Optimal Control of Execution Costs*

Dimitris Bertsimas[†] and Andrew W. Lo[‡]

Abstract

We derive dynamic optimal trading strategies that minimize the expected cost of trading a large block of equity over a fixed time horizon. Specifically, given a fixed block $\overline{S}$ of shares to be executed within a fixed finite number of periods T, and given a price-impact function that yields the execution price of an individual trade as a function of the shares traded and other "state" variables, we obtain the optimal *sequence* of trades as a function of the state variables—closed-form expressions in some cases—that will minimize the expected cost of executing $\overline{S}$ within T periods. We also present best-execution strategies under shortsales constraints, with more general price-impact functions, and for multiple assets with arbitrary correlation, and we provide an analysis of the risks of best execution.

*We thank Robert Ferstenberg for stimulating our interest in this area, and Petr Adamek for excellent research assistance. We are grateful to Dave Cushing, Chris Darnell, Rohit D'Souza, John Heaton, Holly Horrigan, Greg Peterson, Jiang Wang, and seminar participants at Harvard, MIT, and the Society of Quantitative Analysts for helpful comments and discussion. This research was partially supported by the MIT Laboratory for Financial Engineering, an Alfred P. Sloan Research Fellowship, and a Presidential Young Investigator Award DDM-9158118 with matching funds from Draper Laboratory.

[†]Professor of Operations Research, Sloan School of Management, MIT, Cambridge, MA 02142–1347.

[‡]Harris & Harris Group Professor, Sloan School of Management, MIT, Cambridge, MA 02142–1347.

Optimal Control of Execution Costs

Abstract

We derive dynamic optimal trading strategies that minimize the expected cost of trading a large block of equity over a fixed time horizon. Specifically, given a fixed block $\overline{S}$ of shares to be executed within a fixed finite number of periods T, and given a price-impact function that yields the execution price of an individual trade as a function of the shares traded and other "state" variables, we obtain the optimal *sequence* of trades as a function of the state variables—closed-form expressions in some cases—that will minimize the expected cost of executing $\overline{S}$ within T periods. We also present best-execution strategies under shortsales constraints, with more general price-impact functions, and for multiple assets with arbitrary correlation, and we provide an analysis of the risks of best execution.

Contents

1 Introduction

The tremendous growth in equity trading over the past 20 years, fueled largely by the burgeoning assets of institutional investors such as mutual and pension funds, has created a renewed interest in the measurement and management of trading costs.[1] Such costs—often called "execution costs" because they are associated with the execution of investment strategies—include commissions, bid/ask spreads, opportunity costs (of waiting), and price impact (from trading) [see Loeb (1983) and Wagner (1993) for further discussion], and they can have a substantial impact on investment performance. For example, Pérold (1988) observes that a hypothetical or "paper" portfolio constructed according to the Value Line rankings outperforms the market by almost 20% per year during the period from 1965 to 1986, whereas the actual portfolio (the Value Line Fund) outperformed the market by only 2.5% per year, the difference arising from execution costs.[2] This "implementation shortfall" is surprisingly large and underscores the importance of execution cost control, particularly for institutional investors whose trades often comprise a large fraction of the average daily volume of many stocks.

Indeed, recent studies by Chan and Lakonishok (1995) and Keim and Madhavan (1995a-c) show that because the typical institutional investor's trades are so large, they are almost always broken up into smaller trades executed over the course of several days. Chan and Lakonishok call such sequences "packages" and using a sample of 1.2 million transactions of

[1]There is a large and diverse literature on transactions costs in economics and finance. Niehans (1987) provides an excellent review of the transactions costs literature in economics.

The impact of transactions costs on financial decisionmaking has been considered by Aiyagari and Gertler (1991), Bensaid et al. (1992), Boyle and Vorst (1992), Cohen et al. (1981), Davis and Norman (1991), Dumas and Luciano (1991), Epps (1976), Garman and Ohlson (1981), Grossman and Miller (1988), Grossman and Vila (1992), Heaton and Lucas (1994, 1995), Hodges and Neuberger (1989), Litzenberger and Rolfo (1984), Leland (1985), Magill and Constantinides (1976), Tuckman and Vila (1992), and many others.

Empirical studies of transactions costs in financial markets include Berkowitz, Logue, and Noser (1988), Birinyi (1995), Chan and Lakonishok (1993, 1995), Demsetz (1968), Hasbrouck and Schwartz (1988), Huang and Stoll (1995), Keim and Madhavan (1995a-c), Kraus and Stoll (1972), Schwartz and Whitcomb (1988), Stoll (1989, 1993), Tiniç (1972), and Turnbull and White (1995).

And the more practical aspects of transactions cost management have been explored by Arnott and Wagner (1990), Bodurtha and Quinn (1990), Brinson, Hood, and Beebower (1986), Brinson, Singer, and Beebower (1991), Collins and Fabozzi (1991), Gammill and Pérold (1989), Loeb (1983), Pérold (1988), Treynor (1981), Wagner (1975), Wagner and Banks (1992), Wagner and Edwards (1993), and the papers in Sherrerd (1993).

[2]See also Bodurtha and Quinn (1990), Brinson, Hood, and Beebower (1986), Brinson, Singer, and Beebower (1991), Chan and Lakonishok (1993, 1995), Hasbrouck and Schwartz (1988), Keim and Madhavan (1995a-c), Loeb (1983), Treynor (1981), and Turnbull and White (1995).

37 large investment management firms during the period from July 1986 to December 1988, they show that only 20% of the market value of these packages are completed within a day and that over 53% are spread over four trading days or more [Chan and Lakonishok (1995, Table I)].

In this paper we provide an explicit economic rationale for trading in packages by deriving dynamic optimal trading strategies that minimize the expected costs of execution over a fixed time horizon. Specifically, we solve the following problem: given a fixed block $\overline{S}$ of shares to be executed within a fixed finite number of periods T, and given a price-impact function that yields the execution price of an individual trade as a function of the shares traded and other "state" variables, find the optimal *sequence* of trades [as a function of the state variables] that will minimize the expected cost of executing $\overline{S}$ within T periods.[3]

Using stochastic dynamic programming and nonlinear optimization, we obtain explicit closed-form expressions for these optimal trading strategies, which we call "best-execution" strategies, for a variety of price-impact functions. We show that best-execution strategies can sometimes be expressed as the combination of a naive (and not uncommon) strategy—breaking up $\overline{S}$ shares evenly into a package of T trades each of size $\overline{S}/T$—and a correction factor that adjusts each trade size up or down according to additional information, e.g., stock-specific private information, opportunity cost, changing market conditions, etc. In the absence of such information, we derive conditions under which the naive strategy is optimal: an arithmetic random walk for prices with linear price impact and no shortsales constraints [see Section 2.3]. We also show by construction that, apart from these rather restrictive and empirically implausible conditions, the naive strategy is not optimal in general.

We also obtain the expected cost of best execution—the optimal value function which is given implicitly by the Bellman equation—as a by-product of the optimization algorithm, which may serve as a useful benchmark for pricing principal bid and negotiated block transactions. The typical broker/dealer engaging in such transactions will not willingly hold large positions for long, and will seek to trade out of these positions as quickly and as cost-effectively as possible, i.e., they will seek best-execution strategies for their holdings. Of

[3]A complementary set of results concerns the selection of the particular *type* of trade—limit versus market order—in a dynamic context which we do not address in this paper, but which is the focus of Angel (1994), Bernhardt and Hughson (1994), Harris (1994), Hasbrouck and Harris (1992), Kumar and Seppi (1993).

course, risk aversion, adverse selection, and inventory and opportunity costs may change the objective function to be minimized, in which case our benchmark may only be a lower bound on the fair market price of the block transaction. Nevertheless, even in these cases the problem of best execution is still a dynamic optimization problem and our approach is still applicable [although closed-form expressions for best-execution strategies may not be available].

Indeed, our basic approach—described in Sections 2 and 3—can be extended in a number of important ways: (1) trading a portfolio of stocks simultaneously [Section 4]; (2) imposing non-negativity constraints [Section 5]; (3) specifying more general price-impact functions and deriving numerical solutions [Section 6]; and (4) analyzing the risks of best execution [Section 7]. These results comprise a systematic and quantitative approach to controlling execution costs, measuring the liquidity of large-block transactions, and rationalizing within an economic paradigm the kind of trading practices that characterize most institutional equity investors.

2 The Basic Problem

Consider an investor seeking to acquire a large block of $\overline{S}$ shares of some stock over a fixed time interval $[0, T]$.[4] Since it is well-known that the short-term demand curves for even the most actively traded equities are not perfectly elastic, a market order at date 0 for the entire block $\overline{S}$ is clearly not an optimal trading strategy.[5] A more effective strategy would be to break $\overline{S}$ into smaller purchases distributed throughout the interval $[0, T]$, but how should such purchases be parceled out?

The answer depends, of course, on the degree to which a purchase affects the market price, i.e., the reaction or "price-impact" function. Given a particular price-impact function, e.g., the ordered probit model of Hausman, Lo, and MacKinlay (1992), a dynamic optimal trading strategy that minimizes the expected total cost of acquiring $\overline{S}$ in $[0, T]$ may be obtained by

[4] For simplicity, we consider buying $\overline{S}$ only—selling $\overline{S}$ is symmetric and the solution follows directly from our analysis below with only minor changes in notation.

[5] See, for example, Chan and Lakonishok (1993, 1995), Hausman, Lo, and MacKinlay (1992), Holthausen, Leftwich, and Mayers (1987, 1990), and Keim and Madhavan (1995b,c).

stochastic dynamic programming and related techniques in nonlinear optimization.

2.1 Defining Best Execution

To illustrate this approach, suppose that at time 0 the investor begins his program to acquire $\overline{S}$ shares, and this program must be completed by time T. With little loss in generality, let time be measured in discrete intervals of unit length. Since the length of a "period" is arbitrary, it can be set to accommodate even the finest trading decision interval that is of practical relevance. For example, if the decision to acquire $\overline{S}$ is made at the start of the day and the acquisition must be completed by the day's end, setting $T = 13$ yields 30-minute intervals from the 9:30am market open to the 4:00pm market close. If the acquisition is part of an end-of-quarter portfolio rebalancing, the trading horizon may be extended to three or four days, in which case T increases proportionally.[6] Although all of our results are qualitatively independent of both the time horizon and the number of trading periods (with the exception of numerical examples, of course), for concreteness the length of each period should be regarded as some fraction of a single day and related parameters should be calibrated accordingly.

Denote by S_t be the number of shares acquired in period t at price P_t, where $t = 1, 2, \ldots, T$. Then the investor's objective of minimizing execution costs may be expressed as:

$$\underset{\{S_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} P_t S_t \right] \tag{2.1}$$

subject to the constraint

$$\sum_{t=1}^{T} S_t \ = \ \overline{S} \ . \tag{2.2}$$

[6]If 30-minute intervals are deemed too long for a given trading horizon, either because of volatility or opportunity cost, T can be increased and the parameters of the optimization problem rescaled accordingly. Of course, a more natural way to handle the timing issue is to re-cast this problem in continuous time. However, the analysis becomes considerably more complex and closed-form solutions are difficult to obtain [see, for example, Davis and Norman (1990) and Magill and Constantinides (1976)]. Moreover, for empirical applications continuous-time models must be discretized. Therefore, we take a discrete-time approach throughout this paper.

We may also wish to impose shortsales constraints, $S_t \geq 0$, but for expositional convenience we shall ignore these constraints for now and return to them in Section 5.

To complete the statement of the problem, we must specify the impact that a trade of S_t shares has on the execution price P_t, i.e., we require a price-impact function. For simplicity, suppose the price impact of a trade is proportional to its size, so that a purchase of S_t shares may be executed at the prevailing price P_{t-1} plus an impact premium of αS_t, $\alpha > 0$. Then the dynamics or "law of motion" for P_t may be expressed as:

$$P_t \;=\; P_{t-1} \;+\; \alpha S_t \;+\; \epsilon_t \quad , \quad \alpha > 0 \;\; , \;\; \mathrm{E}[\epsilon_t | S_t, P_{t-1}] = 0 \tag{2.3}$$

where ϵ_t is assumed to be a zero-mean independently and identically distributed (IID) random shock, i.e., white noise.[7] Despite the fact that (2.3) has some implausible empirical implications—independent price increments, positive probability of negative prices, percentage price impact that decreases with price, "permanent" price impact, etc.—it provides a concrete illustration of the more general and considerably more complex analysis which we shall present in later sections. Moreover, we shall see in Section 2.3 that (2.3) is precisely the dynamics necessary to render the naive strategy of dividing $\overline{S}$ into T trades each of size $\overline{S}/T$ the optimal one.

The investor's problem is now well-posed: find the *sequence* of trades $\{S_t\}$ that minimizes the expected execution costs $\mathrm{E}_0[\sum_{t=1}^{T} P_t S_t]$, subject to the constraint that $\sum_{t=1}^{T} S_t = \overline{S}$, and given a linear price-impact function incorporated into the law of motion (2.3) for P_t. This is a classical optimal control problem which can be solved by stochastic dynamic programming, and we define the "best-execution strategy" as its solution.

2.2 The Bellman Equation

The basic ingredients for any dynamic programming problem are: the state of the environment at time t, the control variable, the randomness, the cost function, and the law of motion. In our context, the state at time $t = 1, \ldots, T$ consists of the price P_{t-1} realized at

[7] Although we refer to (2.3) as a price-impact function, it is a simple matter to modify it include other components of execution costs such as commissions, a bid/ask spread, and opportunity cost. However, for notational simplicity we shall focus solely on price impact.

the previous period, and W_t the number of shares that remain to be purchased. The state variables summarize all the information the investor requires in each period t to make his decision regarding the control. The control variable at time t is the number of shares S_t purchased. The randomness is characterized by the random variable ϵ_t. The objective is given by (2.1), while the law of motion is given by (2.3) and

$$W_t \;=\; W_{t-1} \;-\; S_{t-1} \qquad , \qquad W_1 = \overline{S} \quad , \quad W_{T+1} = 0, \tag{2.4}$$

where the boundary condition $W_{T+1} = 0$ is equivalent to the constraint that $\overline{S}$ must be executed by period T.

The dynamic programming algorithm is based on the observation that a solution or "optimal control" $\{S_1^*, S_2^*, \ldots, S_T^*\}$ must also optimal for the remaining program at every intermediate date t. That is, for every t, $1<t<T$ the sequence $\{S_t^*, S_{t+1}^*, \ldots, S_T^*\}$ must still be optimal for the remaining program $\mathrm{E}_t[\sum_{k=t}^T P_t S_t]$. This important property is summarized by the Bellman equation (2.5), which relates the optimal value of the objective function in period t to its optimal value in period $t+1$:

$$V_t(P_{t-1}, W_t) \;=\; \underset{S_t}{\mathrm{Min}}\;\mathrm{E}_t\left[\;P_t S_t \;+\; V_{t+1}(P_t, W_{t+1})\;\right]. \tag{2.5}$$

By starting at the end (time T) and applying the Bellman equation (2.5) and the law of motion for P_t (2.3) and W_t (2.4) recursively, the optimal control can be derived as functions of the state variables which characterize the information that the investor must have to make his decision in each period.

Starting at the end ($t = T$), the optimal value function, as a function of the two state variables P_{T-1} and W_T, is given by:

$$V_T(P_{T-1}, W_T) \;=\; \underset{S_T}{\mathrm{Min}}\;\mathrm{E}_T[P_T W_T] \;=\; (P_{T-1} + \alpha W_T)W_T\;. \tag{2.6}$$

Since this is the last period and W_{T+1} must be set to 0, there is no choice but to execute the entire remaining order W_T, hence the optimal trade size S_T^* is simply W_T. Substituting the law of motion (2.3) into $P_T W_T$ yields V_T as a function of P_{T-1} and W_T.

In the next-to-last period $T-1$, the Bellman equation is less trivial:

$$V_{T-1}(P_{T-2}, W_{T-1}) = \underset{S_{T-1}}{\text{Min}} \ \mathrm{E}_{T-1}\left[P_{T-1}S_{T-1} \ + \ V_T(P_{T-1}, W_T) \right] \tag{2.7}$$

$$= \underset{S_{T-1}}{\text{Min}} \ \mathrm{E}_{T-1}\Big[(P_{T-2} + \alpha S_{T-1} + \epsilon_{T-1})S_{T-1} + V_T\Big(P_{T-2} + \alpha S_{T-1} + \epsilon_{T-1}, W_{T-1} - S_{T-1}\Big) \Big] . \tag{2.8}$$

By substituting the right side of (2.6) into (2.8), and substituting for W_T and P_{T-1} using (2.3) and (2.4), respectively, the right side of (2.8) may be cast as an explicit function of S_{T-1} which can be minimized by taking its derivative with respect to S_{T-1} and solving for its zero. This yields:

$$S^*_{T-1} = \frac{W_{T-1}}{2} \tag{2.9}$$

$$V_{T-1}(P_{T-2}, W_{T-1}) = W_{T-1}(P_{T-2} + \frac{3}{4}\alpha W_{T-1}), \tag{2.10}$$

where (2.10) is obtained by substituting S^*_{T-1} into (2.8).

