

OPTIMAL CONTROL OF EXECUTION COSTS FOR PORTFOLIOS

The authors apply stochastic dynamic programming to derive trading strategies that minimize the expected cost of executing a portfolio of securities over a fixed time period. They test their strategies using real-world stock data.

The rapid growth in equity investing, driven by the increasing popularity of mutual funds and defined-contribution retirement plans, has led to a rising concentration of assets among institutional money managers. A typical portfolio manager now oversees a large portfolio of several hundred securities, with individual positions that might constitute a significant fraction of the security's average daily volume. Both active managers and passive indexers must frequently rebalance their portfolios, to include new stock picks, to sell stocks that are out of favor, or to improve the tracking of a given index or benchmark. This generates sizeable orders across many stocks that must be executed within a relatively short time horizon, and that must be executed together so as to maintain the portfolio's risk/reward characteristics. The transaction costs associated with trading such "lists" of securities—often called *execution costs*—can be substantial.

Execution costs have several components: explicit costs such as commissions and bid/ask spreads, and costs that are harder to quantify, such as the *opportunity cost* of waiting and the *price*

impact from trading. Opportunity costs arise because market prices are moving constantly and can move favorably or unfavorably without warning, generating unexpected profits or lost opportunities while a portfolio manager hesitates. Price impact is the typically unfavorable effect on prices that the act of trading creates, not unlike the turbulence that a ship's wake generates. A security's seller will, by the very act of selling, push down the security's price, yielding lower proceeds from the sale, and similarly for the buyer. Moreover, the larger the order, the more heavily the trade affects the price. For portfolios that turn over frequently or have large positions to trade, these costs can significantly hinder the fund's overall performance.¹

Recent studies show that institutional investors often break up their larger trades into smaller "packages" that they execute over the course of several days.²⁻⁵ There is a compelling economic rationale for package trading. Trading is fundamentally a dynamic, path-dependent, stochastic problem. Trading takes time, and the act of trading affects price and price dynamics, which, in turn, affect execution costs. Controlling the execution costs of large blocks of stock must be accomplished by trading over a number of time periods. This was recognized by Dimitris Bertsimas and Andrew Lo, who used stochastic dynamic-programming techniques to

1521-9615/99/\$10.00 © 1999 IEEE

DIMITRIS BERTSIMAS AND ANDREW W. LO

Massachusetts Institute of Technology

PAUL HUMMEL

Long Term Capital Management

This article was edited to *CiSE* house style and space requirements. For the full version, contact Andrew Lo at the address listed for him on p. 23.

derive optimal or *best-execution* strategies.⁶

In this article, we extend the single-asset framework Bertsimas and Lo outlined, to construct best-execution strategies for portfolio problems. We developed a specification for these problems that is both empirically plausible and computationally tractable to implement. The closed-form solution provides insight into the nature of trading portfolios. To quantify the potential cost savings of our strategies, we fit the parameters of our price-impact model using historical data on 25 large-cap New York Stock Exchange (NYSE) stocks.

The portfolio problem

Specifically, we solve this problem:

- given fixed blocks of shares in n stocks, $\bar{s} = [\bar{s}_1 \bar{s}_2 \dots \bar{s}_n]'$, to be purchased within a fixed finite number of periods T , and
- given a set of price dynamics that capture price impact (that is, an individual trade's execution price 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 minimizes the expected execution costs.

Because, as is well-known, the short-term demand curves for even the most actively traded equities are not perfectly elastic,⁷ a market order at date 0 for the entire block \bar{s} is clearly not an optimal trading strategy.

(We follow the common convention that all vectors are column vectors unless they are explicitly transposed, and boldface Roman and Greek letters denote vectors and matrices. For simplicity and without loss in generality, we consider the case of purchasing \bar{s} only. Selling \bar{s} and a combination of buying certain stocks and selling others can easily be accommodated with the appropriate sign conventions [positive numbers for purchases, negative for sales].)

Let $s_t = \{s_{1t}, s_{2t}, \dots, s_{nt}\}$ be the number of shares of each stock acquired in period t at prices $p_t = \{p_{1t}, p_{2t}, \dots, p_{nt}\}$, where $t = 1, \dots, T$. We can express the investor's objective as

$$\text{Min}_{\{s_t\}} E_1 \left[\sum_{t=1}^T p'_t s_t \right] \quad (1)$$

subject to

$$p_t = f(p_{t-1}, x_t, s_t, \epsilon_t)$$

$$x_t = g(x_{t-1}, \eta_t)$$

$$\sum_{t=1}^T s_t = \bar{s} \quad (2)$$

where x_t is a vector of state variables, ϵ_t is vector white noise, and $f(\cdot)$ and $g(\cdot)$ are the state equations or laws of motion that incorporate the price dynamics of p_t , the price impact of trading s_t , and the dynamics of the state variables. We might also wish to impose additional constraints—for example, a no-sales constraint, $s_t \geq 0$ —or other conditions that are placed on the portfolio manager by institutional restrictions, tax considerations, or other aspects of his or her investment process.

(If a portfolio manager is attempting to acquire a block of securities, selling the same securities during the acquisition period is difficult to justify [unless, of course, the manager has extremely accurate negative information regarding the security's price, which is somewhat inconsistent with the original premise that he is a buyer]. Indeed, in many cases, it is illegal because it is considered a violation of the fiduciary trust that portfolio managers have to act in the best interests of their investors.)

The portfolio case contains several interesting features that the single-stock analysis of Bertsimas and Lo and others did not capture. Perhaps the most important feature is the ability to capture cross-stock relations such as the cross-autocorrelations reported by Andrew Lo and Craig MacKinlay.⁸ Price movements in one stock can induce similar movements in the price of another, because of either common factors driving both or linked trading strategies—for example, pairs trading, index arbitrage, and risk arbitrage. In such cases, the price impact of trading a portfolio might be larger than the sum of the price impact of trading the individual stocks separately.

Alternatively, if some stocks are negatively correlated (perhaps because of portfolio substitution effects) or if the portfolio to be executed includes both purchases and sales, the portfolio execution cost might be lower than the sum of the individual stocks' execution costs. This is because of a diversification effect in which trades of one stock lower the price impact of trades in another. Whether execution costs are magnified or mollified in the portfolio case is, of course, an empirical issue that turns on the law of motion for the vector of prices and state variables. In either case, the portfolio setting clearly is considerably more complex than the single-stock case.

The state equations

We now present a specification for the state equations that incorporates a multivariate price-impact function with cross-stock interactions.

Let the execution price p_t be the sum of two components:

$$p_t = \tilde{p}_t + \delta_t, \quad (3)$$

where \tilde{p}_t is a “no-impact” price—the price that would prevail in the absence of any market impact—and δ_t is the impact. A plausible and observable proxy for the no-impact price is the midpoint of the bid/offer spread, although it can be arbitrary so long as the trade size s_t does not affect it. For convenience, and to ensure nonnegative prices, we model \tilde{p}_t as vector-geometric Brownian motion:

$$\tilde{p}_t = \exp(Z_t) \tilde{p}_{t-1}, \quad (4)$$

where Z_t is a diagonal matrix whose diagonal is a normal random vector z_t with mean μ_z and covariance matrix Σ_z . The $\exp(\cdot)$ operator denotes the matrix exponential, which, in this case, reduces to the element-wise exponential of the diagonal entries in Z_t .

For δ_t we set

$$\delta_t = \tilde{P}_t (A \tilde{P}_t s_t + B x_t), \quad (5)$$

where $\tilde{P}_t = \text{diag}[\tilde{p}_{t,i}]$ and $\text{diag}(\cdot)$ is the diagonalization operator that maps its vector argument into a diagonal matrix with the vector as the diagonal. This specification captures the impact of trading s_t shares on the transaction prices p_t . It also implies that as a percentage of the no-impact price, \tilde{p}_t , the price impact is a linear function of the dollar value of the trade and other state variables x_t . The price impact’s form (see Equation 5) differs from the single-stock case in that the percentage price-impact function for each stock i is a linear function, $A \tilde{P}_t s_t$, of the dollar values of the trades of all n stocks, not just of stock i . In the special case where A is diagonal, the portfolio problem reduces to n independent single-stock problems solved by Bertsimas and Lo.⁶

This specification of the dynamics of p_t has several advantages over other specifications (see the “Other specifications” sidebar). First, \tilde{p}_t is guaranteed to be nonnegative, and p_t is guaranteed to be nonnegative under mild restrictions on δ_t . Second, separating the transaction price p_t into the no-impact component \tilde{p}_t and impact component δ_t makes the trade’s price impact temporary. So, the impact affects the current transaction price but does not affect future prices. Third, the

percentage price impact increases linearly with the trade size, which is empirically plausible.^{1,9–11} Fourth, Equation 3 implies a natural decomposition of execution costs, decoupling market-microstructure effects from price dynamics, which is closely related to André Perold’s notion of implementation shortfall.¹² Finally, we shall see in “The dynamic-programming solution” section that Equation 3 admits a closed-form solution in which the best-execution strategy is a simple linear function of the state variables and in which the optimal-value function is quadratic.

