

Mean-Variance Optimal Adaptive Execution

Julian Lorenz* and Robert Almgren**

January 23, 2011

Abstract

Electronic trading of equities and other securities makes heavy use of “arrival price” algorithms, that balance the market impact cost of rapid execution against the volatility risk of slow execution. In the standard formulation, mean-variance optimal trading strategies are static: they do not modify the execution speed in response to price motions observed during trading. We show that substantial improvement is possible by using dynamic trading strategies, and that the improvement is larger for large initial positions.

We develop a technique for computing optimal dynamic strategies to any desired degree of precision. The asset price process is observed on a discrete tree with a arbitrary number of levels. We introduce a novel dynamic programming technique in which the control variables are not only the shares traded at each time step, but also the maximum expected cost for the remainder of the program; the value function is the variance of the remaining program. The resulting adaptive strategies are “aggressive-in-the-money”: they accelerate the execution when the price moves in the trader’s favor, spending parts of the trading gains to reduce risk.

*Work performed when the author was at the Institute of Theoretical Computer Science, ETH Zurich/Switzerland, and was supported by UBS AG. post@julian-lorenz.de

**New York University Courant Institute and Quantitative Brokers, almgren@cims.nyu.edu

1 Introduction

Algorithmic trading considers execution in the market of large transactions within a fixed time period, optimizing some trade-off between risk and reward. In arrival price algorithms, 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 [Perold, 1988]. For instance, the implementation shortfall of a sell program is the initial value of the position minus the dollar amount captured. The implementation shortfall is random, since it depends both on the execution strategy, and on market movements experienced before the trade is complete.

In the simplest model, the expected value of the implementation shortfall is entirely due to the market impact incurred by trading. (Anticipated price drift is potentially important, but in this paper we want to concentrate attention on the balance between risk and reward.) This expected cost is minimized by trading as slowly as possible. Conversely, with deterministic market impact, the variance of the implementation shortfall is entirely due to price volatility, and this variance is minimized by trading rapidly.

In a risk-reward framework, “efficient” strategies minimize risk for a specified maximum level of expected cost or conversely; the set of such strategies is summarized in the “efficient frontier of optimal trading” [Almgren and Chriss, 2000]. In contrast to more mathematically sophisticated utility function formulations, the simple mean-variance approach has the practical advantage that risk and reward are expressed as two real variables that are easily understood and displayed on a two-dimensional picture.

However, using mean and variance in a time-dependent problem presents some subtleties. We may distinguish three different optimization problems, depending on what form of adaptivity we permit.

1. We may require that the number of shares to be sold at each time be specified in advance. Huberman and Stanzl [2005] suggest that a reasonable example of this is insider trading, where trades must be announced in advance. This is the “static” problem of Almgren and Chriss [2000] and of Section 2.3 below, and also the “precommitment” consumption strategies of Strotz [1956].
2. We may require that a *rule* be fixed at the starting time, specifying how the trade quantities will be determined at each future time as a function of market properties observed at that time. The rule is chosen

to optimize a mean-variance criterion evaluated at the initial time, and may not be modified later. For example, Basak and Chabakauri [2010] generalized the static strategies of Strotz [1956] by allowing investment decisions to depend on an additional random variable that is observable but not investable. In our case, the only random market property is asset price, so the question is whether trade decisions are allowed to depend on the realized price process.

3. We may allow the optimal strategy to be recalculated at each time, using a mean-variance criterion determined at that time for the remainder of the program. Strotz [1956] calls this a “consistent planning” strategy. The freedom to modify the strategy may *reduce* its effectiveness, since it must be designed to be appealing to one’s own future self (the “intertemporal tussle” of Strotz [1956]).

For example, consider the well-known problem of Black-Scholes option hedging. Problem 1 can be considered, but its solutions are absurdly far from optimal; the optimal hedge depends essentially on the realized asset price. The solution to problem 3 is the same as the solution to problem 2, exactly because variance is completely eliminated.

For optimal trade execution, for mean-variance optimality with arithmetic Brownian motion, Almgren and Chriss [2000] and Huberman and Stanzl [2005] have shown that the solution to problem 3 is the same as 1. However, by considering a simple rule consisting of a single trajectory modification at one fixed intermediate time, Almgren and Lorenz [2007] showed that the solutions to problem 2 can have better mean-variance properties than the solutions to 3 or 1. This crucially depends on the condition that mean and variance are to be measured at the initial time. These strategies are *not* time-consistent: to obtain the benefit, you must *deny* your future self the opportunity to modify the strategy based on his own preferences.

Whereas Almgren and Lorenz [2007] demonstrated the possibility of improvement by a very simple update strategy which was far from optimal, in this paper we determine the full time-dependent optimal solutions.

Since the nature of the problems and of their solutions are different, it is important to ask which formulation corresponds better to real industrial practice. Indeed, problem 2 corresponds to how trading results are typically reported in practice. Clients of agency trading desks are provided with a post-trade report daily, weekly, or monthly depending on their trading activity. This report shows sample average and standard deviation of execution price

relative to the implementation shortfall benchmark, across all trades executed for that client during the reporting period. Therefore the broker-dealer's goal is to design algorithms that optimize sample mean and variance at the per-order level, so that the post-trade report will be as favorable as possible. As discussed in Almgren and Lorenz [2007], this *ex post* measurement of sample mean and variance across a population corresponds exactly to the probabilistic notion of mean and variance for a single execution.

A problem related to the optimal execution of portfolio transactions is that of mean-variance portfolio optimization in a multiperiod setting (Basak and Chabakauri 2010, Bielecki, Jin, Pliska, and Zhou 2005, Li and Ng 2000, Richardson 1989, Zhou and Li 2000), but the market impact terms significantly complicate the problem. He and Mamaysky [2005] and Forsyth [2011] address the optimal execution problem in a continuous-time setting with a Hamilton-Jacobi-Bellman approach, the former optimizing expected utility, and the latter the mean-variance criterion using a technique proposed by Li and Ng [2000] and by Zhou and Li [2000].

We follow a different approach to determine optimal trading strategies with respect to the specification of risk and reward at the initial time. We give an efficient scheme to obtain fully optimal Markovian trading strategies for the arrival price problem in a discrete time setting which can approximate the continuous time problem arbitrarily closely and yields numerical solutions by solving a series of convex constrained optimization problems.

The improvement through adaptivity is larger for large portfolios, expressed in terms of a preference-free nondimensional parameter that measures the price impact of the trade relative to market volatility. For small portfolios, optimal adaptive trade schedules coincide with optimal static schedules.

Almgren and Lorenz [2007] show that the improvement of a dynamic portfolio strategy over a static portfolio is due to *anti-correlation* between the portfolio return of each period and the expected portfolio return in the remaining periods. After a fortunate price move the investor will try to conserve his realized gains and put less capital at risk in the remainder. In effect, spend any trading gains on market impact to reduce risk. That is, the investor's risk aversion changes in response to past performance, making him *more risk-averse* after *positive performance*. A similar effect was reported by Cvitanic, Lazrak, and Wang [2008] for the Sharpe ratio as a performance measure. Our new optimal trade schedules are "aggressive-in-the-money" (AIM) in the sense of Kissell and Malamut [2005]; a "passive-in-the-money" (PIM) strategy would slow down following a fortunate price change. Also

Basak and Chabakauri [2010] find that the precommitted strategy for the mean-variance portfolio optimization problem decreases the risky investment in good states.

One important reason for using a AIM or PIM strategy would be the expectation of serial correlation in the price process: belief in momentum would lead to a PIM strategy [Almgren and Lorenz, 2006], and see Bertsimas and Lo [1998] for the effect of very short-term correlation. Our strategies arise in a pure random walk model with no serial correlation, using pure classic mean and variance.

Schied and Schöneborn [2009] have shown that AIM or PIM strategies can arise from the structure of a utility function formulation. For functions with increasing absolute risk aversion (IARA), optimal strategies are AIM; with decreasing absolute risk aversion (DARA) optimal strategies are PIM. Only for constant absolute risk aversion, such as an exponential utility function, are optimal strategies static. Our formulation is conceptually similar to a CARA utility function and describes a different mechanism for adaptivity.

The remainder of this paper is organized as follows: In Section 2 we present the market and trading model. In Section 2.2 we review the concept of mean variance efficient strategies and the efficient frontier of trading, including the optimal static trajectories of Almgren and Chriss [2000]. In Section 3, we show how to construct fully optimal adaptive policies by means of a dynamic programming principle for mean-variance optimization. In Section 4 we present an approximation scheme for the dynamic program derived in Section 3, for which we give numerical results in Section 5.

2 Trading Model

Let us start by reviewing the trading model of Almgren and Chriss [2000]. We confine ourselves to sell programs in a single security; the definitions and results for a buy program are completely analogous.

2.1 Market dynamics

We hold a block of X shares of a stock that we want to completely sell by time horizon T . We divide T into N intervals of length $\tau = T/N$, and define discrete times $t_k = k\tau$, $k = 0, \dots, N$. A trading strategy π is a list of stock holdings (x_0, x_1, \dots, x_N) where x_k is the number of shares we plan

to hold at time t_k ; we require $x_0 = X$ and $x_N = 0$. Thus we shall sell $x_0 - x_1$ shares between t_0 and t_1 , $x_1 - x_2$ shares between times t_1 and t_2 and so on. The average rate of trading during the time interval t_{k-1} to t_k is $v_k = (x_{k-1} - x_k)/\tau$.

For a “static” strategy, the trade list $\pi = \{x_k\}$, or equivalently $\pi = \{v_k\}$, is a list of constant numbers. For a “dynamic” strategy, π is a random process adapted to the filtration of the price process as described below. We optimize the strategy by determining the rule by which x_k is computed at time t_{k-1} .

The stock price follows an arithmetic random walk

$$S_k = S_{k-1} + \sigma\tau^{1/2}\xi_k - \tau g(v_k) \quad , \quad k = 1, \dots, N. \quad (1)$$

The ξ_k are i.i.d. random variables taking values in a set Ω and having $\mathbb{E}[\xi_k] = 0$ and $\text{Var}[\xi_k] = 1$. For example, we could take $\Omega = \mathbb{R}$ with each ξ_k standard normal; a binomial tree model would set $\Omega = \{\pm 1\}$.

Let \mathcal{F}_k be the σ -algebra generated by $\{\xi_1, \dots, \xi_k\}$, for $k = 0, \dots, N$. For $k = 1, \dots, N$, \mathcal{F}_{k-1} is the information available to the investor before he makes trade decision x_k , and hence x_k must be \mathcal{F}_{k-1} measurable (x_1 must be a fixed number). The random variable x_k is determined by a function $\Omega^{k-1} \rightarrow \mathbb{R}$. The specification of these functions constitutes the adaptive trading strategy π , but we continue to concentrate attention on x_k as a random variable rather than its determining function.

