Optimal Execution for Portfolio Transitions

Minimizing opportunity cost.

Mark Kritzman, Simon Myrgren, and Sébastien Page

Institutional investors periodically reallocate their portfolios to shift asset mixes or change their investment managers. They face a variety of costs when they undertake these transitions, including commissions, bid-ask spreads, opportunity cost, and market impact (see Perold [1988] and Almgren and Chris [2000]).

Opportunity cost refers to adverse changes in price arising from exogenous market forces, while market impact refers to adverse price movements that occur in response to the purchase and sale of securities. Opportunity cost and market impact are especially interesting because these costs represent the largest share of total cost and because investors influence these costs by how they trade.

We introduce an algorithm for determining the optimal sequence and size of trades that minimize opportunity cost for portfolio transitions. Our algorithm differs from other models intended to minimize opportunity cost because we are concerned with the relative performance of the legacy and the target portfolios rather than the absolute performance of a single security or a group of securities targeted for purchase or sale. Moreover, we assume the investor does not use leverage to execute the transition; hence we constrain the trades to be self-financing.

MARK KRITZMAN
is President and CEO at Windham Capital Management, LLC in Cambridge, MA.
mkritzman@windhamcapital.com

SIMON MYRGREN
is Vice President at State Street Associates in Cambridge, MA.
smyrgren@statestreet.com

SÉBASTIEN PAGE
is Managing Director at State Street Associates in Cambridge, MA.
spage@statestreet.com

ALGORITHM FOR MINIMIZING OPPORTUNITY COST

When investors reallocate portfolios, the single greatest risk they face is that the legacy portfolio will drift down in price relative to the target portfolio or that the target portfolio will drift up in price relative to the legacy portfolio before the transition is complete. The realization of this risk, or opportunity cost, is directly proportional to the tracking error between the legacy and target portfolios. Low tracking error reduces the likelihood of adverse divergent performance, while high tracking error increases it. The optimal sequence and size of trades will constantly minimize tracking error between the legacy and target portfolios as the transition unfolds.

We begin by defining tracking error (squared) between the legacy and the target portfolios as a function of the securities in the legacy and target portfolios and as a function of the proposed trades:[1]

$$TE^2 = \left(w_R^k + w_{MN}\right)\Sigma\left(w_R^k + w_{MN}\right)^T \tag{1}$$

where:

- w_R^k = vector of relative weights of the securities in the legacy and target portfolios after the k-th trade ($w_R^k = w_L^k - w_T^k$);
- w_{MN} = vector of weights associated with selling security M and buying security N; and
- Σ = covariance matrix of the securities in the legacy and target portfolios after the k-th trade.

For example, if we were to consider selling the first security, buying the third security, and maintaining our positions in the second and fourth security, the vector of weights for the proposed trades would have the format as follows:

$$w_{MN} = [-w \quad 0 \quad w \quad 0] \tag{2}$$

Equation (2) indicates that we are selling the same amount of security A as we are buying of security C. This equality ensures that the transition is self-financing; that is, we do not need to use leverage in order to execute any trades, which is a typical constraint faced by transition managers.

We next calculate the partial derivative of tracking error (squared) with respect to the proposed trades:

$$\frac{dTE^2}{dw_{MN}} = 2\Sigma w_R^k + 2\Sigma w_{MN} \tag{3}$$

These derivatives measure the sensitivity to trades representing offsetting purchases and sales. The derivative with the highest value reduces tracking error the most, and the one with the lowest value reduces it the least. We should therefore execute the trade with the highest derivative value first and the one with the lowest derivative value last. But we also need to determine the quantities to be bought and sold for each trade. Assuming a legacy portfolio of securities A and B and a target portfolio of securities C and D, we rewrite Equation (3) as derivatives with respect to the individual security weights, as shown in Equations (4) through (7):

$$dTE_A^2 = \alpha_A + \beta_A w \tag{4}$$

$$dTE_B^2 = \alpha_B + \beta_B w \tag{5}$$

$$dTE_C^2 = \alpha_C + \beta_C w \tag{6}$$

$$dTE_D^2 = \alpha_D + \beta_D w \tag{7}$$

Equations (4) through (7) partition the derivatives into two components, one associated with the remaining securities in the legacy and target portfolios (α), and the other associated with the securities in the proposed trade (β). These equations reveal that tracking error minimization is a linear function of trade size.

