

High Frequency Trading, Algorithmic Buy-Side Execution and Linguistic Syntax

Dan diBartolomeo

Northfield London Investment Seminar

November 2011

Goals for this Talk

- Assert that buy-side investment organizations engage in trading practices that are poorly aligned with underlying portfolio objectives
- Define the set of parameters needed in our approach to algorithmic equity trade execution, and to obtain these inputs by inference from portfolio characteristics
- Show how quantified news can be used to recalibrate the algorithm parameters in near “real time” to improve returns, reduce risks and reduce costs of trading
 - I apologize in advance but we will have to cover a lot of ground before I start talking about the role of quantified news

Introduction

- Algorithmic execution of buy-side orders has steadily gained an ever larger share of trading volume in most equity markets around the world.
- At the same time, the provision of liquidity from high-frequency trading operations has expanded even faster. Taken together, these two developments demand ever-increasing sophistication in execution algorithms.
- Trading is a zero sum game, but the HFT crowd has been making billions of dollars in aggregate annual profits at the expense of the buy-side. The trade execution algorithms available through all major brokers are simply not good enough

Four Steps to Level the Playing Field

- Our first step is a “pre-processor” that aligns the parameters of the execution algorithm to the composition and strategy of the underlying portfolio
- The second element is a “trade scheduling” algorithm that breaks one or more large equity orders (“parent order”) into a series of smaller trades (“child order”) to be executed over time.
- The third ingredient is a second algorithm that “micro-manages” the size, frequency and randomized timing of the execution of child orders
- **Use quantified news to recalibrate both trade scheduling and the micromanagement of the execution in real time!**

A Trade Schedule in Brief

- You can think of a list of undone trades as a long/short portfolio that you need to liquidate
 - You are long things you do have and don't want (sell orders)
 - You are short things you do want and don't have (buy orders)
- Our trade scheduling algorithm is a multi-period mean variance optimization in discrete time
 - We define N periods per day such that each successive period has an expectation of $(100/N)\%$ of the daily volume
- The output trading schedule looks like a spreadsheet where orders are rows and time periods are columns
- The whole schedule can be recalculated in a few seconds if we learn through quantified news that market conditions or other inputs have changed

Trade Scheduling in Words

- **We have two motivations to trade quickly.**
 - The first is opportunity cost or short term alpha. If we are buying a stock, we are presumably doing so because we believe the stock will go up, and we wish to buy before it goes up, not after. If we are selling, we believe the stock price will fall, and we wish to transact before, not after the decline.
 - The second reason to trade quickly is risk. The longer the trades remain unexecuted, the longer our underlying portfolio is different from what we want to hold, and we bear uncertainty around the relative performance of the portfolio we actually hold at each moment in time compared to our desired portfolio.
- **On the other hand, we have a reason to trade slowly**
 - The faster we demand liquidity from the market, the greater our market impact on prices will be. To the extent we cannot do all of our trades in a single execution, our buying (selling) will drive prices up (down) making our subsequent purchases (sales) more expensive.

Trade Scheduling in Math

$$U = E \left[\sum_{t=1}^{N} \left(\sum_{i=1}^{m} w_{it} \alpha_{it} - (1/RAP) \sum_{i=1}^{m} \sum_{j=1}^{m} w_{it} w_{jt} \sigma_i \sigma_j \rho_{ij} \right) - \sum_{i=1}^{m} \left(\text{abs} (w_{it} - w_{it-1}) \right) \kappa_{it} \right]$$

U = expected value of the objective function that we want to maximize

N = number of time periods in the schedule

w_{it} = the weight of security i in the “undone” portfolio at the start of period t

α_{it} = the short term alpha associated with security i during period t

ρ_{ij} = the correlation of security i and security j

σ_i = the kurtosis adjusted standard deviation of returns to security i

κ_{it} = amortized percentage trading cost for security i in period t

RAP = risk acceptance parameter

E = the expectations operator

Complexities in Scheduling

- The optimization problem is maximized subject to the constraint that at the completion of the schedule in N periods the starting weights for the hypothetical next period (w_{n+1}) for all non-cash securities must be zero.
 - The expected value for U will almost always be negative. Our alpha values are reversed from a conventional optimization, as a long position in this “undone” portfolio represents a sell order which is apt to be associated with a negative alpha, and a short portfolio position represents a buy order which is apt to be associated with a positive alpha.
- The multi-period nature of the problem induces path dependence in the solution.
 - If you sold 10,000 shares of stock XYZ in the first period, you don't have those shares any more to sell in the second period, and you must live with some portion of the market impact