(Because Equation 3 implies that price impact is temporary, affecting only p_t and not \tilde{p}_t , the objective function [see Equation 1] separates into two terms. The first is the no-impact cost of execution and the second is the total impact cost. This decomposition is precisely the one Perold proposed in his definition of implementation shortfall, but we apply it to executing \bar{s} . In particular, the first term gives the “paper” return or execution cost, and the sum of the two terms gives the actual cost. So the second term is the implementation shortfall in executing \bar{s} .)

The presence of the vector x_t in Equation 5 captures the potential influence of changing market conditions or private information about the securities on the price impact δ_t . For example, x_t might be the return on the S&P 500 index, a common component in the price of most equities. We model x_t as a vector with r elements, allowing for multiple sources of information to influence execution prices (or several lags of a single state variable).

To complete our specification of the state equation, we must specify the dynamics of x_t . For simplicity, we let

$$x_t = C x_{t-1} + \eta_t, \quad (6)$$

where η_t is vector white noise with mean $\mathbf{0}$ and covariance matrix Σ_η . Because x_t is a vector autoregressive process with one lag (an “AR(1)”), we can capture varying degrees of predictability in information or market conditions. The matrices A and B measure the price impact’s sensitivity to trade size and market conditions. A must be positive definite; B is arbitrary; C must have eigenvalues less than unity in modulus (to ensure stationarity of x_t).

The dynamic-programming solution

We use a stochastic dynamic-programming algorithm to solve the optimal-execution problem

(see Equation 1). We denote by w_t the vector of shares remaining to be bought (or sold) at time t :

$$\begin{aligned} w_t &= w_{t-1} - s_{t-1} \\ w_1 &= \bar{s} \\ w_{T+1} &= \mathbf{0} \end{aligned}$$

The condition $w_{T+1} = \mathbf{0}$ ensures that all \bar{s} shares are executed by time T . The complete statement of the problem is then

$$\text{Min}_{\{s_t\}} E_1 \left[\sum_{t=1}^T p'_t s_t \right],$$

subject to

$$p_t = f(p_{t-1}, x_t, s_t, \epsilon_t)$$

$$x_t = g(x_{t-1}, \eta_t)$$

$$w_t = w_{t-1} - s_{t-1}$$

$$w_1 = \bar{s}$$

$$w_{T+1} = \mathbf{0}$$

Linear-percentage price impact

As with all dynamic-programming solutions, we begin at the end. V_T is the optimal value function at the end of our trading horizon, period T . By definition,

$$\begin{aligned} V_T(\tilde{p}_T, x_T, w_T) &= \text{Min}_{\{s_T\}} E_T[p'_T s_T] = E_T[p'_T w_T] \\ &= [\tilde{P}'_T (e_n + A\tilde{P}'_T w_T + Bx_T)] w_T \end{aligned} \quad (7)$$

Because this is the last period and w_{T+1} must be set to zero, the remaining order w_T must execute. So, the optimal trade size $s_T^* = w_T$. Because $\tilde{P}_T = \tilde{P}'_T$, we can reexpress Equation 7 as

Other specifications

The specification for the state equation proposed in "The state equations" in the main article is only one of many possible specifications. For example, Dimitris Bertsimas and Andrew Lo¹ propose a natural multivariate extension of their "linear price impact" specification in which the state equation is

$$p_t = p_{t-1} + Ax_t + Bx_t + \epsilon_t \quad (A)$$

where A is a positive definite $n \times n$ matrix, B is an arbitrary $n \times m$ matrix, x_t is an m -vector of information variables, and ϵ_t is an n -vector white noise with mean $\mathbf{0}$ and covariance matrix Σ_ϵ .

As before, we assume that x_t follows a stationary AR(1) process. So,

$$x_t = Cx_{t-1} + \eta_t \quad (B)$$

where C is an $m \times m$ matrix with eigenvalues all less than unity in modulus, and η_t is an m -vector white noise with mean $\mathbf{0}$ and covariance matrix Σ_η , and which is independent of ϵ_t .

This specification differs significantly from the linear-percentage price impact in three basic respects. Perhaps the most important difference is that Equation A implies that the price impact has a "permanent" effect on prices, because of the random-walk nature of the price dynamics. Also, the price impact is linear, implying that a 10,000-share trade will have the same dollar impact on a \$1 stock as it will on a \$100 stock. Finally, unless some rather unnatural restrictions are placed on ϵ_t in Equation A, p_t could take on negative values—clearly an unrealistic prospect.

Of course, Bertsimas and Lo consider Equation A primarily for analytic tractability, not because of supporting empirical evidence. Nevertheless, it is instructive to compare this specification with the linear-percentage specification to develop

some intuition for the more practical issues in specifying the state equation. In practice, the state equation must be estimated empirically. For example, several empirical studies seem to point to both permanent and temporary price impact in US equity data.²⁻⁸ However, given the ever-changing nature of financial markets, it is crucial to reestimate the state equation for each application using the most recent data sets available.

References

1. D. Bertsimas and A. Lo, "Optimal Control of Execution Costs: A Study of Government Bonds with the Same Maturity Date," *J. Financial Markets*, Vol. 1, No. 1, 30 Apr. 1998, pp. 1-50.
2. M. Barclay and R. Litzenberger, "Announcement Effects of New Equity Issues and the Use of Intraday Price Data," *J. Financial Economics*, Vol. 21, 1988, pp. 71-100.
3. M. Barclay and J. Warner, "Stealth Trading and Volatility: Which Trades Move Prices?," *J. Financial Economics*, Vol. 34, 1993, pp. 281-306.
4. L. Chan and J. Lakonishok, "Institutional Trades and Intra-Day Stock Price Behavior," *J. Financial Economics*, Vol. 33, 1993, pp. 173-199.
5. L. Chan and J. Lakonishok, "The Behavior of Stock Prices around Institutional Trades," *J. Finance*, Vol. 50, 1995, pp. 1147-1174.
6. R. Holthausen, R. Leftwich, and D. Mayers, "The Effect of Large Block Transactions on Security Prices: A Cross-Sectional Analysis," *J. Financial Economics*, Vol. 19, 1987, pp. 237-267.
7. R. Holthausen, R. Leftwich, and D. Mayers, "Large Block Transactions, the Speed of Response, and Temporary and Permanent Stock-Price Effects," *J. Financial Economics*, Vol. 26, 1990, pp. 71-95.
8. A. Kraus and H. Stoll, "Price Impacts of Block Trading on the New York Stock Exchange," *J. Finance*, Vol. 27, 1972, pp. 569-588.

$$V_T(\tilde{\mathbf{p}}_T, \mathbf{x}_T, \mathbf{w}_T) = \mathbf{e}'_n \tilde{\mathbf{P}}_T \mathbf{w}_T + \mathbf{w}'_T \tilde{\mathbf{P}}_T \mathbf{A}' \tilde{\mathbf{P}}_T \mathbf{w}_T + \mathbf{x}'_T \mathbf{B}' \tilde{\mathbf{P}}_T \mathbf{w}_T,$$

which shows that the optimal value function is linear in \mathbf{x}_T and linear-quadratic in \mathbf{w}_T . By continuing recursively in this fashion and applying Bellman's principle of optimality,¹³ we find that the optimal value function V_{T-k} is

$$\begin{aligned} V_{T-k} &= \text{Min}_{\{\mathbf{s}_{T-k}\}} \mathbb{E}_{T-k}[\mathbf{p}'_{T-k} \mathbf{s}_{T-k} + V_{T-k+1}(\tilde{\mathbf{p}}_{T-k+1}, \mathbf{x}_{T-k+1}, \mathbf{w}_{T-k+1})] \\ &= \mathbf{e}'_n \mathbf{D}_{n,k} \mathbf{e}_n + \mathbf{e}'_r \mathbf{D}_{r,k} \mathbf{e}_r + \mathbf{e}'_n \mathbf{E}_k \mathbf{x}_{T-k} \\ &\quad + \mathbf{x}'_{T-k} \mathbf{F}_k \mathbf{e}_n + \mathbf{x}'_{T-k} \mathbf{G}_k \mathbf{x}_{T-k} + \mathbf{x}'_{T-k} \mathbf{H}_k \mathbf{w}_{T-k} \\ &\quad + \mathbf{w}'_{T-k} \mathbf{J}_k \mathbf{x}_{T-k} + \mathbf{e}'_n \mathbf{K}_k \mathbf{w}_{T-k} + \tilde{\mathbf{w}}'_{T-k} \mathbf{L}_k \mathbf{e}_n \quad (8) \\ &\quad + \mathbf{w}'_{T-k} \mathbf{N}_k \mathbf{w}_{T-k} \end{aligned}$$

This yields the best-execution strategy

$$\tilde{\mathbf{s}}_{T-k}^* = \Delta_{x,k} \mathbf{x}_{T-k} + \Delta_{w,k} \mathbf{w}_{T-k} + \Delta_{1,k} \mathbf{e}_n \quad (9)$$

(For explicit expressions for $\mathbf{D}_{n,k}$, $\mathbf{D}_{r,k}$, \mathbf{E}_k , \mathbf{F}_k , \mathbf{G}_k , \mathbf{H}_k , \mathbf{J}_k , \mathbf{K}_k , \mathbf{L}_k , and \mathbf{N}_k , and for $\Delta_{x,k}$, $\Delta_{w,k}$, and $\Delta_{1,k}$, see the Appendix posted at <http://computer.org/cise>.) The recursion (see Equation 8) and best-execution strategy (see Equation 9) completely characterize the solution to our original problem, and yield the expected best-execution cost, V_{T-k} , as a by-product.