We can rewrite the coefficients for security A for a particular trade such as the sale of security B and the purchase of security D, as below:

$$\alpha_A = \Sigma_{11}w_{R,1}^0 + \Sigma_{12}w_{R,2}^0 + \Sigma_{13}w_{R,3}^0 + \Sigma_{14}w_{R,4}^0 \tag{8}$$

$$\beta_A = -\Sigma_{12} + \Sigma_{14} \tag{9}$$

Finally, to find the optimal trade size, we set the derivatives of the securities to be sold equal to each other, ($dTE_A^2 = dTE_B^2$) and we do the same for the securities intended for purchase. The purpose in setting these derivatives equal to one another is to identify the trade amount beyond which it is optimal to shift from trading one security to the other.[2]

Equations (10) and (11) show these calculations:

$$\alpha_B + \beta_B w = \alpha_A + \beta_A w \qquad w = \frac{\alpha_B - \alpha_A}{\beta_A - \beta_B} \tag{10}$$

$$\alpha_D + \beta_D w = \alpha_C + \beta_C w \qquad w = \frac{\alpha_D - \alpha_C}{\beta_C - \beta_D} \tag{11}$$

Once we determine the optimal trade size for the security to be sold and the security to be purchased, we choose the smaller amount so that we do not trade more than the optimal amount in the other security. The optimal trade size might require a larger sale of a security than the amount owned or a larger purchase of a security than the amount targeted for purchase. It may even require the purchase of a security that ultimately must be sold or the sale of a security targeted for purchase. We are not required to carry out these transactions, however—we could simply trade up to predetermined limits.

Next we rebalance the legacy and target portfolios in accordance with the optimal trades and recalculate tracking error. We find the new derivatives based on the rebalanced portfolio and determine the next optimal trade. We proceed with these calculations until we have determined the sequence and size of trades that shift the legacy portfolio entirely to the target portfolio. All these calculations are performed in advance of the actual execution of the transition, thereby establishing a complete road map for the process.

As the transition unfolds, it is almost certain that opportunities will arise to trade some securities at advantageous prices. Alternatively, the investor may choose to postpone some trades due to unfavorable prices. If the investor departs from the road map, it is a simple matter to recalculate the derivatives and update the plan to account for out-of-sequence trades.

Numerical Example

Let us consider a simple case in which we wish to transition from a legacy portfolio of equal allocations to two securities, A and B, to a target portfolio of equal allocations to two securities, C and D. The portfolios are shown in Exhibit 1.

There are four possible trades to consider:

1. Sell A and buy C.
2. Sell A and buy D.
3. Sell B and buy C.
4. Sell B and buy D.

To implement our trading algorithm, we first estimate the full covariance matrix of the securities in the legacy and target portfolios, which is shown in Exhibit 2. With this information, we calculate the relation between tracking error and the amount traded for each of the four possible trades, which is shown in Exhibit 3.

Exhibit 3 indicates that we should trade our entire positions in securities A and D—that is, sell all of A and buy the entire targeted amount of security D—because this trade reduces tracking error more quickly than the other trades. Exhibit 3 fails to recognize, however, that as we execute this trade the derivatives shift. This point is apparent from Exhibit 4, which shows how the derivatives with respect to the security weights change as the security weights themselves change [see Equations (4) through (7)].

EXHIBIT 1
Legacy and Target Portfolios

Asset	Legacy (%)	Target (%)
A	50	0
B	50	0
C	0	50
D	0	50

EXHIBIT 2
Covariance Matrix

	A	B	C	D
A	0.0040	0.0009	0.0017	0.0012
B	0.0009	0.0029	0.0021	0.0060
C	0.0017	0.0021	0.0046	0.0081
D	0.0012	0.0060	0.0081	0.0266