The coefficient σ is the absolute volatility of the stock, so $\sigma^2\tau$ is the variance of price change over a single time step, and the variance of price change over the entire trading period is σ^2T . The function $g(v)$ represents the rate at which permanent impact is incurred, as a function of the instantaneous average rate of trading $v_k = (x_{k-1} - x_k)/\tau$ during the interval t_{k-1} to t_k .

Temporary market impact is modeled by considering our realized trade price to be less favorable than the “public” price S_k . The effective price per share when selling $x_{k-1} - x_k$ during the interval t_{k-1} to t_k is

$$\tilde{S}_k = S_{k-1} - h(v_k) \quad . \quad (2)$$

Unlike permanent impact $g(v)$, the temporary impact effect $h(v)$ does not affect the next market price S_k .

The total cost of trading, or implementation shortfall, for selling X shares across N steps of length τ , using trading policy $\pi = (x_0, \dots, x_N)$, is the

difference between the initial market value and the final capture of the trade:

$$\begin{aligned} C(X, N, \pi) &= X S_0 - \sum_{k=1}^N (x_{k-1} - x_k) \tilde{S}_k \\ &= \sum_{k=1}^N \left[\tau x_k g(v_k) + \tau v_k h(v_k) - \sigma \tau^{1/2} \xi_k x_k \right]. \end{aligned} \quad (3)$$

$C(X, N, \pi)$ is an \mathcal{F}_N -measurable random variable. Because of our assumption of arithmetic Brownian motion, it is independent of the initial stock price S_0 . Of course C depends on all market parameters as well as the strategy, but we focus on X and N for use in the dynamic programming formulation. When we vary N , we shall assume that the time step τ is held constant, so the time horizon $T = N\tau$ varies.

In general, $g(v)$ may be any convex function, and $h(v)$ may be any function so that $v h(v)$ is convex. We focus on the *linear* case

$$g(v) = \gamma v, \quad h(v) = \eta_0 v.$$

Then, neglecting a constant term involving γ , (3) becomes

$$C(X, N, \pi) = \tau \eta \sum_{k=1}^N v_k^2 - \sigma \tau^{1/2} \sum_{k=1}^N \xi_k x_k \quad (4)$$

with $\eta = \eta_0 - \frac{1}{2}\gamma\tau$. We assume that τ is small enough that $\eta > 0$.

Geometric Brownian motion is a more traditional model than arithmetic. However, for the intraday time scales of interest to us, the difference is negligible, and the arithmetic process is much more convenient. In particular, the expected size of future price changes, as absolute dollar quantities, does not depend on past price changes or the starting price level.

We assume that volatility, as well as the dependence of permanent impact and temporary impact on our trade decisions, are not only non-random and known in advance but are constant. Predictable intraday seasonality can largely be handled by interpreting time t as a “volume time” corresponding to the market’s average rate of trading. Random variations in volatility and liquidity are more difficult to model properly (Almgren 2009, Walia 2006).

As $\tau \rightarrow 0$, that is, as $N \rightarrow \infty$ with T fixed, this discrete trading model converges to a continuous-time process: the exogenous price process is a

Brownian motion, the shares process is an adapted function $x(t)$, and the instantaneous trade rate is $v(t) = -dx/dt$. Hence, the scheme that we will give in the rest of this paper to determine optimal strategies immediately yields a continuous-time scheme as well. The same techniques would also work with nonlinear cost functions, or with a drift term added to the price dynamics, or for multi-asset baskets, though dynamic programming for a multi-asset problem would require a very large state space.

2.2 Efficient frontier of optimal execution

For any trading strategy, the final cost $C(X, N, \pi)$ is a random variable: not only do the price motions ξ_k directly affect our trading gains or losses, but for an adapted strategy the trade list π itself may be different on each realization. An “optimal” strategy will determine some balance between minimizing the expected cost and its variance.

Let

$$\mathcal{D}(X, N) = \left\{ (\pi, \bar{C}) \left| \begin{array}{l} \pi = (x_0, x_1, \dots, x_N) \text{ with } x_0 = X, x_N = 0 \\ x_0 \geq x_1 \geq \dots \geq x_N \\ C(X, N, \pi) \leq \bar{C} \text{ almost surely} \end{array} \right. \right\} \quad (5)$$

be the set of all adapted trading policies that sell X shares in N periods. An agency broker may never buy shares as part of a sell order, so $\{x_k\}$ must be nonincreasing in time. The trade schedule π must be adapted to the filtration \mathcal{F}_k of the stock price process S_k as described above.

The final cost \bar{C} is a \mathcal{F}_N -measurable random variable giving an upper bound for the actual trading cost associated with this strategy for each realization of the stock price. Of course, $(\pi, C(X, N, \pi)) \in \mathcal{D}(X, N)$ for any π , and if $(\pi, \bar{C}) \in \mathcal{D}(X, N)$ then $(\pi, \bar{C} + \delta\bar{C}) \in \mathcal{D}(X, N)$, where $\delta\bar{C}$ is any \mathcal{F}_N -measurable random variable whose values are almost surely nonnegative. We allow the cost bound to be specified separately because below we will interpret it as one of the trader’s control parameters: we allow the trader to deliberately add cost in certain scenarios if that improves his mean-variance trade-off (Section 3.3). This possibility arises from the non-monotonicity of the mean-variance criterion.

For given $E \in \mathbb{R}$, let

$$\mathcal{A}(X, N, E) = \left\{ (\pi, \bar{C}) \in \mathcal{D}(X, N) \mid \mathbb{E}[\bar{C}] \leq E \right\} \quad (6)$$

be the (possibly empty) subset of policies and associated cost bounds, whose expected cost is at most E . Unless otherwise specified, all expectations and variances are unconditional, thus evaluated at time $t = 0$.

In this paper, we concentrate on the mean-variance constrained optimization problem

$$V_{\min}(E) = \inf \left\{ \text{Var}[\bar{C}] \mid (\pi, \bar{C}) \in \mathcal{A}(X, N, E) \right\}. \quad (7)$$

With arithmetic Brownian motion (1), the values $V_{\min}(E)$ and the associated strategies are independent of the initial stock price S_0 . The graph of $V_{\min}(E)$ is the *efficient frontier of optimal trading strategies* [Almgren and Chriss, 2000].

It is not obvious that the infimum in (7) is actually attained by any trading strategy, for fully general problems and for any values of E other than trivial ones. However, in practical applications, the price set Ω is a simple set like $\{\pm 1\}$ or \mathbb{R} , and each x_k is specified by a function on Ω^{k-1} , with a Hilbert space structure inherited from the probability measure on Ω . Problem (7) is a minimization in the Hilbert space of the collection of these functions. For reasonable measures, the functionals $\mathbb{E}[\bar{C}]$ and $\text{Var}[\bar{C}]$ are coercive on this space. Furthermore, in Section 3.3 we show generally that the optimization problem at each step is convex. Then a powerful set of standard techniques [Kurdila and Zabaranin, 2005] are available to demonstrate the existence of minimizers at each step and hence an overall minimizer. We shall therefore proceed under the assumption that minimizing solutions do exist. We also do not assume that minimizing strategies are necessarily unique, and therefore we shall refer to “a” or “every” minimizer.

By varying $E \in \mathbb{R}$ we determine the family of all efficient strategies

$$\mathcal{E}(X, N) = \left\{ (\pi, \bar{C}) \in \mathcal{D}(X, N) \mid \nexists (\tilde{\pi}, \tilde{C}) \in \mathcal{D}(X, N) \text{ such that} \right. \\ \left. \mathbb{E}[\tilde{C}] \leq \mathbb{E}[\bar{C}] \text{ and } \text{Var}[\tilde{C}] < \text{Var}[\bar{C}] \right\}. \quad (8)$$

The domain $\mathcal{A}(X, N, E)$ is empty if the target expected cost E is too small. To determine the minimum possible expected cost, note that the last term in (4) has strictly zero expected value. Minimizing the remaining term $\mathbb{E}[\sum v_k^2]$ gives the nonrandom linear strategy π_{lin} with $x_{k-1} - x_k = X/N$ for $k = 1, \dots, N$. In volume time, this is the popular VWAP profile. The

expectation and variance of the cost of this strategy are

$$E_{\text{lin}}(X, N) = \frac{\eta X^2}{T} , \quad (9)$$

$$V_{\text{lin}}(X, N) = \frac{1}{3} \sigma^2 X^2 T \left(1 - \frac{1}{N}\right) \left(1 - \frac{1}{2N}\right) . \quad (10)$$

Conversely, a strategy having zero variance is the “instantaneous” one that sells the entire position in the first time period: π_{inst} with $x_1 = \dots = x_N = 0$. This yields $V_{\text{inst}}(X, N) = 0$ and

$$E_{\text{inst}}(X, N) = \frac{\eta X^2}{\tau} = N E_{\text{lin}}(X, N) . \quad (11)$$

Thus we may summarize the behavior of (7) as

- For $E < E_{\text{lin}}$, the set $\mathcal{A}(X, N, E)$ is empty and $V_{\text{min}} = +\infty$.
- For $E \geq E_{\text{lin}}$, the set $\mathcal{A}(X, N, E)$ contains at least the linear strategy, with a finite variance. Hence V_{min} exists since it is the infimum of a nonempty set that is bounded below by zero. For $E \geq E_{\text{inst}}$, the minimum possible value $V_{\text{min}} = 0$ is attainable by the instantaneous strategy. Our computational procedure below produces a sequence of strategies whose variance approaches this infimum.

2.3 Static trajectories

An important special set of strategies is determined by restricting the minimum to the subset of $\mathcal{A}(X, N, E)$ consisting of *nonrandom* or “static” strategies. Then (7) becomes a simple numerical minimization

$$\min_{x_1 \leq \dots \leq x_{k-1}} \left\{ \sigma^2 \tau \sum_{k=1}^N x_k^2 \left| \frac{\eta}{\tau} \sum_{k=1}^N (x_{k-1} - x_k)^2 \leq E \right. \right\} \quad (12)$$

with $x_0 = X$ and $x_N = 0$. The strategies are most easily computed by using a Lagrange multiplier or risk aversion, that is, minimizing $\mathbb{E}[\overline{C}] + \lambda \text{Var}[\overline{C}]$. The result is

$$x_j = X \frac{\sinh(\kappa(T - t_j))}{\sinh(\kappa T)}, \quad j = 0, \dots, N , \quad (13)$$

in which the *urgency parameter* is $\kappa \sim \sqrt{\lambda\sigma^2/\eta} + \mathcal{O}(\tau)$. The expected cost and variance can be computed explicitly, and reduce to $(E_{\text{lin}}, V_{\text{lin}})$ and $(E_{\text{inst}}, V_{\text{inst}})$ in the limits $\lambda \rightarrow 0$ and $\lambda \rightarrow \infty$ respectively.