Linear price impact

Under the law of motion (see Equations A and B in the sidebar), Bertsimas and Lo⁶ show that the portfolio problem (see Equation 1) can be solved with Bellman's equation, which yields the following best-execution strategy,

$$\mathbf{s}_{T-k}^* = \left(\mathbf{I} - \frac{1}{2} \mathbf{A}_{k-1}^{-1} \mathbf{A}'\right) \mathbf{w}_{T-k} + \frac{1}{2} \mathbf{A}_{k-1}^{-1} \mathbf{B}'_{k-1} \mathbf{C} \mathbf{x}_{T-k}, \quad (10)$$

and optimal-value function,

$$\begin{aligned} V_{T-k}(\mathbf{p}_{T-k-1}, \mathbf{x}_{T-k}, \mathbf{w}_{T-k}) &= \mathbf{p}'_{T-k-1} \mathbf{w}_{T-k} + \mathbf{w}'_{T-k} \mathbf{A}_k \mathbf{w}_{T-k} \\ &\quad + \mathbf{x}'_{T-k} \mathbf{B}_k \mathbf{w}_{T-k} + \mathbf{x}'_{T-k} \mathbf{C}_k \mathbf{x}_{T-k} + d_k \end{aligned}$$

for $k = 0, \dots, T-1$, where

$$\begin{aligned} \mathbf{A}_k &= \mathbf{A} - \frac{1}{4} \mathbf{A} \mathbf{A}_{k-1}^{-1} \mathbf{A}', & \mathbf{A}_0 &= \mathbf{A} \\ \mathbf{B}_k &= \frac{1}{2} \mathbf{C}' \mathbf{B}_{k-1} (\mathbf{A}'_{k-1})^{-1} \mathbf{A}' + \mathbf{B}', & \mathbf{B}_0 &= \mathbf{B}' \end{aligned}$$

$$\begin{aligned} \mathbf{C}_k &= \mathbf{C}' \mathbf{C}_{k-1} \mathbf{C} - \frac{1}{4} \mathbf{C}' \mathbf{B}_{k-1} (\mathbf{A}'_{k-1})^{-1} \mathbf{B}_{k-1} \mathbf{C}, & \mathbf{C}_0 &= \mathbf{0} \\ d_k &= d_{k-1} + \mathbb{E}[\eta'_{T-k} \mathbf{C}_{k-1} \eta_{T-k}], & d_0 &= 0 \end{aligned}$$

The best-execution strategy (see Equation 10) is qualitatively similar to the optimal single-stock strategy of Bertsimas and Lo.⁶ However, it has one key difference: in the portfolio case, unless the matrix \mathbf{A} is diagonal, the best-execution strategy for one stock will depend on the parameters and state variables of all the other stocks. To see this, observe that the matrix coefficient

$$\left(\mathbf{I} - \frac{1}{2} \mathbf{A}_{k-1}^{-1} \mathbf{A}'\right)$$

multiplying \mathbf{w}_{T-k} in Equation 10 will generally not be a diagonal matrix unless \mathbf{A} is diagonal. Of course, if \mathbf{A} is diagonal, trading in one stock has no price impact on any other stocks (see Equation A in the sidebar). So, the portfolio problem essentially reduces to n independent single-stock problems.

So, whether the portfolio best-execution cost is greater or less than the sum of the individual stocks' best-execution costs depends wholly on the values in \mathbf{A} . This is an empirical issue that we consider in detail later.

Imposing constraints

Most practical applications will have constraints on the kind of execution strategies that institutional investors can follow. For example, if a block of shares is to be purchased within T periods, selling the stock during these T periods is very difficult to justify even if such sales are warranted by the best-execution strategy. (Other common constraints include sector-balance, turnover, tax-motivated, and, in the portfolio case, dollar-balance constraints. This last type of constraint—the portfolio's dollar value at the end of trading must lie within some fixed interval—is one of the most difficult to impose. This is because the constraint is a function of the entire vector of prices, which is stochastic. Bertsimas and Lo devised a probabilistic method of imposing such constraints.⁶) So, in practice, buy programs will almost always have nonnegativity constraints, and sell programs will almost always have nonpositivity constraints. Such constraints are often binding for best-execution strategies, particularly when the information variable has a large effect on price impact.

Why constraints are problematic

Although there are well-known techniques for solving constrained-optimization problems in a static setting, corresponding techniques for dynamic-optimization problems have not yet been developed. To see why this task is difficult, consider the simplest case of imposing nonnegativity constraints $s_t \geq 0$ in the linear percentage price-impact model with only one asset (scalar equations). Without any constraints, the optimal-value function V_{T-k} is quadratic in the state variable w_{T-k} , so the Bellman equation can be easily solved in closed form. But if nonnegativity constraints are imposed, V_{T-k} becomes a piecewise-quadratic function, with 3^k pieces.

To see how this arises, observe that for $k = 0$, the optimal control is $s_T^* = w_T$ and V_T is a quadratic function of w_T . In the next stage, $k = 1$, we calculate the optimal control s_{T-1}^* by minimizing a quadratic function of s_{T-1} subject to the constraints $0 \leq s_{T-1} \leq w_{T-1}$. The solution is

$$s_{T-1}^* = \begin{cases} 0 & \text{if } s_{u,T-1} < 0 \\ s_{u,T-1} & \text{if } 0 < s_{u,T-1} < w_{T-1} \\ w_{T-1} & \text{if } s_{u,T-1} > w_{T-1} \end{cases},$$

where

$$s_{u,T-1} = \frac{1}{\tilde{p}_{T-1}} [\Delta_{w,1} \tilde{p}_{T-1} w_{T-1} + \Delta_{x,1} x_{T-1} + \Delta].$$

This partitions the range of w_{T-1} into three intervals; a different optimal control s_{T-1}^* is over each interval, and a continuous quadratic function of w_{T-1} is within each interval V_{T-1} .

The next stage, $k = 2$, partitions each of these three intervals into another three intervals, each with a different optimal control s_{T-2}^* , and so on. The number of intervals grows exponentially with k . Therefore, even in this simple case, calculating s_{T-k}^* and V_{T-k} exactly is only feasible for a very small number of periods T . (For example, when $T = 20$ there are $3^{20} = 3,486,784,401$ intervals at the last stage of the dynamic program!)

A static-approximation method

Faced with these difficulties, we propose an approximation method to address the optimal control problem with constraints. The dynamic-optimization algorithm we presented for the case without nonnegativity constraints gives the best-execution strategy s_{T-k}^* (see Equation 9) as a function of the state vector $(x_{T-k}, w_{T-k}, \tilde{p}_{T-k})$ at time $T-k$. At time $t = 1$, the expected execution cost V_1 is

$$\begin{aligned} V_1 &= \mathbb{E} \left[\sum_{k=1}^T p'_k s_k \right] \\ &= \mathbb{E} \left[\sum_{k=1}^T [\tilde{P}_k e_n + \tilde{P}_k (A \tilde{P}_k s_k + B x_k)]' s_k \right] \\ &= \mathbb{E} \left[e'_n \tilde{P}_1 s_1 + s'_1 \tilde{P}_1 A' \tilde{P}_1 s_1 + x'_1 B' \tilde{P}_1 s_1 + e'_n \tilde{Z}_2 \tilde{P}_1 s_2 \right. \\ &\quad + s'_2 \tilde{P}_2 \tilde{Z}_2 A' \tilde{Z}_2 \tilde{P}_2 s_2 + (C x_1 + \eta_2)' B' \tilde{Z}_2 \tilde{P}_1 s_2 \\ &\quad + e'_n \tilde{Z}_2 \tilde{Z}_3 \tilde{P}_1 s_3 + s'_3 \tilde{P}_3 \tilde{Z}_3 \tilde{Z}_2 A' \tilde{Z}_2 \tilde{Z}_3 \tilde{P}_1 s_3 \\ &\quad + [C(C x_1 + \eta_2) + \eta_3]' B' \tilde{Z}_2 \tilde{Z}_3 \tilde{P}_1 s_3 \\ &\quad + L \\ &\quad + e'_n \tilde{Z}_2 \tilde{Z}_3 L \tilde{Z}_T \tilde{P}_1 s_T \\ &\quad + s'_T \tilde{P}_1 \tilde{Z}_T \tilde{Z}_{T-1} L \tilde{Z}_2 A' \tilde{Z}_2 \tilde{Z}_3 L \tilde{Z}_T \tilde{P}_1 s_T \\ &\quad + [C^T x_1 + C^{T-1} \eta_2 + L + C \eta_T + \eta_T]' \\ &\quad \left. \times B' \tilde{Z}_2 \tilde{Z}_3 L \tilde{Z}_T \tilde{P}_1 s_T \right] \end{aligned}$$

