OPTIMAL EXECUTION OF PORTFOLIO TRANSACTIONS

JULIAN LORENZ

Portfolio theory delivers insight into optimal asset allocation and optimal portfolio construction in a frictionless market. However, in reality the acquisition (or liquidation) of a portfolio position does not come for free. Transaction costs refer to any costs incurred while implementing an investment decision. Some transaction costs are observable directly: brokerage commissions, fees (e.g. clearing and settlement costs) and taxes. These cost components constitute the overwhelming part of total transaction costs for a retail investor.

However, for an institutional investor, transaction costs are more subtle and stem primarily from *market impact* caused by his trades. Indeed, empirical research shows that for institutional sized transactions, total implementation cost is around 1% for a typical trade and can be as high as 2-3% for very large orders in illiquid stocks [4]. Typically, institutional investors have a portfolio turnover a couple of times a year. Poor execution therefore can erode portfolio performance substantially.

An obvious way to avoid market impact would be to trade more slowly to allow market liquidity recover between trades. Instead of placing a huge 10,000-share order in one big chunk, a trader may break it down in smaller orders and incrementally feed them into the market over time. However, slower trading increases the trader's susceptibility to market volatility and prices may potentially move disadvantageous to the trader. [4] coin this the traders' dilemma: *"Do trade and push the market. Don't trade and the market pushes you."*. Optimal trade schedules seek to balance the market impact cost of rapid execution against the volatility risk of slow execution.

Market Model. The classic market impact model in algorithmic trading is due to Almgren and Chriss [1]. There the execution benchmark is the pre-trade price. The difference between the pre-trade and the post-trade book value of the portfolio (including cash positions) is the implementation shortfall. In the following, we shall restrict the notation to sell programs. Thus, the implementation shortfall of a sell program is the initial value of the position minus the actual amount captured.

In the simplest model the expected value of the implementation shortfall is entirely due to market impact as a deterministic function of the trading rate, integrated over the entire trade. The variance of the implementation shortfall is entirely due to the volatility of the price process S_t. More precisely, in a discrete time setup S_t is modeled as a random walk

$$S_t = S_{t-1} + \xi_t \tag{1}$$

jlorenz@inf.ethz.ch.

for $t = 1, \ldots, T$, where ξ_t are iid random variables. We use an arithmetic model since our interest is in short-term trading (typically less than one day), where the difference between arithmetic and geometric Brownian motion is negligible and the arithmetic process is much more convenient (we neglect the disadvantageous effect that there is a small, yet positive probability of the stock price becoming negative).

On top of this exogenous price process, the activity of our trade program causes *temporary market impact*: our actual trade prices are slightly less favorable than S_t. The effective price per share when we sell v shares at time t is

$$\tilde{S}_t = S_t - h(v) \tag{2}$$

for some impact function $h(v) \geq 0$. The trader has an order of X shares, which must be completed by time $t = T$. A trading strategy is given by $x_1, \ldots, x_{T-1}$, where x_t is the number of shares left to sell at time k (i.e. the trader sells $x_{t-1} - x_t$ at $\tilde{S}_{t-1}$, with $x_0 = X$ and $x_T = 0$). In general x_t may be any non-anticipating random functional of the stock price process S_t.

Static and Adaptive Efficient Strategies. In the mean-variance framework used by [1], "efficient" strategies minimize the variance of the implementation shortfall for a specified maximum level of expected cost or conversely. The set of such strategies is summarized in the "efficient frontier of optimal trading", akin to the well-known Markowitz efficient frontier in portfolio theory. If x_t is fixed independently of the process S_t, then straightforward optimization yields the family of mean-variance efficient static trajectories as $x_t = \sinh\bigl(\kappa(T-t)\bigr)/\sinh\bigl(\kappa T\bigr)$, parametrized by the urgency parameter $\kappa \geq 0$. These execution strategies were determined by Almgren and Chriss [1], and are path-independent (also called static): they do not modify the execution speed in response to price movement during trading.

However, Almgren and Lorenz [3, 5] show that substantial improvement with respect to the mean-variance tradeoff measured at the initial time is possible if one allows path-dependent policies that adapt in response to changes in the price of the asset being traded. [3] construct simple dynamic trading strategies which update exactly once: at some intermediary time they may readjust in response to the stock price movement. Before and after that "intervention" time they may not respond to whether the price goes up or down. Already these simple adaptive strategies improve over the path-independent strategies of [1] with respect to the mean-variance tradeoff measured at the initial time. [5] uses a dynamic programming principle for mean-variance optimization in discrete time to determine optimal Markovian strategies. This technique reduces the determination of optimal dynamic strategies to a series of single-period convex constrained optimization problems.

The relative improvement of a dynamic strategy over a static trade-schedule is larger for large portfolio transactions, expressed in terms of a new preference-free nondimensional parameter μ that measures portfolio size in terms of its ability to move the market relative to market volatility. For small portfolios, $\mu \to 0$, optimal adaptive trade schedules coincide with the optimal static trade schedules of Almgren and Chriss [1].

