JPMorgan

Understanding the Profit and Loss Distribution of Trading Algorithms

Robert Kissell

Vice President, JPMorgan

Robert.Kissell@JPMChase.com

Roberto Malamut, PhD

Vice President, JPMorgan

Roberto.Malamut@JPMChase.com

February 2005

Abstract

With the advent of algorithmic trading it is essential that investors become more proactive in the decision making process to ensure selection of the most appropriate algorithm. Investors need to specify benchmark price, implementation goal, and preferred deviation strategy (i.e., how the optimally prescribed algorithm is to react to changing market conditions or prices). In this paper we describe an analytical process to assess the impact of these decisions on the profit and loss distribution of the algorithm.

Originally published in Institutional Investor, Guide to Algorithmic Trading, Spring 2005.

Introduction

As financial markets have become more competitive, investors have started to turn to algorithmic trading to achieve better execution prices. However, the utilization of algorithmic trading alone does not guarantee better results. Traders need to become more proactive to ensure that the underlying algorithmic strategy is consistent with their investment objectives.

Prior to the selection of the strategy, traders first need to perform pre-trade analysis to assess the suitability of the order and/or trade list for algorithmic trading since not all orders are appropriately handled via algorithms. If algorithmic execution is deemed acceptable for the order, traders subsequently need to address macro- and micro-level issues. Macro level decisions include: specification of desired benchmark price, and implementation goal (e.g., aggressive or passive executions). The latter decision, however, requires in-depth cost analysis to understand the complete consequence of the decision. Micro level decisions include specifying any desired deviation rules. This includes how the algorithm should deviate from the optimally prescribed strategy based upon changing stock prices, market movement, and/or change in index or sector values, as well as changing market conditions such as volume profiles and volatility. Micro-level decisions also include specification of order submission rules such as market or limit order, display size, wait period for new submissions, order revisions, modifications, and/or cancellations. Finally, it is essential traders perform proper post trade analysis to assess the performance of the algorithm and ensure it is consistent with overall investment objectives. Traders who select algorithms without knowledge of potential outcome are ignoring their fiduciary responsibilities to their investors.

To best understand the algorithmic decision making process, however, it is important to understand the basics behind transaction cost management. Seminal transaction cost research is primarily due to Treynor (1981), Perold (1988), Berkowitz, Logue, & Noser (1988), Wagner (1990), and Hasbrouck (1991). More recently, however, Bertsimas & Lo (1996), Almgren & Chriss (1999), and Kissell, Glantz, and Malamut (2004) expanded this work to provide a decision making framework to manage transaction costs. Accordingly, this work now serves as the basis for algorithmic decision making.

Pre-Trade Analysis

Pre-trade analysis provides the necessary data to make informed algorithmic trading decisions. It provides investors with liquidity summaries, cost & risk estimates, as well as trading difficulty and stability measures to determine which orders can be successfully implemented via algorithmic trading and which orders require manual intervention. Not every order is well suited for algorithmic implementation. Pre-trade analysis provides insight into potential risk reduction and hedging opportunities to further improve execution. Finally, pre-trade analysis also provides investors with necessary data to develop views for short-term price movement and market conditions.

After evaluating the suitability of the order for algorithmic trading it is essential traders develop a customized algorithm exactly suited to the specific order. Brokers who do not offer flexibility to customize the algorithmic strategy are unable to provide investors with the most appropriate algorithms.

Macro Level Decisions

Benchmark Price

The first step in the algorithmic decision making framework is specification of the benchmark price. The more common benchmarks can be categorized into pre-, intra-, and post-trade prices. The pre-trade benchmark prices, also commonly referred to as implementation shortfall ("IS") benchmark prices, are those prices that are known before or at the time trading begins. These include the investment decision price, previous night's closing price, opening price, and arrival price (i.e., price at time of order entry). Intraday benchmarks are comprised of those prices that occur during trading such as VWAP, TWAP, and average of open, high, low, and close. Post-trade benchmarks include any price that occurs after or at the end of trading, the most common of which being the day's closing price.

The benchmark price is investor specific and may in fact be different for two investors with identical trade lists. For example, a value manager may desire execution at their decision price (i.e., the price used in the portfolio construction phase), a mutual fund manager may desire execution at the closing price to coincide with valuation of the fund, and a indexer may desire execution that achieves VWAP (e.g., to minimize market impact) as an indication of fair prices for the day. Furthermore, the same investor may specify different benchmark prices for identical orders (i.e., 500,000 shares of MSFT) on different days if the investment objectives have changed.