The static execution strategy is independent of the portfolio size X except for an overall factor, and the expected value and variance of total cost are quadratic in portfolio size. For static strategies, these properties do not hold for nonlinear cost models [Almgren, 2003]. Also, for dynamic strategies, these properties no longer hold.

Almgren and Chriss [2000] as well as Huberman and Stanzl [2005] have shown that dynamic adaptive solutions are the same as the static solutions, if mean and variance are reevaluated at each intermediate time when the new solution is computed. We emphasize that that formulation is a different problem than the one we consider, which optimizes mean and variance measured at the initial time.

3 Optimal Adaptive Strategies

As seen in the previous section, we can construct optimal static trading strategies by solving a straightforward optimization problem. But Almgren and Lorenz [2007] have demonstrated that adaptive strategies can improve over static trajectories, even using only a very simple and restrictive type of adaptivity: only a single update at an intermediate time T^* is allowed, and static trade schedules are used before and after that “intervention” time. We now develop an alternative procedure that lets us compute fully optimal adaptive strategies.

3.1 Dynamic programming

It is alluring to use dynamic programming to determine optimal trading strategies, since this technique works so well for objective functions of the form $\mathbb{E}[u(Y)]$. But dynamic programming for expected values relies on the “smoothing property” $\mathbb{E}[\mathbb{E}[u(Y) | X]] = \mathbb{E}[u(Y)]$. For the square of the expectation in the variance term $\text{Var}[Y] = \mathbb{E}[Y^2] - \mathbb{E}[Y]^2$, there is no immediate analog of this expression, and it is difficult to see how to design an iterative solution procedure. However, with a suitable choice of the value function, mean-variance optimization is indeed amenable to dynamic programming.

For $(\pi, \bar{C}) \in \mathcal{D}(X, N)$ (with $N \geq 2$), $\pi = (X, x_1, \dots, x_{N-1}, 0)$, denote

$$(\pi, \bar{C})_{\xi_1} \in \mathcal{D}(x_1, N-1) \quad \text{the “tail” of the trading strategy } (\pi, \bar{C})$$

for the remaining $N-1$ trading periods *conditional* on the outcome ξ_1 of the first period. This tail strategy has the trade schedule $\pi_{\xi_1} = (x_1, \dots, x_{N-1}, 0)$, and the cost random variable \bar{C}_{ξ_1} . The cost \bar{C} of the total strategy in terms of its tail cost is

$$\bar{C} = \bar{C}_{\xi_1} + \frac{\eta}{\tau}(X - x_1)^2 - \sigma\tau^{1/2}\xi_1 x_1 . \quad (14)$$

An adaptive policy π may use the information ξ_1 from time t_1 onwards, hence in general $(\pi, \bar{C})_{\xi_1}$ indeed depends on the realization of ξ_1 .

The key ingredient in dynamic programming is to write the time-dependent optimization problem on N periods as the combination of a single-step optimization with an optimization on the remaining $N-1$ periods. We must carefully define the parameters of the $(N-1)$ -step problem so that it gives the same solution as the “tail” of the N -step problem.

In Almgren and Chriss [2000], the risk-aversion parameter λ is constant in time and is constant across realizations of the price process $\xi_1 \in \Omega$. For the expected utility function, Schied and Schöneborn [2009] hold constant the analogous parameter α . For our mean-variance formulation, the following Lemma asserts that the “tail” of any initially optimal strategy across N steps is again an optimal strategy across $N-1$ steps, if it is defined to be the minimum-variance solution for an appropriate cost limit. This cost limit is taken to be the expected value of the remainder of the initial strategy, and will be different in each realization.

Lemma 1. *For $N \geq 2$, let $(\pi, \bar{C}) \in \mathcal{E}(X, N)$ be an efficient execution policy $\pi = (X, x_1, \dots, x_{N-1}, 0)$ for (7). Then $(\pi, \bar{C})_{\xi_1} \in \mathcal{E}(x_1, N-1)$ for almost all $\xi_1 \in \Omega$, i.e. $B = \{ \xi_1 \in \Omega \mid (\pi, \bar{C})_{\xi_1} \notin \mathcal{E}(x_1, N-1) \}$ has probability zero.*

Proof. For each $\xi_1 \in B$ (if B is empty the result is immediate), the tail-strategy $(\pi, \bar{C})_{\xi_1}$ is not efficient, hence there exists $(\pi_{\xi_1}^*, \bar{C}_{\xi_1}^*) \in \mathcal{D}(x_1, N-1)$ such that $\mathbb{E}[\bar{C}_{\xi_1}^*] = \mathbb{E}[\bar{C}_{\xi_1}]$ and $\text{Var}[\bar{C}_{\xi_1}^*] < \text{Var}[\bar{C}_{\xi_1}]$.

Define $(\tilde{\pi}, \tilde{C}) \in \mathcal{D}(X, N)$ by replacing the policy for t_1 to t_N in (π, \bar{C}) by $(\pi_{\xi_1}^*, \bar{C}_{\xi_1}^*)$ for all $\xi_1 \in B$ (and identical to π for all other outcomes $\Omega \setminus B$ of ξ_1). Then by construction

$$\tilde{C} \geq C(X, N, \tilde{\pi}) \quad (15)$$

and hence $(\tilde{\pi}, \tilde{C}) \in \mathcal{D}(X, N)$. Also by construction, we have $\mathbb{E}[\tilde{C}] = \mathbb{E}[\overline{C}]$, and conditional on $\xi_1 \in B$, $\text{Var}[\tilde{C} | \xi_1] < \text{Var}[\overline{C} | \xi_1]$. If B has positive probability then $\mathbb{E}[\text{Var}[\tilde{C} | \xi_1 \in B]] < \mathbb{E}[\text{Var}[\overline{C} | \xi_1 \in B]]$ contradicting $(\pi, \overline{C}) \in \mathcal{E}(X, N)$. \square

For use in the dynamic programming construction, we extend the definition (7) for $1 \leq k \leq N$, for $x \geq 0$ and for fixed τ as

$$J_k(x, c) = \inf \left\{ \text{Var}[\overline{C}] \mid (\pi, \overline{C}) \in \mathcal{A}(x, k, c) \right\}. \quad (16)$$

If the cost limit c is below the cost $E_{\text{lin}}(x, k) = \eta x^2 / k\tau$ (9) of the linear strategy, then no admissible solution exists and we set $J_k = \infty$. If $c = E_{\text{lin}}(x, k)$, then the linear strategy gives the minimum variance (10). If the cost limit is above $E_{\text{inst}}(x) = \eta x^2 / \tau$ (11), then instantaneous liquidation is admissible with variance zero and we have $J_k = 0$. Thus

$$J_k(x, c) = \begin{cases} \infty, & c < \eta x^2 / k\tau \\ V_{\text{lin}}(x, k), & c = \eta x^2 / k\tau \\ \text{non-increasing in } c, & \eta x^2 / k\tau \leq c \leq \eta x^2 / \tau \\ 0, & c \geq \eta x^2 / \tau. \end{cases} \quad (17)$$

In particular, for $k = 1$, $E_{\text{inst}} = E_{\text{lin}}$ and so (for $x \geq 0$)

$$J_1(x, c) = \begin{cases} \infty, & c < \eta x^2 / \tau \\ 0, & c \geq \eta x^2 / \tau. \end{cases} \quad (18)$$

The solution to (7) is $V_{\text{lin}}(E) = J_N(X, E)$.

By definitions (5,6,16), the value function $J_k(x, c)$ and the set $\mathcal{E}(x, k)$ are related by

$$(\pi^*, \overline{C}^*) = \underset{(\pi, \overline{C}) \in \mathcal{A}(x, k, c)}{\text{argmin}} \text{Var}[\overline{C}] \implies (\pi^*, \overline{C}^*) \in \mathcal{E}(x, k) \quad (19)$$

and

$$(\pi, \overline{C}) \in \mathcal{E}(x, k) \implies \text{Var}[\overline{C}] = J_k(x, \mathbb{E}[\overline{C}]). \quad (20)$$

In view of the known static solutions, and by inspection of the expressions (17) and (18), it is natural to conjecture that the value function and the cost limit should be proportional to the square of the number of shares:

$J_k(x, c) = x^2 f_k(c/x^2)$. In fact for dynamic strategies this is *not* true, even for linear impact functions, except in the limit of small portfolio size (in a suitable nondimensional sense made clear in Sections 3.2 and 3.4 below).

In the spirit of dynamic programming, we use the efficient frontier for trading over $k - 1$ periods, plus an optimal one-period strategy, to determine the efficient frontier for trading over k periods. The key is to introduce an additional control parameter in addition to the number of shares we trade in the next period. This extra parameter is the expected cost limit for the remaining periods, which we denote by z ; it is a real-valued integrable function $z \in L^1(\Omega; \mathbb{R})$ of the price change $\xi \in \Omega$ on that step.

Theorem 1. *Let the stock price change in the next trading period be $\sigma\tau^{1/2}\xi$ with $\xi \in \Omega$ the random return. Define*

$$\mathcal{G}_k(x, c) = \left\{ (y, z) \in \mathbb{R} \times L^1(\Omega; \mathbb{R}) \mid \mathbb{E}[z(\xi)] + \frac{\eta}{\tau}(x-y)^2 \leq c, 0 \leq y \leq x \right\}. \quad (21)$$

Then for $k \geq 2$,

$$J_k(x, c) = \min_{(y, z) \in \mathcal{G}_k(x, c)} \left(\text{Var}[z(\xi) - \sigma\tau^{1/2}\xi y] + \mathbb{E}[J_{k-1}(y, z(\xi))] \right). \quad (22)$$

Proof. Let ξ be the random price innovation in the first of the remaining k trading periods. For given $x \geq 0$ and $E_{\text{lin}}(x, k) \leq c$, let

$$(\pi^*, \bar{C}^*) = \underset{(\pi, \bar{C}) \in \mathcal{A}(x, k, c)}{\text{argmin}} \text{Var}[\bar{C}] ,$$

That is, π^* is an optimal strategy to sell x shares in k time periods of length τ with expected cost at most c . By (19), we have $(\pi^*, \bar{C}^*) \in \mathcal{E}(X, N)$, and by the definition (16) of J_k we have $J_k(x, c) = \text{Var}[\bar{C}^*]$. Let y be the number of shares held by π^* after the first trading period, so $\pi^* = (x, y, \dots, 0)$.

