

The execution puzzle: How and when to trade to minimize cost

Jim Gatheral



The City University of New York

5th International Financial and Capital Markets Conference
Campos de Jordão, Brazil, August 27th, 2011

Overview of the execution puzzle

Typically, an execution algorithm has three layers:

- The macrotrader
 - This highest level layer decides how to slice the meta-order: when the algorithm should trade, in what size and for roughly how long.
- The microtrader
 - Given a slice of the meta-order to trade (a child order), this layer of the algorithm decides whether to place market or limit orders and at what price level(s).
- The smart order router
 - Given a limit or market order, to which venue should the order be sent?

Impact of proprietary metaorders (from Tóth et al.)

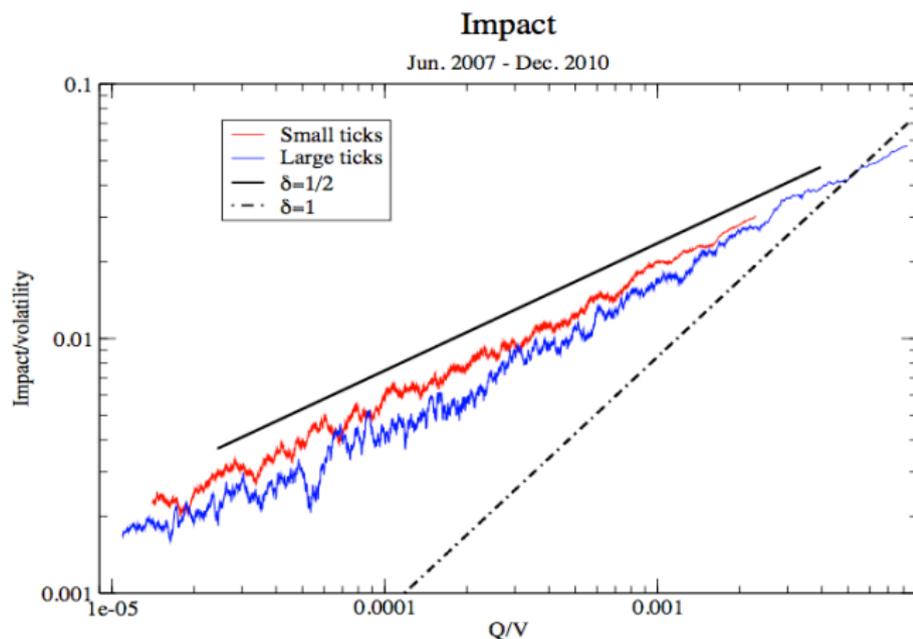


Figure 1: Log-log plot of the volatility-adjusted price impact vs the ratio Q/V

Notes on Figure 1

- In Figure 1 which is taken from [Tóth et al.], we see the impact of metaorders for CFM¹ proprietary trades on futures markets, in the period June 2007 to December 2010.
 - Impact is measured as the average execution shortfall of a meta-order of size Q .
 - The sample studied contained nearly 500,000 trades.
- We see that the square-root market impact formula is verified empirically for meta-orders with a range of sizes spanning two to three orders of magnitude!

¹Capital Fund Management (CFM) is a large Paris-based hedge fund.

Empirically observed stock price path

From Figure 2, we see that

- There is reversion of the stock price after completion of the meta-order.
- Some component of the market impact of the meta-order appears to be permanent.
- The path of the price prior to completion looks like a power law.
 - From [Moro et al.]

$$m_t - m_0 \approx (4.28 \pm 0.21) \left(\frac{t}{T}\right)^{0.71 \pm 0.03} \quad (\text{BME})$$

$$m_t - m_0 \approx (2.13 \pm 0.05) \left(\frac{t}{T}\right)^{0.62 \pm 0.02} \quad (\text{LSE})$$

where T is the duration of the meta-order.

Example of OW optimal strategy

- Consider a Brazilian stock with 14,000 trades per day and a liquidation whose horizon is 1 hour.
 - A rule of thumb is that the order book refreshes after 10-15 trades. So we take the half-life of the order book resilience process to be $20 \times \log 2 \approx 14$ trades.
 - $\log 2 / \rho = \log 2 \times 20$ trades so $\rho = 1/20$ in trade time. One hour has 2,000 trades so $\rho T = 100$.