Table 1: Benchmark Prices

Pre-Trade	Intra	Post
Decision Price	VWAP	Future Close
Previous Close	TWAP	
Opening Price	OHLC	
Arrival Price		

Implementation Goal

The next step in the process is to specify the intended implementation goal. This relates to the level of trading aggressiveness or passiveness. Aggressive trading is associated with higher cost and less risk while passive trading is associated with lower market impact and higher risk. This market phenomenon gives rise to the trader's dilemma:

> *"Trading too aggressive will lead to higher impact cost but trading too passively will lead to higher risk and may result in even more costly trades."*

Therefore, the implementation goal is to solve the trader's dilemma. The solution is found by balancing the tradeoff between cost and risk based at the investor specified level of risk aversion as follows:

$$\underset{\alpha}{Min} \quad Cost(\alpha) + \lambda \cdot Risk(\alpha)$$

where λ is the risk aversion parameter and α is the trading rate (defined as a percentage of volume). Solving for various levels of lambda will result in numerous trading strategies each with different cost & risk. The set of all these points comprises the efficient trading frontier ("ETF") introduced by Almgren & Chriss (1999) and provides quantitative insight

Figure 1: Efficient Trading Frontier for Various Benchmark Prices

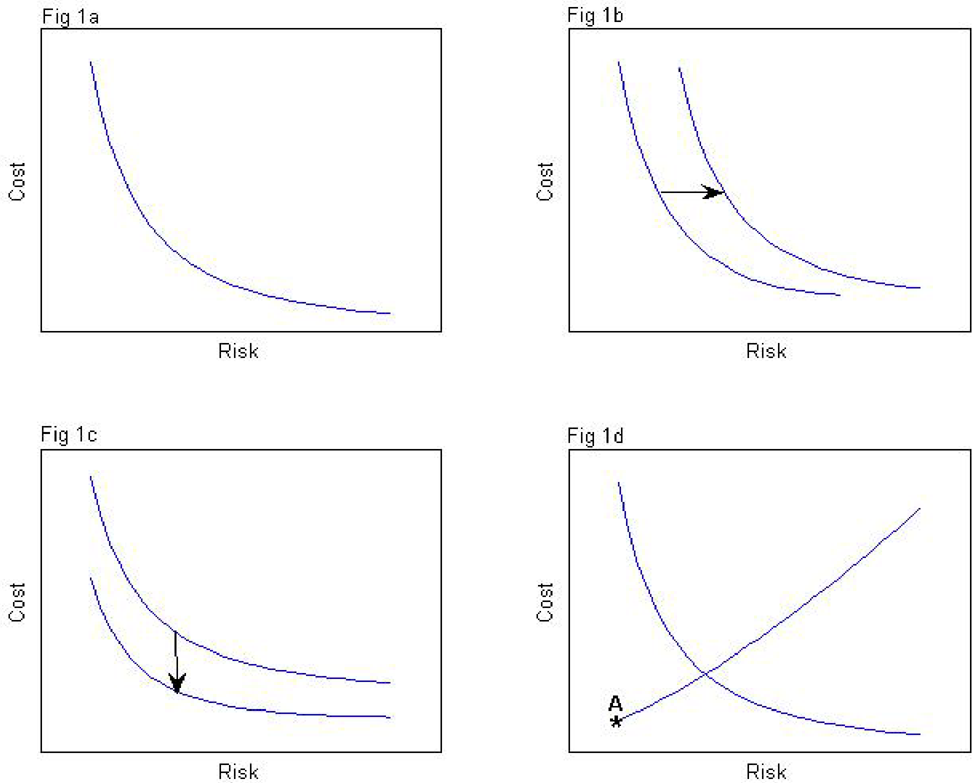

Figure 1: Efficient Trading Frontier. The efficient trading frontier (ETF) is the set of all optimal trading strategies. That is, those strategies with the lowest cost for a given quantity of risk and the least risk for a specified cost. Figure 1a illustrates the ETF for an arrival price benchmark. Notice the aggressive strategies are associated with higher cost but less risk while the passive strategies are associated with lower cost but higher risk. Figure 1b compares the ETF for a pre-trade decision price benchmark (such as the previous day's closing price) to the arrival price ETF. Notice that this ETF is shifted to the right to account for the additional risk of the unknown price change from the closing price to next day's opening price. Figure 1c compares the ETF for a future price benchmark such as the day's closing price to the arrival price ETF. This ETF is shifted downwards to account for a lower benchmark cost since the future price will consist of permanent market impact whereas the order arrival price does not. Figure 1d compares the ETF for a VWAP benchmark to the arrival price ETF. The cost profile for a VWAP execution (or any intra-trade benchmark) is unique in that investors can minimize both benchmark cost and risk by participating with volume over the horizon (denoted by strategy "A"). Therefore, the VWAP ETF is upward sloping and investors can simultaneously minimize benchmark cost and risk.

into cost consequence associated with level of aggressiveness.