Continuing in this fashion, the optimal trades S^*_{T-k} and the optimal value function $V_{T-k}(P_{T-k-1}, W_{T-k})$ may be obtained recursively as:

$$S^*_{T-k} = \frac{W_{T-k}}{k+1} \tag{2.11}$$

$$V_{T-k}(P_{T-k-1}, W_{T-k}) = W_{T-k}(P_{T-k-1} + \frac{k+2}{2(k+1)}\alpha W_{T-k}), \tag{2.12}$$

until we reach the beginning of the program and find:

$$S^*_1 = \frac{W_1}{T} \tag{2.13}$$

$$V_1(P_0, W_1) = W_1(P_0 + \frac{T+1}{2T}\alpha W_1). \tag{2.14}$$

2.3 The Best Execution Strategy

Substituting the initial conditional $W_1 = \overline{S}$ into (2.13) and (2.14) then yields the optimal trade size S_1^* as an explicit function of $\overline{S}$, and the expected best-execution cost V_1 as an explicit function of $\overline{S}$, P_0, and the price-impact parameter α:

$$S_1^* \quad = \quad \frac{\overline{S}}{T} \tag{2.15}$$

$$V_1(P_0, W_1) \quad = \quad \mathrm{E}_1 \left[\, P_t S_t^* \, \right] \;\; P_0 \overline{S} \; + \; \frac{\alpha \overline{S}^2}{2} \left(1 + \frac{1}{T} \right) . \tag{2.16}$$

By forward substitution using (2.15), (2.11), and (2.4), we find that:

$$S_1^* \;=\; S_2^* \;=\; \cdots \;=\; S_T^* \;=\; \frac{\overline{S}}{T} \, . \tag{2.17}$$

The best-execution strategy is simply to divide the total order $\overline{S}$ into T equal "waves" and trade them at regular intervals.

This remarkably simple trading strategy comes from the fact that the price impact αS_t does not depend on either the prevailing price P_{t-1} or the size of the unexecuted order W_t, hence the price-impact function is the same in each period and independent from one period to the next. But since each period's execution cost $P_t S_t$ is a convex (quadratic) function of S_t, the sum of these single-period execution costs will be minimized at the point where the marginal execution costs are equated across all periods. There is no advantage to tilting trades to one period or another—they all offer the same trade-offs to the objective function—hence the trade sizes should be set equal across all periods. Note that in this case the optimal controls $\{S_t^*\}$ are all non-negative hence the non-negativity constraints could have been imposed trivially.

The optimal value function at time 1, $V_1(P_0, W_1)$, gives the expected cost of the best-execution strategy and we see from (2.16) that this cost is the sum of two terms: the no-impact cost $P_0 \overline{S}$ and the cumulative price impact $\frac{\alpha \overline{S}^2}{2}$. Observe that while the impact term is a decreasing function of T—having more time to acquire $\overline{S}$ can never increase the expected cost—the cumulative price impact does not vanish as T increases without bound. This seems counterintuitive since one might expect price impact to become negligible if there is no time limit on completing the purchase. However, observe that our law of motion for P_t (2.3) implies that the price impact αS_t of an individual trade has a *permanent* effect on P_t, hence even infinitesimally small trades will have an impact on next period's price, and the limiting sum of all these infinitesimal trades multiplied by infinitesimally increased prices is finite and non-zero: $\alpha\overline{S}/2$.

These results undescore the importance of the law of motion's specification in determining the total expected cost of executing $\overline{S}$. Of course, (2.3) is empirically implausible for a number of reasons. However, it serves a useful purpose in demonstrating the basic approach to best execution, as well as in rationalizing the rather common practice of parceling a large trade into smaller pieces of equal size and submitting them at regular intervals over some fixed time span. This naive strategy is indeed optimal if the price-impact function and law of motion of P_t is given by (2.3).

In the next section we expand the power and scope of the dynamic we present a closed-form solution for the best-execution strategy under a more complex price-impact function, one which depends both on the trade size and a serially-correlated state variable that proxies for private information or market conditions. With information, the best-execution strategy differs in important ways from the naive strategy $S_t^* = \overline{S}/T$. In particular, the best-execution strategy becomes a nontrivial function of the information variable and can sometimes exhibit counterintuitive trading patterns.

2.4 Linear Price Impact with Information

Suppose that the price-impact function is linear in S_t as in (2.3), but now let X_t denote a serially-correlated state variable which also affects the execution price P_t linearly, hence:

$$P_t = P_{t-1} + \alpha S_t + \gamma X_t + \epsilon_t \quad , \quad \alpha > 0 \tag{2.18}$$

$$X_t = \rho X_{t-1} + \eta_t \quad , \quad X_1 = 1 \quad , \quad \rho \in (-1, 1) \, . \tag{2.19}$$

where ϵ_t and η_t are independent white noise processes with mean 0 and variances σ_ϵ^2 and σ_η^2, respectively.

The presence of X_t in the law of motion for P_t captures the potential impact of changing market conditions or of private information about the security. For example, X_t might be the return on the S&P 500 index, a common component in the prices of most equities. A market rise affects all securities to some degree, and γ measures the sensitivity of this particular security to such market movements.

Alternatively, X_t might represent private information about the security, and γ the importance of that information for P_t. In particular, X_t may denote the outcome of an "alpha model" which incorporates stock-specific analysis to yield an excess return not yet impounded into market prices.

In either case, the impact of X_t on the execution price, and the time series properties of X_t have important implications for the best-execution strategy. Having specified the linear price-impact function with information in (2.19) – (2.19), the best-execution strategy and optimal value function can be obtained by dynamic programming as before, and is given by:

$$S^*_{T-k} = \frac{W_{T-k}}{k+1} + \frac{\rho b_{k-1}}{2a_{k-1}} X_{T-k} \tag{2.20}$$

$$V_{T-k}(P_{T-k-1}, X_{T-k}, W_{T-k}) = P_{T-k-1}W_{T-k} + a_k W^2_{T-k} + b_k X_{T-k} W_{T-k} + c_k X^2_{T-k} + d_k \tag{2.21}$$

for $k = 0, 1, \ldots, T-1$, where:

$$a_k = \frac{\alpha}{2}\left(1 + \frac{1}{k+1}\right), \quad , \quad a_0 = \alpha \tag{2.22}$$

$$b_k = \gamma + \frac{\alpha\rho b_{k-1}}{2a_{k-1}} \quad , \quad b_0 = \gamma \tag{2.23}$$

$$c_k = \rho^2 c_{k-1} - \frac{\rho^2 b^2_{k-1}}{4a_{k-1}} \quad , \quad c_0 = 0 \tag{2.24}$$

$$d_k = d_{k-1} + c_{k-1}\sigma^2_\eta \quad , \quad d_0 = 0\,. \tag{2.25}$$

Since we assume $\alpha > 0$, a_k is positive, and c_k and d_k are negative for all $k > 0$. The sign of b_k can vary, but is positive for all $k > 0$ if α, γ, and ρ are all positive.

In contrast to the case of a linear price-impact function with no information, the best-execution strategy (2.20) varies over time as a linear function of the remaining shares W_{T-k} and the information variable X_{T-k}. In particular, the first term of (2.20) is simply the naive strategy of dividing the remaining shares W_{T-k} at time $T-k$ evenly over the remaining $k+1$ periods.

The second term of (2.20) is an adjustment that arises from the presence of serially-correlated information X_{T-k}. Observe that this term vanishes if $\rho = 0$. When $\rho = 0$ this implies that X_{T-k} is unforecastable, and while X_{T-k} still has an impact on the current execution price, observing X_{T-k} tells us nothing about expected *future* execution prices hence it can no longer affect the best-execution strategy.

If $\rho > 0$ and we assume, without loss of generality, that $\gamma > 0$, then the coefficient multiplying X_{T-k} in (2.20) is also positive, implying that positive realizations of X_{T-k} increases the number of shares purchased at $T-k$, *ceteris paribus*. This may seem counterintuitive at first because a positive X_{T-k} necessarily increases the execution price P_{T-k} by γX_{T-k}, so why trade more?

The answer may be found in the fact that X_{T-k} is positively serially correlated, hence $X_{T-k} > 0$ implies that future realizations are likely to be positive which, in turn, implies additional expected increases in future execution prices.[8] Therefore, although a positive X_{T-k} makes it more costly to purchase shares in period $T-k$, this additional cost is more than offset by the the *sequence* of expected future price increases that arise from positively serially-correlated information.

Alternatively, if $\rho < 0$ so that X_{T-k} exhibits reversals, (2.20) shows that a positive realization of X_{T-k} *decreases* the number of shares purchased, *ceteris paribus*: it is more expensive to trade in period $T-k$ and X_{T-k} is likely to reverse next period making it less expensive to trade then, hence it is optimal to trade less now.

The impact of an increase in X_{T-k} on expected best-execution costs may be measured explicitly by the derivative of the value function V_{T-k} with respect to X_{T-k}:

$$\frac{\partial V_{T-k}}{\partial X_{T-k}} \;=\; b_k W_{T-k} \;+\; 2 c_k X_{T-k} \; . \tag{2.26}$$

Suppose γ and ρ are positive so that b_k is positive. Since c_k is always negative, the impact of an increase in X_{T-k} on the expected best-execution cost depends on whether $b_k W_T$ outweighs

[8] In particular, a one-unit increase in X_{T-k} implies an expected increase in P_{T-k+1} of $\gamma\rho$, an expected increase in P_{T-k+2} of $\gamma\rho^2$, and so on.

$2c_k X_{T-k}$. For empirically plausible parameter values, the first term will generally dominate the second, hence increases in X_{T-k} will typically increase the expected best-execution cost, a sensible implication given that an increase in X_{T-k} increases current and all future expected prices.

It is also not surprising that V_{T-k} is an increasing function of W_{T-k} for empirically plausible parameter values—the larger is the unexecuted portion of the initial block, the higher the expected best-execution cost. In the next section, we provide a numerical example to illustrate the behavior of the best-execution strategy under several simulated scenarios.

2.5 A Numerical Example

Tables 1 and 2 provide illustrative numerical examples of the best-execution strategies under the linear price-impact function with information (2.19) – (2.19) for two simulated realizations of the information variable X_t and pricing shocks ϵ_t. The goal is to minimize the expected execution costs of a 100,000-share purchase over $T = 20$ periods for a stock currently trading at $P_0 = \$50$, given the following parameter values:

$$\begin{aligned}
\alpha &= 5 \times 10^{-5} \\
\gamma &= 5.0 \\
\rho &= 0.50 \\
\sigma_\epsilon^2 &= (0.125)^2 \\
\sigma_\eta^2 &= 0.001 \ .
\end{aligned}$$

To develop some intuition for these parameters, observe that the no-impact cost of acquiring $\overline{S}$ is $100{,}000 \times P_0 = \$5\text{M}$, and the full-impact cost is $100{,}000 \times P(\overline{S}) = 100{,}000 \times \$55 = \$5.5\text{M}$, hence α is calibrated to yield an impact of \$500,000 on a 100,000-share block purchase. From (2.19) it is apparent that $\mathrm{E}[X_t] = 0$ and

$$\sqrt{\mathrm{Var}[\gamma X_t]} \quad = \quad \frac{\gamma \sigma_\eta}{\sqrt{1-\rho^2}} \quad = \quad 0.183$$

hence the standard deviation of the information component is approximately 18 cents (per period). Finally, the standard deviation of ϵ_t is calibrated to be 12.5 cents or one "tick" (per period).

For the first simulated realization the expected best-execution cost is \$5,269,658, which compares favorably with the expected cost of \$5,273,469 for the naive strategy,[9] implying an improvement of 3.8 cents/share on average. However, since the information variable X_t is stochastic, and since the best-execution strategy S_t^* is a function of X_t, the actual execution cost of the best-execution strategy will not typically be equal to its expected value. For this realization, the actual cost of the best-execution strategy is \$5,353,635 and the actual cost of the naive strategy is \$5,368,560, yielding an improvement of 14.9 cents/share.

Of course, while the expected best-execution cost is guaranteed to be lower than the expected cost of all other feasible strategies, there is no assurance that the best-execution strategy will *always* yield the lowest execution cost—Table 2 provides another simulated realization for which the *actual* cost of the naive strategy is lower than that of the best-execution strategy, despite the fact that the best-execution strategy possesses a lower *expected* cost.

Tables 1 and 2 also demonstrate the adaptive nature of the best-execution strategy. For example, Table 1 shows that in period 2, an unusually large trade of 11,278 shares was executed. The columns labelled "$\delta_{w,t}W_t$" and "$\delta_{x,t}X_t$" report the breakdown of this 11,278-share trade into two components, one due to the non-information-related motive for trade, and the other due to the impact of the information contained in the state-variable X_t for the expected price. In period 2, there was an unusually large realization for X_t implying higher expected prices in the future (since X_t is positively autocorrelated), hence the information-based component was large. However, in the very next period, an unusually low realization for X_t occurs, hence the information-based component is negative. Figures 1 and 2 provide a clearer illustration of the adaptive nature of the best-execution strategy.

---

[9]The expected cost of the naive strategy under (2.19) – (2.19) is given by:

$$\mathrm{E}_1\left[\,\sum_{t=1}^{T} P_t\overline{S}/T\,\right] \;=\; P_0\overline{S} \;+\; \alpha\left(\frac{\overline{S}}{T}\right)^2\frac{T(T+1)}{2} \;+\; \frac{\overline{S}}{T}\frac{\gamma X_1}{1-\rho}\left(T-\frac{\rho-\rho^{T+1}}{1-\rho}\right).$$

Performance Summary			
Metric	Naive Strategy ($1000's)	Optimal Strategy ($1000's)	Improvement (cents/share)
Expected Cost	5,273.5	5,269.7	3.8108
Actual Cost	5,368.6	5,353.6	14.9251

Period	P_t	S_t^*	$\delta_{w,t}W_t$	$\delta_{x,t}X_t$	γX_t	$V_t/\$1000$	Cost-To-Go
1	50.197	6,039	5,000	1,039	0.06	5,269.7	5,353.6
2	51.167	11,278	4,945	6,332	0.35	5,006.9	5,050.5
3	51.187	2,204	4,594	–2,390	–0.13	4,386.6	4,473.5
4	51.387	3,660	4,734	–1,074	–0.06	4,278.7	4,360.6
5	51.605	5,453	4,801	652	0.04	4,106.9	4,172.5
6	51.905	5,204	4,758	446	0.03	3,819.7	3,891.1
7	52.182	5,298	4,726	572	0.03	3,553.3	3,621.0
8	52.663	5,400	4,682	718	0.04	3,278.6	3,344.6
9	53.301	6,861	4,622	2,239	0.13	3,016.1	3,060.2
10	53.387	3,787	4,418	–631	–0.04	2,650.2	2,694.5
11	53.749	4,453	4,482	–28	–0.00	2,446.5	2,492.3
12	54.348	7,385	4,485	2,900	0.19	2,226.8	2,252.9
13	55.235	9,460	4,122	5,338	0.36	1,841.6	1,851.6
14	55.874	7,918	3,360	4,558	0.32	1,326.3	1,329.0
15	56.453	5,682	2,600	3,082	0.23	883.9	886.6
16	56.954	5,519	1,984	3,535	0.29	566.7	565.9
17	57.175	111	1,100	–989	–0.09	250.4	251.6
18	57.287	1,095	1,429	–334	–0.04	245.5	245.2
19	57.072	1,026	1,596	–570	–0.11	182.8	182.5
20	57.202	2,167	2,167	0	0.11	123.9	123.9

Table 1: Best-execution strategies for acquiring a $100,000$-share block of stock within 20 periods, for a simulated realization of the information variable $\{X_t\}$ and price shock $\{\epsilon_t\}$, assuming an initial price $P_0 = \$50$ and linear price-impact function $\alpha S_t + \gamma X_t$, where $P_t = P_{t-1} + \alpha S_t + \gamma X_t + \epsilon_t$ and $\alpha = 5 \times 10^{-5}$, $\gamma = 5.0$, $X_t = \rho X_{t-1} + \eta_t$, $X_0 = 0$, $\sigma_\eta^2 = 0.001$, $\sigma_\epsilon^2 = (0.125)^2$, and $\rho = 0.50$. Note: The non-negativity restriction was not imposed.

Figure 1: Best-execution strategy and information-based component for Table 1.