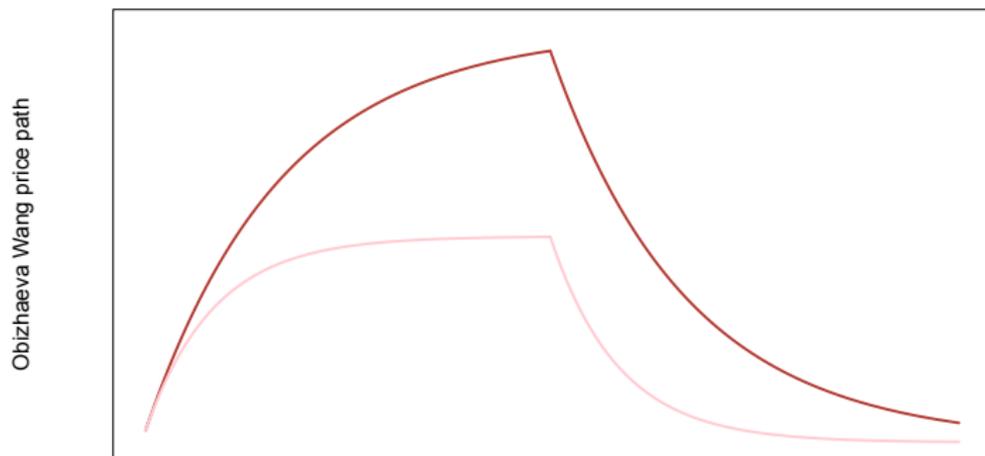
- Recall that the optimal strategy is $u_s = \delta(s) + \rho + \delta(s - T)$. Thus,

$$X = \int_0^T v_s ds = 2 + \rho T = 102$$

- The optimal strategy thus consists of a block trade of relative size one at the beginning, another trade of size one at the end and an interval VWAP of relative size 100.
 - The optimal strategy in this case is very close to VWAP.

Price path in the Obizhaeva-Wang model

Figure 4: The OW average price path is plotted for two different values of ρ .



- The optimal strategy is bucket-like.

The model of Alfonsi, Fruth and Schied

[Alfonsi, Fruth and Schied] consider the following (AS) generalization of the OW model:

- There is a continuous (in general nonlinear) density of orders $f(x)$ above some martingale ask price A_t . The cumulative density of orders up to price level x is given by

$$F(x) := \int_0^x f(y) dy$$

- Executions eat into the order book.
- A purchase of ξ shares at time t causes the ask price to increase from $A_t + D_t$ to $A_t + D_{t+}$ with

$$\xi = \int_{D_t}^{D_{t+}} f(x) dx = F(D_{t+}) - F(D_t)$$

- The order book has exponential resiliency; either the volume impact process E_t or the spread D_t revert exponentially.

Optimal liquidation strategy in the AS model

The optimal strategy in the AS model is

$$v_t = \xi_0 \delta(t) + \xi_0 \rho + \xi_T \delta(T - t).$$

- Just as in the OW model, the optimal strategy consists of a block trade at time $t = 0$, continuous trading at the rate ρ over the interval $(0, T)$ and another block trade at time $t = T$.
- The only difference is that in the AS model, the final block is not the same size as the initial block.

Generalization

- It can be shown that the bucket-shaped strategy is optimal under more general conditions than exponential resiliency.
 - Specifically, if resiliency is a function of the volume impact process E_t (or equivalently the spread D_t) only, the optimal strategy has block trades at inception and completion and continuous trading at a constant rate in-between.
- These conditions may appear quite general but in fact, there are many other models that do not satisfy them.

A transient market impact model

The price process assumed in [Gatheral] is

$$S_t = S_0 + \int_0^t f(v_s) G(t-s) ds + \text{noise} \quad (2)$$

- The instantaneous impact of a trade at time s is given by $f(v_s)$ – some function of the rate of trading.
- A proportion $G(t-s)$ of this initial impact is still felt at time $t > s$.

The square-root model

Consider the following special case of (2) with $f(v) = \frac{3}{4}\sigma \sqrt{v/V}$ and $G(\tau) = 1/\sqrt{\tau}$:

$$S_t = S_0 + \frac{3}{4} \sigma \int_0^t \sqrt{\frac{v_s}{V}} \frac{ds}{\sqrt{t-s}} + \text{noise} \quad (3)$$

which we will call the *square-root process*.