EXHIBIT 3
Tracking Error versus Trade Size

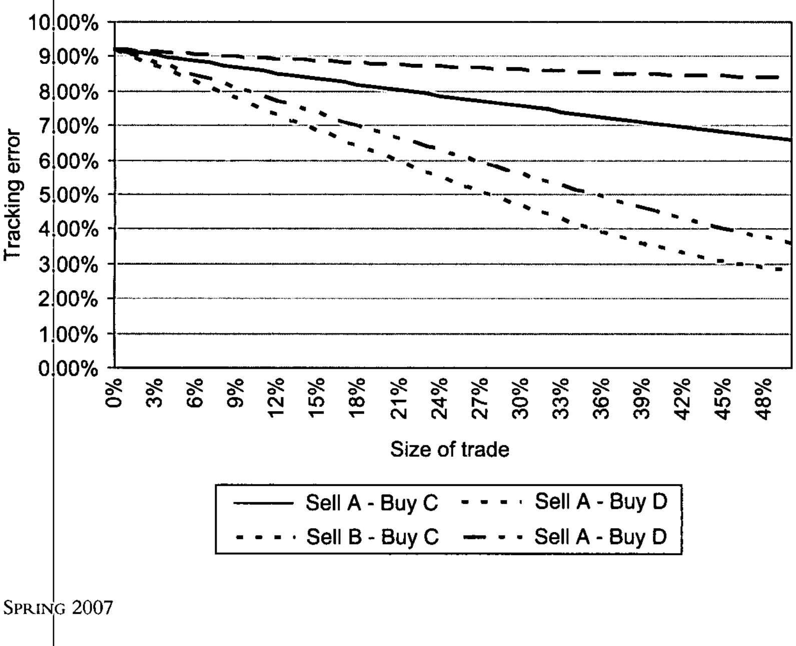

For all trade amounts up to 40.13%, the derivative of A is greater than that of B; hence, up to this threshold we reduce tracking error more quickly by selling security B than security A. The purchase of security D versus C is preferable for nearly the entire trade amount. Nonetheless, in order to preserve the self-financing nature of the transition, we trade only up to 40.13%. Hence, we should reduce our position in security A by 40.13 percentage points from 50% to 9.87% and purchase 40.13% of security D.

EXHIBIT 4
Derivative versus Trade Size (sell A and buy D)

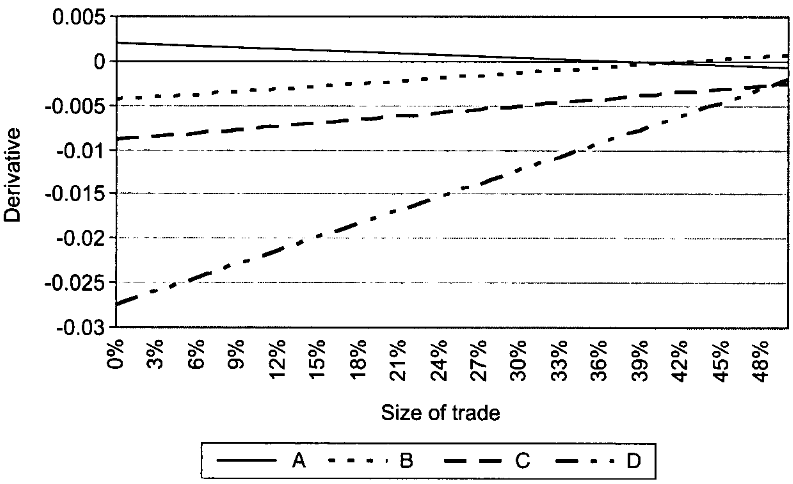

EXHIBIT 5
Tracking Error versus Trade Size

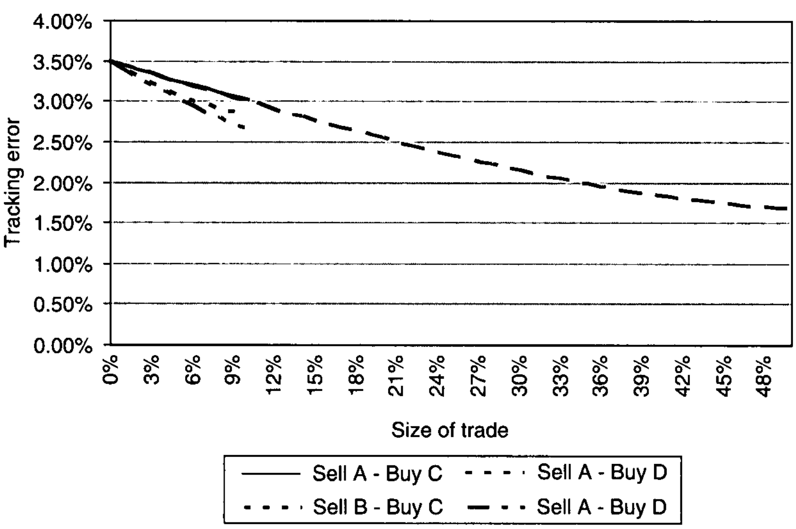

We next reestimate tracking error assuming the portfolios are rebalanced in accordance with the first recommended trade. Exhibit 5 once again shows how tracking error changes as a function of the potential trades that remain, given the new security weights.