Performance Summary			
Metric	Naive Strategy ($1000's)	Optimal Strategy ($1000's)	Improvement (cents/share)
Expected Cost	5,247.4	5,243.5	3.8451
Actual Cost	5,293.8	5,303.3	–9.4413

Period	P_t	S_t^*	$\delta_{w,t}W_t$	$\delta_{x,t}X_t$	γX_t	$V_t/\$1000$	Cost-To-Go
1	50.361	3,565	5,000	–1,435	–0.08	5,243.5	5,303.3
2	50.388	5,216	5,076	141	0.01	5,099.3	5,123.8
3	51.249	8,708	5,068	3,640	0.20	4,847.5	4,860.9
4	52.048	10,122	4,854	5,268	0.30	4,451.2	4,414.7
5	52.532	6,202	4,524	1,677	0.10	3,917.1	3,887.9
6	52.319	2,766	4,413	–1,647	–0.10	3,579.6	3,562.1
7	52.301	1,796	4,530	–2,735	–0.16	3,404.7	3,417.4
8	51.894	207	4,741	–4,533	–0.27	3,292.3	3,323.5
9	51.845	2,981	5,118	–2,137	–0.13	3,273.2	3,312.7
10	51.910	3,101	5,313	–2,211	–0.14	3,107.0	3,158.2
11	52.432	6,829	5,534	1,296	0.08	2,963.6	2,997.2
12	53.211	7,473	5,390	2,083	0.13	2,619.2	2,639.1
13	53.817	7,083	5,129	1,953	0.13	2,239.4	2,241.5
14	54.192	5,712	4,850	861	0.06	1,863.1	1,860.3
15	54.212	4,589	4,707	–118	–0.01	1,552.9	1,550.8
16	54.607	5,172	4,730	441	0.04	1,300.1	1,302.0
17	54.485	2,473	4,620	–2,147	–0.20	1,013.8	1,019.6
18	54.615	3,936	5,336	–1,400	–0.17	876.8	884.8
19	55.105	6,209	6,036	173	0.03	665.3	669.9
20	55.906	5,862	5,862	0	0.44	327.7	327.7

Table 2: Best-execution strategies for acquiring a $100,000$-share block of stock within 20 periods, for a simulated realization of the information variable $\{X_t\}$ and price shock $\{\epsilon_t\}$, assuming an initial price $P_0 = \$50$ and linear price-impact function $\alpha S_t + \gamma X_t$, where $P_t = P_{t-1} + \alpha S_t + \gamma X_t + \epsilon_t$ and $\alpha = 5 \times 10^{-5}$, $\gamma = 5.0$, $X_t = \rho X_{t-1} + \eta_t$, $X_0 = 0$, $\sigma_\eta^2 = 0.001$, $\sigma_\epsilon^2 = (0.125)^2$, and $\rho = 0.50$. Note: The non-negativity restriction was not imposed.

Figure 2: Best-execution strategy and information-based component for Table 2.

Such "whip-sawing" may seem to be a poor way to minimize execution costs, but since ρ is positive, the reversal in X_t that occurred between periods 2 and 3 is unusual. More common is the sequence of positive realizations that occurred between periods 5 and 9, and between 12 and 16. In these more typical sequences, trading "into" expected price increases can be more economical than the natural tendency to wait for a more favorable price.

Of course, the best-execution strategy depends intimately on the specific dynamics of the information variables that enter the price-impact function. In particular, if in our example ρ were negative, the reverse patterns would be observed—a large positive realization of X_t would imply lower prices in the future, hence the information component would cause trading to be shifted from today into the future. Table 3 and Figure 3 illustrate this intuition explicitly. In contrast to the best-execution strategies of Tables 1 and 2, here a large positive realization X_t—in period 5, for example—reduces the trade size for that period, and vice-versa for large negative realizations of X_t. Instead of trading into price rises, the negative autocorrelation in X_t implies that it is optimal to trade against them.

3 Linear Percentage Price Impact

The closed-form best-execution strategies derived in Section 2 were primarily illustrative examples, meant to introduce the overall approach and provide some intuition for the determinants and properties of best-execution strategies in a simple context. From a more practical perspective, those illustrative examples have several important limitations.

In particular, the fact that, aside from price impact, prices P_t are assumed to follow arithmetic random walks in (2.3) and (2.19) implies a positive probability of negative prices. Also, the random walk specification implies that both price impact and information have permanent effects on prices, which contradicts several recent empirical studies [see, for example, Barclay and Litzenberger (1988), Barclay and Warner (1993), Chan and Lakonishok (1993, 1995), and Holthausen, Leftwich, and Mayers (1987, 1990)]. And finally, because of the linearity in(2.3) and (2.19), the *percentage* price impact—as a percentage of the execution price—is a decreasing function of the price level, which is also counterfactual [see, for example, Birinyi (1995), Leinweber (1993, 1994), and Loeb (1983)].

Performance Summary			
Metric	Naive Strategy ($1000's)	Optimal Strategy ($1000's)	Improvement (cents/share)
Expected Cost	5,265.4	5,264.7	0.6
Actual Cost	5,239.8	5,237.8	2.1

Period	P_t	S_t^*	$\delta_{w,t}W_t$	$\delta_{x,t}X_t$	γX_t	V_t/$1000	Cost-To-Go
1	50.255	4,727	5,000	−273	0.04	5,264.7	5,237.8
2	50.508	5,284	5,014	270	−0.04	5,023.5	5,000.2
3	50.578	4,800	4,999	−199	0.03	4,760.2	4,733.4
4	50.416	7,874	5,011	2,863	−0.45	4,474.1	4,490.6
5	50.761	3,383	4,832	−1,449	0.23	4,068.1	4,093.6
6	51.086	5,586	4,929	657	−0.10	3,892.9	3,921.9
7	51.043	6,584	4,882	1,703	−0.27	3,603.6	3,636.5
8	51.436	4,162	4,751	−589	0.09	3,258.7	3,300.4
9	51.813	3,641	4,800	−1,159	0.18	3,059.4	3,086.3
10	51.830	7,486	4,905	2,581	−0.41	2,859.3	2,897.7
11	52.248	3,239	4,647	−1,408	0.23	2,474.9	2,509.7
12	52.463	5,733	4,804	929	−0.15	2,306.0	2,340.5
13	53.117	1,564	4,688	−3,123	0.51	2,019.7	2,039.7
14	53.454	6,192	5,134	1,058	−0.18	1,941.1	1,956.6
15	53.951	3,753	4,957	−1,204	0.20	1,619.8	1,625.6
16	54.173	6,122	5,198	924	−0.16	1,419.4	1,423.1
17	54.440	5,775	4,967	807	−0.14	1,086.5	1,091.5
18	54.744	4,003	4,698	−696	0.14	775.3	777.1
19	55.120	5,182	5,046	136	−0.03	556.1	558.0
20	55.474	4,910	4,910	0	0.11	272.4	272.4

Table 3: Best-execution strategies for acquiring a 100,000-share block of stock within 20 periods, for a simulated realization of the information variable $\{X_t\}$ and price shock $\{\epsilon_t\}$, assuming an initial price $P_0 = \$50$ and linear price-impact function $\alpha S_t + \gamma X_t$, where $P_t = P_{t-1} + \alpha S_t + \gamma X_t + \epsilon_t$ and $\alpha = 5 \times 10^{-5}$, $\gamma = 5.0$, $X_t = \rho X_{t-1} + \eta_t$, $X_0 = 0$, $\sigma_\eta^2 = 0.001$, $\sigma_\epsilon^2 = (0.125)^2$, and $\rho = 0.50$. Note: The non-negativity restriction was not imposed.

Figure 3: Best-execution strategy and information-based component for Table 3.

For these reasons, we propose an alternative to (2.3) and (2.19) which we call the "linear percentage price-impact" function. Specifically, let the execution price at time t be comprised of two components, a no-impact price $\tilde{P}_t$, and the price impact Δ_t:

$$P_t \quad = \quad \tilde{P}_t \; + \; \Delta_t \; . \tag{3.1}$$

The no-impact price may be taken to be the midpoint of the bid/offer spread (although it is, in fact, arbitrary as long as it is unaffected by the trade size S_t). For convenience, we shall adopt the standard geometric random walk specification for $\tilde{P}_t$:

$$\tilde{P}_t \quad = \quad \tilde{P}_{t-1} \exp(Z_t) \qquad , \qquad Z_t \sim \mathrm{N}(\mu, \sigma_z^2) \; . \tag{3.2}$$

The price impact Δ_t captures the effect of trade size S_t on the transaction price (including the bid/offer spread), and we shall assume that as a percentage of the no-impact price $\tilde{P}_t$, it is a linear function of the trade size S_t and X_t where, as before, X_t proxies for private information or market conditions, hence:

$$\Delta_t \quad = \quad (\alpha S_t + \gamma X_t)\tilde{P}_t \tag{3.3}$$

$$X_t \quad = \quad \rho X_{t-1} \; + \; \eta_t \qquad , \qquad \eta_t \sim \mathrm{WN}(0, \sigma_\eta^2) \; . \tag{3.4}$$

Once again we have set X_t to be an AR(1) process which allows us to capture varying degrees of predictability in information or market conditions. The parameters α and γ measure the sensitivity of price-impact to trade size and market conditions.

The problem of best execution then reduces to the following stochastic dynamic program:

$$\underset{\{S_t\}}{\mathrm{Min}}\, \mathrm{E}_1 \left[\sum_{t=1}^{T} P_t S_t \right] \quad = \quad \underset{\{S_t\}}{\mathrm{Min}}\, \mathrm{E}_1 \left[\sum_{t=1}^{T} \tilde{P}_t (1 + \alpha S_t + \gamma X_t) S_t \right] \tag{3.5}$$

$$\quad = \quad \underset{\{S_t\}}{\mathrm{Min}} \left\{ \mathrm{E}_1 \left[\sum_{t=1}^{T} \tilde{P}_t S_t \right] \; + \; \mathrm{E}_1 \left[\sum_{t=1}^{T} \Delta_t S_t \right] \right\} . \tag{3.6}$$

We consider the unconstrained problem here and obtain the best-execution strategy via dynamic programming. The case of nonnegativity constraints is addressed in Section 5 by reducing the problem to a static quadratic optimization problem that can be solved very efficiently.

3.1 The Best Execution Strategy

We assume that at time t, the no-impact price $\tilde{P}_t$, the market condition X_t and the unexecuted portion of the trade W_t are observable. In this subsection we do not impose the nonnegativity constraint. As usual, we begin at the end by characterizing the optimal value function V_T and then proceed recursively:

$$V_T(\tilde{P}_T, X_T, W_T) = \underset{S_T}{\text{Min}}\, \text{E}_T[P_T S_T] = \text{E}_T[P_T W_T] \tag{3.7}$$

$$= \tilde{P}_T(W_T + \gamma X_T W_T + \alpha W_T^2) \tag{3.8}$$

In the next-to-last period, the optimal value function satisfies:

$$V_{T-1} = \underset{S_{T-1}}{\text{Min}}\ \text{E}_{T-1}\left[\ P_{T-1}S_{T-1} + V_T(\tilde{P}_T, X_T, W_T)\right] \tag{3.9}$$

$$= \underset{S_{T-1}}{\text{Min}}\ \Big\{\ \tilde{P}_{T-1}(1 + \alpha S_{T-1} + \gamma X_{T-1})S_{T-1} +$$

$$q\tilde{P}_{T-1}(W_{T-1} - S_{T-1})\Big[\ 1 + \gamma\rho X_{T-1} + \alpha(W_{T-1} - S_{T-1})\ \Big]\Big\} \tag{3.10}$$

where

$$q \equiv \text{E}_{T-1}[\exp(Z_T)] = \exp\left(\mu + \frac{\sigma_z^2}{2}\right). \tag{3.11}$$

The optimal trade size at $T-1$ is then:

$$S^*_{T-1} = \gamma_{x,1}X_{T-1} + \gamma_{w,1}W_{T-1} + \gamma_{1,1} \tag{3.12}$$

where

$$\gamma_{x,1} = \frac{\gamma(q\rho - 1)}{2\alpha(q+1)} \quad , \quad \gamma_{w,1} = \frac{q}{q+1} \quad , \quad \gamma_{1,1} = \frac{q-1}{2\alpha(q+1)} \tag{3.13}$$

Therefore, the optimal value function V_{T-1} is given by:

$$V_{T-1}(\tilde{P}_{T-1}, X_{T-1}, W_{T-1}) \quad = \quad \tilde{P}_{T-1}\Big[\, a_1 \; + \; b_1 X_{T-1} \; + \; c_1 X_{T-1}^2 \; + \\ d_1 X_{T-1} W_{T-1} \; + \; e_1 W_{T-1} \; + \; f_1 W_{T-1}^2 \,\Big] \tag{3.14}$$

where

$$a_1 \quad = \quad \gamma_{1,1}(1 + \alpha\gamma_{1,1}) \; - \; q\gamma_{1,1}(1 - \alpha\gamma_{1,1}) \tag{3.15}$$

$$b_1 \quad = \quad \gamma_{1,1}(\alpha\gamma_{x,1} + \gamma) \; + \; \gamma_{x,1}(1 + \alpha\gamma_{1,1}) \; - \; q\gamma_{x,1}(1 - \alpha\gamma_{1,1}) \; - \\ q\gamma_{1,1}(\gamma\rho - \alpha\gamma_{x,1}) \tag{3.16}$$

$$c_1 \quad = \quad \gamma_{x,1}(\alpha\gamma_{x,1} + \gamma) \; - \; q\gamma_{x,1}(\gamma\rho - \alpha\gamma_{x,1}) \tag{3.17}$$

$$d_1 \quad = \quad \gamma_{w,1}(\alpha\gamma_{x,1} + \gamma) \; + \; \alpha\gamma_{w,1}\gamma_{x,1} \; - \; \alpha q\gamma_{x,1}(1 - \gamma_{w,1}) \; + \\ q(1 - \gamma_{w,1})(\gamma\rho - \alpha\gamma_{x,1}) \tag{3.18}$$

$$e_1 \quad = \quad \gamma_{1,1}\alpha\gamma_{w,1} \; + \; \gamma_{w,1}(1 + \alpha\gamma_{1,1}) \; + \; q(1 - \gamma_{w,1})(1 - \alpha\gamma_{1,1}) \; - \\ \alpha q\gamma_{1,1}(1 - \gamma_{w,1}) \tag{3.19}$$

$$f_1 \quad = \quad \alpha\gamma_{w,1}^2 \; + \; \alpha q(1 - \gamma_{w,1})^2 \; . \tag{3.20}$$

Continuing in this fashion, we arrive at the following recursions:

$$V_{T-k} \quad = \quad \underset{S_{T-k}}{\mathrm{Min}} \; \mathrm{E}_{T-k}\Big[\, P_{T-k} S_{T-k} \; + \; V_{T-k+1}(\tilde{P}_{T-k+1}, X_{T-k+1}, W_{T-k+1}) \,\Big] \tag{3.21}$$

which yields the optimal trade sizes:

$$S_{T-k}^* \quad = \quad \gamma_{x,k} X_{T-k} \; + \; \gamma_{w,k} W_{T-k} \; + \; \gamma_{1,k} \tag{3.22}$$

where

$$\gamma_{x,k} = \frac{q(\rho d_{k-1} - \gamma)}{2(\alpha + q f_{k-1})} \quad , \quad \gamma_{w,k} = \frac{q f_{k-1}}{\alpha + q f_{k-1}} \quad , \quad \gamma_{1,k} = \frac{q e_{k-1} - 1}{2(\alpha + q f_{k-1})} \tag{3.23}$$

hence:

$$V_{T-k}(\tilde{P}_{T-k}, X_{T-k}, W_{T-k}) \quad = \quad \tilde{P}_{T-k}\Big[\, a_k \;+\; b_k X_{T-k} \;+\; c_k X_{T-k}^2 \;+ \\ d_k X_{T-k} W_{T-k} \;+\; e_k W_{T-k} \;+\; f_k W_{T-k}^2 \,\Big] \tag{3.24}$$

where

$$a_k \quad = \quad \gamma_{1,k}(1+\alpha\gamma_{1,k}) \;+\; q(a_{k-1}+\sigma_\eta^2 c_{k-1}) \;-\; q\gamma_{1,k}(e_{k-1}-\gamma_{1,k}f_{k-1}) \tag{3.25}$$

$$b_k \quad = \quad \gamma_{1,k}(\alpha\gamma_{x,k}+\gamma) \;+\; \gamma_{x,k}(1+\alpha\gamma_{1,k}) \;+\; q\rho b_{k-1} \;-\; q\gamma_{1,k}(\rho d_{k-1}-\gamma_{x,k}f_{k-1}) \;- \\ q\gamma_{x,k}(e_{k-1}-\gamma_{1,k}f_{k-1}) \tag{3.26}$$

$$c_k \quad = \quad \gamma_{x,k}(\alpha\gamma_{x,k}+\gamma) \;+\; q\rho^2 c_{k-1} \;-\; q\gamma_{x,k}(\rho d_{k-1}-\gamma_{x,k}f_{k-1}) \tag{3.27}$$

$$d_k \quad = \quad \alpha\gamma_{w,k}\gamma_{x,k} \;+\; \gamma_{w,k}(\alpha\gamma_{x,k}+\gamma) \;+\; q(1-\gamma_{w,k})(\rho d_{k-1}-\gamma_{x,k}f_{k-1}) \;- \\ q\gamma_{x,k}f_{k-1}(1-\gamma_{w,k}) \tag{3.28}$$

$$e_k \quad = \quad \alpha\gamma_{1,k}\gamma_{w,k} \;+\; \gamma_{w,k}(1+\alpha\gamma_{1,k}) \;-\; q\gamma_{1,k}f_{k-1}(1-\gamma_{w,k}) \;+ \\ q(1-\gamma_{w,k})(e_{k-1}-\gamma_{1,k}f_{k-1}) \tag{3.29}$$

$$f_k \quad = \quad \alpha\gamma_{w,k}^2 \;+\; qf_{k-1}(1-\gamma_{w,k})^2 \;. \tag{3.30}$$