It turns out that the square-root process is consistent with the square-root formula for market impact:

$$\frac{c}{X} = \sigma \sqrt{\frac{X}{V}} \quad (4)$$

- Of course, that doesn't mean that the square-root process is the true underlying process!

The optimal strategy under the square-root process

- Because $f(\cdot)$ is concave, an optimal strategy does not exist in this case.
 - It is possible to drive the expected cost of trading to zero by increasing the number of slices and decreasing the duration of each slice.
 - To be more realistic, $f(v)$ must be convex for large v and in this case, an optimal strategy does exist that involves trading in bursts, usually more than two.

Intuition

- The optimal strategy depends on modeling assumptions.
- In the Almgren-Chriss model where there is no price reversion after completion of the meta-order, the optimal strategy is a simple VWAP.
- In other models where there is reversion, the optimal strategy is to make big trades separated in time, perhaps with some small component of continuous trading.

The intuition is easy to see:

The price reversion idea

If the price is expected to revert after completion, stop trading early and start again later after the price has reverted!

The market or limit order decision

- Having decided how to slice the meta-order, should we send market or limit orders?
- Many market participants believe that market orders should only be sent when absolutely necessary – for example when time has run out.
- Conventional wisdom has it that the more aggressive an algorithm is, the more costly it should be.
 - This cannot be true on average. Traders are continuously monitoring whether to send market or limit orders so in equilibrium, market and limit orders must have the same expected cost.
 - Market orders incur an immediate cost of the half-spread but limit orders suffer from adverse selection.

Adverse selection

- Limit orders are subject to adverse selection:
 - If the price is moving towards us, we get filled. We would rather that our order had not been filled. Had we not got the fill, we could have got a better price.
 - If the price is moving away from us, we don't get filled. We need to resubmit at a worse price.
- In general, we regret sending a market order because we have to pay the half-spread.
- In general, we regret sending a limit order because of adverse selection.

The order book signal

- If we know the price is going against us, we should send a market order. Otherwise we should send a limit order.
- In practice, we cannot predict the future; we can compute relative probabilities of future events.
 - One simple idea is to look at the shape of the order book. If there are more bids than offers, the price is more likely to increase than decrease.

The SFGK zero-intelligence model

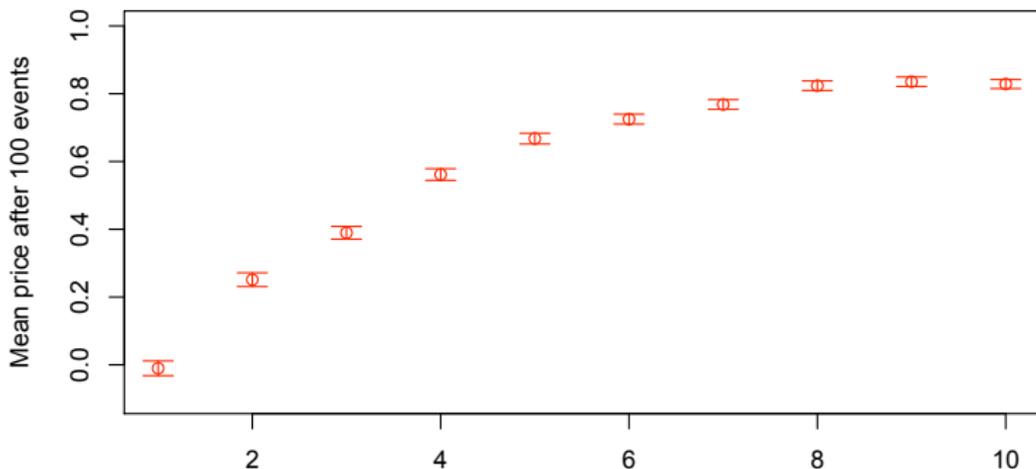
In this model due to Smith, Farmer, Gillemot and Krishnamurthy:

- Limit orders can be placed at any integer price level p where $-\infty < p < \infty$.
 - If worried about negative prices, think of these as being logarithms of the actual price.
- Limit sell orders may be placed at any level greater than the best bid $B(t)$ at time t and limit buy orders at any level less than the best offer $A(t)$.
- Market orders arrive randomly at rate μ .
- Limit orders (per price level) arrive at rate α .
- A proportion δ of existing limit orders is canceled.

Price signal in the ZI simulation

- Even in the ZI model, the shape of the order book allows prediction of price movements.
 - Traders really would need to have zero intelligence not to condition on book shape!