where $\tilde{P}_k = \text{diag}\{p_k\}$ and $\tilde{Z}_k = \exp(Z_k)$. Taking the expectation of the cost function yields

$$\begin{aligned} V_1 &= e'_n \tilde{P}_1 s_1 + s'_1 A' \tilde{P}_1 s_1 + x'_1 B' \tilde{P}_1 s_1 \\ &\quad + e'_n Q \tilde{P}_1 s_1 + s'_1 (A' \bullet R) \tilde{P}_1 s_1 + x'_1 C' B' Q \tilde{P}_1 s_1 \\ &\quad + e'_n Q^2 \tilde{P}_1 s_1 + s'_1 (A' \bullet R \bullet R) \tilde{P}_1 s_1 + x'_1 (C^1)' B' Q^2 \tilde{P}_1 s_1 \\ &\quad + L \\ &\quad + e'_n Q^T \tilde{P}_T s_T + s'_T (A' \bullet R \bullet R \bullet L \bullet R \bullet R) \tilde{P}_T s_T \\ &\quad + x'_T (C^T)' B' Q^T \tilde{P}_T s_T \end{aligned} \quad (11)$$

where Q is a $(n \times n)$ diagonal matrix with entries

$$q_i = \exp\left(\mu_{z,i} + \frac{1}{2} \Sigma_{z,ii}\right),$$

R is an $(n \times n)$ symmetric matrix with elements

$$r_{ij} = \exp\left[\mu_{z,i} + \mu_{z,j} + \frac{1}{2} (\Sigma_{z,ii} + \Sigma_{z,jj} + 2\Sigma_{z,ij})\right],$$

and the matrix dot operator “ \bullet ” denotes an element-wise matrix multiplication—that is, $A \bullet B \equiv [a_{ij} b_{ij}]$.

Equation 11 depends on the entire sequence of controls, $\{s_1, s_2, \dots, s_T\}$, and the observed states at time $t = 1$, \tilde{p}_1 and x_1 . In general, each control variable s_t depends on the state at time t . Under the *static-approximation method*, we will restrict the class of controls to those that s_t depend *only* on the state at time $t = 1$. That is, they depend only on prices \tilde{p}_1 and information vector x_1 .

Table 1. Ticker symbols, CUSIPs, company names, and closing prices on a randomly selected day in 1996 for 25 stocks that constitute the sample portfolio for the empirical implementation of the best-execution strategy.

| Ticker | CUSIP | Company name | Closing price |
|--------|----------|-------------------|---------------|
| AHP | 02660910 | AMER HOME PRODS | 64.0625 |
| AN | 03190510 | AMOCO | 70.5000 |
| BLS | 07986010 | BELLSOUTH | 37.2500 |
| CHV | 16675110 | CHEVRON | 62.6250 |
| DD | 26353410 | DUPONT | 88.9375 |
| DIS | 25468710 | WALT DISNEY | 63.4375 |
| DOW | 26054310 | DOW CHEMICAL | 80.6875 |
| F | 34537010 | FORD MOTOR | 31.3125 |
| FNM | 31358610 | FANNIE MAE | 35.0625 |
| GE | 36960410 | GENERAL ELECTRIC | 90.9375 |
| GM | 37044210 | GM | 48.1250 |
| HWP | 42823610 | HEWLETT PACKARD | 48.8125 |
| IBM | 45920010 | IBM | 25.8750 |
| JNJ | 47816010 | JOHNSON & JOHNSON | 51.4375 |
| KO | 19121610 | COCA COLA | 50.8125 |
| MCD | 58013510 | MCDONALDS | 47.7500 |
| MO | 71815410 | PHILIP MORRIS | 90.1875 |
| MOB | 60705910 | MOBIL | 15.9375 |
| MRK | 58933110 | MERCK & CO | 70.1250 |
| PEP | 71344810 | PEPSICO | 28.3750 |
| PG | 74271810 | PROCTER & GAMBLE | 97.4375 |
| S | 81238710 | SEARS ROEBUCK | 44.8750 |
| T | 00195710 | AT&T | 51.9375 |
| WMT | 93114210 | WAL MART STORES | 26.3125 |
| XON | 30229010 | EXXON | 83.5625 |

Under this approximation, the problem reduces to this quadratic-optimization problem:

$$\begin{aligned}
 \text{Minimize} \quad & e'_1 \bar{P}_1 s_1 + s'_1 A' \bar{P}_1 s_1 + x'_1 B' \bar{P}_1 s_1 + e'_1 Q \bar{P}_1 s_1 \\
 & + s'_1 (A' \cdot R) \bar{P}_1 s_1 + x'_1 C' B' Q \bar{P}_1 s_1 + e'_1 Q' \bar{P}_1 s_1 \\
 & + s'_1 (A' \cdot R \cdot R) \bar{P}_1 s_1 + x'_1 (C')' B' Q' \bar{P}_1 s_1 \\
 & + 1 \\
 & + e'_1 Q' \bar{P}_1 s_T + s'_T (A' \cdot R \cdot R \cdot L \cdot R \cdot R) \bar{P}_1 s_T \\
 & + x'_T (C')' B' Q' \bar{P}_1 s_T \\
 \text{subject to} \quad & \bar{s} = \sum_{t=1}^T s_t \\
 & 0 \leq s_t, \quad t = 1, K, T
 \end{aligned} \tag{12}$$

We solve Equation 12 at time $t = 1$ and find the “optimal” controls s_1^1, \dots, s_T^1 , where the superscript indicates that this is the period-1 solution of Equation 12. However, we only implement the control s_1^1 . After we observe the state vector in

period $t = 2$, we re-solve Equation 12 for time $t = 2$ and find a new set of controls s_2^2, \dots, s_T^2 , but only implement the control s_2^2 . We continue in this fashion, at each step solving a convex quadratic optimization problem that can be handled efficiently using commercially available packages—for example, C-Plex or Minos.

The static-approximation method might not yield adequate approximations in all cases. However, in many of the examples we explored, the technique performs admirably (for example, the empirical analysis in the next section). Of course, deriving accurate bounds on the approximation error in the most interesting cases is difficult because the optimal solutions are unknown for these cases. We hope to explore the theoretical properties of the static-approximation method in future research.

An empirical example

We now implement the best-execution strategies for a hypothetical stock portfolio. Specifically, we estimate the parameters of the linear-percentage model in “The state equations” section for each stock. We then construct several portfolio-rebalancing scenarios and compare the best-execution strategy with a “naive” strategy of trading equal-size lots in each time period.

The data

Our empirical analysis draws on three data sources. The primary source is a proprietary record of trades performed over the NYSE DOT system by the trading desk at Investment Technologies Group (ITG) on every trading day between 2 January and 31 December 1996. Each trade is cataloged with this information: order-submission date and time, order execution date and time, whether it is a buy or sell order, size in shares, execution price, and order type (for example, market order or limit order). We chose the 25 stocks that had the greatest number of market orders over the year-long interval (see Table 1).

Because of our selection rule, our sample consists of companies with large market capitalizations. This ensures that we will have enough data to fit the model and arrive at reasonably accurate estimates of the parameters. But such a sample tends to exhibit a lower-than-average price impact because stocks that trade very frequently are, by definition, very liquid and have much smaller price impact. Such a bias in our sampling procedure by no means invalidates our example's relevance. If we can demonstrate that our best-

execution strategy is beneficial for highly liquid stocks, our approach's value is likely to be even greater for less liquid stocks, where price impact is significantly higher.

The ITG database provides valuable trade information, but we must augment our analysis with NYSE TAQ data to extract quotes prevailing at the time of ITG trades. The TAQ database is a complete history of all trades and quotes on the NYSE, AMEX, and Nasdaq exchanges.

Finally, we use S&P 500 tick data provided by Tick Data Inc. to get intraday levels for the S&P 500 index during 1996.