This set of recursions completely characterize the best-execution strategy $\{S_t^*\}$ and the expected cost of best execution is given by:

$$V_1(\tilde{P}_1, X_1, W_1) \quad = \quad \tilde{P}_1\Big[\, a_{T-1} \;+\; b_{T-1}X_1 \;+\; c_{T-1}X_1^2 \;+ \\ d_{T-1}X_1W_1 \;+\; e_{T-1}W_1 \;+\; f_{T-1}W_1^2 \,\Big] \tag{3.31}$$

$$= \quad \underset{\{S_t\}}{\mathrm{Min}}\, \mathrm{E}_1\Big[\, \sum_{t=1}^{T} P_t S_t \,\Big] \;. \tag{3.32}$$

4 The Portfolio Problem

In the previous section we have focused our attention to the optimal execution of shares of a single stock. In typical applications, however, an investor would like to buy and sell a portfolio of stocks. In particular, an investor would like to buy blocks $\overline{S}_1, \ldots, \overline{S}_k$, and sell blocks $\overline{S}_{k+1}, \ldots, \overline{S}_n$, within a time horizon T. Let $\overline{\mathbf{S}}$ be the vector $[\ \overline{S}_1 \ \cdots \ \overline{S}_n\ ]'$ Let $\mathbf{P}_t$ be the n-vector of prices for the n stocks at time t. Recall that vectors are always column vectors. The symbol $\mathbf{P}_t'$ denotes the transpose of the vector $\mathbf{P}_t$. Moreover, vectors and matrices are written in boldface. Let $\mathbf{S}_t$ be the n-vector of shares to be bought (or sold) at time t. The portfolio problem is then:

$$\underset{\{\mathbf{S}_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} \mathbf{P}_t' \, \mathbf{S}_t \right] \tag{4.1}$$

subject to

$$\begin{aligned} \sum_{t=1}^{T} \mathbf{S}_t &= \overline{\mathbf{S}} && && (4.2) \\ S_{it} &\geq 0 \quad , \quad i = 1, \ldots, k \quad , \quad t = 1, \ldots, T && && (4.3) \\ S_{it} &\leq 0 \quad , \quad i = k+1, \ldots, n \quad , \quad t = 1, \ldots, T \ . && && (4.4) \end{aligned}$$

As it is typical for large mutual funds companies, in addition to nonnegativity constraints for the number of shares bought (or sold) at every period, companies would like to impose constraints that guarantee a certain degree of liquidity. Obviously, since the prices are stochastic and dynamically changing one can only hope for a probabilistic guarantee of liquidity. One possible way to achieve this is to require that

$$P(-\gamma_1 \leq \mathbf{P}_t'\mathbf{S}_t \quad \leq \quad \gamma_2) \geq 1 - \delta. \tag{4.5}$$

i.e., with probability at least $1-\delta$ the total amount needed to buy stocks minus the total amount needed to sell stocks (note that the variables S_{it} are nonnegative for $i = k+1, \ldots, n$). is within an interval $[\gamma_1, \gamma_2]$. These "liquidity" constraints are not the only linking constraint of the stocks. The law of motion for the prices might depend on all the stocks that need to be traded. We will refer to the law of motion of prices as the vector price-impact function, as the price of one stock might depend in general on characteristics of other stocks as well. We will examine two vector price-impact functions that are multidimensional extensions of the models we considered in the single stock case.

4.1 Multivariate Additive Price Impact

In this section we extend the additive price-impact function to multiple assets in order to illustrate how one could use optimization based techniques in order to address the portfolio rebalancing problem. We assume that the price vector has the following law of motion:

$$\mathbf{P}_t \;=\; \mathbf{P}_{t-1} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{X}_t + \boldsymbol{\epsilon}_t \tag{4.6}$$

where $\mathbf{A}$ is an $n \times n$ matrix, $\mathbf{B}$ is an $n \times m$ matrix, $\mathbf{X}_t$ is an m vector of market conditions $\boldsymbol{\epsilon}_t$ are independent random variables. Each $\boldsymbol{\epsilon}_t$ has a multidimensional normal distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{\Sigma}_\epsilon$. A special case is $\mathbf{\Sigma}_\epsilon = \sigma_\epsilon^2\mathbf{I}$, i.e., all components have the same variance and are uncorrelated. The vector of market conditions satisfies an AR(1) process:

$$\mathbf{X}_t \;=\; \mathbf{C}\mathbf{X}_{t-1} + \boldsymbol{\eta}_t \tag{4.7}$$

where $\mathbf{C}$ is an $m \times m$ matrix and $\boldsymbol{\eta}_t$ are independent random variables. Each $\boldsymbol{\eta}_t$ has a multidimensional normal distribution with mean $\mathbf{0}$ and covariance matrix $\mathbf{\Sigma}_\eta$. A special case is $\mathbf{\Sigma}_\eta = \sigma_\eta^2\mathbf{I}$, i.e., all components have the same variance and are uncorrelated.

We relax the non-negativity constraints and apply stochastic dynamic programming. The state at time t is $(\mathbf{P}_{t-1}, \mathbf{W}_t, \mathbf{X}_t)$. Then the dynamic programming equation becomes for $k = 0, \ldots, T-1$:

$$V_{T-k} \;=\; \underset{\mathbf{S}_{T-k}}{\text{Min}}\; \mathrm{E}_{T-k}\left[\, \mathbf{P}'_{T-k}\mathbf{S}_{T-k} + V_{T-k+1}(\mathbf{P}_{T-k}, \mathbf{X}_{T-k+1}, \mathbf{W}_{T-k+1}) \,\right] \tag{4.8}$$

The best execution strategy and optimal value function is then given by (we omit details of the calculation):

$$\mathbf{S}^*_{T-k} = \Big[\mathbf{I} - \frac{1}{2}\mathbf{A}^{-1}_{k-1}\mathbf{A}'\Big]\mathbf{W}_{T-k} + \frac{1}{2}\mathbf{A}^{-1}_{k-1}\mathbf{B}'_{k-1}\mathbf{C}\mathbf{X}_{T-k} \tag{4.9}$$

$$V_{T-k}(\mathbf{P}_{T-k-1}, \mathbf{X}_{T-k}, \mathbf{W}_{T-k}) = \mathbf{P}'_{T-k-1}\mathbf{W}_{T-k} + \mathbf{W}'_{T-k}\mathbf{A}_k\mathbf{W}_{T-k} + \mathbf{X}'_{T-k}\mathbf{B}_k\mathbf{W}_{T-k} + \mathbf{X}'_{T-k}\mathbf{C}_k\mathbf{X}_{T-k} + d_k \tag{4.10}$$

for $k = 0, 1, \ldots, T-1$, where:

$$\mathbf{A}_k = \mathbf{A} - \frac{1}{4}\mathbf{A}\mathbf{A}^{-1}_{k-1}\mathbf{A}', \quad , \quad \mathbf{A}_0 = \mathbf{A} \tag{4.11}$$

$$\mathbf{B}_k = \frac{1}{2}\mathbf{C}'\mathbf{B}_{k-1}(\mathbf{A}'_{k-1})^{-1}\mathbf{A}' + \mathbf{B}' \quad , \quad \mathbf{B}_0 = \mathbf{B}' \tag{4.12}$$

$$\mathbf{C}_k = \mathbf{C}'\mathbf{C}_{k-1}\mathbf{C} - \frac{1}{4}\mathbf{C}'\mathbf{B}_{k-1}(\mathbf{A}'_{k-1})^{-1}\mathbf{B}_{k-1}\mathbf{C} \quad , \quad \mathbf{C}_0 = \mathbf{0} \tag{4.13}$$

$$d_k = d_{k-1} + \mathrm{E}[\eta'_{T-k}\mathbf{C}_{k-1}\eta_{T-k}] \quad , \quad d_0 = 0\,. \tag{4.14}$$

Notice that the optimal execution strategy for the portfolio problem has the same features as the single stock problem. The optimal solution is a linear combination of $\mathbf{W}_t$ and $\mathbf{X}_t$.

4.2 Multivariate Linear-Percentage Price Impact

Analogously to the single asset case we consider the case in which the vector price dynamics for n assets are:

$$\mathbf{P}_t = \tilde{\mathbf{P}}_t\Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{X}_t\Big), \tag{4.15}$$

where $\mathbf{P}_t$ is a $n \times 1$ vector, $\tilde{\mathbf{P}}_t$ is an $n \times n$ matrix, $\mathbf{e}$ is an $n \times 1$ vector of ones, $\mathbf{A}$ is an $n \times n$ matrix (not necessarily diagonal), $\mathbf{S}_t$ is an $n \times 1$ vector of the controls (the vector of number of shares bought or sold at time t) , $\mathbf{B}$ is an $n \times m$ matrix and $\mathbf{X}_t$ is an $r \times 1$ vector of market conditions.

Analogously to the single asset case (see the motivation and the discussion in this section) we assume that that the no-impact matrix $\tilde{\mathbf{P}}_t$ satisfies is specified by a matrix geometric random walk:

$$\tilde{\mathbf{P}}_t = \tilde{\mathbf{P}}_{t-1} \exp(\mathbf{Z}_t) \quad , \tag{4.16}$$

where all the entries of the the matrix $\mathbf{Z}_t$ are independent and identically distributed random variables, each distributed according to a normal distribution with mean μ and variance σ^2.

Moreover, the vector of market information satisfies

$$\mathbf{X}_t = \mathbf{C}\mathbf{X}_{t-1} + \boldsymbol{\eta}_t, \tag{4.17}$$

where $\mathbf{C}$ is an $m \times m$ matrix and $\boldsymbol{\eta}_t$ are independent random variables. Each $\boldsymbol{\eta}_t$ has a multidimensional normal distribution with mean $\mathbf{0}$ and covariance matrix $\boldsymbol{\Sigma}_\eta$. A special case is $\boldsymbol{\Sigma}_\eta = \sigma_\eta^2 \mathbf{I}$, i.e., all components have the same variance and are uncorrelated.

Proceeeding analogously as in the single stock case, we find using dynamic programming that the optimal execution strategy is

$$\mathbf{S}^*_{T-k} = \mathbf{L}_k \mathbf{W}_{T-k} + \mathbf{G}_k \mathbf{X}_{T-k} + \mathbf{c}_k \tag{4.18}$$

where the parameters $\mathbf{L}_k$, $\mathbf{G}_k$ and $\mathbf{c}_k$ are calculated recursively (we omit the formulae as they do not provide particular insights).

5 Imposing Non-Negativity Constraints

In the previous sections we did not impose non-negativity constraints on the control variables. In this section we propose methods to address such constraints. In Section 5.1 we find the optimal execution strategy for a particular price impact function. The technique does not generalize to the other models we considered so far. For this reason we propose in Section 5.2 an approximation algorithm that reduces the problem to a static nonlinear optimization problem that can be solved efficiently. We further generalize the proposed method to the portfolio problem with non-negativity and liquidity constraints in Section 5.3.

5.1 A Multiplicative State-Dependent Price Impact Function.

Let X_t be a multiplicative shock to the slope α of the linear price-impact function, i.e., the price impact of a trade of S_t shares is $\alpha X_t S_t$. The state variable X_t attempts to model "market conditions." In this sense, the price impact of a trade of S_t shares is proportional to the market conditions. We assume that X_t is observable at time t and $X_1 = 1$.

In this case, the investor's optimization problem becomes:

$$\underset{\{S_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} P_t S_t \right]$$

subject to the constraints:

$$\sum_{t=1}^{T} S_t = \overline{S} \qquad , \qquad S_t \geq 0$$

with state-equations:

$$P_t = P_{t-1} + \alpha S_t X_t + \epsilon_t \tag{5.1}$$

$$\log X_t = \log(X_{t-1}) + \eta_t \qquad , \qquad X_1 = 1 \tag{5.2}$$

$$W_t = W_{t-1} - S_{t-1} \qquad , \qquad W_t = \overline{S} \ , \ W_{T+1} = 0 \tag{5.3}$$

where ϵ_t and η_t are assumed to be independent normally distributed random variables with mean zero and variance σ_ϵ^2 and σ_η^2 respectively. Then,

$$X_t = e^{\eta_t} X_{t-1} \ . \tag{5.4}$$

Let $\rho = \mathrm{E}[e^{\eta_t}] = e^{\sigma_\eta^2/2}$.

Unlike the linear price-impact function (2.19) – (2.19) of Section 2.4, in this case it is the *logarithm* of X_t which follows a random walk. The reason for this specification is to prevent the price-elasticity αX_t from changing sign—in addition to creating non-convexities that make the dynamic optimization difficult if not impossible, sign-changes in the price-elasticity are economically nonsensical.

The state variables at time t are: (P_{t-1}, X_t, W_t), i.e., the prevailing price P_{t-1} of the previous period, the market condition X_t and the unexecuted portion of the original order W_t. The control is clearly the number of shares S_t traded in period t.

Once again, we apply Bellman's equation recursively to characterize the best-execution strategy, beginning at the end:

$$V_T(P_{T-1}, X_T, W_T) = \underset{S_T}{\mathrm{Min}}\ \mathrm{E}_T[P_T S_T] = (P_{T-1} + \alpha W_T X_T) W_T \tag{5.5}$$

$$S_T^* = W_T\ . \tag{5.6}$$

For $t = T-1, \ldots, 1$ the value function satisfies the Bellman equation:

$$V_t(P_{t-1}, X_t, W_t) = \underset{S_t \geq 0}{\mathrm{Min}}\ \mathrm{E}_t \left[P_t S_t + V_{t+1}(P_t, X_{t+1}, W_{t+1}) \right] \tag{5.7}$$

We can next characterize the optimal policy for $k = T-1, \ldots, 1$ as follows:

$$S_{T-k}^* = a_k W_{T-k} \tag{5.8}$$

$$V_{T-k}(P_{T-k-1}, X_{T-k}, W_{T-k}) = P_{T-k-1} W_{T-k} + \alpha b_k X_{T-k} W_{T-k}^2 \tag{5.9}$$

where the numbers a_k, b_k are computed recursively as follows:

$$(a_0, b_0) = (1, 1)$$

and for $k = 1, \ldots, T-1$:
If $b_{k-1} \geq 0$

$$(a_k, b_k) = \begin{cases} (1 - \frac{1}{2\rho b_{k-1}}, 1 - \frac{1}{4\rho b_{k-1}}) & \text{if } 0 \leq 1 - \frac{1}{2\rho b_{k-1}} \leq 1; \\ (0, \rho b_{k-1}) & \text{if } 1 - \frac{1}{2\rho b_{k-1}} < 0; \\ (1, 1) & \text{if } 1 - \frac{1}{2\rho b_{k-1}} > 1. \end{cases} \tag{5.10}$$

If $b_{k-1} < 0$ and $b_{k-1}\rho \leq 1$, then $(a_k, b_k) = (0, b_{k-1})$.
If $b_{k-1} < 0$ and $b_{k-1}\rho > 1$, then $(a_k, b_k) = (1, 1)$.

We show by induction on k that the optimal cost is indeed given by (5.9). For $k = 0$, from (5.6) and (5.6) we obtain (**??**) and (5.9) hold. Assuming that (**??**) and (5.9) hold for

k, we will show that they hold for $k+1$. Applying Bellman's equation we obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1}) \;=$$

$$\underset{0 \le S_{T-k-1} \le W_{T_k-1}}{\text{Min}} \quad \mathrm{E}_{T-k-1}\left[\; P_{T-k-1}S_{T-k-1} \;+\; V_{T-k}(P_{T-k-1}, X_{T-k}, W_{T-k}) \;\right].$$

Using the induction hypothesis we substitute (5.9) to the above equation and obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1}) \;=$$

$$\underset{0 \le S_{T-k-1} \le W_{T_k-1}}{\text{Min}} \quad \mathrm{E}_{T-k-1}\left[\; P_{T-k-1}S_{T-k-1} \;+ P_{T-k-1}W_{T-k} \;+\; \alpha b_k X_{T-k} W_{T-k}^2 \;\right].$$

Using the laws of motion (5.2), (5.3), (5.3) we obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1}) \;=$$

$$\underset{0 \le S_{T-k-1} \le W_{T_k-1}}{\text{Min}} \quad \mathrm{E}_{T-k-1}\Big[(P_{T-k-2} + \alpha S_{T-k-1}X_{T-k-1})W_{T-k-1} + \alpha b_k \rho X_{T-k-1}(S_{T-k-1} - W_{T-k-1})^2$$

Notice that the function to be minimized is a quadratic function of S_{T-k-1}. The function is convex if $\alpha b_k \rho X_{T-k-1} \ge 0$, i.e., if $b_k \ge 0$, since $\alpha,\ \rho,\ X_{T-k-1} \ge 0$; otherwise the function is concave.

Suppose that $b_k \ge 0$. The function attains its minimum at

$$S_{T-k-1} = (1 - \frac{1}{2b_k\rho})W_{T-k-1}.$$

If $0 \le 1 - \frac{1}{2b_k\rho} \le 1$, then $0 \le S_{T-k-1} \le W_{T-k-1}$ and hence this is the optimal solution, i.e.,

$$S^*_{T-k-1} = a_{k+1}W_{T-k-1}, \text{ with } a_{k+1} = 1 - \frac{1}{2b_k\rho}.$$

If $1 - \frac{1}{2b_k\rho} < 0$, then $S^*_{T-k-1} = 0$, i.e., $a_{k+1} = 0$.