After computing the ETF investors select the most appropriate optimal strategy based upon their investment objective. For example, informed traders with expectations regarding future price movement are likely to select an aggressive strategy (e.g., POV=30%) with higher cost but more certainty surrounding expected transaction prices. Indexers are likely to select a passive strategy (e.g., POV = 5%) or a risk neutral strategy to reduce cost. Some investors may select a strategy that balances the tradeoff between cost and risk depending upon their level of risk aversion, and others may elect to participate with volume throughout the day (e.g., VWAP strategy).

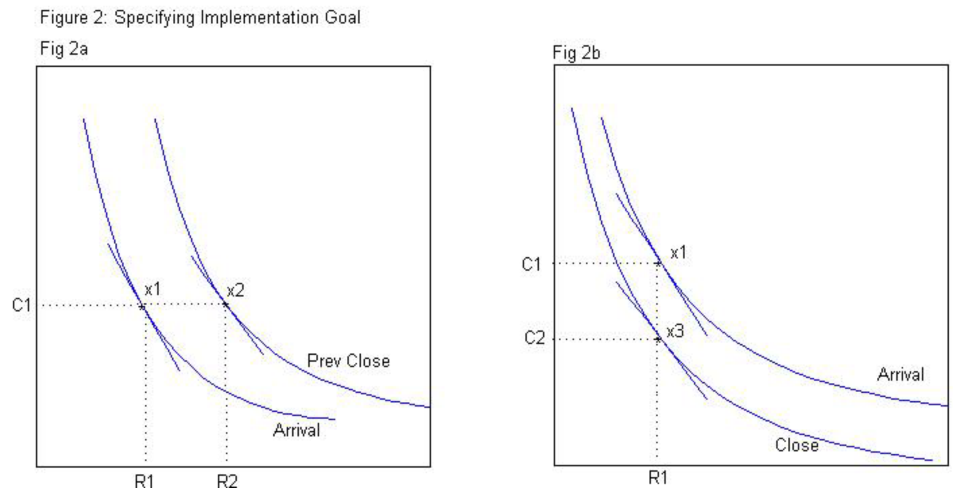

Figure 2: Specifying Implementation Goal. Figure 2a illustrates how an investor with an implementation goal to achieve a cost C_1 will have different execution strategies for an arrival price and previous close (decision price) benchmark. Since there is less risk associated with the arrival price benchmark investors will execute following strategy X_1 to achieve cost C_1 with lower risk R_1. Investors with the previous closing price benchmark are subject to higher risk due to potential overnight price movement and will execute following X_2 to achieve cost C_1 with higher risk R_2 ($R_1<R_2$). Figure 2b illustrates how different benchmark prices and different implementation goals can result in identical execution strategies. Here investors with different benchmark prices (arrival and day's close) as well as different desired costs can result in identical strategies X_1 and X_2 both with risk R_1.

To solve the optimization problem and construct the ETF, however, it is necessary to formulate the cost and risk of the execution. To address this need we first estimate the expected average execution price at commencement of trading as follows[1]:

$$\overline{P}_0(\alpha) = P_0 + f(X,\alpha) + g(X) + \varepsilon(\alpha)$$

where, X is the number of shares to trade, P_0 is the market price at time of order entry, $f(X,\alpha)$ is temporary impact cost due to the liquidity demands of traders, $g(X)$ is the permanent impact cost due to the information leakage of the order, $\varepsilon(\alpha)$ is price volatility. The cost function is computed as follows:

$$\begin{aligned} Cost(\alpha) &= E_0\left(\overline{P}(\alpha) - P_b\right) \\ Risk(\alpha) &= \sigma(\varepsilon(\alpha)) \end{aligned}$$

We next provide the derivation of Cost and Risk for the arrival, previous night's closing price, future closing price, and VWAP benchmarks.

Arrival Price

$$\begin{aligned} Cost(\alpha) &= E_0\left(\overline{P}_0(\alpha) - P_0\right) \\ &= P_0 + f(X,\alpha) + g(X) + E[\varepsilon(\alpha)] - P_0 \\ &= f(X,\alpha) + g(X) \end{aligned}$$

with $E[\varepsilon(\alpha)] = 0$ and $Var[\varepsilon(\alpha)] = \sigma^2(\alpha)$.

The corresponding risk of the execution is:

$$\Re(\alpha) = \sigma(\varepsilon(\alpha))$$

The ETF for the arrival price benchmark

[1] For simplicity we express expected execution price only in terms of market impact. In actuality, the expected execution price will consist of both market impact and price appreciation (momentum, trend, alpha) cost.

is constructed using various values for the trading rate α (i.e., percentage of volume) and is shown in figure 1a. The chart illustrates that more aggressive trading will incur higher cost but less risk, and more passive trading incurs lower cost but a higher quantity of risk. The ETF allows traders to analyze and then select their preferred optimal trading strategy (i.e., the strategy with the lowest cost for a given level of risk and the least risk for a specified cost).