Another Complexity

- Our last challenge is that the matrix of trading cost values, (κ_{it}) is a complex function of the amount each stock that is to be traded in a given period.
 - It is also path dependent because we assume that some portion of market impact is permanent and accumulates across multiple periods.
 - We also take into account that the market impact of trading a particular security will affect the price of other securities. For example, buying a large amount of Ford stock will push up the price. To the extent that investors base their valuation of GM and Toyota at partially on the price of Ford, the expected prices of GM and Toyota are also changed. This means that a combined order to “buy Ford, buy GM” will have a different expected cost than a combined order to “buy Ford, sell GM” executed in the same period.
- News will substantially effect both issues

Micromanagement of Executions

- Once an optimal trade schedule has been created we need to further sub-divide the shares to be transacted in each period into a series of child orders typically of a few hundred shares each
- Our micromanagement process randomizes time between child executions, and should adjust trading speed in response to real time market conditions such as price changes and **NEWS**
- Most Wall Street algorithms simplistically respond to price changes
 - If the current price is favorable relative to the decision price, you speed up. If the current price is unfavorable relative to the decision price, you slow down

Trade Scheduling Live Results

- Our trade schedule algorithm was implemented by Instinet (with their then existing micromanagement)
- Schmidt (2007) compared the effectiveness of our algorithm to VWAP and “participation trading” over more than 20,000 program trades done by Instinet
 - The multi-period algorithm dominated the others for all trades larger 6% of ADV, while results were inconclusive for trades less than 2% of ADV
 - For the other algorithms to be perceived as better, traders must have either extreme alpha expectations (over 50% annually) or be extremely risk averse
 - <http://www.northinfo.com/Documents/246.Pdf>

Calibration to Investor Preferences and Initial Beliefs

- The “pre-processor” tool that uses our risk models to analyze the portfolio from which the orders arise
- The analysis performed provides three scalar values and one vector as inputs for trading algorithms.
- No other information on the underlying portfolio is disclosed into the trading process.
- Both the pre-processor and the trade schedule computation engine sit at the buy-side client site so as to remove any potential for inadvertent leakage of position information.

Inferring Risk Aversion

- The first thing we need to understand how aggressive or conservative the underlying portfolio actually is
- If you have an explicit alpha forecast for the portfolio we can just calculate the slope of the mean variance efficient frontier as you lever the portfolio slightly
- If no explicit alpha forecast is available, we can make an inference of risk acceptance (RAP) from only the portfolio tracking error. Our approach is adopted from Wilcox (2003), and results in a simple rule of thumb that derives the RAP from underlying portfolio risk level
- This input is not recalibrated based on news during executions

Inferring Initial State Return Beliefs

- To trade efficiently we need to know what return beliefs are driving the trade decisions
- Once we know the risk and risk tolerance of a given portfolio, **a set of implied alpha forecasts must exist that will make the underlying portfolio mean-variance optimal**, as first described in Sharpe (1974)
- For long-only portfolios an adjustment is typically required to the basic implied alpha calculation. See diBartolomeo (2008)
- We apply an alpha decay rate and work backward to a vector of “short horizon” alpha estimates for the period of the trade execution. See Grinold and Stuckelman (1993) for the decay calculation assumptions
- **Adjustments for news will play a big role**

Inferring “Style” of the Portfolio

- The last initial parameter is a measure of the implicit assumptions of the underlying portfolio with respect to serial correlation in security returns.
- Relevant risk factors are used to characterize the portfolio.
 - Momentum oriented portfolios, it is expected that stocks that have gone up will keep going up and stocks that have gone down will keep going down (positive serial correlation).
 - A value oriented investor is generally assuming that stocks that have gone down will come back up and those that have risen will subsequently decline (negative serial correlation). For more details, see diBartolomeo (2003)
- This parameter as adjusted for “news” will dominate the calibration of the micromanagement process

Why We Need News in Addition to Observing Prices to Recalibrate

- Financial markets are driven by the arrival of information in the form of “news” (truly unanticipated) and the form of “announcements” that are anticipated with respect to time but not with respect to content.
- The time intervals it takes markets to absorb and adjust to new information ranges from minutes to days. Generally much smaller than a month, but up to and often larger than a day. That’s why US markets were closed for a week at September 11th.
- GARCH and other trend related models don’t work well on announcements
 - Market participants anticipate announcements
 - Volume and volatility dry up as investors wait for outcomes
 - Reduce volatility into the announcement and boost it after the announcement, so they are wrong twice

Interpreting the Content of News

- Several papers have examined the relative market response to “news” and “announcements”
 - Ederington and Lee (1996)
 - Kwag Shrieves and Wansley(2000)
 - Abraham and Taylor (1993)
- Jones, Lamont and Lumsdaine (1998) show a remarkable result for the US bond market
 - Total returns for long bonds and Treasury bills are not different if announcement days are removed from the data set
- Brown, Harlow and Tinic (1988) provide a framework for asymmetrical response to “good” and “bad” news
 - Good news increases projected cash flows, bad news decreases
 - All new information is a “surprise”, decreasing investor confidence and increasing discount rates
 - Upward price movements are muted, while downward movements are accentuated

