAP strategy is relatively independent of market moves. Define by $\bar{p}$ the VWAP prevailing in the market during the life of a particular trade. The shortfall between the executed price and the market's VWAP can be approximated using the following formula,

$$\text{sf} = \frac{10,000}{\bar{p}} \sum_{i=0}^{N-1} (w_i - v_i) p_i \qquad \text{bp}, \tag{1}$$

where N is the number of time periods, T is the life of the trade, w_i is the fraction of the daily market volume traded during the period $[t_0 + i\frac{T}{N}, t_0 + (i+1)\frac{T}{N}]$, v_i is the fraction of the traded shares executed during the period $[t_0 + i\frac{T}{N}, t_0 + (i+1)\frac{T}{N}]$, and p_i is the prevailing market price during the period $[t_0 + i\frac{T}{N}, t_0 + (i+1)\frac{T}{N}]$. This formula tells us that if the fraction of the order traded during a certain period equals the fraction of the total volume traded during that period, the shortfall is going to be zero *independently* of the price moves. If one knew ex-ante how the volume pattern is going to look like at the end of the day, one could achieve a near-zero VWAP shortfall by trading according to that pattern, regardless of price volatility. It is the inevitable deviations between w_i and v_i that result in the VWAP shortfall.

Second, the average shortfall over a large number of trades is zero. Assume that one trades all the shares at the beginning of the trade, i.e. $v_0 = 1$ while $v_1 = \ldots = v_{N-1} = 0$. In this case, for a sell (buy) trade the shortfall is[3] $\frac{\bar{p}-p_0}{\bar{p}}$ ($\frac{p_0-\bar{p}}{\bar{p}}$). Note that the buy and sell shortfalls offset each other. Therefore, if we assume that the side of our trade is random, the average shortfall is zero. This argument can be extended to show that the average shortfall is zero, for any curve $v_0, v_1, \ldots v_{N-1}$ used for trading.

Although the VWAP shortfall is zero on average, the dispersion of shortfall values does have an impact on the profitability of a trading strategy. Consider a trading strategy that has a theoretical Sharpe ratio of r if executed precisely at market VWAP. In reality, some practical VWAP algorithm is used, and a non-zero shortfall is incurred. Since the average shortfall is zero, the average return of the strategy remains the same. But the standard deviation of the shortfall increases the denominator of the Sharpe ratio. For example, if the standard deviation of the shortfall is approximately equal to the volatility of the strategy per bet, the strategy's actual Sharpe ratio is going be about 70% of r. This example and the fact that no VWAP algorithm can systematically outperform the market demonstrates that the trader/algorithm objective should be to minimize the deviation from the market VWAP independent of their side. The most convenient measure of dispersion is the standard deviation of the algorithm's shortfall.

[3] We use the sign convention that positive shortfalls are unfavorable.