Previous Night's Closing Price

The ETF for the previous night's close (denoted as P_d) is computed in a similar manner to the arrival price as follows:

$$\begin{aligned} Cost(\alpha) &= E_0\left(\overline{P}_0(\alpha) - P_d\right) \\ &= P_0 + f(X,\alpha) + g(X) + E[\varepsilon(\alpha)] - P_d \end{aligned}$$

But since $P_0 = P_d + \xi$ where ξ is the price change from previous close to the time of order entry (unknown in advance) with $E[\xi] = 0$ and $Var[\xi] = \sigma^2(\xi)$ we have:

$$\begin{aligned} Cost(\alpha) &= P_0 + f(X,\alpha) + g(X) + E[\varepsilon(\alpha)] - (P_0 - \xi) \\ &= f(X,\alpha) + g(X) \end{aligned}$$

and,

$$\Re(\alpha) = \sqrt{\sigma^2(\varepsilon(\alpha)) + \sigma^2(\xi)}$$

Notice that in this case the associated timing risk of the trade is larger than for the arrival price which is due to the fact that there is potential for price movement in the stock price from the time of the investment decision to the time of the order entry. For analytical purposes, the higher timing risk causes a shift to the right in the efficient trading frontier. This is shown in figure 1b. Therefore, two investors with the identical level of risk aversion but one is using the arrival price benchmark and the other is using the previous night's closing benchmark will specify different algorithmic strategies altogether. These strategies, however, will have the same expected cost but different timing risk. This is depicted in figure 2a. Notice that strategies X1 & X2 have the same expected cost C1 but different risk. The arrival price benchmark has a lower risk R1 than the previous closing price benchmark R2 because it is subject to incremental overnight risk. Therefore, for identical levels of risk aversion the previous night's closing price benchmark will be more passive than the arrival price algorithm.

Future Closing Price

To compute the cost profile for a future closing price it is necessary to estimate the expected future closing price P_T. By definition, all future prices will include the permanent market impact of the order (recall that for simplicity we assume no price appreciation over the trading period) but not temporary impact cost.

Therefore, we have:

$$P_T = P_0 + g(X))$$

The cost profile is thus computed as follows:

$$\begin{aligned} Cost(\alpha) &= E_0\left(\overline{P}_0(\alpha) - P_T\right) \\ &= P_0 + f(X,\alpha) + g(X) + \varepsilon(\alpha) - (P_0 + g(X)) \\ &= f(X,\alpha) \end{aligned}$$

with corresponding timing risk as follows:

$$\Re(\alpha) = \sigma(\varepsilon(\alpha))$$

Notice that for the future closing price benchmark the expected cost is equal to temporary market impact only, and it is less than the arrival price benchmark but has the same quantity of timing risk. This causes a downward shift in the ETF (illustrated in figure 1c). The magnitude of the downward shift is equal to the quantity of permanent market impact. The consequence of the future price benchmark on the algorithmic trading decision will result in the same strategy but with lower expected cost. This is illustrated in figure 2b. Notice that strategies X1 and X3 have different expected costs, C1 & C3 respectively, but for the same level of risk aversion, these strategies will have the same quantity of risk R1. It is important to reiterate here that the risk term R1 denotes the uncertainty corresponding to the expected cost estimate. That is, the uncertainty surrounding the expected execution price and the benchmark price.

VWAP Benchmark

The ETF associated with the VWAP benchmark (or any of the intra-day benchmarks) has a much different shape than any of the pre-trade or post-trade benchmarks and is an increasing function with respect to risk. This is because the intra-day benchmark prices include both temporary & permanent impact cost, and the timing risk calculation is based upon all prices and trades over the entire trading period rather than from a specific point in time. For intraday benchmark prices, traders minimize both cost (i.e., expected gain/loss to the benchmark) and uncertainty of the expected gain/loss (timing risk) by participating with volume. Any deviation from this strategy will result in higher risk. This is depicted in figure 1d with the only "rational" optimal strategy denoted by "A."

Micro Level Decisions

Strategy Deviation Rules

It is important that investors specify their preferences for the algorithmic deviation rules. This defines how the algorithm should deviate from the originally prescribed optimal strategy and implementation goal. The more common deviation rules include changing the execution rate (to be more or less aggressive) or POV rate based upon changing market volumes, stock price movement (momentum), overall market movement, or possibly based on changes in index values such as the S&P500 or based on changes in the stock's specific sector index. Additionally, investors may also choose to change the execution strategy based on changing market conditions such as volume profiles due to special events such as new announcements, or fed indicators, as well as changes in volatility.

