AutomatedTrading

Terrence Hendershott

Keywords: Liquidity, algorithms, automation, market microstructure, market making, limit order book, price impact, market fragmentation

Abstract:

Technology changes the way securities are traded. Both liquidity suppliers and liquidity demanders use computer algorithms to improve and automate their trading. Liquidity demanders attempt to minimize the price impact of a large order by dividing their order into small parts and executing it over time across multiple markets. Liquidity demanders do this to disguise the magnitude of their trading interest to prevent other market participants from capitalizing on their future trading demands. Liquid suppliers condition their quotes -- buying and selling price schedules -- on all available information. Technology allows more information to be processed more quickly. This should improve the efficiency of prices and allow for liquidity to be supplied at lower costs. While humans create and parameterize the algorithms used to trade, the actual trading process involves fewer and fewer humans.

Technological change has revolutionized the way financial assets are traded. Investors place orders via computer rather than speaking to a broker on the phone. Trading floors have largely been replaced by electronic trading platforms [13]. The nature of order execution has changed dramatically as well, as many market participants now employ algorithmic trading (AT), commonly defined as the use of computer algorithms to manage the trading process. Starting from near zero in the 1990s, by 2007 AT is responsible for roughly 1/3 of trading volume in the U.S and is expected to account for perhaps half of trading volume by 2010. The intense activity generated by algorithms threatens to overwhelm exchanges and market data providers, forcing significant upgrades to their infrastructures.

Before algorithmic trading, a pension fund manager wanting to buy 100,000 shares of IBM might have hired a broker-dealer to search for a counterparty to execute the entire quantity at once in a block trade. Alternatively, that institutional investor might have hired a broker to quietly "work" the order (possibly on the floor of the New York Stock Exchange), using his judgment and discretion to buy a little bit here and there over the course of the trading day to keep from driving the IBM share price up too far (see [14] for evidence of large orders being broken up). As trading became more electronic, it became easier and cheaper to replicate the human trader with a computer program doing algorithmic trading. Now virtually every large broker-dealer offers a suite of algorithms to its institutional customers to help them execute orders in a single stock, in pairs of stocks, or in baskets of stocks. Algorithms typically determine the timing, price, and quantity of orders, dynamically monitoring market conditions across different securities and trading venues, reducing market impact by optimally (and possibly randomly) breaking large orders into smaller pieces, and closely tracking benchmarks such as the volume-weighted average price (VWAP) over the execution interval.

Algorithms for Executing Large Orders

How to execute trades over a trading horizon involves complex dynamic optimization problems to determine order size, order frequency, and order type. Bertsimas and Lo [3] study optimal execution strategies in the presence of temporary price impacts. Conditional on the constraint of completing the entire transaction by a fixed date, orders are broken into pieces so as to minimize cost. Almgren and Chriss [1] extend this by considering the risk that arises from breaking up orders and slowly executing them. The basic Almgren and Chriss setup begins by assuming an initial holding of $x_0 = X$ units of a security that must by completely liquidated before time T. Dividing T into N intervals of length $\tau = T / N$ and defining the discrete times $t_k = k\ \tau$, for $k = 0, \ldots , N$, they solve for the optimal holdings, x_k, at each time t_k. Equivalently this can be formulated as the series to trades to accomplish this liquidation with $n_k = x_k - x_{k-1}$ being the number of units sold between times t_{k-1} and t_k. Capturing the risk involved in the trading process (see [8] for the relationship between execution risk and investment risk) requires specifying a price process for the security:

$$S_k = S_{k-1} + \sigma \tau^{1/2} \xi_k - \tau\, g\left(\frac{n_k}{\tau}\right)$$

for $k = 0; \ldots, N$. Here σ represents the volatility of the asset, the ξ_k are draws from independent random variables each with zero mean and unit variance, and the permanent price impact of the trade, *g(v)*, is a function of the average rate of trading $v = n_k / \tau$ during the interval t_{k-1} and t_k. In addition there may be a temporary price impact function, *h(v)*, which captures the temporary drop in average price per share caused by trading at average rate *v* during a time interval. Thus, the actual price per share received on sale *k* is

$$\tilde{S}_k = S_{k-1} - h\left(\frac{n_k}{\tau}\right),$$

where *h(v)* does not impact the subsequent market price S_k. *g(v)* and *h(v)* can be chosen to reflect the trader's beliefs about market dynamics and market microstructure.