Figure 6: With one share at best offer, future price change vs size at best bid.



Microprice

- The relationship between the imbalance in the order book and future price movements is sometimes described in terms of the *microprice*.
 - This can be thought of as a fair price, usually between the bid and the ask.
- As an example, in the context of the zero-intelligence model, Cont and Larrard derived the following asymptotic expression:

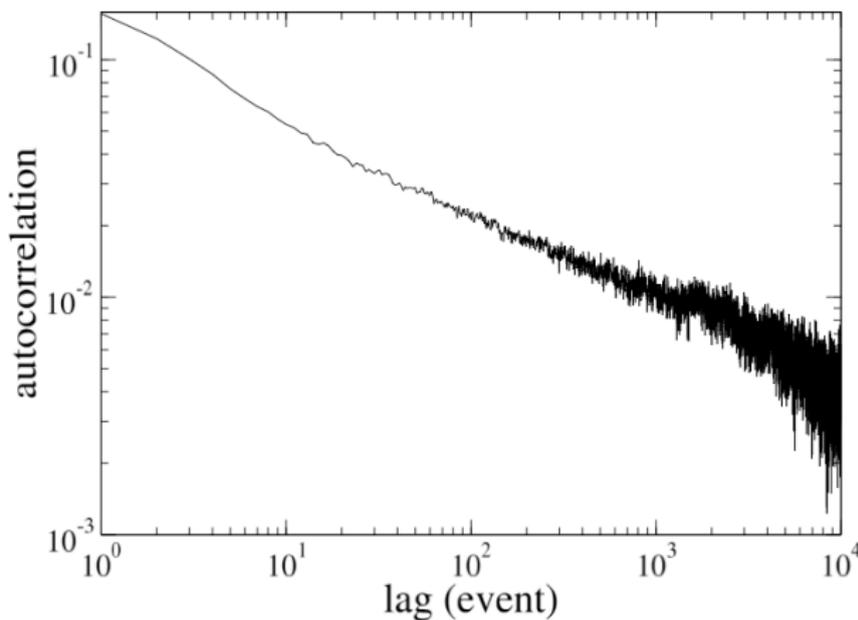
Proposition 2 of Cont & Larrard

The probability $\phi(n, p)$ that the next price is an increase, conditioned on having n orders at the bid and p orders at the ask is:

$$\phi(n, p) = \frac{1}{\pi} \int_0^\pi dt \left(2 - \cos t - \sqrt{(2 - \cos t)^2 - 1} \right)^p \frac{\sin n t \cos \frac{t}{2}}{\sin \frac{t}{2}}$$

Autocorrelation of order signs

Figure 7: Autocorrelation function of the time series of signs of Vodafone market orders in the period May 2000 – December 2002, a total of 580,000 events (from [Bouchaud, Farmer, Lillo])



The market/ limit order decision

- We have two (related) signals:
 - From the current state of the order book, we can predict the sign of the next price change.
 - From recent order flow history, we can judge whether active meta-orders are more on the buy side or on the sell side.
- The practical recipe is therefore to:
 - Send market orders when the market is going against us, limit orders otherwise.
 - Trade more when others want to trade with us, less when there are fewer counterparties.

Order routing

- Having optimally scheduled child orders and for each such child order, having decided whether to send a market or a limit order, where should the order be sent?
- In Brazil, the answer is straightforward; there is currently only one market – the BOVESPA.
- In the US, there are currently approximately 13 lit venues and over 40 dark venues.
 - On the one hand, it would be prohibitively complicated and expensive to route to all of them.
 - On the other hand, by not routing to a particular venue, the trader misses out on potential liquidity, and all things being equal, will cause incur greater market impact.
- Most traders have to use a smart order routing (SOR) algorithm provided by a dealer.

The Almgren and Harts (AH) algorithm

- The idea of this algorithm is that the more hidden quantity is detected in a given venue, the more hidden quantity there is likely to be.
 - This is a characteristic of distributions with fatter tails than exponential.
 - Empirically, we find that order sizes are power-law distributed in which case this assumption would definitely be justified.
- For simplicity, let's focus on the sale of stock.
- If hidden quantity w is detected (by selling more than the visible quantity) on a particular venue, the current estimate of hidden liquidity is increased by w .
- If no hidden quantity quantity is detected on a venue, the existing estimate is decremented by a factor ρ .