In specifying the appropriate deviation strategy, it is important to have a complete understanding of how the deviation rule impacts the cost distribution. For example, many price-based scaling rules (adjusting the execution rate based on price movement) are designed to result in lower cost on average, but it is accompanied with an increase in tail risk and diminished possibility for large gains. Therefore, in times of adverse momentum, the deviation strategy may be more costly than a strategy without any specified deviation rule. Additionally, it is possible to develop deviation rules that further minimize the potential for large losses with an increase in potential for gains, but this comes at an increased cost. Regardless of the specified deviation rule, it is essential that traders understand the impact of the decision on the cost distribution.

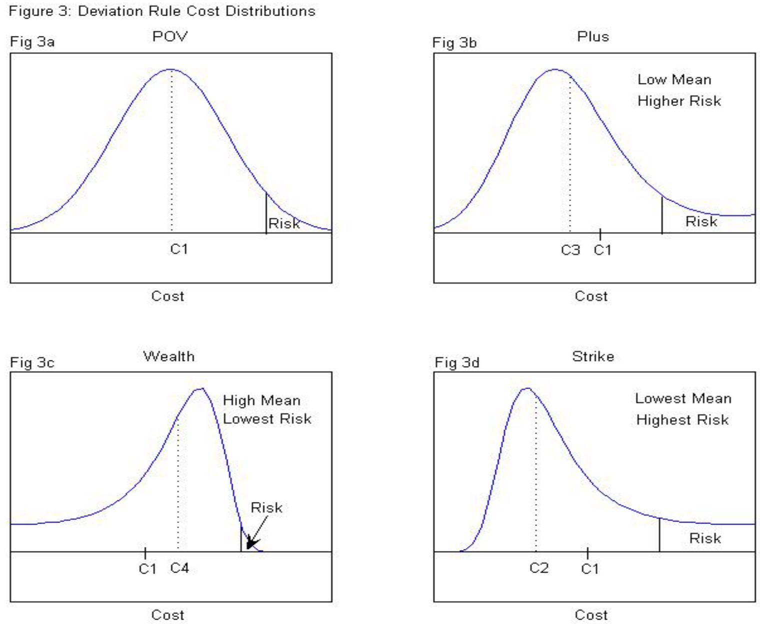

Figure 3: Deviation Rule Cost Distributions. Figure 3 illustrates the trading cost distribution associated with alternative deviation strategies (i.e., how the algorithm will adapt to changing market conditions such as price trends). Figure 3a illustrates a standard trading cost distribution for a specified initial trading rate expressed as a percentage of volume ("POV"). In this situation the expected cost is C_1 with the potential of unfavorable prices denoted as "Risk" in the right hand tail. Figure 3b illustrates the distribution for a "Plus" strategy that maximizes the likelihood that the realized price will be more favorable than the benchmark price. In this situation the algorithm will be aggressive in times of favorable price and passive in times of adverse price movement to limit excess market impact cost. The overall result will be a lower cost on average $C_3<C_1$ but higher potential for unfavorable prices if the adverse trend continues. This higher risk is denoted as "Risk" in the right hand tail. Additionally, the "Plus" algorithm limits the potential to realize better prices if the favorable trend persists. Figure 3c illustrates the distribution for a "Wealth" strategy that becomes aggressive in times of adverse price movement to limit losses and passive in times of favorable prices to take advantage of the better prices. The overall result is a higher cost on average $C_4>C_1$ caused by being aggressive in times of adverse price movement but the benefit is lower risk of high costs and higher potential for gains. Figure 3d illustrates the cost distribution for a "Strike" strategy that will continuously adapt to changing market conditions such that expected realized price will be equal to the benchmark price. In this situation the algorithm will be most aggressive in times of favorable prices to lock into the benchmark price and passive in times of adverse price movement in order to limit unnecessary market impact cost. The overall result here is lower cost on average $C_2<C_1$ but with increased exposure to unfavorable prices if the adverse trend continues. The "strike" deviation strategy also limits the potential for better price if the favorable trend persists.

To best illustrate this point we will discuss the impact of deviation rules on the cost distributions resulting from price based scaling schemes[2]. For example, one price-based scaling scheme may be to become more aggressive (i.e., increase the POV rate) when prices are more favorable than the benchmark and become less aggressive (i.e., decrease the POV rate) when prices are less favorable than the benchmark. An alternative deviation scheme is to become more passive when prices are favorable and more aggressive when prices are unfavorable to limit potentially higher costs and continue to participate with the better prices with the

[2] In addition to stock price movement, price-based scaling schemes can also be developed based on market movement, index values, sector values, or industry values, or any other momentum measure.

hope of even better execution prices.