Note that all the potential trades that include securities A and D are constrained to 9.87%, which is what remains to be completed in these securities after the initial trade of 40.13%. The trade that now reduces tracking error the most is to sell security B and to purchase security D.

Exhibit 6 again shows the derivatives of tracking error with respect to each security, given the new portfolio weights. The derivative with respect to security B is higher than that of security A for the entire potential trade up to 9.87%, and the derivative with respect to security D is lower than that of security C for the entire potential trade as well. Thus, we should sell 9.87% of security B and purchase 9.87% of security D, which results in a full allocation of 50% to security D.

The only security available for purchase now is security C, but we must decide whether it is preferable to sell security A or B. By proceeding in this fashion, we determine that we should sell 17.93% of security B. We continue in this fashion until we determine the optimal sequence and size of trades required to divest the legacy portfolio in its entirety and to acquire the target portfolio. This sequence of trades is shown in Exhibit 7.

It is important to note that all these calculations are carried out in advance of any trades, thus providing a road map for the transition. As the transition unfolds, the investor will likely deviate from this road map as market conditions warrant, but it is a simple matter to revise the calculations along the way in order to update the plan.

Our Algorithm versus the Industry Approach

The standard approach of investors who specialize in portfolio transition is to strive for sector neutrality; that is, to minimize the differences in sector weights between the legacy and target portfolios as the transition unfolds.

EXHIBIT 6

Derivative versus Trade Size (sell B and buy D)

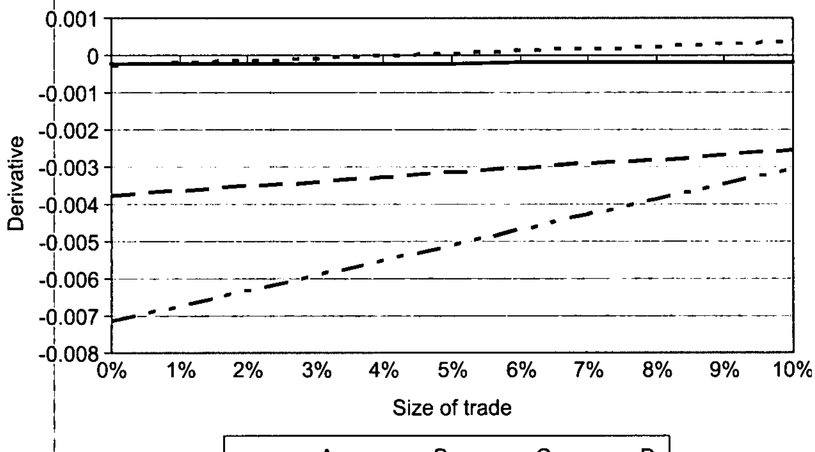

EXHIBIT 7

Optimal Trades Without Market Impact

Sells	Buys	Size
A	D	40.13%
B	D	9.87%
B	C	17.93%
A	C	9.87%
B	C	22.20%

EXHIBIT 8

Tracking Error: Optimal versus Sector-Based Trades

Fraction Completed	Optimal Trades	Sector-Based Trades	Difference
0%	4.94%	4.94%	0.00%
20%	2.97%	4.35%	-1.38%
40%	1.92%	3.78%	-1.86%
60%	1.12%	2.60%	-1.48%
80%	0.52%	1.42%	-0.89%
100%	0.00%	0.00%	0.00%

To test the efficacy of our approach, we compare the evolution of tracking error of a hypothetical transition using the sector approach with the evolution of tracking error of the same transition using our algorithm. In this example, we assume the investor wishes to reallocate from a hypothetical growth portfolio to a hypothetical value portfolio.[3]

We estimate the covariance matrix from one year of monthly returns beginning January 31, 2005, and ending January 31, 2006. We first calculate the optimal sequence and size of trades using our algorithm. Then, in order to apply the sector approach, we trade the same amount of each sector as the calculated optimal trade size at that step, beginning with the sector that has the greatest difference in weights and ending with the sector for which the weights are most similar. Within sectors we start by trading the security with the largest holding and proceed to the security with the smallest holding.

Exhibit 8 shows the tracking error between the legacy and target portfolios for both approaches at 20% increments of completion, and Exhibit 9 shows in more detail how tracking error evolves with the fraction completed for both strategies.