Conditional distribution of quantity: Power-law case

Suppose that the distribution of order sizes Q is power-law so that

$$\Pr(Q > n) = \frac{C}{n^\alpha}$$

Assuming the conditional probability that hidden quantity is greater than n given that n slices have already been observed is given by

$$\begin{aligned}\Pr(Q \geq (n+1) | Q \geq n) &= \frac{\Pr(Q \geq (n+1))}{\Pr(Q \geq n)} \\ &= \left(\frac{n}{n+1}\right)^\alpha \\ &\rightarrow 1 \text{ as } n \rightarrow \infty\end{aligned}$$

If the distribution of Q is power-law, the more quantity you observe, the more likely it is that there is more quantity remaining.

Simulation results

- An algorithm can only be tested by experiment or simulation.
 - The data used for model estimation comes from particular choices of algorithm and we can't predict what would have been if these algorithms had chosen to act differently.
- In simulations, the GKNW algorithm outperformed two other obvious choices of algorithm:
 - Equal allocation across venues
 - A *bandit* algorithm that begins with equal weights. If there is any execution at a particular venue, that venue's weight is increased by a factor $\alpha = 1.05$.

A philosophical question

Question

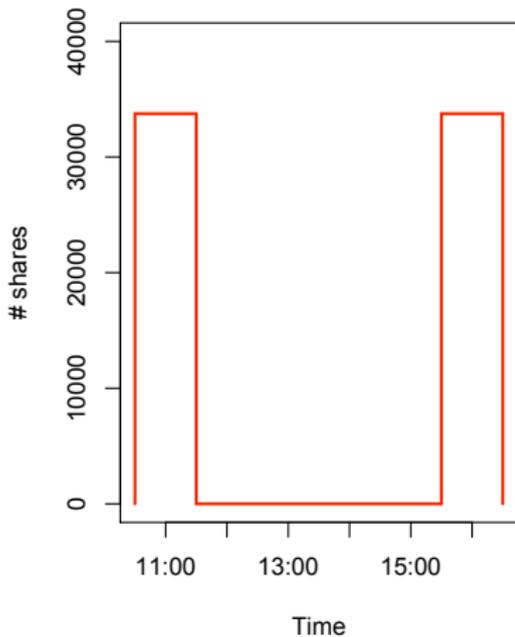
Why should it be so complicated to trade stock?

One answer goes something like this:

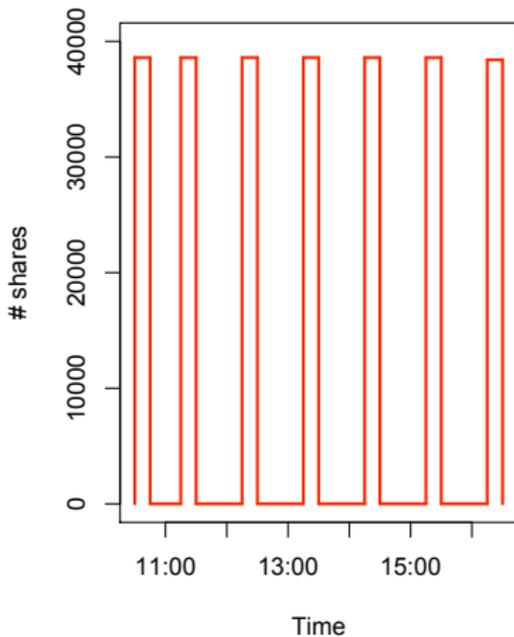
- Changes in market structure together with technological innovation have massively reduced trading costs.
- Nevertheless, some market participants achieve significantly lower costs.
 - This requires either substantial investment in technology and trading expertise or
 - careful selection of broker algorithms.
- Note however that algorithm performance is very hard to assess ex-post. Ideally, randomized experiments are required.

Stock trading schedules

Bucket-like schedule



Quasi-optimal schedule



Potential cost savings from smart order routing

- A naïve estimate would be to use the square-root market impact formula, changing the denominator V to reflect the potential increase in liquidity from routing to extra venues.
 - If the potential liquidity accessed is doubled, costs should be decreased by $1 - 1/\sqrt{2} \approx 30\%$ according to this simple computation!
 - This could be one explanation for the multiplicity of trading venues in the US.
- We expect actual savings to be much less than this because the different venues are all connected and information leaks from one venue to the other.