There are three (3) main categories of price-based scaling techniques: i) Strike (associated with the lowest mean and highest risk compared to the static POV strategy and alternative price-based scaling schemes), ii) Plus (associated with a lower mean and higher risk than the static POV strategy), and iii) Wealth (associated with a higher mean but lower risk compared to the static POV strategy)[3]. These distributions are shown in figures 3a – 3d. Each technique adjusts the algorithm in a different manner to become more or less aggressive based upon changing prices and market conditions. But before investors/traders specify the preferred deviation strategy, it is essential to have a clear understanding of the consequences of the "rule" on the cost distribution. These are described below:

Strike:

$$\underset{\alpha(t)}{Min} \quad E_t\left(\overline{P}_t(\alpha_t) - P_b\right)^2$$

The strike deviation rule dynamically adjusts the POV rate to minimize the quadratic cost function without regards to risk. Here, the algorithm will become more aggressive in times of favorable prices and less aggressive in times of unfavorable prices. For example, for a buy order the strike algorithm will become more aggressive when prices are less than the specified benchmark price ("in-the-money") and less aggressive when prices are higher than the benchmark ("out-of-the-money"). This deviation scheme provides lower cost on average but comes at the expense of increased risk on the cost side and reduced potential to achieve very favorable prices because the order will likely be completed prior to these prices arising. This cost distribution for the strike is shown in figure 4a. The expected cost C2 is less than that for the static POV strategy with expected cost C1. It is lower on average because the algorithm will take advantage of favorable prices when they arise. The distribution, however, is skewed to the cost side. This is because by becoming more passive when prices are unfavorable there is a higher chance of realizing extremely unfavorable prices if adverse price trends persist. Similarly, this deviation rule does not provide potential opportunity to achieve better prices if favorable trends persist because it is likely the order will already be completed by the time these prices arise. Therefore, the cost curve is also truncated on the benefit side.

Plus:

$$\underset{\alpha(t)}{Max} \quad \frac{E_t\left(\overline{P}_t(\alpha_t) - P_b\right)}{\Re(\alpha_t)}$$

The "plus" deviation rule dynamically adjusts the POV rate to maximize the likelihood that the algorithm will outperform the specified benchmark price. Notice that this formulation is identical to that of maximizing the Sharpe ratio. The "plus" deviation rule behaves similar to the "strike" algorithm by becoming more aggressive when "in-the-money" and less aggressive when "out-of-the-money." But unlike the "strike", the "plus" is risk sensitive (notice the risk term in the denominator) so it does not expose investors to the same degree of fat tail risk. It does, however, provide opportunity to achieve somewhat better prices if favorable trends persist. Unfortunately, this comes at a slightly higher cost on average $C_3 > C_2$. Figure 4b compares the distribution of the "plus" deviation rule to the static POV rule. By continuously adapting to

[3] Price-based scaling is only one type of strategy deviation rule. Other rules can be developed based on volumes and volume profiles, volatility, news and information, etc.

