The Capacity of Liquidity-Demanding Equity Strategies

VITALY SERBIN, PETER M. BULL, AND HOWARD ZHU

VITALY SERBIN
is head of the Portfolio Analytics Group at Investment Technology Group, Inc., in Boston, MA.
vitaly.serbin@itg.com

PETER M. BULL
is a portfolio manager and analyst at Realindex Investments in Sydney, Australia.
peter.bull@realindexinvestments.com.au

HOWARD ZHU
is an analytical product specialist at MSCI Inc., in Hong Kong.
howard.zhu@mscibarra.com

In light of the significant growth of the U.S. mutual fund industry in recent years, with over $10 trillion under management as of October 2006 (Maxey [2006]), an understanding of the sources of fund performance has never been more critical for investors, regulators, and investment managers. Beyond this growth in assets, the issue of capacity becomes increasingly important as traditionally passive products, such as ETFs, become structured on higher turnover rules or indices that no longer imply passive trading (Authers [2006]). These developments coincide with the ever-increasing mobility of investor capital through global markets in pursuit of higher returns (Rüffer and Stracca [2006]). As the term implies, "excess liquidity" exceeds available strategies' capacity to provide risk-adjusted returns above a certain threshold.

Perold and Salomon [1991] were the first to focus both academicians and practitioners on the effect of the amount of assets under management on fund performance. Most authors document the inverse relationship between a fund's size and its net return. The dominant explanation in the literature attributes diminishing returns to size to rising transaction/turnover costs. Chan et al. [2004] showed that fund excess returns decrease with the amount of capital employed. While higher turnover is often necessary to utilize the informational advantages that a superior manager might have (e.g., Wermers [2000]), Perold and Salomon [1991], and Vangelisti [2006] showed that performance deteriorates as fund managers must turn over larger volumes of stock, incurring explicit (commission) and implicit (market-impact, front-running, and so forth) transaction costs.

Our contribution to the literature is twofold. First, we quantify market impact costs on an individual stock level, thus providing much better granularity regarding their source. Combining stock-specific estimates of trading costs with stock-specific alphas provides a more realistic way to analyze optimal capacity. We believe the industry is primed for this type of approach because the existing literature has thus far assumed highly simplified mechanisms for transaction costs. For example, Bogle [1994] multiplied the fund's turnover rate by two and then by 0.6% to estimate the lower bound of the microstructure drag-down on performance, and Vangelisti [2006] investigated fund capacity from a transaction cost perspective; the underlying model for transaction costs, however, remains a simple linear function of median daily volume (MDV).[1]

We calculate cost estimates using the ITG Agency Cost Estimator (ACE) model.[2] The input variables for this model include the relative order size (number of shares in proportion to MDV), time-weighted five-day average spread, and forward-looking daily volatility for each stock. The cost curves are also conditioned on the individual stock's liquidity level, split- and dividend-adjusted close-to-close returns, and intraday volume profile, as well as the trading intensity of specific orders. The ITG Cost Curves, although a commercial product, relies on a transparent methodology that is described in an ACE white paper (ITG [2007]). We believe that a different price-impact model should lead to the same qualitative conclusions, as long as the model recognizes the concave relationship between unit costs and volume and correctly accounts for different liquidity, volatility, order urgency, and spread dynamics. In addition, the parameters in a good price-impact model should be calibrated to match empirical trading costs. The ACE model used in this article was calibrated using 2004–2005 execution data from more than 80 large institutional clients. Making the transaction cost analysis stock specific allows us to explore the relationship between assets under management (AUM) and alpha across different segments of stocks and, potentially, across different markets.

Our second contribution is to formulate the capacity problem as an optimization problem in which users must accept the trade-off between higher AUM and lower net alpha subject to a monthly turnover constraint. In order to quantify that trade-off, we make the alpha-generating process transparent and strategy independent. Any strategy with a given information coefficient, IC, (see Grinold [1989]) will produce the same set of alphas.[3] We also implement a simple way of making alphas time sensitive and, therefore, turnover sensitive. Lifting the veil of opaqueness covering the portfolio formation process allows us to see how alphas, trading costs, and turnover combine to form an investment strategy, but it also allows for a reasonable degree of separation between their effects. As a consequence, our conclusions are not investment strategy dependent, and we are able to avoid a pitfall common to several existing studies. For example, the analysis of Vangelisti [2006] revolved around a proprietary investment strategy, but this fact is never made explicit.[4] The main problem with putting the portfolio construction mechanism into a black box is that alphas, AUM, and turnover become impossible to separate. Perold and Salomon [1991, Table II], however, assumed that transaction costs were independent of alpha; they thus failed to recognize the relationship between portfolio alpha and turnover.