Recalibrating Alpha Expectations with Quantified News

- Sentiment scores describing the content of news are obvious sources of directional forecasts.
 - See several chapters in [The Handbook of New Analytics in Finance](#) (Wiley, 2010)
- The abnormal volumes of news cause market participants to believe alpha values decay more quickly as market participants focus on short term events
 - Our information becomes “old news” faster
- Price trends strengthen on higher volumes of news
 - Liquidity declines (more on this coming up)
 - More positive serial correlation is observed
 - Value strategies work less effectively
- On average effects are stronger on the downside

Recalibrating Risk Expectations with Quantified News

- We can use the flow of news itself to estimate changes in market risk conditions
 - If there are an average of twenty stories a day on business news wires about MSFT, and today there are three hundred you can assume something is “up” and risk has increased
 - The text of news articles can be scanned for words and phrases that have positive or negative connotations
 - We can condition on this information just like option implied volatility
- See Mitra, Mitra and diBartolomeo (2008)
<http://www.northinfo.com/documents/313.pdf>

Recalibrating Risk Expectations with Quantified News

- Continue to use the existing risk models
 - Estimated from lower frequency return observations avoiding many statistical complexities
- Use information (i.e. news flow) that is not part of the risk model to adjust various components of the risk forecast to short-term conditions
 - Think of a vector of scaling factors that has one element for each important aspect of the risk model
 - Each element has a default value of one
 - Bayesian framework weights the new information with the default value to come to a scaling
 - We can use the same risk model for trading as for portfolio management, improving transparency and communication
- Will also impact expected security correlations
 - See diBartolomeo (1999) for math

Recalibrating Trading Cost Expectations with News

- Standard market impact models have two parameters, δ and π
 - The δ is usually security specific and in many models is proportional to expected volatility
 - The π exponent is assumed be either one half or one in most models.
- As news flows increase liquidity declines
 - Liquidity providers feel the increased risk and decrease capital available
 - Market impact costs increase
 - The “permanent” portion of market impact becomes longer lasting

Conclusions

- Most buy-side execution algorithms are not particularly good
 - If they were, the HFT crowd would not be so profitable
 - A lot of the problem is poor communication between portfolio managers and equity traders
- We are proponents of more sophisticated buy-side execution algorithms that balance short term alpha, risk and trading costs in an optimal fashion, conditional on inputs that must be inferred statistically from the original portfolio
- Every aspect of an algorithmic execution process can be improved by incorporating quantified news metrics to recalibrate an execution program once it has begun

References

- Ederington and Lee, "Creation and Resolution of Market Uncertainty: The Importance of Information Releases, Journal of Financial and Quantitative Analysis, 1996
- Kwag, Shrieves and Wansley, "Partially Anticipated Events: An Application to Dividend Announcements", University of Tennessee Working Paper, March 2000
- Abraham and Taylor, "Pricing Currency Options with Scheduled and Unscheduled Announcement Effects on Volatility", Managerial and Decision Science 1993
- Jones, Charles M., Owen Lamont and Robin L. Lumsdaine. "Macroeconomic News And Bond Market Volatility," Journal of Financial Economics, 1998, v47(3,Mar), 315-337.
- Brown, Keith C., W. V. Harlow and Seha M. Tinic. "Risk Aversion, Uncertain Information, And Market Efficiency," Journal of Financial Economics, 1988, v22(2), 355-386.

References

- Schmidt, Thorsten. "Implied Risk Acceptance Parameters in Institutional Equity Trades", Proceedings of Northfield Seminar, <http://www.northinfo.com/documents/246.pdf>, June 2007
- Wilcox, Jarrod. "Harry Markowitz and the Discretionary Wealth Hypothesis, *Journal of Portfolio Management*, 2003
- Sharpe, William F. "Imputing Expected Security Returns From Portfolio Composition," *Journal of Financial and Quantitative Analysis*, 1974, v9(3), 462-472.
- diBartolomeo, Dan. "Measuring Investment Skill Using the Effective Information Coefficient", *Journal of Performance Measurement*, Fall 2008.
- Grinold, Richard C. and Mark Stuckelman. "The Value-Added/Turnover Frontier," *Journal of Portfolio Management*, 1993,v19(4), 8-17.
- diBartolomeo, Dan. "Growth/Value/Momentum Returns as a Function of the Cross-Sectional Dispersion of Stock Returns", Northfield Research Conference, 2003, <http://www.northinfo.com/documents/350.pdf>.

References

- Wilcox, Jarrod W. "Better Risk Management," *Journal of Portfolio Management*, 2000, v26(4,Summer), 53-64.
- diBartolomeo, Dan. "Optimization of Composite Assets Using Implied Covariance Matrices", Northfield Working Paper, 1999.