The strategy π^* may be understood as consisting of two parts: First, the number of shares to be sold in the first period, $x - y$. This is a deterministic variable, and may not depend on the next period price change ξ . Second, the strategy for the remaining $k - 1$ periods. When the trader proceeds with this $(k - 1)$ -period strategy, the outcome of ξ is known, and the strategy may depend on it. Conditional on ξ , let $(\pi^*, \bar{C}^*)_\xi$ be the $(k - 1)$ -period tail-strategy.

By Lemma 1, $(\pi^*, \bar{C}^*)_\xi \in \mathcal{E}(y, k-1)$ for almost all realizations $\xi \in \Omega$. Thus, there exists $z \in L^1(\Omega; \mathbb{R})$ such that using (20) we have for each ξ

$$\begin{aligned}\mathbb{E}[\bar{C}_\xi^*] &= z(\xi) \\ \text{Var}[\bar{C}_\xi^*] &= J_{k-1}(y, z(\xi)) .\end{aligned}$$

Since $(\pi^*, \bar{C}^*)_\xi \in \mathcal{E}(y, k-1)$, we must have

$$z(\xi) \geq E_{\text{lin}}(y, k-1) \quad (23)$$

(the minimal expected cost is achieved by the linear profile π_{lin}). With (14), we conclude

$$\begin{aligned}\mathbb{E}[\bar{C}^* | \xi] &= z(\xi) + \frac{\eta}{\tau}(x-y)^2 - \sigma\tau^{1/2}\xi y , \\ \text{Var}[\bar{C}^* | \xi] &= J_{k-1}(y, z(\xi)) ,\end{aligned}$$

and by the law of total expectation and total variance

$$\begin{aligned}\mathbb{E}[\bar{C}^*] &= \mathbb{E}[z(\xi)] + \frac{\eta}{\tau}(x-y)^2 , \\ \text{Var}[\bar{C}^*] &= \text{Var}[z(\xi) - \sigma\tau^{1/2}\xi y] + \mathbb{E}[J_{k-1}(y, z(\xi))] .\end{aligned}$$

The pair (π^*, \bar{C}^*) is an optimal solution to

$$\min_{(\pi, \bar{C}) \in \mathcal{A}(x, k, c)} \text{Var}[\bar{C}] ,$$

and so indeed $(z(\xi), y)$ must be such that they are an optimal solution to

$$\min_{(z, y)} \text{Var}[z(\xi) - \sigma\tau^{1/2}\xi y] + \mathbb{E}[J_{k-1}(y, z(\xi))]$$

$$\text{s.t. } \mathbb{E}[z(\xi)] + \frac{\eta}{\tau}(x-y)^2 \leq c$$

$$0 \leq y \leq x \quad (24)$$

$$E_{\text{lin}}(y, k-1) \leq z(\xi) . \quad (25)$$

The constraint (24) comes from our requirement that (π^*, \bar{C}^*) must be a pure sell-program. Since $\mathcal{A}(x, k, c) = \emptyset$ for $c < E_{\text{lin}}(x, k)$ in (16), $J_{k-1}(y, z(\xi)) = \infty$ for $z(\xi) < E_{\text{lin}}(y, k-1)$ and thus the constraint (25) never becomes binding. Thus, the result (21,22) follows. \square

Thus an optimal strategy (π^*, \bar{C}^*) for k periods is defined by $(y, z(\xi))$: the number of shares $x - y$ to sell in the first period, and the $(k - 1)$ -period tail strategy specified by $\mathbb{E}[\bar{C}_\xi^*] = z(\xi)$ (that is, we commit ourselves that if we see ξ in the first period, then we sell the remaining y shares using the mean-variance optimal strategy with expected cost $z(\xi)$ and variance $J_{k-1}(y, z(\xi))$). The expectation and variance of (π^*, \bar{C}^*) are then given by

$$\mathbb{E}[\bar{C}^*] = \mathbb{E}[z(\xi)] + \frac{\eta}{\tau}(x - y)^2, \quad (26)$$

$$\text{Var}[\bar{C}^*] = \text{Var}[z(\xi) - \sigma\tau^{1/2}\xi y] + \mathbb{E}[J_{k-1}(y, z(\xi))] . \quad (27)$$

As noted in (23), not all $z(\xi)$ are possible, since the minimal possible expected cost of the tail strategy $\mathbb{E}[\bar{C}_\xi^*]$ is the expected cost of a linear profile.

In terms of (π, \bar{C}) , we have the following recursion for any optimal strategy: For given (x, k, c) , let $(\pi_k^*(x, c), \bar{C}_k^*(x, c)) \in \mathcal{A}(x, k, c)$ be an optimal strategy for (16), and let $(y, z(\xi))$ be an optimal one-step control in (22). Then we have (for $c \geq \eta x^2/(k\tau)$)

$$\begin{aligned} \pi_k^*(x, c) &= \left(y, \pi_{k-1}^*(y, z(\xi)) \right) \\ \bar{C}_k^*(x, c) &= \bar{C}_{k-1}^*(y, z(\xi)) + \frac{\eta}{\tau}(x - y)^2 - \sigma\tau^{1/2}\xi y, \end{aligned} \quad (28)$$

for $k \geq 2$, and $\pi_1^*(x, c) = x$ and $\bar{C}_1^*(x, c) = \max\{\eta x^2/\tau, c\}$.

The dynamic program (18, 21, 22) indeed allows for $\bar{C} > C(X, N, \pi)$ as stated in the definition (5). Recall that \bar{C} and $C(X, N, \pi)$ are random variables that give a cost for each realization of the stock price process. Thus, $\bar{C} > C(X, N, \pi)$ in (5) means that the trader can incur extra costs (by giving away money) in some realizations of the stock price, if that improves his mean-variance tradeoff. In terms of the dynamic program (18, 21, 22) this is essentially accomplished by the trader's choice of the last period tail strategy: the trader has a certain number x of shares left to sell with actual cost of $\eta x^2/\tau$; there is no decision variable in reality. However, in the specification of the dynamic program (18, 21, 22), which is in line with the definition (7) of \mathcal{E} , the trader additionally specifies $\bar{C} \geq \eta x^2/\tau$; the difference $\bar{C} - \eta x^2/\tau$ is the money that the trader is giving away.

3.2 Nondimensionalization

The optimization problem (7) (respectively the one-step optimization problem (21, 22)) depends on five dimensional constants: the initial shares X , the total time T (or the time step τ in conjunction with the number of steps N), the absolute volatility σ and the impact coefficient η . To simplify the structure of the problem, it is convenient to define scaled variables.

We measure shares x relative to the initial position X . We measure impact cost c and its limit z relative to the total dollar cost that would be incurred by liquidating X shares in time T using the linear strategy; the per-share cost of this strategy is $\eta v = \eta X/T$ so the total cost is $\eta X^2/T$. We measure variance J_k relative to the variance (in squared dollars) of holding X shares across time T with absolute volatility σ . The standard deviation of the price per share is $\sigma\sqrt{T}$, so the standard deviation of dollar value is $\sigma\sqrt{T}X$ and the variance scale is $\sigma^2 T X^2$.

We denote nondimensional values by a caret $\hat{\cdot}$, so we write

$$x = X\hat{x}, \quad c = \frac{\eta X^2}{T}\hat{c}, \quad z = \frac{\eta X^2}{T}\hat{z} \quad \text{and} \quad J_k(x, c) = \sigma^2 T X^2 \hat{J}_k\left(\frac{x}{X}, \frac{c}{\eta X^2/T}\right).$$

Then $\hat{X} = \hat{x}_0 = 1$, so the trading strategy is $\hat{\pi} = (1, \hat{x}_1, \dots, \hat{x}_{N-1}, 0)$.

The one-period value function is

$$\hat{J}_1(\hat{x}, \hat{c}) = \begin{cases} \infty, & \hat{c} < N\hat{x}^2 \\ 0, & \hat{c} \geq N\hat{x}^2 \end{cases} \quad (29)$$

and in Theorem 1 we have the scaled set of admissible controls

$$\hat{\mathcal{G}}_k(\hat{x}, \hat{c}) = \left\{ (\hat{y}, \hat{z}) \in \mathbb{R} \times L^1(\Omega; \mathbb{R}) \mid \mathbb{E}[\hat{z}(\xi)] + N(\hat{x} - \hat{y})^2 \leq \hat{c}, \quad 0 \leq \hat{y} \leq \hat{x} \right\} \quad (30)$$

and the dynamic programming step

$$\hat{J}_k(\hat{x}, \hat{c}) = \min_{(\hat{y}, \hat{z}) \in \hat{\mathcal{G}}_k(\hat{x}, \hat{c})} \left(\text{Var}[\mu\hat{z}(\xi) - N^{-1/2}\xi\hat{y}] + \mathbb{E}[\hat{J}_{k-1}(\hat{y}, \hat{z}(\xi))] \right). \quad (31)$$

The nondimensional ‘‘market power’’ parameter

$$\mu = \frac{\eta X}{\sigma T^{3/2}} = \frac{\eta X/T}{\sigma\sqrt{T}} \quad (32)$$

was identified by Almgren and Lorenz [2007] as a preference-free measure of portfolio size. The numerator is the per-share price impact that would

be caused by liquidating the portfolio linearly across the available time; the denominator is the amount that the price would move on its own due to volatility in the same time. As noted by Almgren and Lorenz [2007], for realistic trade sizes μ will be substantially smaller than one.

The nondimensional version (29,30,31) of the optimization problem now depends only on two nondimensional parameters: the time discretization parameter N and the new market power parameter μ . Especially for numerical treatment, this reduction is very useful. From now on, we shall drop the nondimensionalization mark $\hat{\cdot}$, assuming that all variables have been nondimensionalized.

3.3 Convexity

We now show that the optimization problem at each step is convex, and that the value function J_k is a convex function of its two arguments.

We need the following lemma which is proved by an easy modification of the argument in Boyd and Vandenberghe [2004, sect. 3.2.5].

Lemma 2. *Let $f(v)$ and $h(u, v)$ be real-valued convex functions on vector spaces V and $U \times V$ respectively, possibly taking the value $+\infty$. Then $g : U \mapsto \mathbb{R}$ defined by*

$$g(u) = \inf_{v \in V} \{ f(v) \mid h(u, v) \leq 0 \}$$

is convex.

Now we are ready to prove the convexity of $J_k(x, c)$ and the dynamic programming step.

Theorem 2. *The optimization problem (30,31) is convex. The value function $J_k(x, c)$ is convex for $k \geq 1$.*

Proof. We proceed by induction. Clearly, $J_1(x, c)$ in (29) is convex, since it is finite on the convex domain $\{(x, c) \mid c \geq Nx^2\} \subseteq \mathbb{R}^2$.

The optimization problem (30,31) is of the form described in Lemma 2, with the identifications $u = (x, c)$ and $v = (y, z)$, and $h(u, v)$, $f(v)$ given by the functions appearing on the right side of (30) and (31) respectively. Thus we need only show that these functions are convex.