The estimation procedure

Our estimation procedure consists of three steps. First, we estimate the parameters μ_z and Σ_z of the no-impact price dynamics (see Equation 4) for each stock. Given the geometric-Brownian-motion specification, we know that the continuously compounded returns z_{it} are independently and identically (IID) normal random variates:

$$z_{i,t} = \log\left(\frac{\tilde{p}_{i,t}}{\tilde{p}_{i,t-1}}\right) \sim \mathcal{N}(\mu_i, \sigma_i^2) \quad (13)$$

for each stock i , where $i = 1, \dots, 25$ and $\mathcal{N}(\mu_i, \sigma_i^2)$ is the normal distribution with mean μ_i and variance σ_i^2 . The no-impact price is taken to be the midpoint of the prevailing bid and offer prices at time t (hence the need for quotes):

$$\tilde{p}_{it} = \frac{\tilde{p}_{it}^b + \tilde{p}_{it}^a}{2} \quad (14)$$

where \tilde{p}_{it}^b and \tilde{p}_{it}^a are the bid and ask prices for stock i at time t . For each of our 25 stocks, we collect TAQ quotes at every half hour over the course of the 1996 trading year and calculate the midpoint to construct the no-impact prices, \tilde{p}_t . Thus, the time index, t , ranges over half hours, $t = 1, 2, \dots, N_b$, where N_b is the total number of half hours in the 1996 trading year (approximately 250 days times 13 periods per day).

We then form log returns according to Equation 13 and discard any overnight returns to eliminate non-synchronous trading effects. This gives us a sample of 2,069 observations of z_t during the 1996 calendar year from which we can estimate μ_z and Σ_z in the standard way. Table 2 summarizes the results (to conserve space, we report estimates only for the first five stocks of Table 1). The drift and volatility are expressed in percent per year; we scale them up from the half-hourly units by assuming each of the 250 trading days per year consists of 13 half-hour trading intervals. The drift and volatility estimates are consistent with intuition and agree reasonably well with other data sources such as BARRA.

Our second task is to estimate the parameters of the market-information process in Equation 6. The variable x_t captures the potential impact of changing market conditions or private information about the security. For example, we could construct a short-term excess-returns model for this purpose. In our example, x_t denotes the half-hourly returns on the S&P 500 index, a common factor that influences the prices of most securities. For this specification, the AR(1) coefficients, C , and the covariance matrix of the noise, Σ_η , reduce to scalar quantities, C and σ_η . Using the S&P 500 tick data from 2 January to 31 December 1996, we construct the returns x_t , where t is the same time index used previously. We rescale the returns by subtracting out the mean and dividing by standard deviation. This leaves us with a zero-mean, unit-standard-deviation process:

$$\tilde{x}_t = \frac{x_t - \mu_x}{\sigma_x} \quad (15)$$

Assuming $|C| < 1$, we can rewrite Equation 15 as


$$\tilde{x}_t = C\tilde{x}_{t-1} + \eta_t$$

The maximum likelihood estimator of the AR(1) coefficient C is

AutoSignal™

*A timely breakthrough
in cutting-edge
signal analysis*

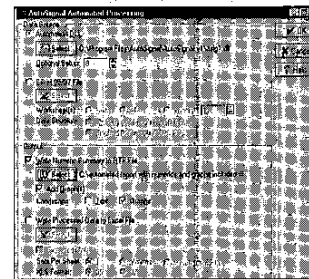
Unlike any other tool, AutoSignal has an easy-to-use interface that requires no programming to filter, process and analyze your complex signals with interactive graphics and detailed numerical reports.



Signal analysis doesn't get any easier — no programming is required! Every step of your analysis is automated.

Quickly locate your signal components using one of 19 built-in spectral analysis procedures:

- Six FFT techniques and 30 data tapering windows
- Fourteen Auto Regressive algorithms
- Complex exponential modeling
- Minimum variance methods
- Eigenanalysis frequency estimation
- Wavelets



Save precious research time with the production facility by batch processing up to 256 data sets or connect directly to your hardware.

Visit our Web site:
www.spss.com/software/science/AutoSignal

(888) 867-8660

SPSS Science
233 S. Wacker Dr., 11th Floor
Chicago, IL 60606-6307
+1 (312) 651-3000, +49 (0) 2104 1393
Fax: +312 651 3268, +49 (0) 2104 13064
Distributors worldwide:
www.spss.com/cedr

SPSS Science Software GmbH
Schimmelshausen 25
40699 Erkrath, Germany
+49 (0) 2104 9340
Fax: +49 (0) 2104 93410
France: +0800 90 37 55

Table 2. Parameter estimates and correlations for the no-impact price process \tilde{p}_i for five stocks, using 2,069 half-hourly observations from 2 January to 31 December 1996. The first and second rows give the annual drift and volatility parameters (percent/year) scaled up from half-hourly estimates by assuming 250 trading days with 13 half-hour periods per day. The last five rows report the correlation coefficients for the half-hourly returns of the five stocks.

| | AHP | AN | BLS | CHV | DD |
|----------------|-------|--------|--------|--------|-------|
| $\hat{\mu}$ | 0.200 | -0.193 | -0.387 | -0.167 | 0.283 |
| $\hat{\sigma}$ | 0.268 | 0.196 | 0.283 | 0.222 | 0.225 |
| AHP | 1.000 | 0.196 | 0.284 | 0.226 | 0.345 |
| AN | 0.196 | 1.000 | 0.173 | 0.408 | 0.283 |
| BLS | 0.284 | 0.173 | 1.000 | 0.259 | 0.328 |
| CHV | 0.226 | 0.408 | 0.259 | 1.000 | 0.314 |
| DD | 0.345 | 0.283 | 0.328 | 0.314 | 1.000 |

$$\hat{C} = \frac{\frac{1}{T_1} \sum_{t=2}^{T_1} \tilde{x}_t \tilde{x}_{t-1}}{\frac{1}{T_2} \sum_{t=1}^{T_2} \tilde{x}_t^2} \quad (16)$$

To avoid nonsynchronous trading effects, we discard all products $\tilde{x}_t \tilde{x}_{t-1}$ that straddle an overnight period in Equation 16's numerator. Similarly, we exclude overnight-return terms from Equation 16's denominator. The constants, T_1 and T_2 , are the number of terms that are included in calculating the numerator and denominator and are 1,902 and 2,078. Our estimate of \hat{C} is 0.0354. (Not surprisingly, the level of serial correlation in the S&P 500 index is quite low. If not, profitable trading strategies would be possible that would quickly drive the predictability of index returns back to a low level.)

Given \hat{C} , the maximum-likelihood estimator for the standard deviation of η_t is

$$\sigma_\eta = \sqrt{1 - \hat{C}^2}.$$

Our estimate is 0.999. The parameters \hat{C} and $\hat{\sigma}_\eta$ fully characterize the AR(1) process that describes the S&P 500 returns.

Our final task is to estimate the parameters A and B of the price-impact equation (see Equation 5). We can recast the vector equation as 25 separate linear regressions by rearranging terms:

$$\frac{p_{it} - \tilde{p}_{it}}{\tilde{p}_{it}} = \tilde{p}_{it} s_i a'_i + x_i b'_i,$$

where a_i and b_i are the i th rows of A and B . This expression shows that the percentage price impact to the i th security is a linear function of the

dollar volume we intend to trade in the i th security, the dollar volumes that we and others are currently trading in the other 24 stocks, and the S&P 500 return over the preceding half hour.

Obviously, trading in stock i should have a price impact on p_{it} . But less obvious is the role that trading in other stocks might play in determining the price impact on p_{it} . Such cross-effects have several economic sources. One stock might be a close substitute for another, so a high price impact for one would imply the same for the other. Another motivation is that, in a market with sharply rising (or falling) prices and high volume, the overall market impact will likely increase as liquidity providers demand higher premiums above posted quotes for large market orders.

To estimate A and B , we use a combination of ITG proprietary data, TAQ data, and SPX tick data. For each executed market order in a given stock i , the ITG database gives complete information about the market order except for the prevailing quote. We search the TAQ database to find the quote. As before, we form the no-impact price, \tilde{p}_{it} , as the average of the bid and offer (see Equation 14) and then construct the dependent variables,

$$\frac{p_{it} - \tilde{p}_{it}}{\tilde{p}_{it}},$$

for each trade. The ITG database provides one independent variable—namely, the dollar volume of stock i : $\tilde{p}_{it} s_{it}$.

A difficulty arises in constructing other dollar-volume-related independent variables (that is, $\tilde{p}_{jt} s_{jt}$ for $i \neq j$). The ITG data is too sparse to find nearby trades in time, so we must turn to the TAQ data to resolve this observability problem. One possible solution is to use the nearest TAQ trade that occurs before an ITG market order. Unfortunately, the time alignment of the TAQ and ITG data sets can be imprecise because of recording lags by either party. To reduce the impact of this type of error, we define a proxy for the closest trade by forming a 30-second window before each market order and computing an average dollar volume within it. Although this averaging procedure tends to smooth the data and reduces its information content, it ensures that temporal sequencing is not violated.

Specifically, we find all N_k trades in stock j that occur within that window. Each trade is executed at price p_{ik} , where $k = 1, \dots, N_k$. Trades that are executed above or at the midpoint of their quotes are classified as buys, and the rest as sells. We then compute an average dollar volume

within the window for stock j as

$$\tilde{p}_{jt^s} = \frac{1}{N_k} \sum_{k=1}^{N_k} \tilde{p}_{jk^s}.$$

Finally, for the S&P 500 return, x_t , we split the trading day into 13 half-hour intervals and compute the return in the half hour before the trade.