Combining these and summing across periods yields the proceeds of the liquidation:

$$\sum_{k=1}^{N} n_k \tilde{S}_k = XS_0 + \sum_{k=1}^{N}\left(\sigma \tau^{1/2} \xi_k - \tau\, g\left(\frac{n_k}{\tau}\right)\right)x_k - \sum_{k=1}^{N} n_k h\left(\frac{n_k}{\tau}\right),$$

with the first term being the initial market value, the middle term including both the volatility and the permanent price impacts, and the final term being the temporary price impacts. Calculating the expected cost of trading, *C(x)*, as the difference between the initial value and the proceeds of the sales gives:

$$C(x) = \sum_{k=1}^{N} x_k g\left(\frac{n_k}{\tau}\right) + n_k h\left(\frac{n_k}{\tau}\right).$$

The variance of *C(x)* is:

$$V(C(x)) = \sigma^2 \tau \sum_{k=1}^{N} x_k^2 .$$

Using this setup Almgren and Chriss calculate the mean-variance efficient frontier for trading/liquidation strategies. Furthermore, assuming the objective is to minimize the function $C(x) - \lambda\, V(C(x))$, they calculate the optimal strategies for various values of λ. As long as the parameters are fixed and known in advance, Almgren and Chriss show that the optimal trading strategy can be calculated in advance and does not need to be updated dynamically. The calculation of such strategies may still be complex depending on the functional form of the permanent and transitory price impact functions *g(v)* and *h(v)*. Some typical impact functions include: power functions, weighted sums of linear and non-linear functions, and exponential decay functions (see eqf18/002 for detail about a specific functional form for the permanent and temporary impact functions). The significant computer power often required to find the optimum under complex price impact functions makes implementation via algorithms attractive.

Numerous features have been added to the above model that significantly increase the complexity of the calculations and algorithms needed to find and implement the optimal trading strategies. Some of the most common enhancements include portfolio considerations and modeling the information that motivates the trade. Obizhaeva and Wang [18] optimize assuming that liquidity does not replenish immediately after it is taken but only gradually over time. Algorithms can be improved by modeling variation within and cross days to account for changing market conditions and predictable price movements [2]. Kissell and Malamut [15] discuss how trading should be sped up or slowed down conditional on beliefs about prices changes mean reverting or continuing. Many brokers build models with these considerations into their AT products that they sell to their clients.

The above optimal execution approaches either abstract away from the market mechanism or assume liquidity is taken via market orders. Implementation strategies that allow for more passive trading strategies enable the algorithm/trader to choose the type and aggressiveness of each order/component of the larger transaction. Cohen et al. [6] and Harris [11] focus on the simplest choice: market order versus limit order. If a trader chooses a non-marketable limit order, the aggressiveness of the order is determined by its limit price ([10] and [19]). How aggressively a limit order should be priced depends on how price affects the time to execution [17]. For assets that trade on multiple venues, determining where to send orders adds further complexity.

Market Making Algorithms

Many observers think of algorithms from the standpoint of institutional buy-side investor. But there are other important users of algorithms. Some hedge funds and broker-dealers supply liquidity using algorithms, competing with designated market-makers and other liquidity suppliers.