Our trading algorithm reduces tracking error and hence opportunity cost more quickly than the sector-based approach. We can determine precisely how much we reduce opportunity cost by comparing the integrals under each curve in Exhibit 9. The integral under the sector-based approach is 2.93%, and for the optimal approach it is 1.77%. Hence the optimal approach yields a 40% improvement over the industry norm.

We acknowledge that these results are sample-specific, but we believe this particular sample is broadly representative, given the point we want to illustrate.

CONSIDERATION OF MARKET IMPACT

We introduce market impact as a second step in our trading process after we determine the sequence and size of trades that minimize opportunity cost. We are required to consider opportunity cost and market impact sequentially rather than simultaneously because we consider only pair trades. If instead we solve simultaneously for both the optimal combination of a security to sell and to buy and the optimal size of this trade, we would have to perform an optimization for every possible combination of securities intended for sale and purchase.

If, for example, the legacy and target portfolios include 250 securities each, we would have to perform 62,500 optimizations. Moreover, we would have to devise an optimization algorithm that would determine when it is optimal to switch to another pair of securities to buy and sell.[4]

In light of persuasive empirical evidence, we model market impact as a function of the square root of the ratio of trade size to average daily volume:[5]

$$MI(s_m) = \kappa_m \sqrt{\frac{s_m}{adv_m}} \tag{12}$$

where:

κ_m = a constant that is determined empirically;
s_m = number of shares traded in security m; and
adv_m = average daily volume traded in security m.

We modify Equation (12) to convert cost from an absolute amount to a percentage of the portfolio's value, and to reflect the fact that the security positions change throughout the transition:

$$MI\left(w_{R,m}^k\right) = -\frac{\kappa_m}{MV_p w_{R,m}^k} \sqrt{\frac{s_m^0 w_{R,m}^k}{w_{R,m}^0 adv_m}} \tag{13}$$

where:

$w_{R,m}^k$ = relative weight in security m at the k-th step of the transition, and
MV_p = market value of the portfolio.

Next we multiply market impact expressed as a percentage cost by the portfolio weight of each security, which gives the market impact cost associated with trading w_m^k in asset m at step k:

$$MICost\left(w_m^k, w_{R,m}^k\right) = w_m^k MI\left(w_{R,m}^k\right) \tag{14}$$

We next add this term to our function that minimizes tracking error (squared) in order to account for market impact:

$$C(w_{MN}) = TE^2 + MICost\left(w_{MN}, w_R^k\right)$$
$$C(w_{MN}) = \left(w_R^k + w_{MN}\right)\Sigma\left(w_R^k + w_{MN}\right)^T + w_{MN} MI\left(w_R^k\right) \tag{15}$$

Because every trade involves only two securities, one from the legacy portfolio and one from the target portfolio, the market impact term affects only the derivatives for these two securities. Taking the derivatives and performing the same simple algebraic manipulations as before, we find the solutions for the optimal trade size that minimizes both opportunity cost and market impact:

$$w = \frac{\alpha_B - \alpha_A + MI(w_{R,B}^k)}{\beta_A - \beta_C}$$
$$w = \frac{\alpha_D - \alpha_C + MI(w_{R,D}^k)}{\beta_C - \beta_D} \tag{16}$$

These equations assume we have already determined it is optimal to sell a particular amount of security B and

EXHIBIT 9
Tracking Error versus Fraction Completed

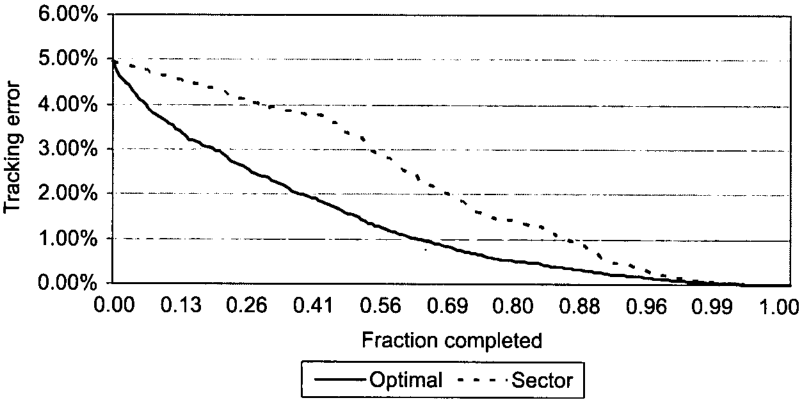

EXHIBIT 10
Assumptions for Market Impact Analysis

	A	B	C	D
Shares	400	200	250	800
ADV	850	1100	650	3700
Price	$25.00	$50.00	$40.00	$12.50
MV	$10,000	$10,000	$10,000	$10,000
Shares/ADV	47.06%	18.18%	38.46%	21.62%
Kappa	10.00			