We now have a complete set of data with which to estimate the parameters of our model's price-impact portion. We performed the regressions in SAS; they contained no intercept term because the price impact should be zero if no stocks are being traded. Table 3 summarizes the first five of the 25 regressions. For each regression, the table reports the parameter estimates for the 26 regressors—the lagged returns for the 25 stocks and the S&P 500 lagged return—and their t -statistics. R^2 and the sample size appear at the bottom of each column. (We performed diagnostics on the residuals to test for the presence of heteroskedasticity and autocorrelation. The Durbin-Watson test indicated low levels of positive serial correlation, with statistics ranging from 1.12 to 1.69 for the 25 regressions. The test of first and second moment specification indicated a very weak pres-

ence of heteroskedasticity, because the p -values were, in general, very low.)

To develop some intuition for the coefficients, consider the estimated price impact for American Home Products in \hat{A} caused by trading in AHP, which is 4.97×10^{-10} , according to Table 3. If we traded a 100,000-share block of AHP at its beginning-of-year price of \$64.0625 with no impact, our total cost would be $100,000 \times \$64.0625 = \$6,406,250$. But according to Table 3, the full-impact cost would be

$$100,000 \times (\tilde{p}_t + \delta_t) = 100,000 \times (\$64.0625 + \delta_t)$$

$$\delta_t = \tilde{p}_t (4.97 \times 10^{-10} \times \tilde{p}_t \times 100,000) = 0.203969$$

$$100,000 \times (\tilde{p}_t + \delta_t) = \$6,426,647$$

which implies a price impact of approximately 20 cents per share (ignoring the other factors in the regression). This estimated price impact is unacceptably high; no professional trader would submit such a large order except in the most desperate circumstances.

Further inspection of the regression diagnostics shows that R^2 ranges from 0.052 to 0.440 for the 25 regressions, indicating that the regressions have varying degrees of explanatory power. How-

The
MIT
Press

Using MPI

Portable Parallel Programming with the
Message Passing Interface
second edition

William Gropp, Ewing Lusk, and Anthony Skjellum

Using MPI-2

Advanced Features of the Message Passing Interface

William Gropp, Ewing Lusk, and Rajeev Thakur

The Message Passing Interface (MPI) specification is widely used for solving significant scientific and engineering problems on parallel computers. There exist more than a dozen implementations on computer platforms ranging from IBM SP-2 supercomputers to clusters of PCs running Windows NT or Linux ("Beowulf" machines). The initial MPI Standard document, MPI-1, was recently updated by the MPI Forum. The new version, MPI-2, contains both significant enhancements to the existing MPI core and new features.

Using MPI is a completely up-to-date version of the authors' 1994 introduction to the core functions of MPI.

It adds material on the new C++ and Fortran 90 bindings for MPI throughout the book. It contains greater discussion of datatype extents, the most frequently misunderstood feature of MPI-1, as well as material on the new extensions to basic MPI functionality added by the MPI-2 Forum in the area of MPI datatypes and collective operations.

Using MPI-2 covers the new extensions to basic MPI. These include parallel I/O, remote memory access operations, and dynamic process management. The volume also includes material on tuning MPI applications for high performance on modern MPI implementations.

To order call 800-356-0343 or (617) 625-8569. Prices higher outside US and subject to change without notice.

Scientific and Engineering Computation series
Using MPI, 2nd edition: 350 pp., \$32.50 paper
Using MPI-2: 275 pp., \$32.50 paper
Set: \$55

<http://mitpress.mit.edu>

Table 3. Coefficients of the unconstrained price-impact regressions for five stocks, based on market orders from 2 January to 31 December 1996. All coefficients have been multiplied by 10^{10} except the SPX coefficients, which have been multiplied by 10^5 . The last two rows contain the sample size T and R^2 coefficients; the t -statistics appear in parentheses below the coefficients.

| Variable | AHP | AN | BLS | CHV | DD |
|----------|------------------|------------------|------------------|------------------|------------------|
| SPX | 3.74 (0.92) | -2.29 (0.43) | 0.27 (0.04) | 11.73 (1.94) | 2.38 (0.56) |
| AHP | 4.97 (3.36) | -2.04 (-2.11) | -1.89 (-1.12) | 0.80 (0.69) | -1.78 (-1.65) |
| AN | -1.49 (-1.62) | 5.86 (5.96) | -2.27 (-1.49) | 1.38 (1.14) | -1.89 (-1.81) |
| BLS | 4.46 (2.05) | 0.31 (0.18) | 0.85 (0.29) | 4.37 (2.20) | -1.84 (-1.07) |
| CHV | 2.02 (1.38) | -0.25 (-0.32) | 4.90 (2.54) | 8.72 (2.84) | 3.24 (3.30) |
| DD | -1.08 (-0.81) | -0.70 (-0.62) | 5.31 (3.17) | -1.29 (0.89) | 7.21 (5.53) |
| DIS | 2.35 (2.26) | 3.04 (3.23) | 0.01 (0.03) | -1.54 (-1.28) | 1.17 (1.35) |
| DOW | -1.12 (-1.12) | 0.12 (0.27) | -3.53 (-2.04) | -2.43 (-1.89) | -2.52 (-3.44) |
| F | 1.75 (0.70) | 0.19 (0.15) | 0.01 (0.02) | -1.01 (-0.47) | 2.19 (1.16) |
| FNM | -5.10 (-1.44) | -4.87 (-2.63) | 4.59 (1.19) | 2.33 (0.95) | 2.81 (1.23) |
| GE | 0.14 (0.12) | 0.78 (1.30) | 0.11 (0.12) | 0.96 (1.13) | -0.29 (-0.67) |
| GM | -3.74 (-2.32) | -0.67 (-0.99) | -0.30 (-0.50) | -0.16 (-0.21) | -3.30 (-4.80) |
| HWP | -1.20 (-0.74) | 1.45 (1.28) | -5.55 (-2.32) | 3.78 (2.95) | 3.26 (4.30) |
| IBM | -0.64 (-0.68) | 0.51 (0.80) | -0.01 (-0.05) | -0.17 (-0.16) | -0.50 (-1.80) |
| JNJ | 0.81 (0.40) | 1.60 (2.20) | 3.88 (2.15) | 1.37 (1.43) | 1.65 (1.26) |
| KO | -0.58 (0.53) | -0.20 (-0.28) | 2.43 (1.46) | -0.36 (-0.52) | -0.30 (-0.56) |
| MCD | 2.56 (1.20) | -5.83 (-3.75) | -5.80 (1.56) | -3.34 (-2.96) | -4.00 (-2.82) |
| MO | 0.17 (0.28) | -1.78 (-1.53) | 0.69 (0.77) | 3.06 (1.93) | -0.01 (-0.05) |
| MOB | 0.70 (0.69) | -1.36 (1.32) | -2.02 (-1.16) | -0.48 (-0.56) | -0.32 (-0.40) |
| MRK | -1.20 (-2.33) | 0.65 (0.70) | 0.31 (0.24) | -3.23 (-2.06) | 1.03 (1.37) |
| PEP | -0.55 (-0.32) | 5.71 (5.82) | 2.01 (0.83) | -0.37 (-0.17) | 1.12 (1.62) |
| PG | -0.28 (-0.24) | -2.15 (-1.99) | 0.91 (0.50) | 2.92 (2.19) | 2.85 (2.47) |
| S | -0.01 (-0.04) | -2.69 (-1.95) | -9.44 (-2.80) | 0.21 (0.08) | 1.43 (0.96) |
| T | 1.78 (2.63) | 0.47 (0.05) | -0.49 (-0.32) | -1.49 (-2.82) | 0.45 (0.82) |
| WMT | -1.58 (0.89) | -0.41 (-0.24) | 3.89 (1.17) | -8.31 (-2.35) | 3.69 (1.83) |
| XON | 1.66 (2.56) | -1.02 (-1.98) | -0.35 (-0.28) | 0.20 (0.38) | 1.04 (2.02) |
| T | 726 | 887 | 768 | 791 | 759 |
| R^2 | 0.11 | 0.23 | 0.08 | 0.09 | 0.21 |

ever, 30% (193 of 625) of the t -statistics are significant at the 5% level, implying the importance of factors other than own-stock trading in determining price impact. Also, 18 of the 25 own-stock price-impact terms—that is, \hat{a}_{ii} (the i th diagonal entry of \hat{A})—are statistically significant. These terms should be the most dominant in determining price impact, and our regression confirms this conjecture. Nevertheless, cross-stock effects are significant.

Consider, for example, the AHP regression: while the own-price effect has a coefficient of 4.97, coefficients for BLS and FNM are 4.46 and -5.10. That these two cross-stock coefficients have opposite signs underscores the portfolio approach's importance for minimizing execution costs. Because of significant cross-stock price-impact effects, the expected cost of executing a portfolio is not simply the sum of the expected values of executing each security in isolation.