If $1-\frac{1}{2b_k\rho}>1$, then $S^*_{T-k-1}=W_{T-k-1}$, i.e., $a_{k+1}=1$.
In all cases it is straightforward to verify that

$$V_{T-k-1}(P_{T-k-2},X_{T-k-1},W_{T-k-1})=P_{T-k-2}W_{T-k-1}\ +\ \alpha b_{k+1}X_{T-k-1}W^2_{T-k-1},$$

with $b_{k+1}=a_{k+1}+\rho(1-a_{k+1})^2$.
If $a_{k+1}=1$, we obtain $b_{k+1}=1$.
If $a_{k+1}=0$, then $b_{k+1}=\rho$.
Finally if $a_{k+1}=1-\frac{1}{2b_k\rho}$, then $b_{k+1}=1-\frac{1}{4b_k\rho}$. In all cases (5.10) is satisfied.

Suppose that $b_k<0$. Then, the function to be minimized is concave and therefore, its minimum is attained at either $S_{T-k-1}=0$ or $S_{T-k-1}=W_{T-k-1}$. Therefore,

$$V_{T-k-1}(P_{T-k-2},X_{T-k-1},W_{T-k-1})\ =$$

$$\mathrm{Min}[P_{T-k-2}+W_{T-k-1}+\alpha b_k\rho X_{T-k-1}W^2_{T-k-1},P_{T-k-2}W_{T-k-1}+\alpha X_{T-k-1}W^2_{T-k-1}].$$

Hence, if $b_k\rho\leq 1$, then $S^*_{T-k-1}=0$ and

$$V_{T-k-1}(P_{T-k-2},X_{T-k-1},W_{T-k-1})=P_{T-k-2}+W_{T-k-1}+\alpha\rho b_kX_{T-k-1}W^2_{T-k-1},$$

i.e., $a_{k+1}=0$ and $b_{k+1}=b_k$.
If however, $b_k\rho>1$, then $S^*_{T-k-1}=W_{T-k-1}$ and

$$V_{T-k-1}(P_{T-k-2},X_{T-k-1},W_{T-k-1})=P_{T-k-2}W_{T-k-1}+\alpha X_{T-k-1}W^2_{T-k-1},$$

i.e., $a_{k+1}=1$ and $b_{k+1}=1$. This completes the induction.

By forward substitution using the state equations, explicit expressions for the optimal controls $S^*_2,\ldots,S^*_T$ may be obtained:

$$\begin{aligned}S^*_1 &= a_{T-1}\overline{S}\\ S^*_2 &= a_{T-2}(1-a_{T-1})\overline{S}\end{aligned}$$

$$
\begin{aligned}
S_3^* &= a_{T-3}(1 - a_{T-2} - a_{T-1})\overline{S} \\
&\vdots \\
S_k^* &= a_{T-k}(1 - a_{T-k-1} - \cdots - a_{T-1})\overline{S} \\
&\vdots \\
S_{T-1}^* &= a_1(1 - a_2 - \cdots - a_{T-1})\overline{S} \\
S_T^* &= \overline{S} - \sum_{t=1}^{T-1} S_t^*
\end{aligned}
$$

5.2 A static optimization approximation algorithm

In previous sections we were able to find a closed form formula for the optimal execution strategy in the absence of non-negativity constraints. The source of the closed form solutions is the quadratic nature of the cost-to-go function $V_t(.)$ as a function of W_t. If we impose non-negativity constraints, however, the optimal cost function is no longer quadratic. It turns out that for the additive and linear percentage price impact functions under non-neaativity constraints, the cost-to-go function $V_t(.)$ is a piece-wise quadratic function of W_t. However, the numbewr of pieces increases exponentially in the number of periods. This limits our ability to solve the problem exactly using an exact dynamic programming algorithm.

For this reason, we propose a different approximate approach. We illustrate the approach using the additive price impact function as an example first.

Additive Price Impact, revisited We address the optimal control problem under the additive price impact function in Section 2.4. The investor's problem is formulated as follows

$$
\underset{\{S_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} P_t S_t \right]
$$

subject to the constraints

$$
\sum_{t=1}^{T} S_t = \overline{S} \qquad , \qquad S_t \geq 0 .
$$

At time $t = 1$, the prevailing price P_0, the market condition X_1 and the quantity of shares $\overline{S}$ are known. Let us calculate the $\mathrm{E}_1 [P_t]$. Notice that from the dynamics of the price (2.19) we obtain that

$$P_t = P_0 \ + \ \alpha \sum_{i=1}^{t} S_i \ + \ \gamma \sum_{i=1}^{t} X_i \ + \ \sum_{i=1}^{t} \epsilon_i.$$

Notice that $E_1[X_i] = X_1 \rho^{i-1}$.

We now make **the approximation** that at time $t = 1$ the controls S_t are constants (as opposed to random variables). Then, by taking expectations we obtain that

$$E_1[P_t] = P_0 \ + \ \alpha \sum_{i=1}^{t} S_i \ + \ \gamma X_1 \frac{1-\rho^t}{1-\rho}.$$

Therefore, under this approximation the problem becomes

$$\underset{\{S_t\}}{\text{Min}} \ \sum_{t=1}^{T} \Big(P_0 \ + \ \alpha \sum_{i=1}^{t} S_i \ + \ \gamma X_1 \frac{1-\rho^t}{1-\rho} \Big) S_t \tag{5.11}$$

subject to the constraints

$$\sum_{t=1}^{T} S_t \ = \ \overline{S} \qquad , \qquad S_t \ \geq \ 0 \ .$$

Notice that the above problem is a static nonlinear (quadratic) optimization problem.

We first solve the unconstrained problem:

$$\underset{\{S_t\}}{\text{Min}} \ \sum_{t=1}^{T} \Big(P_0 \ + \ \alpha \sum_{i=1}^{t} S_i \ + \ \gamma X_1 \frac{1-\rho^t}{1-\rho} \Big) S_t$$

subject to

$$\sum_{t=1}^{T} S_t \ = \ \overline{S}.$$

We form the Lagrangean

$$L = L(S_1, \ldots, S_T, \lambda) = \sum_{t=1}^{T} \Big(P_0 \ + \ \alpha \sum_{i=1}^{t} S_i \ + \ \gamma X_1 \frac{1-\rho^t}{1-\rho} \Big) S_t - \lambda (\sum_{t=1}^{T} S_t - \overline{S}).$$

The optimal solution satisfies the first order conditions:

$$\frac{\partial L}{\partial S_t} = 0, \; t = 1, \ldots, T,$$

$$\frac{\partial L}{\partial \lambda} = 0.$$

These equations lead to :

$$P_0 + \gamma X_1 \frac{1-\rho^t}{1-\rho} + \alpha \overline{S} + \alpha S_t = \lambda$$

$$\sum_{t=1}^{T} S_t \;\; = \;\; \overline{S}.$$

Solving this system of equations we obtain for $t = 1, \ldots, T$

$$S_t^* = \frac{\overline{S}}{T} + \gamma X_1 \Big(\frac{\rho^t}{\alpha(1-\rho)} - \frac{\rho(1-\rho^T)}{T\alpha(1-\rho)^2} \Big),$$

$$\lambda = P_0 + \alpha \overline{S} \Big(1 + \frac{1}{T} \Big) + \frac{\gamma X_1}{1-\rho} \Big(1 - \frac{\rho(1-\rho^T)}{T(1-\rho)} \Big).$$

Notice that the optimal control is qualitatively the same (although not identical) with the optimal control found in Section 2.4 using stochastic dynamic programming. In particular, S_1^* is a linear combination of $\overline{S}$ and the market condition X_1. Moreover, the optimal control divides the trade in T equal parts (the term $\frac{\overline{S}}{T}$), but it adjusts the trade to account for the market conditions. This is exactly what the optimal policy was. The only difference, is that the coefficient of X_1 is different. Moreover, under the static approach the optimal static solution is found explicitly.

Also, the optimal value function is given by the expression:

$$V_1 = \sum_{t=1}^{T} \Big(P_0 \; + \; \alpha \sum_{i=1}^{t} S_i^* \; + \; \gamma X_1 \frac{1-\rho^t}{1-\rho} \Big) S_t^*.$$

Finally the sensitivity of the optimal execution cost to changes in $\overline{S}$ is simply:

$$\frac{\partial V_1}{\partial \overline{S}} = \lambda.$$

At time $t = 1$ only the control S_1^* will be implemented. At time $t = 2$, the new market condition X_2 will be observed. Moreover, there are now $T - 1$ periods remaining and the number of shares remaining to be bought is $\overline{S} - S_1^*$. We can then solve the same type of static optimization problem as before but with $T-1$ periods. The stochastic dynamic programming method finds directly the optimal trade as a function of the unexecuted portion of the trade W_t and the market condition. Notice that under the optimal

The static optimization approach, however, is particularly useful in how to address the optimal control problem (5.11) with non-negativity constraints. We first observe that the objective function of problem (5.11):

$$V(S_1, \ldots, S_T) = \sum_{t=1}^{T} \Big(P_0 \; + \; \alpha \sum_{i=1}^{t} S_i \; + \; \gamma X_1 \frac{1-\rho^t}{1-\rho} \Big) S_t$$

is convex. It is straightforward to see that

$$\frac{\partial^2 V(S_1, \ldots, S_T)}{\partial^2 S_t} = 2$$

$$\frac{\partial^2 V(S_1, \ldots, S_T)}{\partial S_t \partial S_j} = 1, \quad j \neq t.$$

Therefore the matrix of the second derivatives (the Hessian of the function $V(.)$) is semidefinite, which implies that the function $V(.)$ is convex.

Therefore, problem (5.11) is a convex quadratic problem with linear constraints. While a closed form solution is not available, there are several algorithms (see Bertsekas (1995b)), as well as commercially available software (for example MINOS, CPLEX, etc.) to solve effectively very large scale problems of this type (with thousands of variables S_t).

Linear Percentage Price Impact

We next apply the static approach to the linear percentage price impact function considered in Section 3. The investor's problem without nonnegativity constraints is formulated as follows

$$\underset{\{S_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} \tilde{P}_t (1 + \alpha S_t + \gamma X_t) S_t \right]$$

subject to

$$\sum_{t=1}^{T} S_t = \overline{S}.$$

The no-impact price is

$$\tilde{P}_t = \tilde{P}_1 \exp(\sum_{i=2}^{t} Z_i).$$

Therefore,

$$E_1[P_t] = \tilde{P}_1 \exp\Big((\mu + \frac{\sigma_z^2}{2})(t-1) \Big) = \tilde{P}_1 q^{t-1},$$

with $q = \exp(\mu + \frac{\sigma_z^2}{2})$. Moreover,

$$E_1[X_t] = X_1 \rho^{t-1}.$$

Therefore, the investor's problem is equivalent to the static optimization problem:

$$\underset{\{S_t\}}{\text{Min}} \ \sum_{t=1}^{T} \tilde{P}_1 q^{t-1} (1 + \alpha S_t + \gamma X_1 \rho^{t-1}) S_t$$

subject to

$$\sum_{t=1}^{T} S_t = \overline{S}.$$

Introducing a Lagrange multiplier λ to the constraint we can proceed as before. The solution is:

$$S_t^* = \frac{q^{T-t}(1-q)}{1-q^T}\overline{S} + \frac{\gamma}{2\alpha}\Big(\frac{q^{T-t}(1-q)(1-\rho^T)}{(1-q^T)(1-\rho)} - \gamma\rho^{t-1}\Big)X_1 + \frac{1}{2\alpha}\Big(\frac{Tq^{T-t}(1-q)}{1-q^T} - 1\Big).$$

$$\lambda = \tilde{P}_1\frac{q^{T-1}(1-q)}{1-q^T}\Big(2\alpha\overline{S} + T + \gamma\frac{1-\rho^T}{1-\rho}X_1\Big).$$

The optimal control S_1^* is a linear combination of the number of shares $\overline{S}$ to be bought, the market condition X_1 and a constant. Note that this is qualitatively (but not identical) the same with the optimal policy obtained by dynamic programming.

If we impose nonnegativity constraints on S_t, the investor's problem becomes:

$$V_1 = \underset{\{S_t\}}{\text{Min}} \sum_{t=1}^{T} \tilde{P}_1 q^{t-1}(1 + \alpha S_t + \gamma X_1 \rho^{t-1})S_t$$

subject to

$$\sum_{t=1}^{T} S_t = \overline{S}, \quad S_t \geq 0.$$

Under the assumption that $\alpha > 0$, the objective function is clearly a convex quadratic function of S_t. In addition, the objective function is separable. Although one could use a commercially available software package to find the optimal solution, we can exploit the separability and convexity of the objective function to propose a sublinear ($O(\log \mathrm{T})$) algorithm for the problem. We introduce the Lagrange multiplier λ and consider the dual problem:

$$Z(\lambda) = \underset{\{S_t\}}{\text{Min}} \sum_{t=1}^{T} \tilde{P}_1 q^{t-1}(1 + \alpha S_t + \gamma X_1 \rho^{t-1})S_t - \lambda(\sum_{t=1}^{T} S_t - \overline{S}) \qquad (5.12)$$

subject to

$$S_t \geq 0.$$

Because of the convexity of the objective function it is well known (see for example Bertsekas

(1995b)) that

$$\operatorname*{Max}_{\lambda} Z(\lambda) = V_1.$$

Therefore it suffices to solve the dual problem. Notice, however that the problem (5.12) can be solved separately for each variable S_t. As a function of the Lagrange multiplier λ the optimal solution is simply

$$S_t(\lambda) = \mathrm{Max}(0, \frac{\lambda - \gamma \tilde{P}_1 X_1 (\rho q)^{t-1}}{2\alpha q^{t-1} \tilde{P}_1}).$$

Let $u_t = \gamma \tilde{P}_1 X_1 (\rho q)^{t-1}$. Then, $S_t(\lambda) = \frac{\lambda - u_t}{2\alpha q^{t-1} \tilde{P}_1}$ if and only if $\lambda \geq u_t$. Without loss of generality let us assume that $\gamma X_1 > 0$ and $\rho q > 1$. Then the numbers u_t are sorted as follows:

$$u_1 \leq u_2 \leq \ldots \leq u_T.$$

If $u_i \leq \lambda \leq u_{i+1}$, then $S_t(\lambda) = \frac{\lambda - u_t}{2\alpha q^{t-1} \tilde{P}_1}$ for $t = 1, \ldots, i$ and $S_t(\lambda) = 0$ for $t = i+1, \ldots, T$. In this case we denote the objective function $Z(\lambda)$ as $Z_i(\lambda)$:

$$Z(\lambda) = Z_i(\lambda) = \tilde{P}_1 \sum_{t=1}^{T} q^{t-1} + \tilde{P}_1 \sum_{t=1}^{i} (\alpha q^{t-1} \frac{\lambda - u_t}{2\alpha q^{t-1} \tilde{P}_1} + \gamma X_1 \rho^{t-1}) \frac{\lambda - u_t}{2\alpha q^{t-1} \tilde{P}_1} - \lambda (\sum_{t=1}^{i} \frac{\lambda - u_t}{2\alpha q^{t-1} \tilde{P}_1} - \overline{S}).$$

Notice that this is a convex quadratic function of λ that needs to be maximized in the interval $\lambda \in [u_i, u_{i+1}]$. Clearly the optimum is obtained in either u_i or u_{i+1}. Since $Z_i(u_{i+1}) = Z_{i+1}(u_{i+1})$, we need to find the maximum

$$V_1 = \mathrm{Max}\, Z(\lambda) = \operatorname*{Max}_{i=1,\ldots,T} Z_i(u_i).$$

Note that λ^* is the u_i that maximizes $Z(\lambda)$. Because the function $Z(\lambda)$ is concave we can find the optimal u_i using binary search in $O(\log T)$ steps. Once λ^* has been found, then the optimal control is simply $S_t(\lambda^*)$.

5.3 A static optimization for the portfolio problem

We next illustrate how to incorporate non-negativity and liquidity constraints approximately in the portfolio problem.