For each k , the constraint function in (30) is convex in (x, c, y, z) , since the expectation operator is linear and the quadratic term is convex. In (31),

the second term in the objective function is convex in (y, z) , since J_{k-1} is assumed convex and expectation is linear.

The first term in (31) may be written $\text{Var}[w(\xi)]$ where $w(\xi)$ depends linearly on y and $z(\xi)$. And it is easy to see that $\text{Var}[w]$ is convex in w . (Indeed, this is certainly true for random variables w having $\mathbb{E}[w] = 0$ since then $\text{Var}[w] = \int w^2$. For general w , one can write $w(\xi) = \bar{w} + u(\xi)$ where \bar{w} is constant and $u(\xi)$ has mean zero; then $\text{Var}[w] = \int u^2$ and convexity follows.) The result follows. \square

Now let us return to the definitions (5,6,7), especially the constraint “ $C(X, N, \pi) \leq \bar{C}$ ” in (5). This effectively allows the trader to destroy money: he may report a cost that is higher than the costs actually incurred by the trade strategy. This is counterintuitive, but the trader may want to make use of it due to a rather undesirable property of the mean-variance criterion: A mean-variance optimizer can reduce his variance by making positive outcomes less so. Of course this also reduces his mean benefit, but depending on the parameters the tradeoff may be advantageous. As it is well-known, mean-variance comparison does not necessarily respect stochastic dominance.

If we want to bar the trader from making use of this peculiarity, we replace (5) by

$$\mathcal{D}'(X, N) = \left\{ (\pi, \bar{C}) \left| \begin{array}{l} \pi = (x_0, x_1, \dots, x_N) \text{ with } x_0 = X, x_N = 0 \\ x_0 \geq x_1 \geq \dots \geq x_N \\ C(X, N, \pi) = \bar{C} \end{array} \right. \right\} \quad (33)$$

Now the random variable \bar{C} gives the *exact* actual cost of the trade schedule π at all times. We change the definitions (6,7) of \mathcal{A} and \mathcal{E} accordingly, replacing \mathcal{D} by \mathcal{D}' , and denote these new sets \mathcal{A}' and \mathcal{E}' . We define

$$J'_k(x, c) = \min_{(\pi, \bar{C}) \in \mathcal{A}'(x, k, c)} \text{Var}[\bar{C}] \quad . \quad (34)$$

With these definitions, $\mathcal{D}'(X, 1)$ is the single-element set $\mathcal{D}'(X, 1) = \{(\pi_{\text{inst}}, E_{\text{inst}}(X))\}$ where $\pi_{\text{inst}} = (X, 0)$ is the immediate liquidation of X shares and $E_{\text{inst}}(X)$ its cost.

It can be shown that Lemma 1 also holds for \mathcal{E}' .

Lemma 3. *For $N \geq 2$, let $(\pi, \bar{C}) \in \mathcal{E}'(X, N)$ with $\pi = (X, x_1, \dots, x_{N-1}, 0)$. Then $B = \{a \in \Omega \mid (\pi, \bar{C})_{\xi_1=a} \notin \mathcal{E}'(x_1, N-1)\}$ has probability zero.*

The proof follows the proof of Lemma 3 word by word, with \mathcal{D} and \mathcal{E} replaced by \mathcal{D}' and \mathcal{E}' , respectively; only (15) changes to $\tilde{C} = C(X, N, \tilde{\pi})$, which indeed implies $(\tilde{\pi}, \tilde{C}) \in \mathcal{D}'(X, N)$ accordingly.

Using Lemma 3, we can then argue along the lines of the proof of Theorem 1 to obtain the (nondimensionalized) dynamic program

$$\mathcal{G}'_k(x, c) = \left\{ (y, z) \in \mathbb{R} \times L^1(\Omega; \mathbb{R}) \left| \begin{array}{l} \mathbb{E}[z] + N(x - y)^2 \leq c \\ z \leq Ny^2 \text{ a.e.} \\ 0 \leq y \leq x \end{array} \right. \right\} \quad (35)$$

and

$$J'_k(x, c) = \min_{(y, z) \in \mathcal{G}'_k(x, c)} \left(\text{Var}[\mu z - N^{-1/2} \xi y] + \mathbb{E}[J_{k-1}(y, z)] \right), \quad (36)$$

with $J'_1(x, c) = J_1(x, c)$ unchanged. The additional constraint “ $z(\xi) \leq Ny^2$ a.e.” in (35) comes from the fact that $z(\xi)$ specifies the $(k-1)$ -period tail strategy $(\pi^*, \bar{C}^*)_\xi$ by means of $\mathbb{E}[\bar{C}^*] = z(\xi)$ and not all $z(\xi)$ correspond to a valid $(\pi^*, \bar{C}^*)_\xi$: the maximal expected cost of a $(k-1)$ -period tail strategy is the cost of immediate liquidation $\eta y^2 / \tau$ (or Ny^2 in nondimensionalized variables). In Theorem 1, this upper bound does not apply because by means of $(\pi_{\text{inst}}, c) \in \mathcal{E}(y, k-1)$ for *all* $c \geq E_{\text{inst}} = Ny^2$ the trader may give away the extra cash.

Obviously, we have $J'_k(x, c) \geq J_k(x, c)$ for all $k \geq 1$. Contrary to (30,31), the optimization problem (35,36) is *not* a convex optimization problem because the additional constraint breaks the convexity of the set $\mathcal{G}'_k(x, c)$ and Lemma 2 is not applicable.

In the following, we shall continue to work with the value function $J_k(x, c)$ as defined in Section 3.1.

3.4 Small portfolios

Almgren and Lorenz [2007] observed that for small portfolios the optimal adaptive and optimal static efficient frontier coincide. We now prove that this indeed holds also for general strategies. We denote by J_k^0 the value function when $\mu = 0$. Note that this case is perfectly natural in the nondimensional form (30,31). But in the original dimensional form, from the definition (32), this requires $\eta \rightarrow 0$, $X \rightarrow 0$, $\sigma \rightarrow \infty$, or $T \rightarrow \infty$, all of which pose conceptual problems for the model. We are not able to prove rigorously that the solution

of the problem for $\mu = 0$ is the same as the limit of the solutions for positive μ , but we conjecture that this is true.

Theorem 3. *For $\mu = 0$, the optimal policy of (30,31) is path-independent (static) and the efficient frontier coincides with the static efficient frontier.*

Proof. For $\mu = 0$, (31) becomes

$$J_k^0(x, c) = \min_{(y,z) \in \mathcal{G}_k(x,c)} \left(N^{-1}y^2 + \mathbb{E}[J_{k-1}^0(y, z(\xi))] \right) . \quad (37)$$

Inductively, we now show that for $k \geq 1$ (defining $x_k = 0$ to shorten notation)

$$J_k^0(x, c) = \min_{x_1 \geq \dots \geq x_{k-1}} \left\{ \frac{1}{N} \sum_{j=1}^{k-1} x_j^2 \mid (x - x_1)^2 + \sum_{j=2}^k (x_{j-1} - x_j)^2 \leq \frac{c}{N} \right\} \quad (38)$$

for $c \geq Nx^2/k$, and $J_k^0(x, c) = \infty$ otherwise. For $k = 1$, (38) reduces to

$$J_1^0(x, c) = 0 \quad \text{for } c \geq Nx^2, \quad \text{and } J_1^0(x, c) = \infty \quad \text{for } c < Nx^2 ,$$

and by definition (29) indeed $J_1^0(x, c) = J_1(x, c)$. For $k \geq 2$, suppose that (38) holds for $k - 1$. $J_{k-1}^0(x, c)$ is convex by Lemma 2. Thus, for any nonconstant $z(\xi)$, Jensen's inequality implies $J_{k-1}^0(y, \mathbb{E}[z]) \leq \mathbb{E}[J_{k-1}^0(y, z)]$. i.e. there exists a constant nonadaptive optimal control $z(\xi) \equiv z$. Thus, (37) becomes

$$J_k^0(x, c) = \min_{y,z} \left\{ N^{-1}y^2 + J_{k-1}^0(y, z) \mid z + N(x - y)^2 \leq c \right\} .$$

After undoing the nondimensionalization, for $k = N$ the optimization problem (38) is exactly problem (12) for the static trajectory. Hence, for $\mu = 0$ the adaptive efficient frontier does coincide with the static efficient frontier. \square

The theorem also holds for the variant (35, 36), where we restrict the trader from giving away money. The reason is that for $c \geq Ny^2$, $J_{k-1}^0(y, c) = 0$ since then $x_1 = \dots = x_{k-1} = 0$ in (38) is admissible. Hence, the constraint $z \leq Ny^2$ in (35, 37) will in fact never become binding.

For $\mu > 0$, improvements over static strategies come from introducing an anticorrelation between the two terms inside the variance in (31). This reduces the overall variance, which we can trade for a reduction in expected cost. Thus, following a positive investment return, we decrease our cost limit for the remaining part of the program.

4 Approximation of Optimal Control

The dynamic programming principle presented in Section 3 gives a method to obtain fully optimal mean-variance strategies for (7). However, the optimization problem (30,31) that we have to solve in each of the N recursive steps is generally very hard, since it requires determining an optimal control function $z : \Omega \rightarrow \mathbb{R}$ and Ω may be of infinite cardinality (e.g. $\Omega = \mathbb{R}$ and $\xi_i \sim \mathcal{N}(0, 1)$). In some situations, for instance for the classical Markowitz portfolio problem in a multiperiod setting, a closed-form analytical solution may be obtained [Lorenz, 2008, chap. 5]. There, this is due to an anti-correlation argument; the optimal control function $z : \Omega \rightarrow \mathbb{R}$ is perfectly anti-correlated to portfolio return in the current period. Unfortunately, this argument cannot be used for (21,22) because of the additional constraint of the type $z \geq E_{\text{lin}}(x, k, \tau)$; since the stock price change ξ may take arbitrarily large positive and negative values, a lower bounded z can never be *perfectly* correlated to ξ . Nonetheless, we will see in Theorem 4 below that z is indeed correlated to ξ , yet not perfectly.

In this section, we show how approximate solutions can be obtained, which converge to the full optimal solution as $N \rightarrow \infty$.