Although some regressions have low explanatory power, recall that we have proposed a rather naive specification for these regressions, omitting many other variables that proprietary traders and other professional portfolio managers have at their disposal. But even with our naive specification, we still achieve R^2 's as high as 0.440 (for Merck, not shown in Table 3), which is quite substantial, considering the data's high frequency.

No-arbitrage constraints

One additional aspect of the estimation procedure must be considered: whether or not the parameter estimates yield a well-posed optimization problem (see Equations 1 and 2). In particular, for certain parameter values, the optimization problem is not convex, so the objective function can be made arbitrarily negative. The economic interpretation for such circumstances is an *arbitrage opportunity* (also known as a "free lunch"), a situation in which riskless profits can be manufactured out of thin air. Ordinarily, this would be a welcome state of affairs for investment professionals. In this case, the arbitrage is more likely a spurious side effect of sampling variation in our parameter estimates.

To avoid these false-arbitrage opportunities, a no-arbitrage restriction should be imposed on the estimation procedure. For the linear-percentage price-impact model, we accomplish this by constraining both \hat{A} and $\hat{A} \cdot R$ to be positive definite matrices. This, in turn, involves estimating a constrained linear-regression model. Table 4 reports the results of such a procedure. The two most significant differences between Tables 3 and 4 are

that the latter shows lower R^2 's and the higher significance of the own-stock coefficients. The former is not surprising, because any constraint is bound to decrease the regression's explanatory power, although the decline is rather small for AN and DD. The higher significance of own-stock coefficients follows from the definition of positive definiteness—that is, $x'Ax > 0$ for all vectors x . The diagonal elements of A are coefficients of squared terms of the x values in the matrix product $x'Ax$. Therefore, by making the squared terms sufficiently large relative to the cross-terms, we arrive at a positive definite matrix.

As Table 4 shows, the cross-effects are also affected by the no-arbitrage constraint, highlighting its significance in the portfolio context. For example, in AHP's case, the coefficients of BLS and FNM are now smaller (3.49 and 0.01) than in Table 3 (4.46 and -5.10), where the no-arbitrage constraint has not been imposed. However, the coefficients of MCD and WMT become larger, increasing to 3.38 and 3.70 from 2.56 and 1.58 in Table 3.

Table 5 shows the ratio of the total sum of squared errors of the constrained regression to the unconstrained regression for all 25 stocks. The increase in squared errors is approximately 5% overall, a rather modest increase that provides some support for imposing the restriction. More important, if the no-arbitrage condition were not imposed, the dynamic-optimization algorithms described earlier might yield nonsensical results.

The empirical results of Tables 3 to 5 suggest that the state equations necessary for our dynamic-optimization algorithm can be estimated reasonably accurately and that a portfolio approach to execution-cost control has significant benefits.

Monte Carlo analysis

Having calibrated the state equation (see "The estimation procedure" section) for the linear percentage case (see "The state equations" section), we now investigate the performance of the best-execution strategy through Monte Carlo simulations. Specifically, we consider minimizing the execution costs of purchasing \bar{s} shares of each of the 25 stocks in Table 1 over T periods. This occurs under the price dynamics (see Equations 3 to 6) where A and B are the estimates \hat{A} and \hat{B} from the constrained regression (see the previous section) and C , μ_x , and Σ_{η} are as we estimated in "The estimation procedure" section. We assume that the baseline covariance matrix of the no-impact price is Σ_x and that the initial no-impact prices are the prices in Table 1 (closing prices se-

Table 4. Coefficients of the constrained price-impact regressions for five stocks, based on market orders from 2 January to 31 December 1996. All coefficients have been multiplied by 10^{10} except for the SPX coefficients, which have been multiplied by 10^5 . The last two rows contain sample size T and R^2 coefficients; t -statistics are in parentheses below the coefficients.

| Variable | AHP | AN | BLS | CHV | DD |
|----------|------------------|------------------|------------------|------------------|------------------|
| SPX | 3.74 (0.90) | 2.28 (0.41) | 0.26 (0.04) | 11.70 (1.89) | 2.38 (0.54) |
| AHP | 12.40 (8.17) | -1.69 (-1.70) | -1.99 (-1.15) | 1.04 (0.88) | -1.07 (-0.96) |
| AN | -1.32 (-1.40) | 10.10 (9.98) | -1.96 (-1.25) | 1.09 (0.88) | -1.63 (-1.53) |
| BLS | 3.49 (1.56) | 0.83 (0.48) | 14.40 (4.82) | 2.26 (1.11) | -2.45 (-1.39) |
| CHV | 2.09 (1.39) | -0.17 (-0.22) | 2.34 (1.19) | 21.20 (6.72) | 2.84 (2.83) |
| DD | -0.93 (-0.68) | 0.66 (0.57) | 6.18 (3.60) | 1.55 (1.05) | 11.70 (8.72) |
| DIS | 1.98 (1.85) | 2.73 (2.84) | -0.30 (-0.09) | -2.59 (-2.09) | 0.92 (1.04) |
| DOW | -0.79 (-0.77) | 0.18 (0.40) | -3.88 (-2.19) | -0.23 (-0.18) | -1.59 (-2.13) |
| F | -0.23 (-0.09) | 0.66 (0.50) | 3.29 (0.82) | -0.65 (-0.29) | 4.24 (2.19) |
| FNM | 0.01 (0.00) | -0.92 (-0.48) | -0.43 (-0.11) | 3.27 (1.31) | 0.39 (0.17) |
| GE | 0.77 (0.66) | -0.45 (-0.73) | 0.53 (0.59) | 0.03 (0.04) | -0.41 (-0.92) |
| GM | -2.83 (-1.72) | -0.86 (-1.23) | 0.23 (0.37) | 0.28 (0.37) | -2.99 (-4.25) |
| HWP | -1.66 (-1.00) | 0.53 (0.46) | -2.73 (-1.11) | 0.73 (0.55) | 2.35 (3.03) |
| IBM | -0.09 (-0.09) | 1.00 (1.53) | -0.77 (-0.47) | -1.33 (-1.28) | -0.24 (-0.84) |
| JNJ | -1.39 (-0.67) | 1.26 (1.70) | 0.13 (0.07) | 2.07 (2.11) | 2.40 (1.78) |
| KO | 1.11 (0.99) | 0.61 (0.83) | 0.69 (0.40) | -0.42 (-0.60) | 1.44 (2.58) |
| MCD | 3.38 (1.55) | -3.73 (-2.34) | -0.24 (-0.06) | -2.48 (-2.15) | -1.08 (-0.74) |
| MO | -0.63 (-1.02) | -0.84 (-0.71) | 1.87 (2.05) | -1.75 (-1.08) | 0.25 (0.22) |
| MOB | 0.36 (0.34) | -0.12 (-0.11) | -0.98 (-0.55) | -0.92 (-1.03) | 0.41 (0.48) |
| MRK | -0.75 (-1.41) | 0.55 (0.58) | -0.29 (-0.22) | -0.18 (-0.11) | 1.30 (1.67) |
| PEP | -0.85 (-0.48) | 1.75 (1.74) | 0.53 (0.21) | -1.93 (-0.90) | -0.28 (-0.39) |
| PG | -1.48 (-1.25) | -1.94 (-1.75) | 0.89 (0.47) | 3.51 (2.56) | 2.70 (2.29) |
| S | 0.01 (0.00) | -1.76 (-1.25) | -6.90 (-2.00) | 1.13 (0.43) | -0.30 (-0.20) |
| T | 1.84 (2.67) | 0.70 (0.77) | -1.18 (-0.74) | -0.20 (-0.36) | 0.59 (1.04) |
| WMT | 3.70 (2.02) | 2.01 (1.15) | 2.36 (0.69) | -3.65 (-1.01) | 0.96 (0.46) |
| XON | 1.45 (2.18) | -0.14 (-0.26) | 0.77 (0.61) | -0.16 (-0.29) | 0.09 (0.17) |
| T | 726 | 887 | 768 | 791 | 759 |
| R^2 | 0.06 | 0.19 | 0.03 | 0.04 | 0.17 |

Table 5. Ratios of the sum of squared residuals of the unconstrained and constrained price-impact regressions for 25 stocks, based on market orders from 2 January to 31 December 1996.

| Ticker | Ratio | Ticker | Ratio | Ticker | Ratio | Ticker | Ratio | Ticker | Ratio |
|--------|--------|--------|--------|--------|--------|--------|--------|--------|--------|
| AHP | 1.0501 | DIS | 1.0501 | GM | 1.0501 | MCD | 1.0501 | PG | 1.0500 |
| AN | 1.0498 | DOW | 1.0507 | HWP | 1.0416 | MO | 1.0499 | S | 1.0500 |
| BLS | 1.0500 | F | 1.0500 | IBM | 1.0508 | MOB | 1.0479 | T | 1.0281 |
| CHV | 1.0501 | FNM | 1.0497 | JNJ | 1.0501 | MRK | 1.0496 | WMT | 1.0500 |
| DD | 1.0501 | GE | 1.0500 | KO | 1.0494 | PEP | 1.0500 | XON | 1.0516 |