The multivariate additive price impat function

We consider again the portfolio problem in Section 4.1. Solving the recursion (4.6) we obtain:

$$\mathbf{P}_t = \mathbf{P}_0 + \mathbf{A}\sum_{i=1}^{t}\mathbf{S}_i + \mathbf{B}\sum_{i=1}^{t}\mathbf{C}^{i-1}\mathbf{X}_1 + \mathbf{B}\sum_{k=2}^{t}\mathbf{C}^{t-k}\boldsymbol{\epsilon}_k + \sum_{k=1}^{t}\boldsymbol{\eta}_k.$$

Under the approximation that the controls $\mathbf{S}_t$ are constant at time $t = 1$, then the random variable

$$Z_t = \mathbf{P}_t'\mathbf{S}_t,$$

has a normal distribution with mean

$$E[Z_t] = \Big(\mathbf{P}_0 + \mathbf{A}\sum_{i=1}^{t}\mathbf{S}_i + \mathbf{B}\sum_{i=1}^{t}\mathbf{C}^{i-1}\mathbf{X}_1\Big)'\mathbf{S}_t, \tag{5.13}$$

and variance

$$Var[Z_t] = E\Big[\mathbf{S}_t'(\mathbf{P}_t - E[\mathbf{P}_t])(\mathbf{P}_t - E[\mathbf{P}_t])'\mathbf{S}_t\Big] =$$

$$E\Big[\mathbf{S}_t'\Big(\mathbf{B}\sum_{k=2}^{t}\mathbf{C}^{t-k}\boldsymbol{\epsilon}_k + \sum_{i=k}^{t}\boldsymbol{\eta}_k\Big)\Big(\mathbf{B}\sum_{k=2}^{t}\mathbf{C}^{t-k}\boldsymbol{\epsilon}_k + \sum_{i=k}^{t}\boldsymbol{\eta}_k\Big)'\mathbf{S}_t\Big] =$$

$$= \mathbf{S}_t'\Big(\mathbf{B}\sum_{k=2}^{t}\mathbf{C}^{t-k}\boldsymbol{\Sigma}_\epsilon(\mathbf{C}')^{t-k} + \boldsymbol{\Sigma}_\eta t\Big)\mathbf{S}_t. \tag{5.14}$$

Therefore, constraint (4.5) can be written as

$$\Phi\Big(\frac{\gamma_2 - E[Z_t]}{\sqrt{Var[Z_t]}}\Big) - \Phi\Big(\frac{-\gamma_1 - E[Z_t]}{\sqrt{Var[Z_t]}}\Big) \geq 1 - \delta,$$

where $\Phi(x)$ is the cumulative distribution function of a standard normal. Therefore, the optimal control problem for n assets becomes:

$$\operatorname*{Min}_{\{\mathbf{S}_t\}} \left[\sum_{t=1}^{T} \Big(\mathbf{P}_1 + \mathbf{A} \sum_{i=1}^{t} \mathbf{S}_i + \mathbf{B} \sum_{i=1}^{t} \mathbf{C}^{i-1} \mathbf{X}_1 \Big)' \, \mathbf{S}_t \right] \tag{5.15}$$

subject to

$$\sum_{t=1}^{T} \mathbf{S}_t = \overline{\mathbf{S}} \tag{5.16}$$

$$\Phi\Big(\frac{\gamma_2 - E[Z_t]}{\sqrt{Var[Z_t]}} \Big) - \Phi\Big(\frac{-\gamma_1 - E[Z_t]}{\sqrt{Var[Z_t]}} \Big) \geq 1 - \delta, \quad t = 1, \ldots, T \tag{5.17}$$

$$S_{it} \geq 0, \quad i = 1, \ldots, k, \quad t = 1, \ldots, T. \tag{5.18}$$

$$S_{it} \leq 0, \quad i = k+1, \ldots, n, \quad t = 1, \ldots, T. \tag{5.19}$$

where $E[Z_t]$ and $Var[Z_t]$, given in (5.13) and (5.14) respectively, are quadratic functions of the control variables. The nonlinear optimization problem (5.15) can be solved using a commercial software package (such as MINOS).

In order to obtain some insight about the structure of the solution we will be able to solve in closed form the portfolio problem where we relax the liquidity constraints (5.18) and constraints (5.19). Forming the Lagrangean, taking the first order conditions, which are necessary and sufficient if the matrix $\mathbf{A}$ is positive semidefinite and solving the resulting equations we obtain:

$$\mathbf{S}_t = \frac{1}{T}\overline{\mathbf{S}} + \frac{1}{T}\mathbf{A}^{-1}\mathbf{B} \sum_{k=1}^{T} (T+1-k)\mathbf{C}^{k-1}\mathbf{X}_1 - \mathbf{A}^{-1}\mathbf{B} \sum_{k=1}^{t} \mathbf{C}^{k-1}\mathbf{X}_1. \tag{5.20}$$

As in the exact solution, the optimal static solution without non-negativity and liquidity constraints is a linear combination of two terms: A term proportional to the vector of shares

$\overline{\mathbf{S}}$ we want to buy and sell and a term proportional to the market conditions $\mathbf{X}_1$.

Multivariate linear-percentage price impact

We next apply the approximate static approach for the portfolio problem in Section 4.2.

Solving the recursion (4.15) we obtain:

$$\mathbf{P}_t = \tilde{\mathbf{P}}_1 \exp\Big(\sum_{k=1}^{t-1} \mathbf{Z}_k\Big) \Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{C}^{t-1}\mathbf{X}_1 + \mathbf{B}\sum_{k=2}^{t} \mathbf{C}^{t-k}\boldsymbol{\epsilon}_k\Big).$$

Therefore, under the appraoximation that $\mathbf{S}_t$ are constants,

$$E[\mathbf{P}_t] = \tilde{\mathbf{P}}_1 \exp\Big((\mathbf{M} + \frac{1}{2}\mathbf{G})(t-1)\Big)\Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{C}^{t-1}\mathbf{X}_1\Big),$$

where $\mathbf{M}$ is an $n \times n$ matrix with all the entries equal to μ and $\mathbf{G}$ is an $n \times n$ matrix with all entries equal to σ_z^2. Let $\mathbf{Q} = \exp\Big((\mathbf{M} + \frac{1}{2}\mathbf{G})\Big)$. Then

$$E[\mathbf{P}_t] = \tilde{\mathbf{P}}_1 \mathbf{Q}^{t-1}\Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{C}^{t-1}\mathbf{X}_1\Big).$$

Note that the random variable

$$U_t = \mathbf{P}_t'\mathbf{S}_t = \mathbf{S}_t'\mathbf{P}_t,$$

has a normal distribution with mean

$$E[U_t] = \mathbf{S}_t'\tilde{\mathbf{P}}_1\, \mathbf{Q}^{t-1}\Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{C}^{t-1}\mathbf{X}_1\Big). \tag{5.21}$$

and variance that can be computed after somewhat involved calculations as a fourth degree polynomial of the variables $\mathbf{S}_t$ (we omit the calculations as they are involved and not very insightful). Therefore, constraint (4.5) can be written as

$$\Phi\Big(\frac{\gamma_2 - E[U_t]}{\sqrt{Var[U_t]}}\Big) - \Phi\Big(\frac{-\gamma_1 - E[U_t]}{\sqrt{Var[U_t]}}\Big) \geq 1 - \delta,$$

where $\Phi(x)$ is the cumulative distribution function of a standard normal. Therefore, the optimal control problem for n assets in this case becomes:

$$\underset{\{\mathbf{S}_t\}}{\text{Min}} \left[\sum_{t=1}^{T} \mathbf{S}_t' \tilde{\mathbf{P}}_1 \, \mathbf{Q}^{t-1} \Big(\mathbf{e} + \mathbf{A}\mathbf{S}_t + \mathbf{B}\mathbf{C}^{t-1}\mathbf{X}_1\Big)\right] \tag{5.22}$$

subject to

$$\sum_{t=1}^{T} \mathbf{S}_t = \overline{\mathbf{S}} \tag{5.23}$$

$$\Phi\Big(\frac{\gamma_2 - E[U_t]}{\sqrt{Var[U_t]}}\Big) - \Phi\Big(\frac{-\gamma_1 - E[U_t]}{\sqrt{Var[U_t]}}\Big) \geq 1 - \delta, \quad t = 1, \ldots, T \tag{5.24}$$

$$S_{it} \geq 0, \quad i = 1, \ldots, k, \quad t = 1, \ldots, T. \tag{5.25}$$

$$S_{it} \leq 0, \quad i = k+1, \ldots, n, \quad t = 1, \ldots, T. \tag{5.26}$$

where $E[U_t]$, given in (5.21), is a quadratic function of the control variables and $Var[U_t]$ is a fourth degree polynomial. The nonlinear optimization problem (5.22) can be solved using a commercial software package (such as MINOS).

We again solve in closed form the problem in which we relax the liquidity constraints (5.25) and constraints (5.26). Forming the Lagrangean, taking the first order conditions, which are necessary and sufficient if the matrix $\mathbf{A}$ is positive semidefinite and solving the resulting equations we obtain:

$$\mathbf{S}_t^* = \frac{1}{2}\mathbf{A}^{-1}\mathbf{Q}^{-(t-1)}(\mathbf{I} - \mathbf{Q}^{-1})(\mathbf{I} - \mathbf{Q}^{-T})^{-1}\Big(2\mathbf{A}\overline{\mathbf{S}} + \mathbf{B}\sum_{k=1}^{T} C^{k-1}\mathbf{X}_1 + T\mathbf{e}\Big) -$$

$$-\frac{1}{2}\mathbf{A}^{-1}\mathbf{B}C^{t-1}\mathbf{X}_1 - \frac{1}{2}\mathbf{A}^{-1}\mathbf{e}.$$

As in the exact solution, the optimal solution without the non-negativity and liquidity constraints is a linear combination of two terms: A term proportional to the vector of shares $\overline{\mathbf{S}}$ we want to buy and sell and a term proportional to the market conditions $\mathbf{X}_1$ and a constant.

6 More General Price-Impact Functions

In this section we describe the general approach to minimizing expected execution costs. Despite the fact that closed-form solutions may not exist for many cases of empirical interest, in almost all cases it is possible to obtain numerical solutions to well-posed dynamic programming problems. We begin with the same objective function:

$$\underset{\{S_t\}}{\text{Min}} \ \mathrm{E}_1 \left[\sum_{t=1}^{T} P_t S_t \right]$$

subject to the same constraints:

$$\sum_{t=1}^{T} S_t \ = \ \overline{S} \qquad , \qquad S_t \ \geq \ 0$$

but with a more general set of state-equations:

$$P_t \ = \ f_t(P_{t-1}, \mathbf{X}_t, S_t, \epsilon_t) \tag{6.1}$$

$$\mathbf{X}_t \ = \ g_t(\mathbf{X}_{t-1}, \eta_t) \tag{6.2}$$

$$W_t \ = \ W_{t-1} \ - \ S_{t-1} \qquad W_1 = \overline{S} \qquad , \qquad W_{T+1} = 0 \ . \tag{6.3}$$

The price-impact function is now incorporated into the law of motion for P_t, which is a general nonlinear and possibly time-varying function f_t of P_{t-1}, a state-vector $\mathbf{X}_t$, the control S_t, and a random shock ϵ_t. Since $\mathbf{X}_t$ is a vector of arbitrary dimension, it can accommodate several factors as in Hausman et al. (1992), and the fact that g_t is a function of only a single lag is not a restrictive as it may seem because additional lags may be captured by increasing the dimension of $\mathbf{X}_t$. Moreover, since g_t is not restricted to be linear, it is possible to capture complex dynamic behavior of the state variables parsimoniously.

The Bellman equation and optimal trade size at the end of the fixed horizon is given by:

$$
\begin{aligned}
V_T(P_{T-1}, \mathbf{X}_T, W_T) \quad &= \quad \underset{S_T}{\text{Min}} \ \text{E}_T[P_T S_T] \quad = \quad \text{E}_T[f_T(P_{T-1}, \mathbf{X}_T, W_T, \epsilon_T) W_T] && (6.4) \\
S_T^* \quad &= \quad W_T \ . && (6.5)
\end{aligned}
$$

In period $T-1$, the Bellman equation becomes:

$$
\begin{aligned}
V_{T-1}(P_{T-2}, \mathbf{X}_{T-1}, W_{T-1}) \quad &= \quad \underset{S_{T-1}}{\text{Min}} \ \text{E}_{T-1}\Big[\ P_{T-1} S_{T-1} \ + \ V_T(P_{T-1}, \mathbf{X}_T, W_T) \ \Big] && (6.6) \\
&= \quad \underset{S_{T-1}}{\text{Min}} \ \text{E}_{T-1}\Big[\ f_{T-1}(P_{T-2}, \mathbf{X}_{T-1}, S_{T-1}, \epsilon_{T-1}) S_{T-1} \ + \\
&\qquad\qquad V_T(f_{T-1}(\cdot), g_T(\cdot), W_{T-1} - S_{T-1}) \ \Big] \ . && (6.7)
\end{aligned}
$$

Performing the one-period minimization in (6.7)—subject to any additional constraints that might be present—yields the optimal trade size S_{T-1}^* as a function of the state variables P_{T-2}, $\mathbf{X}_{T-1}$, and W_{T-1}, hence:

$$
S_{T-1}^* \quad = \quad h_{T-1}(P_{T-2}, \mathbf{X}_{T-1}, W_{T-1}) \tag{6.8}
$$

Proceeding recursively in this fashion, the optimal value function and optimal trade size at time $T-k$ is then:

$$
\begin{aligned}
V_{T-k}(P_{T-k-1}, \mathbf{X}_{T-k}, W_{T-k}) \quad &= \quad \underset{S_t}{\text{Min}} \ \text{E}_{T-k}\Big[\ P_{T-k} S_{T-k} \ + \\
&\qquad\qquad V_{T-k+1}(P_{T-k}, \mathbf{X}_{T-k+1}, W_{T-k+1}) \ \Big] && (6.9) \\
&= \quad \underset{S_{T-k}}{\text{Min}} \ \text{E}_{T-k}\Big[\ f_{T-k}(P_{T-k-1}, \mathbf{X}_{T-k}, S_{T-k}, \epsilon_{T-k}) S_{T-k} \ + \\
&\qquad\qquad V_{T-k+1}(f_{T-k}(\cdot), g_{T-k+1}(\cdot), W_{T-k} - S_{T-k}) \ \Big] && (6.10) \\
S_{T-k}^* \quad &= \quad h_{T-k}(P_{T-k-1}, \mathbf{X}_{T-k}, W_{T-k}) && (6.11)
\end{aligned}
$$

and when we reach the starting point of the trading horizon:

$$
V_1(P_0, \mathbf{X}_1, W_1) \quad = \quad \underset{S_1}{\text{Min}} \ \text{E}_1\Big[\ P_1 S_1 \ + \ V_2(P_1, \mathbf{X}_2, W_2) \ \Big] \tag{6.12}
$$

$$
\begin{aligned}
&= \quad \underset{S_1}{\text{Min}}\ \mathrm{E}_1\Big[\ f_1(P_0,\mathbf{X}_1,S_1,\epsilon_1)S_1 + \\
&\qquad\qquad V_2(f_1(\cdot),g_2(\cdot),W_1-S_1)\ \Big] \qquad (6.13)
\end{aligned}
$$

$$
S_1^* \quad = \quad h_1(P_0,\mathbf{X}_1,W_1) \quad , \quad W_1 \ = \ \overline{S} \qquad (6.14)
$$

the initial conditions will enable us to obtain the entire sequence of optimal trade sizes as functions of the state variables:

$$
\begin{aligned}
S_1^* &= h_1(P_0,\mathbf{X}_1,\overline{S}) \\
S_2^* &= h_2(P_1,\mathbf{X}_2,\overline{S}-S_1^*) \\
&\vdots \\
S_k^* &= h_k(P_{k-1},\mathbf{X}_k,\overline{S}-\sum_{t=1}^{k-1} S_t^*) \qquad (6.15) \\
&\vdots \\
S_{T-1}^* &= h_{T-1}(P_{T-2},\mathbf{X}_{T-1},\overline{S}-\sum_{t=1}^{T-1} S_t^*) \\
S_T^* &= \overline{S} \ - \ \sum_{t=1}^{T-1} S_t^* \ .
\end{aligned}
$$

Moreover, the optimal value function at the start of the trading horizon, $V_1(P_0,\mathbf{X}_1,\overline{S})$, is an estimate of the best-execution costs, and depends only on the initial price, the initial state-vector, and the total number of shares to be traded.

In general, depending on the particular dynamics of the price and the market conditions we assume, the problem of calculating the optimal control at time $T-k$ (Equation (6.11)), might be intractable as a closed form expression for the cost-to-go functions $V_{T-k+1}(.,.,.)$ is not available. For this reason, we propose next two approaches that address the computational difficulties of the problem.

6.1 A Discretization Approach

Faced with the need to propose an efficient method to calculate the optimal controls and also motivated by some practical considerations we propose in this section a discretization scheme adapted to the dynamic programming algorithm we proposed earlier.

We discretize the possible prices P_t as a multiple of a constant δ. A natural choice is to select $\delta = \frac{1}{8}$, as prices in the NYSE typically move by this amount. We also assume that the horizon T is such that the range of possible prices is finite. Let K be the number of possible prices.

Moreover, we assume that the number of shares the investor buys at each period is a multiple of a fixed quantity s. Depending on the initial number of shares $\overline{S}$ that need to be bought, a typical value for s might be 100. This is also an attractive feature, as in practice, investors buy shares in round lots. This implies that at every time t the unexecuted portion of the trade W_t takes values that are multiples of s. Let $J = \frac{\overline{S}}{s}$ be the number of blocks that need to be bought initially.

Finally, we also discretize the market conditions into a discrete number of states. For example we assume that $\mathbf{X}_t$ takes a finite number N of possible values. In this way, for every time t, the cost-to-go function $V_t(P_{t-1}, \mathbf{X}_t, W_t)$ needs to be evaluated for $K\ J\ N$ possible values. Note that these values need to be stored only temporarily, as they will be used in step $t-1$ of the dynamic programming algorithm, but not at step $t-2$ (at this step the values $V_{t-1}(P_{t-2}, \mathbf{X}_{t-1}, W_{t-1})$ are needed). As a result, the total memory requirements are $O(K\ J\ N)$.