Figure 4: Comparison of Price-Based Scaling Cost Distributions

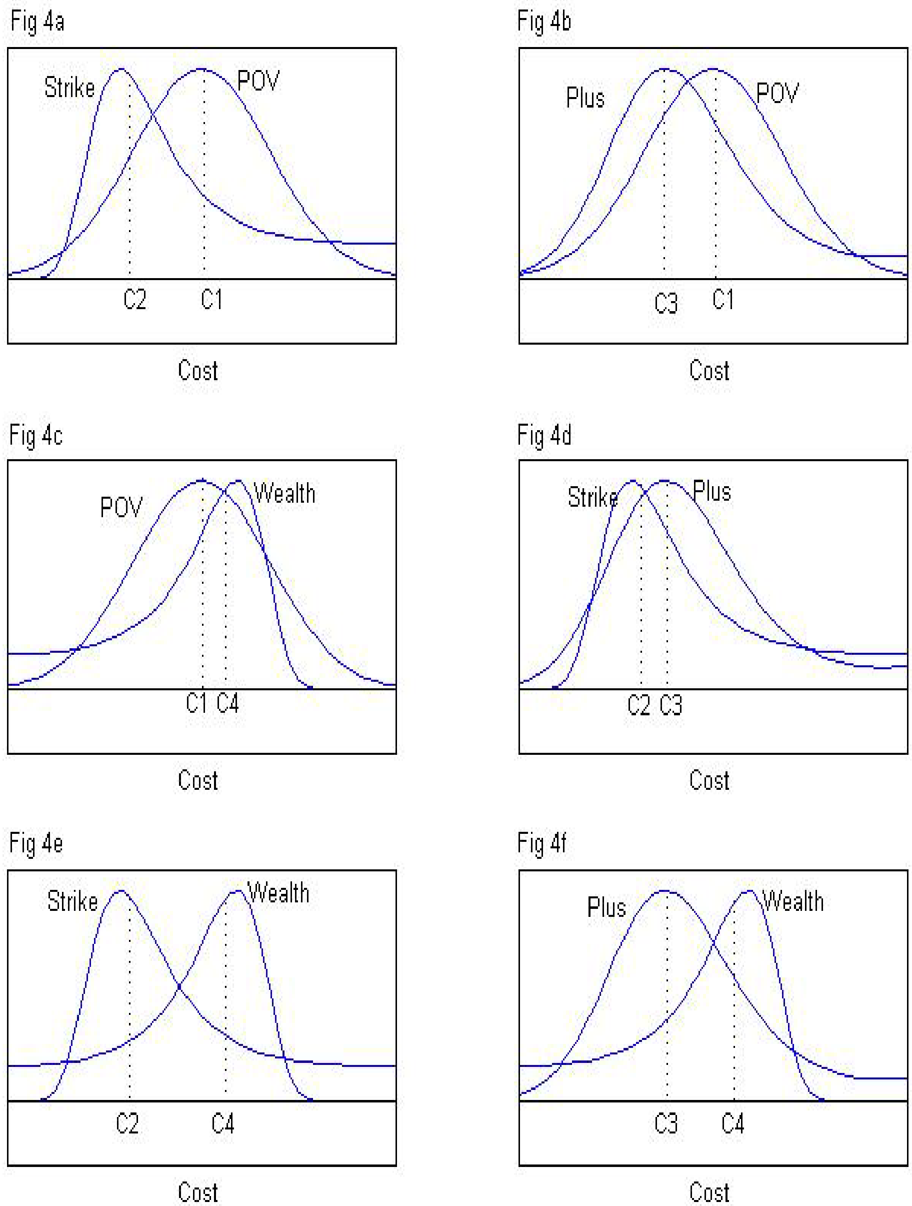

Figure 4: Comparison of Price-Based Scaling Cost Distributions. Figure 4a compares the "Strike" algorithm to the POV strategy. The Strike algorithm will incur a lower cost on average but exposes the trade to great risk of unfavorable prices and limits the potential to participate with favorable prices. Figure 4b compares the "Plus" strategy to the POV strategy. The Plus strategy incurs also lower cost on average but with increased risk of unfavorable prices and decreased potential for participating in better prices if favorable trends continue. Figure 4c compares the "Wealth" strategy to the POV strategy. The Wealth strategy limits the potential to incur unfavorable prices and increases potential to participate with better prices in times of favorable trends, but this comes at a slightly higher cost on average than a POV strategy. Figures 4d-4e compare the cost distribution of each of the deviation strategies. It illustrates to investors the change in costs and shift in distribution for the different deviation algorithms so that investors can choose the algorithm most inline with overall investment objective.

market conditions, it will realize lower costs on average $C_3<C_1$. However it does increase risk of higher potential costs and does not allow for the full array of better prices, although it does provide additional protection over the "strike" as well as opportunity for additional gains. Figure 4d compares the "plus" algorithm to the "strike" algorithm.

Wealth:

$$\underset{\alpha(t)}{Max} \quad \log E_t\left\{H\left[\overline{P}_t(\alpha_t)\right]\right\}$$

The "wealth" deviation rule is most closely related to the traditional economic style risk aversion utility function H introduced by Pratt (1964) and later formalized by Arrow (1971), and it is intended to maximize investor wealth in presence of uncertainty. This deviation rule behaves in the opposite manner of "Strike" and "Plus" by becoming more passive in times of favorable prices and more aggressive in times of unfavorable prices as a means of limiting potential large losses if adverse conditions persist. The cost distribution of the "wealth" deviation rule is illustrated in figure 3c. Notice that it is skewed to the left and truncated on the right showing the larger potential for large gains and ample protection from adverse market conditions. But it does come at an increase in expected cost on average with $C_4 > C_1$. Figure 4e and 4f compare the "wealth" cost distribution to the "strike" and "plus" cost distribution curves respectively.

A final point with regards to these price based scaling deviation rules is that they are designed for individual order execution. For list trading, it is important that deviation rules be specified based on how the adjustments affect the cost & risk profile of the entire trade list. For example, investors selecting the "strike" deviation rule for hedged list (i.e., buy-sell beta neutral list) will find that the algorithms execute orders "in-the-money" and hold onto orders "out-of-the-money" which will result in higher costs (if trends do not reverse) than an algorithmic strategy without any price-based scaling schemes. For list trading, price-based scaling and deviation rules need to be determined based on the consequence to the entire trade list. In list trading situations, an appropriately specified price based scaling scheme should only deviate when doing so would result in lower cost, better prices, and overall reduced total trade list risk. Some examples of proper trade list deviation strategies are described in Malamut (2002) and Kissell & Glantz (2005).