4.1 Step-function approximation

For simplicity, we assume $\Omega = \mathbb{R}$. Suppose we restrict the space of admissible controls z to the space of step functions rather than all measurable functions. More precisely, we partition the real line into n intervals I_1, \dots, I_n with

$$I_1 = (-\infty, a_1), \quad I_2 = [a_1, a_2), \dots, I_{n-1} = [a_{n-1}, a_n), \quad I_n = [a_n, \infty) \quad (39)$$

for $a_1 < \dots < a_n$. For given $(z_1, \dots, z_n) \in \mathbb{R}^n$ we define the step function $z : \mathbb{R} \rightarrow \mathbb{R}$ by $z(\xi) = z_j$ for $\xi \in I_j$. For large n , this approaches a continuous dependence $z(\xi)$. Let

$$p_i = \mathbb{P}[\xi \in I_i] \quad , \quad E_i = \mathbb{E}[\xi | \xi \in I_i] \quad \text{and} \quad V_i = \text{Var}[\xi | \xi \in I_i] \quad .$$

Let $j = j(\xi)$ be the indicator random variable $j \in \{1, \dots, n\}$ such that $j = j$ if and only if $\xi \in I_j$.

Straightforward calculation yields

$$\text{Var} \left[\mu z_j - \frac{\xi y}{\sqrt{N}} \right] = \sum_{i=1}^n p_i \left(\mu z_i - \frac{E_i y}{\sqrt{N}} - \mu \sum_{j=1}^n p_j z_j \right)^2 + \frac{y^2}{N} \sum_{i=1}^n p_i V_i \quad .$$

Defining $\bar{z} = \sum_{i=1}^n p_i z_i$ and

$$E(y, z) = N(x - y)^2 + \bar{z}, \quad (40)$$

$$V(y, z) = \sum_{i=1}^n p_i \left\{ \left(\mu(z_i - \bar{z}) - \frac{E_i y}{\sqrt{N}} \right)^2 + \frac{y^2 V_i}{N} + \tilde{J}_{k-1}(y, z_i) \right\} \quad (41)$$

the optimization problem (30,31) reads

$$\tilde{J}_k(x, c) = \min_{(y, z_1, \dots, z_n) \in \mathbb{R}^{n+1}} \left\{ V(y, z) \left| \begin{array}{l} E(y, z) \leq c \quad (C1) \\ Ny^2/(k-1) \leq z_i \quad \forall i \quad (C2) \\ 0 \leq y \leq x \quad (C3) \end{array} \right. \right\} \quad (42)$$

in $\text{dom } J_k = \{(x, c) \mid x \geq 0, c \geq Nx^2/k\}$, with $\tilde{J}_1(x, c) = J_1(x, c)$ in (29).

Thus, we have to solve an optimization problem with $n+1$ variables in each step. Since $E(y, z)$ and $V(y, z)$ are convex in (y, z) , as the general dynamic program (31, 30) the approximate problem (42) constitutes a convex optimization problem as well. Hence, by Lemma 2 all approximate value functions $\tilde{J}_k(x, z)$ are convex functions in (x, z) .

The corresponding approximation of (35, 36) for $J'(x, c)$ is

$$\tilde{J}'_k(x, c) = \min_{(y, z_1, \dots, z_n) \in \mathbb{R}^{n+1}} \left\{ V(y, z) \left| \begin{array}{l} E(y, z) \leq c \quad (C1) \\ Ny^2/(k-1) \leq z_i \quad \forall i \quad (C2) \\ 0 \leq y \leq x \quad (C3) \\ Ny^2 \geq z_i, i = 1 \dots n \quad (C4) \end{array} \right. \right\}, \quad (43)$$

with the additional constraint (C4). $\tilde{J}'_1(x, c) = J'_1(x, c)$ remains unchanged.

4.2 Aggressiveness in-the-money (AIM)

In this section we shall prove that for step-function controls z as described in Section 4.1 any optimal control $z(\xi)$ is positively correlated to the stock price return. In terms of the controls z_1, \dots, z_n this means that $z_1 \leq \dots \leq z_n$ since by definition (39) of the intervals I_j for $i < j$ the control z_i corresponds to price changes ξ that are smaller than those for z_j .

The interpretation is that if the stock price goes up, we sell faster (higher expected cost z_i for the remainder). We obtain a AIM strategy (aggressive in the money), which burns part of the windfall trading gains to sell faster and reduce the risk for the time left, as observed by Almgren and Lorenz [2007].

Theorem 4. *Let $\mu > 0$ and $3 \leq k \leq N$. For $Nx^2/k < c < Nx^2$ any optimal control for $\tilde{J}_k(x, c)$ in (40,41,42) must satisfy*

$$y > 0 \quad \text{and} \quad z_1 \leq z_2 \leq \cdots \leq z_k . \quad (44)$$

For $c \geq Nx^2$, the optimal value is attained by the unique optimal control $y = z_1 = \cdots = z_k = 0$ (immediate liquidation) and for $c = Nx^2/k$ it is attained by the unique optimal control $y = (k-1)x/k$ and $z_1 = \cdots = z_k = x^2(k-1)N/k^2$ (linear profile).

Proof. It is easy to see by induction that $\tilde{J}_k(x, c) \geq 0$, and that for $c \geq E_{\text{inst}} = Nx^2$ the minimum value $\tilde{J}_k(x, c) = 0$ is attained only for $z_1 = \cdots = z_k = 0$ and $y = 0$. For $c = Nx^2/k$, the only point that satisfies (C1)–(C3) in (42) is $y = (k-1)x/k$ and $z_1 = \cdots = z_k = x^2(k-1)N/k^2$, which indeed corresponds to the linear strategy. For $c < Nx^2$, suppose $y = 0$. Then, by (40) we have $E(y, z) \geq Nx^2$, a contradiction.

Now suppose $z_s > z_r$ for $r > s$. Let

$$\bar{z} = (p_r z_r + p_s z_s)/(p_r + p_s) \quad \text{and} \quad \delta = z_r - z_s .$$

Then $\delta < 0$, and

$$z_r = \bar{z} + \delta p_s/(p_r + p_s) \quad \text{and} \quad z_s = \bar{z} - \delta p_r/(p_r + p_s) . \quad (45)$$

Let $A = \mu \sum_{j=1}^n p_j z_j = (p_r + p_s)\mu\bar{z} + \mu \sum_{j \neq r,s} p_j z_j$, and

$$\Delta = V(y, \tilde{z}_1, \dots, \tilde{z}_n) - V(y, z_1, \dots, z_n) \quad (46)$$

with $\tilde{z}_i = z_i$ for $i \notin \{r, s\}$ and $\tilde{z}_r = \tilde{z}_s = \bar{z}$. Since $E(y, \tilde{z}_1, \dots, \tilde{z}_n) = E(y, z_1, \dots, z_n)$, the control $(y, \tilde{z}_1, \dots, \tilde{z}_n)$ satisfies (C1) in (42). Since z satisfies (C2) in (42), i.e. $z_r \geq y^2 N/(k-1)$, we have $\bar{z} > z_r \geq y^2 N/(k-1)$, and hence \tilde{z} also satisfies (C2).

We shall prove that $\Delta < 0$, contradicting the optimality of (y, z_1, \dots, z_n) . To shorten notation, let $J(x, c) = \tilde{J}_{k-1}(x, c)$. Since $J(x, c)$ is convex,

$$p_r J(y, z_r) + p_s J(y, z_s) \geq (p_r + p_s) J\left(y, \frac{p_r z_r + p_s z_s}{p_r + p_s}\right) = (p_r + p_s) J(y, \bar{z}) .$$

Hence,

$$\Delta \leq \sum_{i=r,s} p_i \left[(\mu\bar{z} - yE_i N^{-1/2} - A)^2 - (\mu z_i - yE_i N^{-1/2} - A)^2 \right] .$$

Using (45), $E_r > E_s$, $\delta < 0$ and $y > 0$, we obtain $\Delta < 0$. \square




Figure 1: Backwards optimization in the binomial framework. If we already have the efficient frontier to sell from time t_j onwards, at time t_{j-1} we need to determine an optimal control (y, z_+, z_-) : we sell $x - y$ shares between t_{j-1} and t_j , and commit ourselves to trading strategies for the remainder chosen from the set of efficient strategies for t_j to t_N depending on whether the stock goes up or down. If the stock goes up (down), we follow the efficient strategy with expected cost z_+ (z_-). The choice $z_+ = z_- = z_0$ would lead to a path-independent strategy. By choosing $z_+ > z_-$ (and y) optimally, we can reduce the variance of the whole strategy measured at t_{j-1} . Instead of the square shaped point on the static frontier at time t_{j-1} , we obtain a point on the improved frontier (blue dashed line).

5 Numerical Example

5.1 Binomial Model

The simplest approximation scheme of the type outlined in Section 4 are step functions

$$z(\xi) = \begin{cases} z_+ & \text{for } \xi \geq 0 \\ z_- & \text{for } \xi < 0 \end{cases} .$$

Assuming that the price innovations ξ are independent Gaussian random variables $\xi \sim \mathcal{N}(0, 1)$, straightforward calculation yields

$$E_+ = \mathbb{E}[\xi | \xi \geq 0] = \sqrt{2/\pi}, \quad E_- = \mathbb{E}[\xi | \xi < 0] = -E_+ ,$$

$$V_+ = \text{Var}[\xi | \xi \geq 0] = 1 - 2/\pi, \quad V_- = \text{Var}[\xi | \xi < 0] = 1 - 2/\pi ,$$

and $p_+ = \mathbb{P}[\xi \geq 0] = 1/2$, $p_- = \mathbb{P}[\xi < 0] = 1/2$.

That is, the adaptivity of our trading strategy is restricted to reacting to the observation whether the stock price goes up or down during each of the N trading periods. As $N \rightarrow \infty$, this simple model of adaptivity will converge to the fully adaptive optimal strategy, similar to a binomial tree converging to Brownian motion. Figure 1 illustrates this binomial framework.

5.2 Numerical Results

For numerical computations, we discretize the state space of the value functions $J_k(x, c)$. The figures presented in this section were generated for $T = 1$, $N = 50$ time steps ($\tau = 1/50$) with $N_x = 250$ grid points for the relative portfolio size $x \in [0, 1]$ and $N_c = 100$ in the cost dimension (N_c points on the frontier for each value of x). Starting with $J_1(x, c)$, we successively determine $J_2(x, c), \dots, J_N(x, c)$ by means of (42), using interpolated values from the grid data of the previous value function. We use standard direct search methods to find the optimal control (y, z_+, z_-) . Since the optimization problem is convex (Section 3.3), local minima are global optimal solutions. For each level of x we have to trace an efficient frontier. The function value (9) for the linear strategy at the upper-left end of the frontier is readily available; from there, we work towards the right (increasing c) and compute optimal controls for each c by taking the optimal controls for the point $c - h$ (where h is the discretization along the cost dimension) as the starting point for the iteration. Note that the optimal control for $c - h$ is indeed a feasible starting point for the optimization problem with maximal cost c .