In order to have a feeling for the computational requirements let us consider an example. Assume that 50,000 shares, currently at \$60, need to be bought in the next 20 periods. We assume that within this time period, the price will be within \$55 and \$65, meaning that there are $K = 80$ possible prices to consider (multiples of $\frac{1}{8}$). If we choose $s = 500$, i.e., we will buy shares in multiples of 500, then $J = 100$. We also discretize the market conditions in $N = 10$ states. Then at each step we need to compute the cost-to-go function for $K\ J\ N = 80,000$ values of the parameters. As there are 20 periods, in total 1.6 million evaluations of the cost-to-go function are needed.

With the speed of today's workstations this computation can be performed in seconds. For example, if we assume that the computation of each cost-to-go function $V_t(P_{t-1}, X_t, W_t)$ takes 10^{-6} second, then the total computation will be 1.6 seconds.

On the other hand, if we want to have a more refined discretization, the computational requirements increase drastically. For example, if instead of considering blocks of 500 shares we consider blocks of 100 shares, and instead of using 10 states for the discretization of the market conditions, we use 100 states, then the computational requirements increase by 50 times.

6.2 Approximate Dynamic Programming

If the required discretization is excessive, we propose to use approximate dynamic programming as follows (see Bertsekas (1995a)).

We let the state vector at time t $\mathbf{Y}_t = (P_{t-1}, \mathbf{X}_t, W_t)$ We first find $V_T(\mathbf{Y}_T)$. Note that no minimization is required to obtain this value. We approximate this function with

$$\hat{V}_T(\mathbf{Y}_T) = \mathbf{Y}'_T\mathbf{Q}_T\mathbf{Y}_T + \mathbf{b}'_T\mathbf{Y}_T,$$

where the matrix $\mathbf{Q}_T$ and the vector $\mathbf{b}_T$ are selected so as to minimize:

$$\int_{\mathbf{y}_T} \left(V_T(\mathbf{Y}_T) - \hat{V}_T(\mathbf{Y}_T) \right)^2 d\mathbf{Y}_T.$$

In terms of the approximating matrix $\mathbf{Q}_T$ and vector $\mathbf{b}_T$ this is a least square problem that can be solved in closed form.

At step $T-k$ we find

$$V_{T-k}(\mathbf{Y}_{T-k}) = \underset{S_{T-k}}{\text{Min}} \; \mathrm{E}_{T-k}\Big[f_{T-k}(P_{T-k-1}, \mathbf{X}_{T-k}, S_{T-k}, \epsilon_{T-k})S_{T-k}$$

$$+ \hat{V}_{T-k+1}(f_{T-k}(\cdot), g_{T-k+1}(\cdot), W_{T-k} - S_{T-k}) \Big].$$

This minimization can be easily performed because the approximating cost-to-go function is quadratic. After finding $V_{T-k}(\mathbf{Y}_{T-k})$, we approximate the function with the quadratic function

$$\hat{V}_{T-k}(\mathbf{Y}_{T-k}) = \mathbf{Y}'_{T-k}\mathbf{Q}_{T-k}\mathbf{Y}_{T-k} + \mathbf{b}'_{T-k}\mathbf{Y}_{T-k},$$

where the matrix $\mathbf{Q}_{T-k}$ and the vector $\mathbf{b}_{T-k}$ are selected so as to minimize:

$$\int_{\mathbf{Y}_{T-k}} \left(V_{T-k}(\mathbf{Y}_{T-k}) - \hat{V}_T(\mathbf{Y}_{T-k}) \right)^2 d\mathbf{Y}_{T-k}.$$

One natural question is why we used a quadratic approximation of the cost-to-go function. For all the particular price-impact functions we examined the cost-to-go function was always quadratic. In addition, this approximation leads to a significant reduction of the dimensions. If the vector of market conditions is composed of two scalars, then the total state vector $(P_{t-1}, \mathbf{X}_t, W_t)$ is composed of four scalars. Therefore, the dimension of the matrix $\mathbf{Q}_t$ is 4×4 and the vector $\mathbf{b}_t$ has dimension 4×1. Moreover, the minimization that needs to be performed in order to obtain the optimal control S_t^* can be solved very easily when the cost to go is quadratic.

7 The Risks of Best Execution

As Table 2 forcefully illustrates, the best-execution strategy minimizes *expected* execution costs, not *actual* execution costs. Of course, without perfect foresight, minimizing actual execution costs is impossible. But in the same way that investors evaluate the risks surrounding a given expected return, we can also evaluate the uncertainty around our expected best-execution costs. In this section we address the question of assessing the risk of best execution. In addition, we propose an algorithm to control the risk of best execution.

7.1 Assessing the Risk of Best Execution

There are at least four sources of uncertainty surrounding the expected best-execution costs V_1: (1) the realized cost is typically different from the expected cost, and although over time and trades the average of the former should equal the latter, this may be small comfort if the "tracking errors" are wild and unpredictable; (2) the expected cost V_1 is itself a function of random initial conditions (X_1 in particular), and will vary from program to program even if the same block of shares were to be executed each time and at the same initial price; (3) estimation errors of the parameter estimates for the price-impact function and state-variable dynamics will be propagated recursively through Bellman's equation to V_1, so that even small uncertainties in the parameter values may lead to large uncertainty in the expected execution cost; and (4) the price-impact function and state-variable dynamics may suffer from the kind of nonstationarities and time-variation that plague all economic models—factors that affect price-impact last year need not be relevant next year, and vice-versa.

Signal/Noise Ratio	γ	ρ	Mean V_1-Error in Cents/Share	Std. Dev. V_1-Error in Cents/Share	Mean Expected Performance in Cents/Share (SE)	Mean Actual Performance in Cents/Share (SE)
0.0683	1.00	0.25	−0.64	34.47	0.0175 (0.0000)	0.0136 (0.0085)
0.6144	3.00	0.25	−0.48	45.95	0.1576 (0.0004)	0.1870 (0.0335)
1.7067	5.00	0.25	−1.00	61.94	0.4374 (0.0011)	0.3761 (0.0715)
0.0853	1.00	0.50	−0.80	35.99	0.1627 (0.0005)	0.1708 (0.0344)
0.7680	3.00	0.50	−1.97	53.55	1.4576 (0.0043)	1.1389 (0.1478)
2.1333	5.00	0.50	−11.24	68.59	4.0067 (0.0110)	1.9290 (0.2747)

Table 4: Monte Carlo simulation of the performance of the best-execution strategy for acquiring a 100,000-share block of stock within 20 periods, for 1,000 simulated sample paths of the state-variable X_t, assuming an initial price $P_0 = \$50$ and price-impact function $\alpha S_t + \gamma X_t$, where $\alpha = 5 \times 10^{-5}$, $X_t = \rho X_{t-1} + \eta_t$, $X_0 = 0$, $\sigma_\eta^2 = 0.001$, $\sigma_\epsilon^2 = (0.125)^2$, and γ and ρ vary as indicated. Each row corresponds to a separate and independent simulation of 1,000 sample paths. Note: No nonnegativity restriction was imposed.

To address the first and second sources of uncertainty, we simulate 1,000 realizations of $\{X_t\}$ and $\{P_t\}$ for the linear price-impact function with information in Section 2.4, and for each realization or "sample path" we calculate the expected best-execution cost V_1 and the actual cost that the best-execution strategy generated. The variance of V_1 over the

1,000 sample paths measures of the uncertainty induced by the initial condition X_1, and the variance of the difference between V_1 and actual costs measures the uncertainty surrounding V_1 itself.

For example, the entry in the last row of column 6 of Table 4 shows that across the 1,000 realizations of V_1, each calculated from an entire 20-period sample path of $\{X_t, P_t\}$, the mean is 4.0067 cents/share. Recall that under these same parameter values, Tables 1 and 2 expected execution costs of 3.8108 and 3.8451, respectively. The standard error of 0.0110 reported in parentheses under the mean 4.0067 implies that with 95 percent probability, the theoretical mean of V_1 will lie within 2×0.0110 of 4.0067.

With respect to the tracking error of V_1, columns 4 and 5 of Table 4 show that while the mean tracking error is sometimes small, it does fluctuate substantially. For example, under the parameter values for the first row, the mean tracking error is -0.64 cents/share, but the standard deviation of the tracking error is 34.47 cents/share. One implication of this uncertainty is the need for a larger number of simulations—with only 1,000 sample paths, the standard error of the mean tracking error is given by the standard deviation estimator divided by $\sqrt{1{,}000}$ and under the parameter values of the first row, this standard error is 1.09 cents/share which exceeds the mean in magnitude.

7.2 Controlling the Risk of Best Execution

In this section we address the problem of controlling the risk of optimal execution first for a single asset and then for multiple assets. For a single asset we have focused so far in minimizing the expected cost of execution subject to the constraint that we trade the total number of shares $\overline{S}$ over a period T and nonnegativity constraints. We have not imposed any requirement on the variability of the execution cost. In this section we examine how to introduce these considerations. Our treatment is approximate as we use the static optimization approach which, by its nature, is an approximate method.

We consider the additive price-impact function. From the dynamics (2.19) of the additive price-impact function we obtain

$$P_t = P_0 \; + \; \alpha \sum_{i=1}^{t} S_i \; + \; \gamma \sum_{i=1}^{t} \rho^{i-1} X_1 \; + \; \sum_{i=1}^{t} \epsilon_i .$$

The cost of execution is

$$V = \sum_{t=1}^{T} P_t S_t.$$

It will be convenient to introduce vectors in order to make the calculation more transparent. We introduce the vector of controls $\mathbf{S} = (S_1, \ldots, S_T)'$ and the vector of prices (which are random variables) $\mathbf{P} = (P_1, \ldots, P_T)'$. Then

$$V = \mathbf{S}'\mathbf{P},$$

$$\mathbf{P} = P_0\mathbf{e} \; + \; \alpha\mathbf{R}\mathbf{S} \; + \; \gamma\mathbf{R}\boldsymbol{\eta}_0 X_1 \; + \; \gamma\mathbf{R}_1\boldsymbol{\eta} + \boldsymbol{\epsilon},$$

where $\mathbf{e}$ is a $T \times 1$ vector of ones,

$$\mathbf{R} = \begin{bmatrix} 1 & 0 & 0 & \ldots & 0 \\ 1 & 1 & 0 & \ldots & 0 \\ 1 & 1 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & 1 & 1 & \ldots & 1 \end{bmatrix}.$$

$$\mathbf{R}_1 = \begin{bmatrix} 1 & 1 & 1 & \ldots & 1 \\ 1 & 1 & 0 & \ldots & 0 \\ 1 & \rho & 1 & \ldots & 0 \\ 1 & \rho^2 & \rho & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 1 & \rho^{T-2} & \rho^{T-3} & \ldots & 1 \end{bmatrix}.$$

$$\boldsymbol{\eta}_0 = (1, \rho, \rho^2, \ldots, \rho^{T-1})',$$

$$\boldsymbol{\eta} = (0, \eta_2, \eta_3, \ldots, \eta_T)',$$

$$\boldsymbol{\epsilon} = (\epsilon_1, \epsilon_2, \epsilon_3, \ldots, \epsilon_T)'.$$

Notice that η_1 is not defined.

Using this notation it is straightforward to calculate the first two moments of the execution cost. It is

$$E[V] = \mathbf{S}'E[\mathbf{P}] = \alpha\mathbf{S}'\mathbf{R}\mathbf{S} \ + \ \mathbf{S}'\Big(P_0\mathbf{e} \ + \ \gamma\mathbf{R}\boldsymbol{\eta}_0 X_1\Big),$$

$$\mathrm{Var}[V] = E[\mathbf{S}'(\mathbf{P} - E[\mathbf{P}]) \ (\mathbf{P} - E[\mathbf{P}])'\mathbf{S}] =$$

$$\mathbf{S}'E\Big[(\gamma\mathbf{R}_1\boldsymbol{\eta} + \boldsymbol{\epsilon}) \ (\gamma\mathbf{R}_1\boldsymbol{\eta} + \boldsymbol{\epsilon})'\Big]\mathbf{S} = \mathbf{S}'\Big(\gamma^2\sigma_\eta^2\mathbf{R}_1\mathbf{I}_1\mathbf{R}_1' + \sigma_\epsilon^2\mathbf{I}\Big)\mathbf{S},$$

where

$$\mathbf{I}_1 = \begin{bmatrix} 0 & 0 & 0 & \ldots & 0 \\ 0 & 1 & 0 & \ldots & 0 \\ 0 & 0 & 1 & \ldots & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \ldots & 1 \end{bmatrix},$$

has zeros in the first row and column and ones in the remaining $t-1$ positions of the diagonal. Therefore, the first two moments of the execution cost are quadratic functions of the controls $\mathbf{S}$:

$$E[V] = \mathbf{S}'\mathbf{K}\mathbf{S} + \mathbf{S}'\mathbf{c}$$

$$\mathrm{Var}\ [V] = \mathbf{S}'\mathbf{L}\mathbf{S},$$

where the matrices $\mathbf{K}, \mathbf{L}$ and the vector $\mathbf{c}$ are completely specified from the data in the problem:

$$\mathbf{K} = \alpha\mathbf{R}, \quad \mathbf{c} = P_0\mathbf{e} \ + \ \gamma\mathbf{R}\boldsymbol{\eta}_0 X_1, \quad \mathbf{L} = \gamma^2\sigma_\eta^2\mathbf{R}_1\mathbf{I}_1\mathbf{R}_1' + \sigma_\epsilon^2\mathbf{I}.$$

Let us pause for a moment and attempt to minimize the variance of the execution cost independent of the expected execution cost. The problem can be formulated as

$$\underset{\{\mathbf{S}\}}{\text{Min}} \ \mathbf{S}'\mathbf{L}\mathbf{S} \tag{7.1}$$

subject to

$$\mathbf{S}'\mathbf{e} = \overline{S}$$

$$\mathbf{S} \geq \mathbf{0}.$$

This is a convex quadratic programming problem that can be solved easily using standard packages. Let us call the optimal value of the is quadratic programming problem σ_V^2. This represents the smallest possible risk of execution costs.

If we ignore for a moment the nonnegativity constraints, we can solve the problem in closed form as follows:

$$\mathbf{S}^* = \overline{S}\frac{\mathbf{L}^{-1}\mathbf{e}}{\mathbf{e}'\mathbf{L}^{-1}\mathbf{e}}.$$

Notice that the optimal solution is not $S_t = \frac{\overline{S}}{T}$, as one might expect on intuitive grounds. The reason is that the covariance matrix $\mathbf{K}$ is not proportional to the identity matrix (in this case the solution $S_t = \frac{\overline{S}}{T}$ would minimize the variance of the execution cost), as there are off-diagonal correlation terms.

Suppose that in addition to minimizing the expected execution cost we would like to impose the constraint

$$\text{Var}\ [V] \leq \theta$$

and then start increasing θ from σ_V^2 to higher values. What would be the effect? Consider the following optimization problem:

$$\underset{\{\mathbf{S}\}}{\text{Min}} \;\; \mathbf{S}'\mathbf{K}\mathbf{S} + \mathbf{S}'\mathbf{c} \tag{7.2}$$

subject to

$$\mathbf{S}'\mathbf{L}\mathbf{S} \leq \theta$$

$$\mathbf{S}'\mathbf{e} = \overline{S}$$

$$\mathbf{S} \geq \mathbf{0}.$$

As the optimal solution value depends on θ we denote the optimal solution value as $V(\theta)$. The function $V(\theta)$ represents the optimal expected execution cost subject to the constraint that the variance of the total cost is at most θ. By varying θ we can control the risk of optimal execution effectively.

Problem (7.2) is a convex optimization problem that can be solved efficiently using standard packages.

The above discussion on controlling the risk of execution costs can be easily generalized to the other models we considered. There is only a minor difference. Similar to the additive model, as we have seen in earlier sections, for both the multiplicative and the linear percentage price-impact functions, the expected execution cost is a quadratic function of the control variables $\mathbf{S}$. However, the variance in these cases is a fourth (as opposed to quadratic) degree polynomial of the controls. The resulting optimization problem is slightly more complicated, but it can be solved numerically very easily.

Finally, the above considerations naturally generalize for the multiple assets case. Again we can calculate the variance of the execution cost of the total portfolio. For the vector additive price-impact function, the variance is a quadratic function of the controls, while for the vector linear percentage case it is again a fourth degree polynomial of the controls.

8 Conclusion

We developed a methodology based on nonlinear optimization and stochastic dynamic programming to optimally control execution costs. Our approach incorporates sophisticated measures of price impact—measures that vary with current market conditions—and yet is conceptually straightforward to implement. The solution of the stochastic optimization problem yields both optimal trade sizes and optimal timing as functions of market conditions, hence it is a truly dynamic and adaptive trading strategy that ensures "best execution" in an expected-valued sense over a fixed trading window. In addition, the current approach takes into account several constraints that are typical in practical settings: no short-sales, round lots, fixed completion date liquidity constraints and controlling risk. We examined several price-impact functions and showed that our approach leads to explicit solutions that provide insight on the optimal execution strategy as well as numerical algorithms to incorporate practical considerations. Moreover, the optimal value function yields a measure of the minimum average cost of the portfolio rebalancing, and may be used as a benchmark for pricing principal bids and other block transactions. We intend to explore these issues in our continuing research.