Order Submission Rules

Order Submission Rules refer to the actual market pricing schemes (e.g., market or limit order), share quantities, wait period between order submissions, revisions, and cancellation. The more common pricing rules include market and limit orders (all variations) as well as floating prices that are pegged to a reference price such as the bid, ask, or midpoint and change with the reference price. Varying these order types allows the algorithm to adhere to the optimally prescribed strategy by executing aggressively (i.e., market orders) and/or passively (i.e., limit orders) when needed. The use of an algorithm when submitting marketable orders affords a degree of anonymity and does not highlight the type of order being entered.

In most situations it is appropriate to combine limit, market, floats and reserve orders. For example, suppose the specified macro level optimal strategy is a POV rate of 15%. Here a micro-level algorithm may submit limit orders to the market for execution better than the mid-quote for as long as the actual POV rate is consistent with 15% of market volume, but once the algorithm starts lagging behind the specified rate it would submit appropriately sized and spaced market orders to be more aggressive. A reserve order could also be used to automatically replenish limit orders at favorable prices. More advanced micro pricing strategies utilize real-time data, prices and quotes, and recent trading activity to forecast very short-term price trends while providing probabilistic estimates surrounding the likelihood that a limit order will execute within a certain period of time. For example, a limit order model will provide probabilistic estimates regarding the likelihood that specified number of shares will be executed within a desired amount of time. When the likelihood of completion is too low, investors would be better served via a market order rather than the desired limit order.

Post Trade Analysis

Algorithmic post trade analysis is a two part process that consists of cost measurement and algorithm performance analysis. First, cost is measured as the difference between the actual realized execution price and the specified benchmark price. This allows investors to critique the accuracy of the trading cost model to improve future cost estimates and macro strategy decisions, and it provides managers with higher quality price information to improve investment decisions. Second, algorithmic performance is analyzed to assess the ability of the algorithm to adhere to the optimally prescribed strategy, its ability to achieve fair and reasonable prices, and determine if the algorithm deviates from the optimally specified strategy in an appropriate manner. Investors must continuously perform post-trade analysis to ensure brokers are delivering as advertised and question those executions that are out of line with pre-trade cost estimates.

Summary

Algorithmic trading has recently become a popular vehicle for efficient and low cost executions. However, proper specification of algorithmic rules requires investors to become more proactive during implementation by taking greater control of their execution decisions. Accordingly, investors are required to specify decision on the macro level, i.e., specification of benchmark price and implementation goal, as well as on the micro level, i.e., appropriate deviation schemes (e.g., price-based scaling, volume patterns, volatility, etc.) and order submission rules. It is essential that traders not only develop customized algorithms so that expected transaction costs (e.g., market impact and timing risk) are consistent with their overall investment objectives, but that they also analyze and compare alternative algorithms to determine the most appropriate algorithm. Moreover, it requires investors to select brokers who can best customize algorithms with investor implementation and investment goals. Only then will investors have highest chances to achieve their investment goals.

References

Arrow, K.J., “Essays in the Theory of Risk-Bearing,” North Holland, Amsterdam, 1971.

Almgren, R., and N. Chriss, (1999). "Value Under Liquidation," J. Risk 12(12), 61-63.

Almgren, R., and N. Chriss, (2000). "Optimal Execution of Portfolio Transactions." J. Risk 3(2), 5-39.

Berkowitz, S., D. Logue, and E. Noser (1988). “The Total Cost of Transactions on the NYSE.” Journal of Finance 41, pg 97-112.

Bertsimas, D., and A. Lo, (1998). "Optimal Control of Liquidation Costs," Journal of Financial Markets, Vol 1.

Hasbrouch, J. (1991). “Measuring the information content of stock trades,” Journal of Finance 46, page 179-207.

Kissell, R. and M. Glantz (2003). "Optimal Trading Strategies," AMACOM, Inc., NY, 2003.

Kissell, R., M. Glantz, and R. Malamut, (2004). “A Practical Framework for Estimating Market Impact and Achieving Best Execution,” Finance Research Letters I pg. 35-46.

Kissell, R., and M. Glantz (2005). “Managing Trading Risk.” Forthcoming Elsevier Science.

Malamut, R. (2002), “Multi-Period Optimization Techniques for Trade Scheduling,” QWAFAFEW, April 2002.

O’Hara, M. (1995). Market Microstructure Theory. Cambridge, MA: Blackwell 1995

Perold, A.F., (1988). "The Implementation Shortfall: Paper versus Reality," Journal of Portfolio Management, Spring 1988.

Pratt, J.W., “Risk Aversion in the Small and the Large,” Econometrica, Jan-Apr, 1964, page 122-136.

Treynor, J. L., (1981). "What Does It Take to Win the Trading Game," Financial Analysts Journal, January/February 1981, pg. 55-60.

Wagner, W. (1990). “Transaction Costs: Measurement and Control,” Financial Analysts Journal, September/October 1990.