Figure 5 shows the set of efficient frontiers at the initial time $t = 0$ for the entire initial portfolio (relative portfolio size $x = 1$) for different values of the market power $0 \leq \mu \leq 0.15$. (Recall the discussion about the order of magnitude for μ in Section 3.2.) The horizontal axis is the expectation of total cost, and the vertical axis its variance. We scale both expectation and variance by their values for the linear trajectories (9). The two blue marks on the frontier for $\mu = 0.15$ correspond to optimal adaptive strategies with the same mean, but lower variance (below the black mark) and same variance, but lower mean (to the left of the black mark) as the static strategy corresponding to the black mark. The inset shows the cost distributions associated with these three strategies (trading costs increase to the right of the x -axis). Those distributions were determined by applying these three strategies to 10^5 randomly sampled paths of the stock price: for each sample path (ξ_1, \dots, ξ_N) the final cost C_N is obtained by sequentially

applying the optimal one-step controls $(y_k^*(x, c), z_{k\pm}^*(x, c))$ associated with $J_k(x_k, c)$ (respectively, their interpolated values over the discretized state space),

$$\begin{aligned} x_{i+1} &= y_i^*(x_i, c_i) \\ c_{i+1} &= \begin{cases} z_{i+}^*(x_i, c_i) & \xi_{i+1} \geq 0 \\ z_{i-}^*(x_i, c_i) & \xi_{i+1} < 0 \end{cases} \\ C_{i+1} &= C_i + \mu N(x_i - x_{i+1})^2 - x_{i+1}\xi_{i+1}/\sqrt{N} \end{aligned}$$

with $x_0 = 1$, $C_0 = 0$ and the initial limit $c_0 = E$ for the expected cost. C_i is measured in units of $\sigma\sqrt{T}X$, the standard deviation of the initial portfolio value due to the stock price volatility across the trading horizon.

The adaptive cost distributions are slightly skewed, suggesting that mean-variance optimization may not give the best possible solutions. Figure 6 shows four static and adaptive cost distributions along the frontier. In the upper left corner (near the linear strategy), the adaptive cost distributions are almost Gaussian (Point #1); indeed, for high values of V adaptive and static strategies coincide. As we move down the frontiers (towards less risk-averse strategies), the skewness first increases (Point #2). Interestingly, as we move further down, where the improvement of the adaptive strategy becomes larger, the adaptive distributions look more and more Gaussian again (Point #3 and #4). All adaptive distributions are strictly preferable to their reference static strategy, since they have lower probability of high costs and higher probability of low costs. Table 1 compares the semi-variance, value-at-risk (VaR) and conditional value-at-risk (CVaR) [see Artzner, Delbaen, Eber, and Heath, 1999, for instance] for the four distribution pairs shown in Figure 6. For Gaussian random variables, mean-variance is consistent with expected utility maximization as well as stochastic dominance (see for instance Bertsimas, Lauprete, and Samarov [2004], Levy [1992]). As the adaptive distributions are indeed not too far from Gaussian, we can expect mean-variance to give reasonable results.

To illustrate the behavior of adaptive policies in the binomial framework, Figure 4 shows trajectories for two sample paths of the stock price in a small instance with $N = 4$. The inset shows the underlying binomial decision tree.

Figure 7 shows optimal adaptive and static trading for $N = 50$. The static strategy is chosen such that it has the same expected cost (yet higher variance) as the adaptive policy.




Figure 2: Left plot shows $J_2(x, c)$ for different values of $\mu \in \{0.25, 0.5, 0.75, 1.0\}$ as a function of c (and $x = 1$). The black line is $J_2'(x, c)$. The right plot shows $\|J_2 - J_2'\|_1$ as a function of μ .

As can be seen, the optimal adaptive policies are indeed “aggressive in the money”. If the stock price goes up, we incur unexpectedly smaller total trading cost and react with selling faster (burning some of the gains), whereas for falling stock price, we slow down trading.

The numerical results presented so far were obtained using the value function J_N , (21, 22). Let us briefly discuss the results for the value function definition J'_N , (35, 36) (respectively \tilde{J}' in (43)). As mentioned there, $J'_N(x, c) \geq J_N(x, c)$. In the specification of J_N (respectively, \mathcal{E}) the trader can reduce his variance by destroying money in order to make a positive outcome less so (see discussion in Section 3.3). In the specification of J'_N (respectively, \mathcal{E}') this is not possible. The numerical results show that while this effect is important for small values of N and large values of μ (see Figure 2), it diminishes very rapidly as N increases (see Figure 3). In fact, the value of $\mu = 2$ in Figure 3 is rather large (recall our discussion for the order of μ in Section 3.2), and for realistic values of μ the difference is even smaller. That is, while the specification of J_N and $\mathcal{E}(X, N)$ allow for the undesirable peculiarity that the trader gives away money in positive outcomes, our numerical results show that this effect does not play a big role in practice. The specification of J_N has the advantage that the associated dynamic program is convex, which makes the numerical optimization easier.

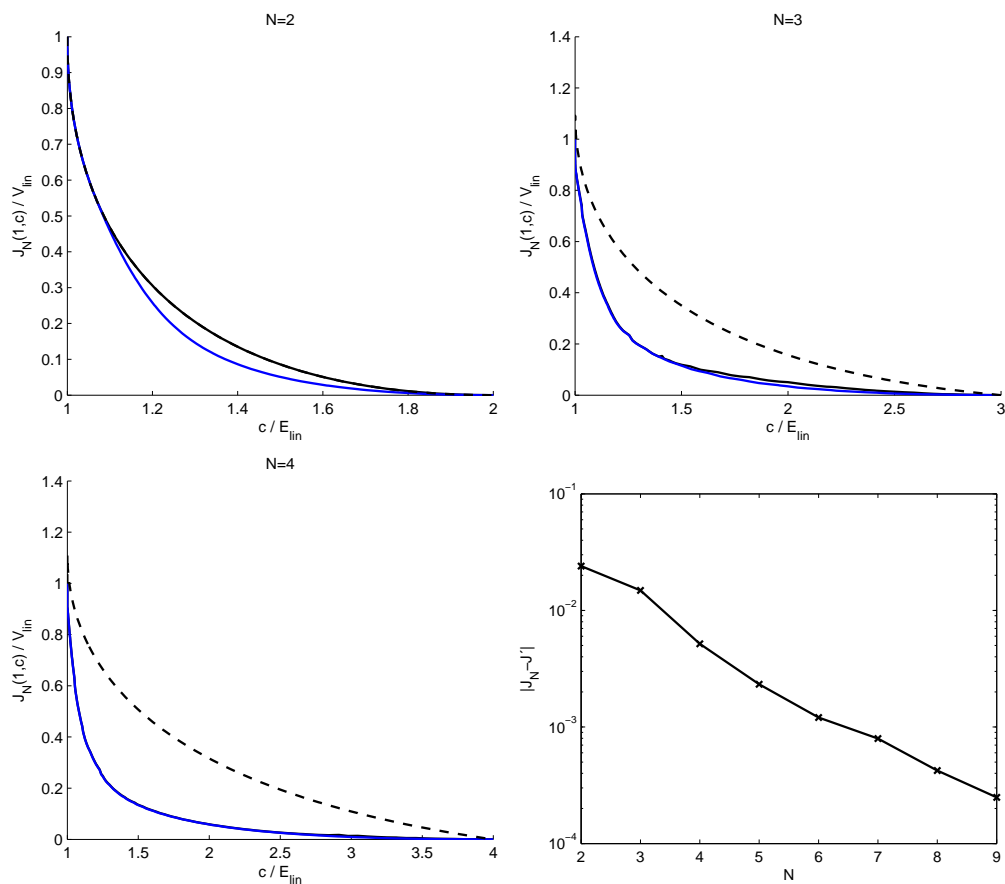


Figure 3: The first three plots show $J_N(x, c)$ (blue solid line) vs. $J'_N(x, c)$ (black solid line) for $N = 2, 3, 4$ and $\mu = 2$. For $N = 2$, the two curves are clearly separated. For $N = 3$ there is only a visible difference for larger values of c , and for $N = 4$ the two curves are almost identical. The black dashed line is the static efficient frontier (i.e. $\mu = 0$), which coincides for $N = 2$ with $J'_2(x, c)$ in the first plot. The last plot shows $\|J_N - J'_N\|_1$ as a function of N .

	#1		#2		#3		#4	
	static	adapt	static	adapt	static	adapt	static	adapt
$\mathbb{E}[\cdot]$	1.68	1.52	2.96	2.27	7.00	3.92	13.00	7.09
$\text{Var}[\cdot]$	5.98	5.98	3.19	3.19	1.20	1.20	0.44	0.44
$\text{SVar}[\cdot]$	3.01	3.35	1.61	1.89	0.60	0.68	0.22	0.22
$\text{VaR}_{5.0\%}$	5.70	5.85	5.90	5.58	8.80	5.83	14.09	8.17
$\text{VaR}_{2.5\%}$	6.47	6.77	6.46	6.38	9.14	6.34	14.29	8.41
$\text{VaR}_{1.0\%}$	7.37	7.91	7.12	7.35	9.54	7.02	14.54	8.73
$\text{VaR}_{0.5\%}$	7.98	8.62	7.56	8.00	9.82	7.56	14.70	8.97
$\text{VaR}_{0.1\%}$	9.23	10.18	8.48	9.45	10.38	8.57	15.04	9.43
$\text{CVaR}_{5.0\%}$	6.72	7.09	6.65	6.67	9.25	6.56	14.36	8.51
$\text{CVaR}_{2.5\%}$	7.43	7.91	7.15	7.39	9.57	7.06	14.54	8.75
$\text{CVaR}_{1.0\%}$	8.23	8.87	7.71	8.26	9.93	7.71	14.77	9.06
$\text{CVaR}_{0.5\%}$	8.76	9.52	8.14	8.87	10.20	8.19	14.91	9.28
$\text{CVaR}_{0.1\%}$	9.94	10.96	8.97	10.16	10.69	9.19	15.20	9.80

Table 1: Statistics for the adaptive and static cost distribution functions shown in Figure 6, obtained by Monte Carlo simulation (10^5 sample paths). For the random variable C , the total cost in units of E_{lin} , the value-at-risk VaR_β is defined by $\mathbb{P}[C \geq \text{VaR}_\beta(C)] = \beta$, and the conditional-value-at-risk $\text{CVaR}_\beta(C) = \mathbb{E}[C | C \geq \text{VaR}_\beta(C)]$. Thus, low values for VaR and CVaR are desirable.

6 Conclusion

By a suitable application of the dynamic programming principle, we show how one can obtain fully optimal adaptive policies for the optimal execution problem considered by Almgren and Chriss [2000]. The optimal adaptive policies significantly improve over static trade schedules, with the improvement being larger for large portfolios. As already observed by Almgren and Lorenz [2007], those adaptive policies are “aggressive in the money”: If the price moves in your favor, then spend those gains on market impact costs by trading faster for the remainder of the program. If the price moves against you, then reduce future costs by trading more slowly. This kind of correlating the profit in the current period to the expected profit in the remaining time can also be observed for the classical portfolio optimization problem in a multiperiod setting [Lorenz, 2008, Richardson, 1989].