Table 6. Monte Carlo simulations of optimal execution strategies with short-sales constraints for purchasing \bar{s} shares of all 25 stocks in the portfolio over T periods. Each row corresponds to an independent simulation experiment consisting of 10,000 IID replications. Costs are in cents per share. T is the number of execution periods. \bar{s} is the amount of shares to purchase. k_v is the scaling of the volatility. s^* is the average execution cost for the optimal strategy with no short-sale constraints. s_c^* is the average execution cost for the optimal strategy with short-sale constraints. \bar{s}/T is the naive strategy's average cost. Standard errors are in parentheses below the main entries.

| T | \bar{s} | k_v | s^* | s_c^* | \bar{s}/T |
|-----|-----------|-------|-----------------|-----------------|-----------------|
| 20 | 10,000 | 1 | -8.05 (0.99) | 0.87 (0.78) | 1.53 (0.47) |
| 20 | 20,000 | 1 | -2.83 (0.59) | 0.21 (0.45) | 1.86 (0.47) |
| 20 | 30,000 | 1 | -0.87 (0.50) | 0.44 (0.41) | 2.19 (0.47) |
| 20 | 50,000 | 1 | 1.09 (0.46) | 1.35 (0.42) | 2.85 (0.47) |
| 20 | 100,000 | 1 | 3.71 (0.46) | 3.69 (0.45) | 4.50 (0.47) |
| 20 | 200,000 | 1 | 7.50 (0.46) | 7.48 (0.46) | 7.80 (0.47) |
| 20 | 100,000 | 0.5 | 3.04 (0.32) | 3.04 (0.32) | 3.26 (0.33) |
| 20 | 100,000 | 2 | 3.79 (0.62) | 3.75 (0.61) | 3.45 (0.66) |
| 20 | 100,000 | 4 | 4.79 (0.84) | 4.77 (0.84) | 4.05 (0.93) |
| 5 | 100,000 | 1 | 13.72 (0.26) | 13.55 (0.26) | 13.63 (0.26) |
| 10 | 100,000 | 1 | 7.08 (0.34) | 6.91 (0.34) | 7.19 (0.34) |
| 15 | 100,000 | 1 | 4.95 (0.41) | 4.81 (0.40) | 5.20 (0.41) |
| 20 | 100,000 | 1 | 3.71 (0.46) | 3.69 (0.45) | 4.50 (0.47) |

lected from a random trading day in 1996).

To gauge the sensitivity of execution costs to the model's parameters, we vary the time horizon, T , the number of shares traded, \bar{s} (assumed the same for each stock), and the no-impact price volatility. We modify the price volatility by scaling the variances by a constant while keeping the correlation structure fixed. The results for 10,000 replications are in Table 6, which reports the expected execution cost in cents per share. The table also lists the standard error of the estimate for the best-execution strategy (s^*), the strategy under no-sales constraints computed through the static-approximation method (s_c^*), and the naive strategy (\bar{s}/T).

Some general patterns emerge from the simulations. First, as T increases, execution costs fall. Because we can spread the trading over more time periods, and because we have the flexibility to be more patient and wait for particularly opportune times to trade, expected costs decline. Second, as \bar{s} decreases, execution costs also decrease. With small-enough trade sizes, the expected price impact is negative! This is because the price impact consists of two terms: the impact of shares traded, s_t , which is quadratic in the share size, and the impact of information, which is linear in the information variable, x_t . When we trade small-enough quantities, the quadratic term is negligible and the information term dominates. Our strategy optimally uses information so as to trade when trading is least expensive. For sufficiently significant pieces of information, trading can be quite profitable (not a new insight to proprietary traders). Finally, increasing volatility seems to increase execution costs slightly.

In all but two cases, the optimal strategy outperforms the naive on average. In the two anomalous cases, the confidence interval of the difference between the two strategies is so wide that this outcome could easily have occurred purely by chance. If we increased the number of repli-

cations to 100,000, these two anomalies would no doubt disappear.

Another anomalous result is that for some simulations, the constrained strategy's execution cost is less than the unconstrained strategy's cost. Although the point estimates are indeed reversed in these cases, the sampling variation is so great (consider their standard errors) that making accurate inferences about their relative magnitudes is difficult. Indeed, the differences are not statistically significant. For the cases we consider, the no-sales constraint seems to have relatively little impact on the best-execution strategy's performance, except in cases with negative execution costs. To achieve negative execution costs, the no-sales constraints must be violated, so imposing them increases the costs dramatically.

Of course, these conclusions are highly portfolio- and time-period-specific. Similar analyses should be conducted case by case to determine the value added by the best-execution strategy in a given context.

The remaining challenge is to integrate these best-execution strategies directly into the investment process, which requires solving the portfolio optimization problem subject to transactions costs. This is a formidable challenge that is both theoretically and computationally intensive, and we plan to turn to these problems in future research. ■

Acknowledgments

This research is part of the MIT Laboratory for Financial Engineering's Transactions Costs Project. We are grateful to Investment Technology Group, the National Science Foundation (Grants DMI-9610486 and SBR-9709976), and the sponsors of the LFE for financial support. We thank Dave Cushing, Chris Darnell, Robert Ferstenberg, Rohit D'Souza, John Heaton, Leonid Kogan, Bruce Lehmann, Greg Peterson, Jiang Wang, and seminar participants at Columbia, Harvard, the London Business School, MIT, Northwestern, the NYSE Conference on Best Execution, the Society of Quantitative Analysts, and Yale for helpful comments and discussion.

References

1. T. Loeb, "Trading Cost: The Critical Link between Investment Information and Results," *Financial Analysts J.*, Vol. 39, 1983, pp. 39-44.
2. L. Chan and J. Lakonishok, "The Behavior of Stock Prices around Institutional Trades," *J. Finance*, Vol. 50, 1995, pp. 1147-1174.
3. D. Keim and A. Madhavan, "Anatomy of the Trading Process: Empirical Evidence on the Behavior of Institutional Traders," *J.*

Financial Economics, Vol. 37, No. 3, Mar. 1995, pp. 371-398.

4. D. Keim and A. Madhavan, "The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects," *Rev. Financial Studies*, Vol. 9, No. 1, Spring 1996, pp. 1-36.
5. D. Keim and A. Madhavan, "Execution Costs and Investment Performance: An Empirical Analysis of Institutional Equity Trades," working paper, School of Business Administration, Univ. of Southern California, Los Angeles, 1995.
6. D. Bertsimas and A. Lo, "Optimal Control of Execution Costs," *J. Financial Markets*, Vol. 1, No. 1, 30 Apr. 1998, pp. 1-50.
7. A. Shleifer, "Do Demand Curves for Stocks Slope Down?" *J. Finance*, Vol. 41, 1986, pp. 579-590.
8. A. Lo and C. MacKinlay, "When Are Contrarian Profits Due to Stock Market Overreaction?" *Rev. Financial Studies*, Vol. 3, No. 2, Summer 1990, pp. 175-205.
9. L. Birinyi, *What Does Institutional Trading Cost?* Birinyi Associates, Greenwich, Conn., 1995.
10. D. Leinweber, "Using Information from Trading in Trading and Portfolio Management," *Execution Techniques, True Trading Costs, and the Microstructure of Markets*, K. Sherrerd, ed., Assoc. for Investment Management and Research, Charlottesville, Va., 1993.
11. D. Leinweber, "Careful Structuring Reins in Transaction Costs," *Pensions and Investments*, 25 July 1994, p. 19.
12. A. Perold, "The Implementation Shortfall: Paper versus Reality," *J. Portfolio Management*, Vol. 14, No. 3, Spring 1988, pp. 4-9.
13. D. Bertsekas, *Dynamic Programming and Optimal Control*, Vol. 1, Athena Scientific, Belmont, Mass., 1995.

Dimitris Bertsimas is the Boeing Professor of Operations Research at the MIT Sloan School of Management. He studies the theory and practice of optimization, as well as performance, analysis, and control of large-scale stochastic systems. He received his PhD in operations research and applied mathematics from MIT. Contact him at the MIT Sloan School of Management and Operations Research Center, Massachusetts Inst. of Technology, E53-363, Cambridge, MA 02142-1347; dbertsim@mit.edu.

Paul Hummel is a portfolio strategist at Long Term Capital Management. His responsibilities include trading in risk arbitrage and special situations; designing automated spread-trading, volume-weighted-average-price, and options-hedging strategies; and research support for short-term, high-frequency trading strategies. He received his MS in mechanical engineering from Cornell University and his MBA from the MIT Sloan School of Management. Contact him at Portfolio Strategist, Long Term Capital Management, One E. Weaver St., Greenwich, CT 06831; phummel@aol.com.

Andrew W. Lo's biography appears on p. 23.