Appendix

In this section we present the formal derivation of the optimal policy for the linear-percentage-price-impact function. We will prove that (2.20) using induction on k. For $k = 0$, the optimal control is $S_T^* = W_T$ and the value function is

$$V_T(P_{T-1}, X_T, W_T) = \underset{S_T}{\mathrm{Min}}\, E_T\ [P_T S_T] = (P_{T-1}\ +\ \alpha S_T + \gamma X_T)S_T =$$

$$P_{T-1}W_T\ +\ \alpha W_T^2 + \gamma X_T W_T,$$

i.e., $a_0 = \alpha, b_0 = \gamma, c_0 = 0, d_0 = 0$. Assuming that (2.21) holds for k, we will prove it for $k+1$. Applying Bellman's equation we obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1})\ =\ \underset{S_{T-k-1}}{\mathrm{Min}}\ \mathrm{E}_{T-k-1}\left[\ P_{T-k-1}S_{T-k-1} + V_{T-k}(P_{T-k-1}, X_{T-k}, W_{T-k})\right].$$

Using the induction hypothesis we substitute (2.21) to the above equation and obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1})\ =$$

$$\underset{S_{T-k-1}}{\mathrm{Min}}\ \mathrm{E}_{T-k-1}\left[\ P_{T-k-1}S_{T-k-1}\ + P_{T-k-1}W_{T-k}\ +\ a_k W_{T-k}^2\ +\ b_k X_{T-k}W_{T-k}\ +\ c_k X_{T-k}^2\ +\ d_k\ \right].$$

Using the laws of motion (2.19), (2.19), (5.3) we obtain:

$$V_{T-k-1}(P_{T-k-2}, X_{T-k-1}, W_{T-k-1})\ =\ \underset{S_{T-k-1}}{\mathrm{Min}}\ \left[\ (P_{T-k-2} + \alpha S_{T-k-1} + \gamma X_{T-k-1})W_{T-k-1} + \right.$$

$$\left. a_k(S_{T-k-1} - W_{T-k-1})^2\ +\ b_k\rho X_{T-k-1}(W_{T-k-1} - S_{T-k-1})\ +\ c_k(\rho^2 X_{T-k-1}^2 + \sigma_\eta^2)\ +\ d_k\ \right]. \quad (8.1)$$

Notice that the function to be minimized is a quadratic function of S_{T-k-1}. Moreover, from the induction hypothesis $a_k \geq 0$ (if $\alpha \geq 0$) and therefore, the function is convex. It is straightforward to verify that the optimal solution is attained at

$$S^*_{T-k-1} \quad = \quad (1 - \frac{\alpha}{2a_k})W_{T-k-1} + \frac{\rho b_k}{2a_k}X_{T-k-1}.$$

Substituting to (8.1) we obtain (2.21) for $k+1$. Notice that a_k satisfies the recursion:

$$a_k = \alpha\Big(1 - \frac{\alpha}{4a_{k-1}}\Big), \quad a_0 = \alpha.$$

Solving the recursion we obtain that

$$a_k = \frac{\alpha}{2}\Big(1 + \frac{1}{k+1}\Big).$$

Therefore, the optimal control becomes

$$S^*_{T-k-1} \quad = \quad \frac{W_{T-k-1}}{k} + \frac{\rho b_k}{2a_k}X_{T-k-1},$$

and the induction is complete.

References

Aiyagari, R. and M. Gertler, "Asset Returns With Transactions Costs And Uninsured Individual Risk", *Journal of Monetary Economics* 27, 311–332.

Amihud, Y. and H. Mendelson, 1980, "Dealership Markets: Market Making with Uncertainty", *Journal of Financial Economics* 8, 31–54.

Amihud, Y. and H. Mendelson, 1987, "Trading Mechanisms and Stock Returns: An Empirical Investigation", *Journal of Finance* 42, 533–553.

Angel, J., 1994, "Limit Versus Market Orders", Working Paper No. FINC-1377-01-293, School of Business Administration, Georgetown University.

Arnott, R. and W. Wagner, 1990, "The Measurement And Control of Trading Costs", *Financial Analyst Journal* 46, 73–80.

Bagehot, W. (a.k.a. Jack Treynor), 1971, "The Only Game in Town", *Financial Analysts Journal* 27, 12–14.

Barclay, M. and R. Litzenberger, 1988, "Announcement Effects Of New Equity Issues And The Use Of Intraday Price Data", *Journal of Financial Economics* 21, 71–100.

Barclay, M. and J. Warner, 1993, "Stealth Trading And Volatility: Which Trades Move Prices?", *Journal of Financial Economics* 34, 281–306.

Beebower, G. and W. Priest, 1980, "The Tricks of the Trade", *Journal of Portfolio Management* 6, 36–42.

Bensaid, B., Lesne, J., Pages, H. and J. Scheinkman, 1992, "Derivative Asset Pricing With Transaction Costs", *Mathematical Finance* 2, 63–86.

Benston, G. and R. Hagerman, 1974, "Determinants of Bid-Asked Spreads in the Over-the-Counter Market", *Journal of Financial Economics* 1, 353–364.

Berkowitz, S., D. Logue and E. Noser, Jr., 1988, "The Total Cost Of Transactions On The NYSE", *Journal of Finance* 43, 97–112.

Bernhardt, D. and E. Hughson, 1994, "Splitting Orders", Working Paper, Department of Economics, Queen's University.

Bertsekas, D., 1995a, *Dynamic Programming and Optimal Control*, Vol. I. Belmont: Athena Scientific.

Bertsekas, D., 1995b, *Nonlinear Programming*, Belmont: Athena Scientific, to appear.

Bertsekas, D. and S. Shreve, 1978, *Stochastic Optimal Control: The Discrete Time Case.* New York: Academic Press.

Birinyi, L., 1995, *What Does Institutional Trading Cost?* Greenwich, CT: Birinyi Associates, Inc.

Bodurtha, S. and T. Quinn, 1990, "Does Patient Program Trading Really Pay?" *Financial Analysts Journal* 46, 35–42.

Boyle, P. and T. Vorst, 1992, "Option Replication In Discrete Time With Transaction Costs", *Journal of Finance* 47, 271–294.

Branch, B. and W. Freed, 1977, "Bid-Asked Spreads on the AMEX and the Big Board", *Journal of Finance* 32, 159–163.

Brennan, M. and T. Copeland, 1988, "Stock Splits, Stock Prices, and Transaction Costs", *Journal of Financial Economics* 22, 83–101.

Brinson, G., Hood, R. and G. Beebower, 1986, "Determinants of Portfolio Performance", *Financial Analysts Journal* 42, 39–44.

Brinson, G., Singer, B. and G. Beebower, 1991, "Determinants of Portfolio Performance II: An Update", *Financial Analysts Journal* 47, 40–48.

Bronfman, C., 1991, "From Trades to Orders on the NYSE: Pitfalls in Inference Using Transactions Data", working paper, Department of Finance and Real Estate, College of Business and Public Administration, University of Arizona, Tucson, AZ.

Chan, K., Christie, W. and P. Schultz, 1993, "Market Structure and the Intraday Evolution of Bid-Ask Spreads for NASDAQ Securities", working paper, Owen Graduate School of Management, Vanderbilt University.

Chan, L, and J. Lakonishok, 1993, "Institutional Trades and Intra-Day Stock Price Behavior", *Journal of Financial Economics* 33, 173–199.

Chan, L. and J. Lakonishok, 1995, "The Behavior of Stock Prices Around Institutional Trades", *Journal of Finance* 50, 1147–74.

Cohen, K., Maier, S., Schwartz, R. and D. Whitcomb, 1981, "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread", *Journal of Political Economy* 89, 287–305.

Cohen, Kalman, Steven Maier, Robert Schwartz, and David Whitcomb, 1986, *The Microstructure of Securities Markets*. Englewood Cliffs, NJ: Prentice-Hall.

Collins, B. and F. Fabozzi, 1991, "A Methodology for Measuring Transaction Costs", *Financial Analysts Journal* 47, 27–36.

Copeland, T. and D. Galai, 1983, "Information Effects on the Bid-Ask Spread", *Journal of Finance* 38, 1457–1469.

Cuneo, L. and W. Wagner, 1975, "Reducing the Cost of Stock Trading", *Financial Analysts Journal* 26, 35–44.

Davis, M. and A. Norman, 1990, "Portfolio Selection with Transactions Costs", *Mathematiacs of Operations Research* 15, 676–713.

Demsetz, H., 1968, "The Cost of Transacting", *Quarterly Journal of Economics* 82, 33–53.

Epps, T. W., 1976, "The Demand For Brokers' Services: The Relation Between Security Trading Volume And Transaction Cost", *Bell Journal of Economics* 7, 163–196.

Easley, D. and M. O'Hara, 1987, "Price, Trade Size, and Information in Securities Markets", *Journal of Financial Economics* 19, 69–90.

Easley, D. and M. O'Hara, 1990, "The Process of Price Adjustment in Securities Markets", working paper, Johnson Graduate School of Management, Cornell University, Ithaca, NY.

Gammill, J. and A. Perold, 1989, "The Changing Character Of Stock Market Liquidity", *Journal of Portfolio Management* 15, 13–18.

Garman, M., 1976, "Market Microstructure", *Journal of Financial Economics* 3, 257–275.

Glosten, L., 1987, "Components of the Bid-Ask Spread and the Statistical Properties of Transaction Prices", *Journal of Finance* 42, 1293–1307.

Glosten, L. and L. Harris, 1988, "Estimating the Components of the Bid/Ask Spread", *Journal of Financial Economics* 21, 123–142.

Glosten, L. and P. Milgrom, 1985, "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders", *Journal of Financial Economics* 14, 71–100.

Grossman, S. and J. Vila, 1992, "Optimal Dynamic Trading With Leverage Constraints", *Journal of Financial and Quantitative Analysis* 27, 151–168.

Harris, L., 1990, "Estimation of Stock Variances and Serial Covariances From Discrete Observations", *Journal of Financial and Quantitative Analysis* 25, 291–306.

Harris, L., 1991, "Stock Price Clustering and Discreteness", *Review of Financial Studies* 4, 389–415.

Harris, L., 1994, "Minimum Price Variations, Discrete Bid-Ask Spreads, and Quotation Sizes", *Review of Financial Studies* 7, 149–178.

Harris, L., 1994, "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems", Working Paper No. 94–8, School of Business Administration, University of Southern California.

Harris, L. and J. Hasbrouck, 1992, "Market vs. Limit Orders: The SuperDot Evidence on Order Submission Strategy", NYSE Working Paper No. 92-02.

Harris, L., Sofianos, G. and J. Shapiro, 1990, "Program Trading and Intraday Volatility", Working Paper #90–03, New York Stock Exchange, New York, NY.

Hasbrouck, J., 1988, "Trades, Quotes, Inventories, and Information", *Journal of Financial Economics* 22, 229–252.

Hasbrouck, J. and T. Ho, 1987, "Order Arrival, Quote Behavior, and the Return-Generating Process", *Journal of Finance* 42, 1035–1048.

Hasbrouck, J. and R. Schwartz, 1988, "Liquidity and Execution Costs in Equity Markets", *Journal of Portfolio Management* 14, 10–16.

Hausman, J., Lo, A. and C. MacKinlay, 1992, "An Ordered Probit Analysis of Transaction Stock Prices", *Journal of Financial Economics* 31, 319–379.

Ho, T. and H. Stoll, 1980, "On Dealership Markets Under Competition", *Journal of Finance* 35, 259–267.

Ho, T. and H. Stoll, 1981, "Optimal Dealer Pricing Under Transactions and Return Uncertainty", *Journal of Financial Economics* 9, 47–73.

Holthausen, R., Leftwich, R. and D. Mayers, 1987, "The Effect of Large Block Transactions on Security Prices: A Cross-Sectional Analysis", *Journal of Financial Economics* 19, 237–267.

Holthausen, R., Leftwich, R. and D. Mayers, 1990, "Large Block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects", *Journal of Financial Economics* 26, 71–95.

Huang, R. and H. Stoll, 1995, "Dealer versus Auction Markets: A Paired Comparison of Execution Costs on NASDAQ and the NYSE", Working Paper 95-16, Financial Markets Research Center, Owen Graduate School of Management, Vanderbilt University.

Keim, D. and A. Madhavan, 1995a, "The Anatomy of the Trading Process", *Journal of Financial Economics* 37, 371–398.

Keim, D. and A. Madhavan, 1995b, "The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects", to appear in *Review of Financial Studies.*

Keim, D. and A. Madhavan, 1995c, "Execution Costs and Investment Performance: An Empirical Analysis of Institutional Equity Trades", working paper, School of Business Administration, University of Southern California.

Kraus, A. and H. Stoll, 1972, "Price Impacts of Block Trading on the New York Stock Exchange", *Journal of Finance* 27, 569–588.

Kumar, P. and D. Seppi, 1993, "Limit and Market Orders with Optimizing Traders", working paper, Graduate School of Industrial Administration, Carnegie-Mellon University.

Leamer, E., 1978, *Specification searches.* New York: John Wiley and Sons.

Lee, C. and M. Ready, 1991, "Inferring Trade Direction From Intraday Data", *Journal of Finance* 46, 733–746.

Leinweber, D., 1993, "Using Information From Trading in Trading and Portfolio Management", in K. Sherrerd, ed.: *Execution Techniques, True Trading Costs, and the Microstructure of Markets.* Charlottesville, VA: Association for Investment Management and Research.

Leinweber, D., 1994, "Careful Structuring Reins In Transaction Costs", *Pensions and Investments* July 25, 19.

Lo, A. and C. MacKinlay, 1990c, "Data-Snooping Biases in Tests of Financial Asset Pricing Models", *Review of Financial Studies* 3, 431–468.

Loeb, T., 1983, "Trading Cost: The Critical Link Between Investment Information and Results", *Financial Analysts Journal* 39, 39–44.

Madhavan, A. and S. Smidt, 1991, "A Bayesian Model of Intraday Specialist Pricing", *Journal of Financial Economics* 30, 99–134.

Magill, M. and G. Constantinides, 1976, "Portfolio Selection with Transactions Costs", *Journal of Economic Theory* 13, 245–263.

Morrison, D., 1976, *Multivariate Statistical Methods*, 2nd edition. New York: MacMillan Publishing Company.

New York Stock Exchange Fact Book: 1994 Data, New York Stock Exchange, April 1995.

Niehans, J., 1987, "Transaction Costs", in J. Eatwell, M. Milgate, and P. Newman, eds., *The New Palgrave: A Dictionary of Economics.* London, U.K.: The MacMillan Press, Ltd.

Pérold, A., 1988, "The Implementation Shortfall: Paper Versus Reality", *Journal of Portfolio Management* 14, 4–9.

Randles, R. and D. Wolfe, 1979, *Introduction to the Theory of Nonparametric Statistics.* New York: John Wiley and Sons.

Schwartz, R. and D. Whitcomb, 1988, *Transaction Costs and Institutional Investor Trading Strategies.* Monograph Series in Finance and Economics 1988–2/3, New York: Salomon Brothers Center for the Study of Financial Institutions, New York University.

Sherrerd, K., ed., 1993, *Execution Techniques, True Trading Costs, and the Microstructure of Markets.* Charlottesville, VA: Association for Investment Management and Research.

Stoll, H., 1978, "The Supply of Dealer Services in Securities Markets", *Journal of Finance* 33, 1133–1151.

Stoll, H., 1989, "Inferring the Components of the Bid-Ask Spread: Theory and Empirical Tests", *Journal of Finance* 44, 115–134.

Stoll, H., 1993, *Equity Trading Costs.* Charlottesville, VA: Association for Investment Management and Research.

Tiniç, S., 1972, "The Economics of Liquidity Services", *Quarterly Journal of Economics* 86, 79–93.

Treynor, J., 1981, "What Does It Take to Win the Trading Game?", *Financial Analysts Journal* 37, 55–60.

Tuckman, B. and J. Vila, 1992, "Arbitrage With Holding Costs: A Utility-Based Approach", *Journal of Finance* 47, 1283–1302.

Turnbull, A. and R. White, 1995, "Trade Type and the Costs of Making Markets", working paper, Western Business School University of Western Ontario.

Wagner, W., 1993, "Defining and Measuring Trading Costs", in K. Sherrerd, ed.: *Execution Techniques, True Trading Costs, and the Microstructure of Markets.* Charlottesville, VA: Association for Investment Management and Research.

Wagner, W. and M. Banks, 1992, "Increasing Portfolio Effectiveness Via Transaction Cost Management", *Journal of Portfolio Management* 19, 6–11.

Wagner, W. and M. Edwards, 1993, "Best Execution", *Financial Analyst Journal* 49, 65–71.

Wood, R., McInish, T. and K. Ord, 1985, "An Investigation of Transactions Data for NYSE Stocks", *Journal of Finance* 40, 723–738.