EXHIBIT 11
Optimal Trades With Market Impact

Sells	Buys	Size
A	D	35.78%
A	D	4.34%
B	D	7.87%
B	D	2.00%
B	C	3.32%
B	C	14.61%
A	C	9.87%
B	C	22.20%

to buy the same amount of security D. The market impact component simply determines how to split this trade into smaller components.

We apply a set of arbitrary assumptions shown in Exhibit 10 to determine the new set of optimal trades, taking market impact into consideration, as shown in Exhibit 11.

CONCLUSION

When investors change their asset mix or switch one asset manager for another, they incur transition costs. The dominant components of these costs are opportunity cost and market impact. Opportunity cost, which is the single greatest risk for a transition, is directly related to the tracking error between the legacy and target portfolios.

We propose an algorithm for determining the sequence and size of trades that minimize tracking error as the transition unfolds. Our algorithm ensures that the trades are self-financing so that there is no need to resort to leverage in order to execute the transition. Using a realistic example of a portfolio transition, we demonstrate that our algorithm exposes the transition to about 40% lower tracking error than the industry norm, which is to minimize sector differences as the transition is executed.

We also augment our algorithm to account for market impact. We do not consider market impact simultaneously with opportunity cost, because that is not a tractable approach, given that we are restricted to self-financing pair trades. Rather, we first determine the optimal trades on the basis of opportunity cost, and then as a second step introduce a market impact function to determine how to partition these trades into smaller units.

Our algorithm to reduce opportunity cost depends only on the covariance matrix of the securities in the legacy and target portfolios. It does not consider assumptions about expected return. Nonetheless, we can adapt the algorithm to incorporate investor views about security returns, in which case we would shift from minimizing tracking error to maximizing expected utility for a given risk aversion coefficient.

Finally, our algorithm is easy to apply interactively. Although it produces a complete road map for portfolio transitions in advance of execution, it is easy to recalculate the approach at any point during the transition in the event market conditions warrant occasional departures from prescribed trades.

ENDNOTES

The authors thank Justin Balogh, Neil Constable, Cel Kulasekaran, Ross McLellan, Anya Obizhaeva, and especially Edouard Stirling for helpful comments.

[1]Our algorithm for portfolio transition is motivated by Sharpe [1987]. Sharpe shows that by calculating the derivatives of quadratic utility with respect to portfolio weights and then continually shifting the portfolio weights from the security with the lowest derivative to the one with the highest derivative, one arrives at the mean-variance efficient portfolio when all the derivatives are equal. We apply this insight to derive our algorithm for minimizing tracking error as a function of the sequence and size of trades.

[2]In other words, we optimize the buy and sell decisions independently, but we select the smaller trade, because if we were to select the larger trade part of it would be suboptimal.

[3]These portfolios are based upon Citigroup's S&P 500 Growth Index and Value Index, where we have chosen the top one-third of the securities in each index ranked by market capitalization.

[4]If we were determined to address market impact simultaneously with our algorithm to minimize opportunity cost, we could impose a size constraint on our algorithm. We prefer, however, to model market impact explicitly and to introduce it as a second step to our algorithm.

[5]See, for example, Barra's "Market Impact Handbook" [1997], Grinold and Kahn [1999], and Loeb [1983].

REFERENCES

Almgren, R., and N. Chriss. "Optimal Execution of Portfolio Transactions." *Risk*, 3 (2) 2000.

Grinold, R.C., and R.N. Kahn. *Active Portfolio Management.* New York: McGraw-Hill, 1999.

Loeb, T.F. "Trading Costs: The Critical Link Between Investment Information and Results." *Financial Analysts Journal*, 39(3) (1983), pp. 39-44.

"Market Impact Handbook." Barra, 1997. Available at http://www.barra.com/Newsletter/nl168/mim4-168.asp.

Perold, A. "The Implementation Shortfall: Paper versus Reality." *The Journal of Portfolio Management*, Spring 1988.

Sharpe, W. "An Algorithm for Portfolio Improvement." *Advances in Mathematical Programming and Financial Planning*, 1 (1987), pp. 155-169.