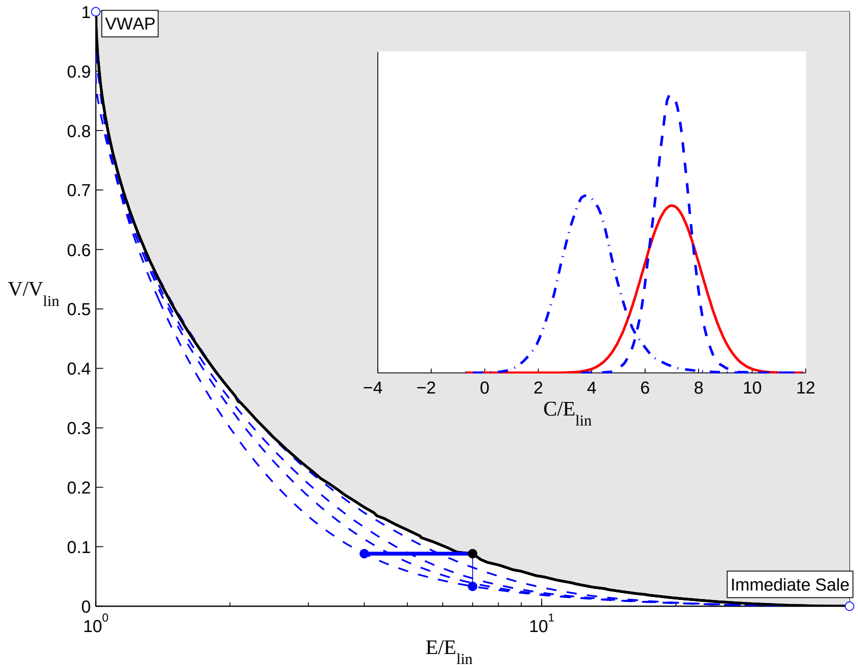

FIGURE 1. Adaptive efficient frontiers for different values of the market power μ and $T = 50$. The grey shaded region is the set of values accessible to static trading trajectories and the black line is the static efficient frontier, which is also the limit $\mu \to 0$. The blue dashed curves are the improved frontiers, with the improvement increasing with μ. The inset shows the cost distributions corresponding to the three points marked on the frontiers.

The optimal adaptive trade schedules are "aggressive-in-the-money": if the price moves in our favor in the early part of the trading, then we reduce risk by accelerating the remainder of the program, spending parts of the windfall gains on higher market impact costs. If the price moves against us, then we reduce future costs by trading more slowly, despite the increased exposure to risk of future fluctuations.

Further Work. The model described in this note allows many extensions. Trading of multiple securities in one joint execution program may be considered ("program trading"), where the different asset price processes and price impact functions are correlated. Other possible extensions are different types of price impact, or different asset price processes e.g. price momentum. Almgren and Lorenz [2] discuss one way to model price momentum resulting from the daily trading cycle: investment decisions are often made overnight or in the early morning and then left to be executed during the day, causing intraday momentum, and we can use price observations during the day to estimate the momentum and to determine optimal dynamic trading strategies.

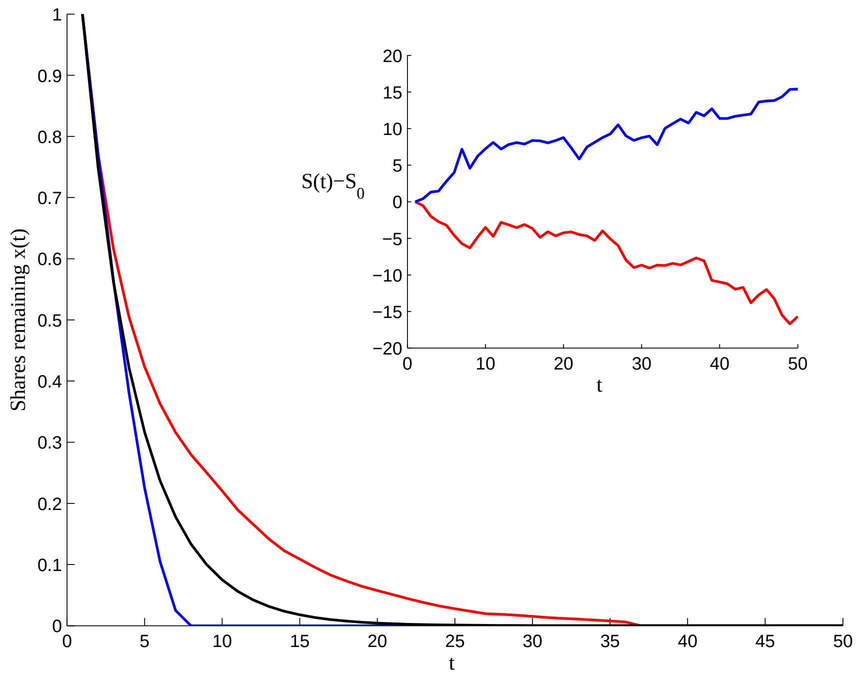

Figure 2. Sample trading trajectories of the adaptive strategy (sell program) for the point on the frontier in Figure 1, having the same variance but lower expected cost than the static trajectory (solid black line). The blue trajectory belongs to a rising stock price path and accelerates trading. The red trajectory corresponds to a falling stock price and trades more slowly.

References

[1] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *J. Risk*, 3(2):5–39, 2000.

[2] Robert Almgren and Julian Lorenz. Bayesian adaptive trading with a daily cycle. *J. Trading*, 1(4):38–46, 2006.

[3] Robert Almgren and Julian Lorenz. Adaptive arrival price. *Algorithmic Trading III*, pages 59–66, 2007.

[4] Robert Kissell and Morton Glantz. *Optimal Trading Strategies*. Amacom, 2003.

[5] Julian Lorenz. *Optimal Trading Algorithms*. PhD thesis, ETH Zurich, 2008.

E-mail address: `jlorenz@inf.ethz.ch`