Figure 4: Optimal adaptive trading for $N = 4$ time steps, illustrating the binomial adaptivity model. The blue trajectory corresponds to the rising stock price path, and sells faster than red trajectory (falling stock price path). The inset shows the schematics of the stock price on the binomial tree. x_1 at $t = 1$ is the same for all adaptive trajectories because x_1 is determined at $t = 0$ with no information available. Only from $t = 1$ onwards the adaptive trajectories split.

References

- R. Almgren and N. Chriss. Optimal execution of portfolio transactions. *J. Risk*, 3(2):5–39, 2000.
- R. Almgren and J. Lorenz. Bayesian adaptive trading with a daily cycle. *J. Trading*, 1(4), 2006.
- R. Almgren and J. Lorenz. Adaptive arrival price. In B. R. Bruce, editor, *Algorithmic Trading III*, pages 59–66. Institutional Investor, 2007.
- R. F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Appl. Math. Fin.*, 10:1–18, 2003.
- R. F. Almgren. Optimal trading in a dynamic market. Preprint, 2009.

- P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent measures of risk. *Math. Finance*, 9:203–228, 1999.
- S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Rev. Financial Studies*, 23(8):2970–3016, 2010.
- D. Bertsimas and A. W. Lo. Optimal control of execution costs. *J. Financial Markets*, 1:1–50, 1998.
- D. Bertsimas, G. J. Lauprete, and A. Samarov. Shortfall as a risk measure: Properties, optimization and applications. *J. Econom. Dynam. Control*, 28(7):1353–1381, 2004.
- T. R. Bielecki, H. Jin, S. R. Pliska, and Z. Y. Zhou. Continuous-time mean-variance portfolio selection with bankruptcy prohibition. *Math. Finance*, 15(2):213–244, 2005.
- S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- J. Cvitanić, A. Lazrak, and T. Wang. Implications of the Sharpe ratio as a performance measure in multi-period settings. *J. Econom. Dynam. Control*, 32:1622–1649, 2008.
- P. A. Forsyth. A Hamilton-Jacobi-Bellman approach to optimal trade execution. *Appl. Numer. Math.*, 61(2):241–265, 2011.
- H. He and H. Mamaysky. Dynamic trading policies with price impact. *J. Econom. Dynam. Control*, 29(5):891–930, 2005.
- G. Huberman and W. Stanzl. Optimal liquidity trading. *Review of Finance*, 9(2):165–200, 2005.
- R. Kissell and R. Malamut. Understanding the profit and loss distribution of trading algorithms. In B. R. Bruce, editor, *Algorithmic Trading*, pages 41–49. Institutional Investor, 2005.
- A. J. Kurdila and M. Zabrankin. *Convex Functional Analysis*. Birkhäuser Verlag, Basel–Boston–Berlin, 2005.
- H. Levy. Stochastic dominance and expected utility: Survey and analysis. *Management Science*, 38(4):555–593, 1992.

- D. Li and W.-L. Ng. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Math. Finance*, 10(3):387–406, 2000.
- J. Lorenz. *Optimal Trading Algorithms*. PhD thesis, ETH Zurich, 2008.
- A. F. Perold. The implementation shortfall: Paper versus reality. *J. Portfolio Management*, 14(3):4–9, 1988.
- H. R. Richardson. A minimum variance result in continuous trading portfolio optimization. *Management Science*, 35(9):1045–1055, 1989.
- A. Schied and T. Schöneborn. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance Stochast.*, 13(2):181–204, 2009.
- R. H. Strotz. Myopia and inconsistency in dynamic utility maximization. *Rev. Economic Studies*, 23(3):165–180, 1956.
- N. Walia. Optimal trading: Dynamic stock liquidation strategies, 2006. Senior thesis, Princeton University.
- X. Y. Zhou and D. Li. Continuous-time mean-variance portfolio selection: A stochastic LQ framework. *Appl. Math. Optim.*, 42:19–33, 2000.




Figure 5: Adaptive efficient frontiers (in the binomial model described in Section 5.1) for different values of the market power $\mu \in \{0.025, 0.05, 0.075, 0.15\}$, and $N = 50$. The expectation and variance of the total trading cost are normalized by their values of a linear trajectory (VWAP) and plotted in a semilogarithmic scale. 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 $\mu \rightarrow 0$. The blue dashed curves are the improved efficient frontiers, with the improvement increasing with μ . The inset shows the distributions of total cost corresponding to the three points marked on the frontiers, determined by applying these three strategies to 10^5 randomly sampled paths of the stock price.

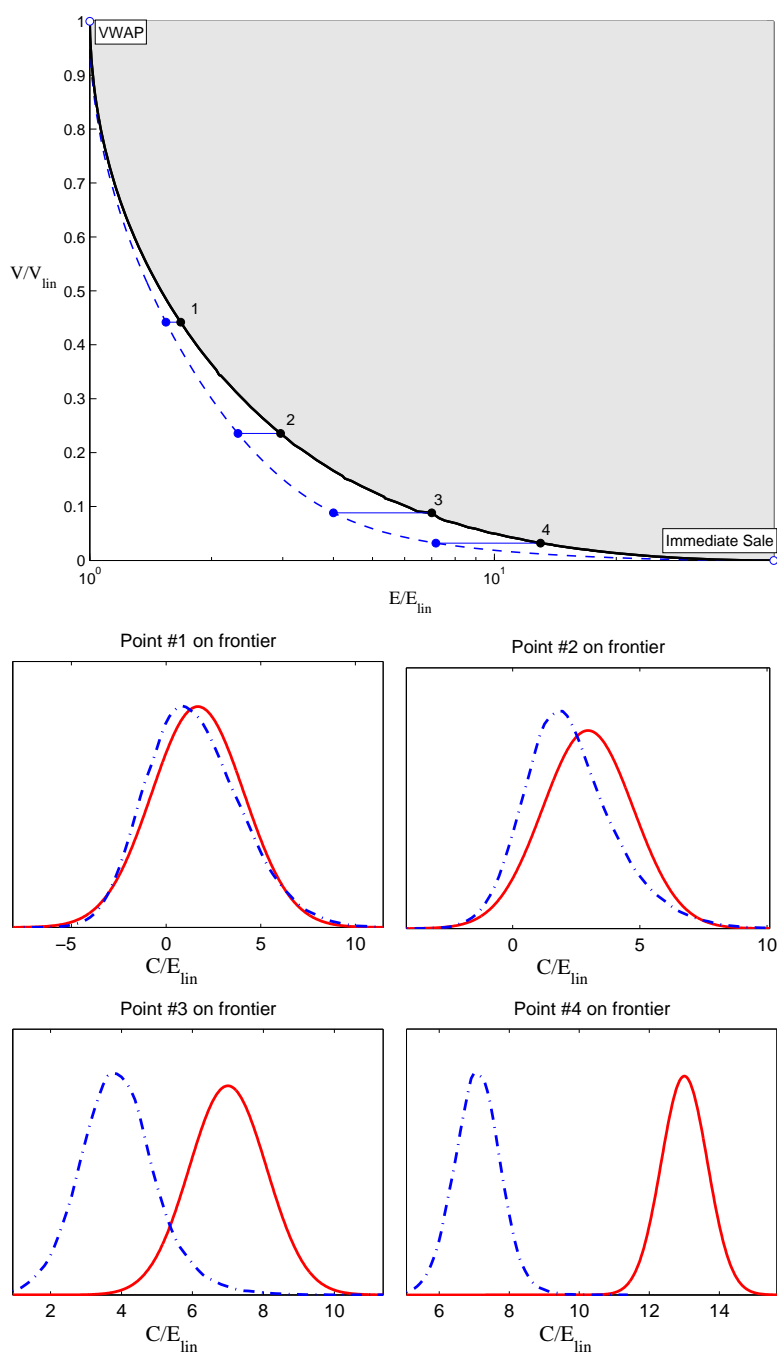


Figure 6: Distributions of total cost corresponding to four points on the frontier (for $\mu = 0.15$ and $N = 50$), determined by applying the corresponding strategies to 10^5 randomly sampled paths of the stock price. Table 1 gives their VaR and CVar values for different levels of confidence. Especially for those values of V where the improvement is large, the adaptive distributions are very close to Gaussian. The static distribution functions are always Gaussian.

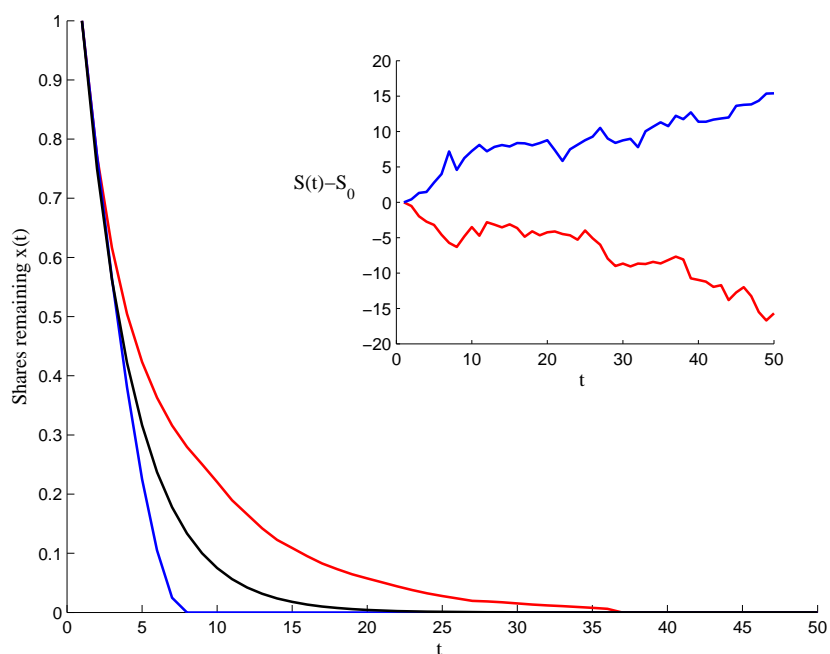


Figure 7: Optimal adaptive strategy for the point on the efficient frontier in Figure 5, having the same variance but lower expected cost than the static trajectory (solid black line), computed using 50 time steps. Specific trading trajectories are shown for two rather extreme realizations of the stock price process. The blue trajectory incurs impact costs that are comparable to the static trajectory, but has trading gains because it holds more stock as the price rises. The red trajectory has lower impact costs because of its slower trade rate, but it has trading losses because of the price decline. The mean and variance of the adaptive strategy cannot be seen in this picture, because they are properties of the entire ensemble of possible realizations.