Most models take the traditional view that one set of traders provides liquidity via quotes or limit orders and another set of traders initiates a trade to take that liquidity – for either informational or liquidity/hedging motivations. Liquidity supply involves posting firm commitments to trade. These standing orders provide free trading options to other traders. Using standard option pricing techniques Copeland and Galai [7] value the cost of the option granted by liquidity suppliers. The arrival of public information renders existing orders stale and can put the free trading option into the money. Biais et al. [4] provide a general theoretical framework for a liquidity supply price schedule $q_t(p)$ at time t:

$$\underset{q_t(p)}{\text{Max}}\, EU(C_t + I_t v_t + (v_t - p)q_t(p) \mid H_t),$$

where $U(x)$ is the utility function, C_t is cash, I_t is the inventory of the security, v_t is the expected fundamental value, and H_t is the information set available. The use of algorithms lowers the cost of calculating and managing this price schedule (quotes) in the presence of constant changes in the information set. Any time an announcement, trade, or quote change occurs in any security, a liquidity provider may want to revise their prices and quantities. Foucault, Roell, and Sandas [9] study the equilibrium level of effort that liquidity suppliers should expend in monitoring the market to avoid this picking off risk. Improvements in technology allow algorithms to inexpensively perform such monitoring. In this way AT may be a way to implement Black's [5] idea of limit orders indexed to a market index.

The widespread adoption of algorithms by traders supplying and demand liquidity has led some observers to liken the situation to an arms' race [16]. While evidence points to algorithmic trading improving liquidity [12], the growing dominance of algorithms in trading makes their impact worthy of continued study.

References

[1] Almgren, R, Chriss, N. (2000) Optimal execution of portfolio transactions. *Journal of Risk* **3**, 5-39.

[2] Almgren, R, Lorenz, J. (2006) Bayesian adaptive trading with a daily cycle. *Journal of Trading* **1**, 38-46.

[3] Bertsimas, D, Lo, AW. (1998) Optimal control of execution costs. *Journal of Financial Markets* **1**, 1-50.

[4] Biais, B, Glosten, L, Spatt, C. (2005) Market microstructure: a survey of microfoundations, empirical results and policy implications. *Journal of Financial Markets* **8**, 217-264.

[5] Black, F. (1995) Equilibrium exchanges. *Financial Analysts Journal* **51**, 23-29.

[6] Cohen, K, Maier, S, Schwartz, R, Whitcomb, D. (1981) Transaction costs, order placement strategy and existence of the bid-ask apread. *Journal of Political Economy* **89**, 287-305.

[7] Copeland, TE, Galai, D. (1983) Information effects on the bid-ask spread. *Journal of Finance* **38**, 1457-1469.

[8] Engle, RF, Ferstenberg, R. (2007) Execution risk. *Journal of Portfolio Management* **33**, 34-45

[9] Foucault, T, Roell, A, Sandas, P. (2003) Market making with costly monitoring: An analysis of the SOES controversy. *Review of Financial Studies* **16**, 345-384.

[10] Griffiths, M, Smith, B, Turnbull, DA, White, R. (2000) The costs and determinants of order aggressiveness. *Journal of Financial Economics* **56**, 65-88.

[11] Harris, L. (1998) Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions, and Instruments* **7**, 1-76.

[12] Hendershott, T, Jones, C, Menkveld A. (2008) Does algorithmic trading improve liquidity? Technical Report, UC Berkeley.

[13] Jain, PK. (2005) Financial market design and the equity premium: Electronic versus floor trading. *Journal of Finance* **60**, 2955-2985.

[14] Keim, D, Madhavan, A. (1995). Anatomy of the trading process: Empirical evidence on the behavior of institutional traders. *Journal of Financial Economics* **37**, 371-398.

[15] Kissell, R, Malamut, R. (2006) Algorithmic decision making framework. Journal of Trading 1, 12-21.

[16] Leinweber, D. (2007) Algo vs. Algo. *Institutional Investor's Alpha*, 45-51.

[17] Lo, WA, MacKinlay, AC, Zhang, J. (2002) Econometric models of limit-order executions. *Journal of Financial Economics* **65**, 31-71.

[18] Obizhaeva, A, Wang, J. (2005) Optimal trading strategy and supply/demand dynamics. MIT working paper.

[19] Ranaldo, A. (2004) Order aggressiveness in limit order book markets. *Journal of Financial Markets* **7**, 53-74.