The work of Kahn and Shaffer [2005] is the closest to ours in terms of motivation and methodology. In their framework, increasing assets under management requires a decrease in turnover and is accompanied by declining alpha (both gross and net). Kahn and Shaffer concluded that net alpha is not very sensitive to increases in capacity, provided that the money manager takes good care to control turnover. The main difference between our approach and that of Kahn and Shaffer is that we use a bottom-up approach for measuring costs and alpha, they used a top-down approach. More specifically, we simulate the next-month alpha for each stock and use a stock-specific estimate of market-impact cost for the same month. This approach allows us to relax two of the crucial assumptions of Kahn and Shaffer. First, as noted above, we account for the possibility that alpha and liquidity might be related. This assumption is more consistent with the evidence that high-alpha (e.g., high-momentum) stocks might be less liquid and thus more expensive to trade (Korajczyk and Sadka [2004]). Second, by using market-impact curves for the same period in which we measure alpha, we relax the assumption that trading costs are independent of the investment strategy. If competitors follow similar investment strategies, then stocks with high next-month alpha might become more expensive to trade. If many investors buy or sell a limited universe of "hot" stocks, spreading trades over more stocks, as Kahn and Shaffer [2005] suggested, might not be sufficient. Besides the limited potential for reducing costs, investing in more stocks might result in deterioration of gross alpha. In addition, many investment strategies, such as merger or index arbitrage, do not lend themselves to expanding the investable universe (Coval and Stafford [2005]).

The remainder of the article proceeds as follows. The next section presents capacity as a strategy-dependent problem and relates it to the underlying issue of liquidity. The section following that describes our simulation methodology and the corresponding results. The final section concludes.

THE MARKET-IMPACT PERSPECTIVE

In this section, we introduce the general framework for subsequent discussion. We demonstrate that the level of transaction costs and, therefore, optimal capacity is a strategy-specific phenomenon. We also provide a simple example to illustrate that optimal capacity could not be determined for each fund in isolation, since many money managers follow similar investment strategies at the same time. Discussion in this section is mostly conceptual, while in the next section we quantify the relationship between expected portfolio return, turnover rate, and optimal capacity.

We use the following key assumptions throughout the article: 1) the nominal return, or *paper return*, of a strategy is known or satisfactorily estimated a priori, and 2) trading volume in investable assets increases proportionately to AUM, so that the turnover rate is held constant while the strategy is implemented. We also hold constant other strategy inputs such as the commission rate. Under this framework, the performance drag that accompanies higher AUM is due solely to increased marginal market-impact costs. A forecast of the market-impact cost of a volume-weighted average price (VWAP) strategy provides a good definition of a stock's average (across trading venues) liquidity and, therefore, is appropriate for our analysis.[5] We can extrapolate these daily market-impact cost estimates and adjust them for leverage, cash drag, and so forth to complete the net expected return calculation for individual funds.

We believe that capacity should be defined only in terms of the capacity *of an investment strategy*. Acknowledging this statement leads us to realize that portfolio managers themselves might not always know the optimal capacity of their strategy. For example, the capacity of convertible arbitrage strategies might vary depending on the supply of convertible bonds.

We will illustrate how the choice of a particular investment strategy can affect a fund's optimal capacity. Suppose we have four trade lists representing distinct size-segments of the U.S. equity universe: equal- and market-cap weighted large-cap stocks (S&P 500, or SP500) as well as equal- and market-cap-weighted small-cap stocks (Russell 2000, or R2000). Not surprisingly, trading an equal-weighted R2000 list costs noticeably more than an index-weighted list, because more dollar volume is pushed into the less liquid names. Exhibit 1 illustrates the capacity implications of the volume-participation cost consistency and Exhibit 2 shows the dollar amounts traded for selected average percent MDV breakpoints. We obtained the market-impact costs depicted in Exhibits 1, 3, and 4 from ITG Cost Curves, calculated using the ACE methodology. For more details on the methodology, see ITG [2007]. Investment strategies and structures are routinely influenced, if not driven, by these types of liquidity characteristics.[6]

EXHIBIT 1

Daily Market-Impact Costs for U.S. Indices Trade Lists, $100 Million to $5 Billion

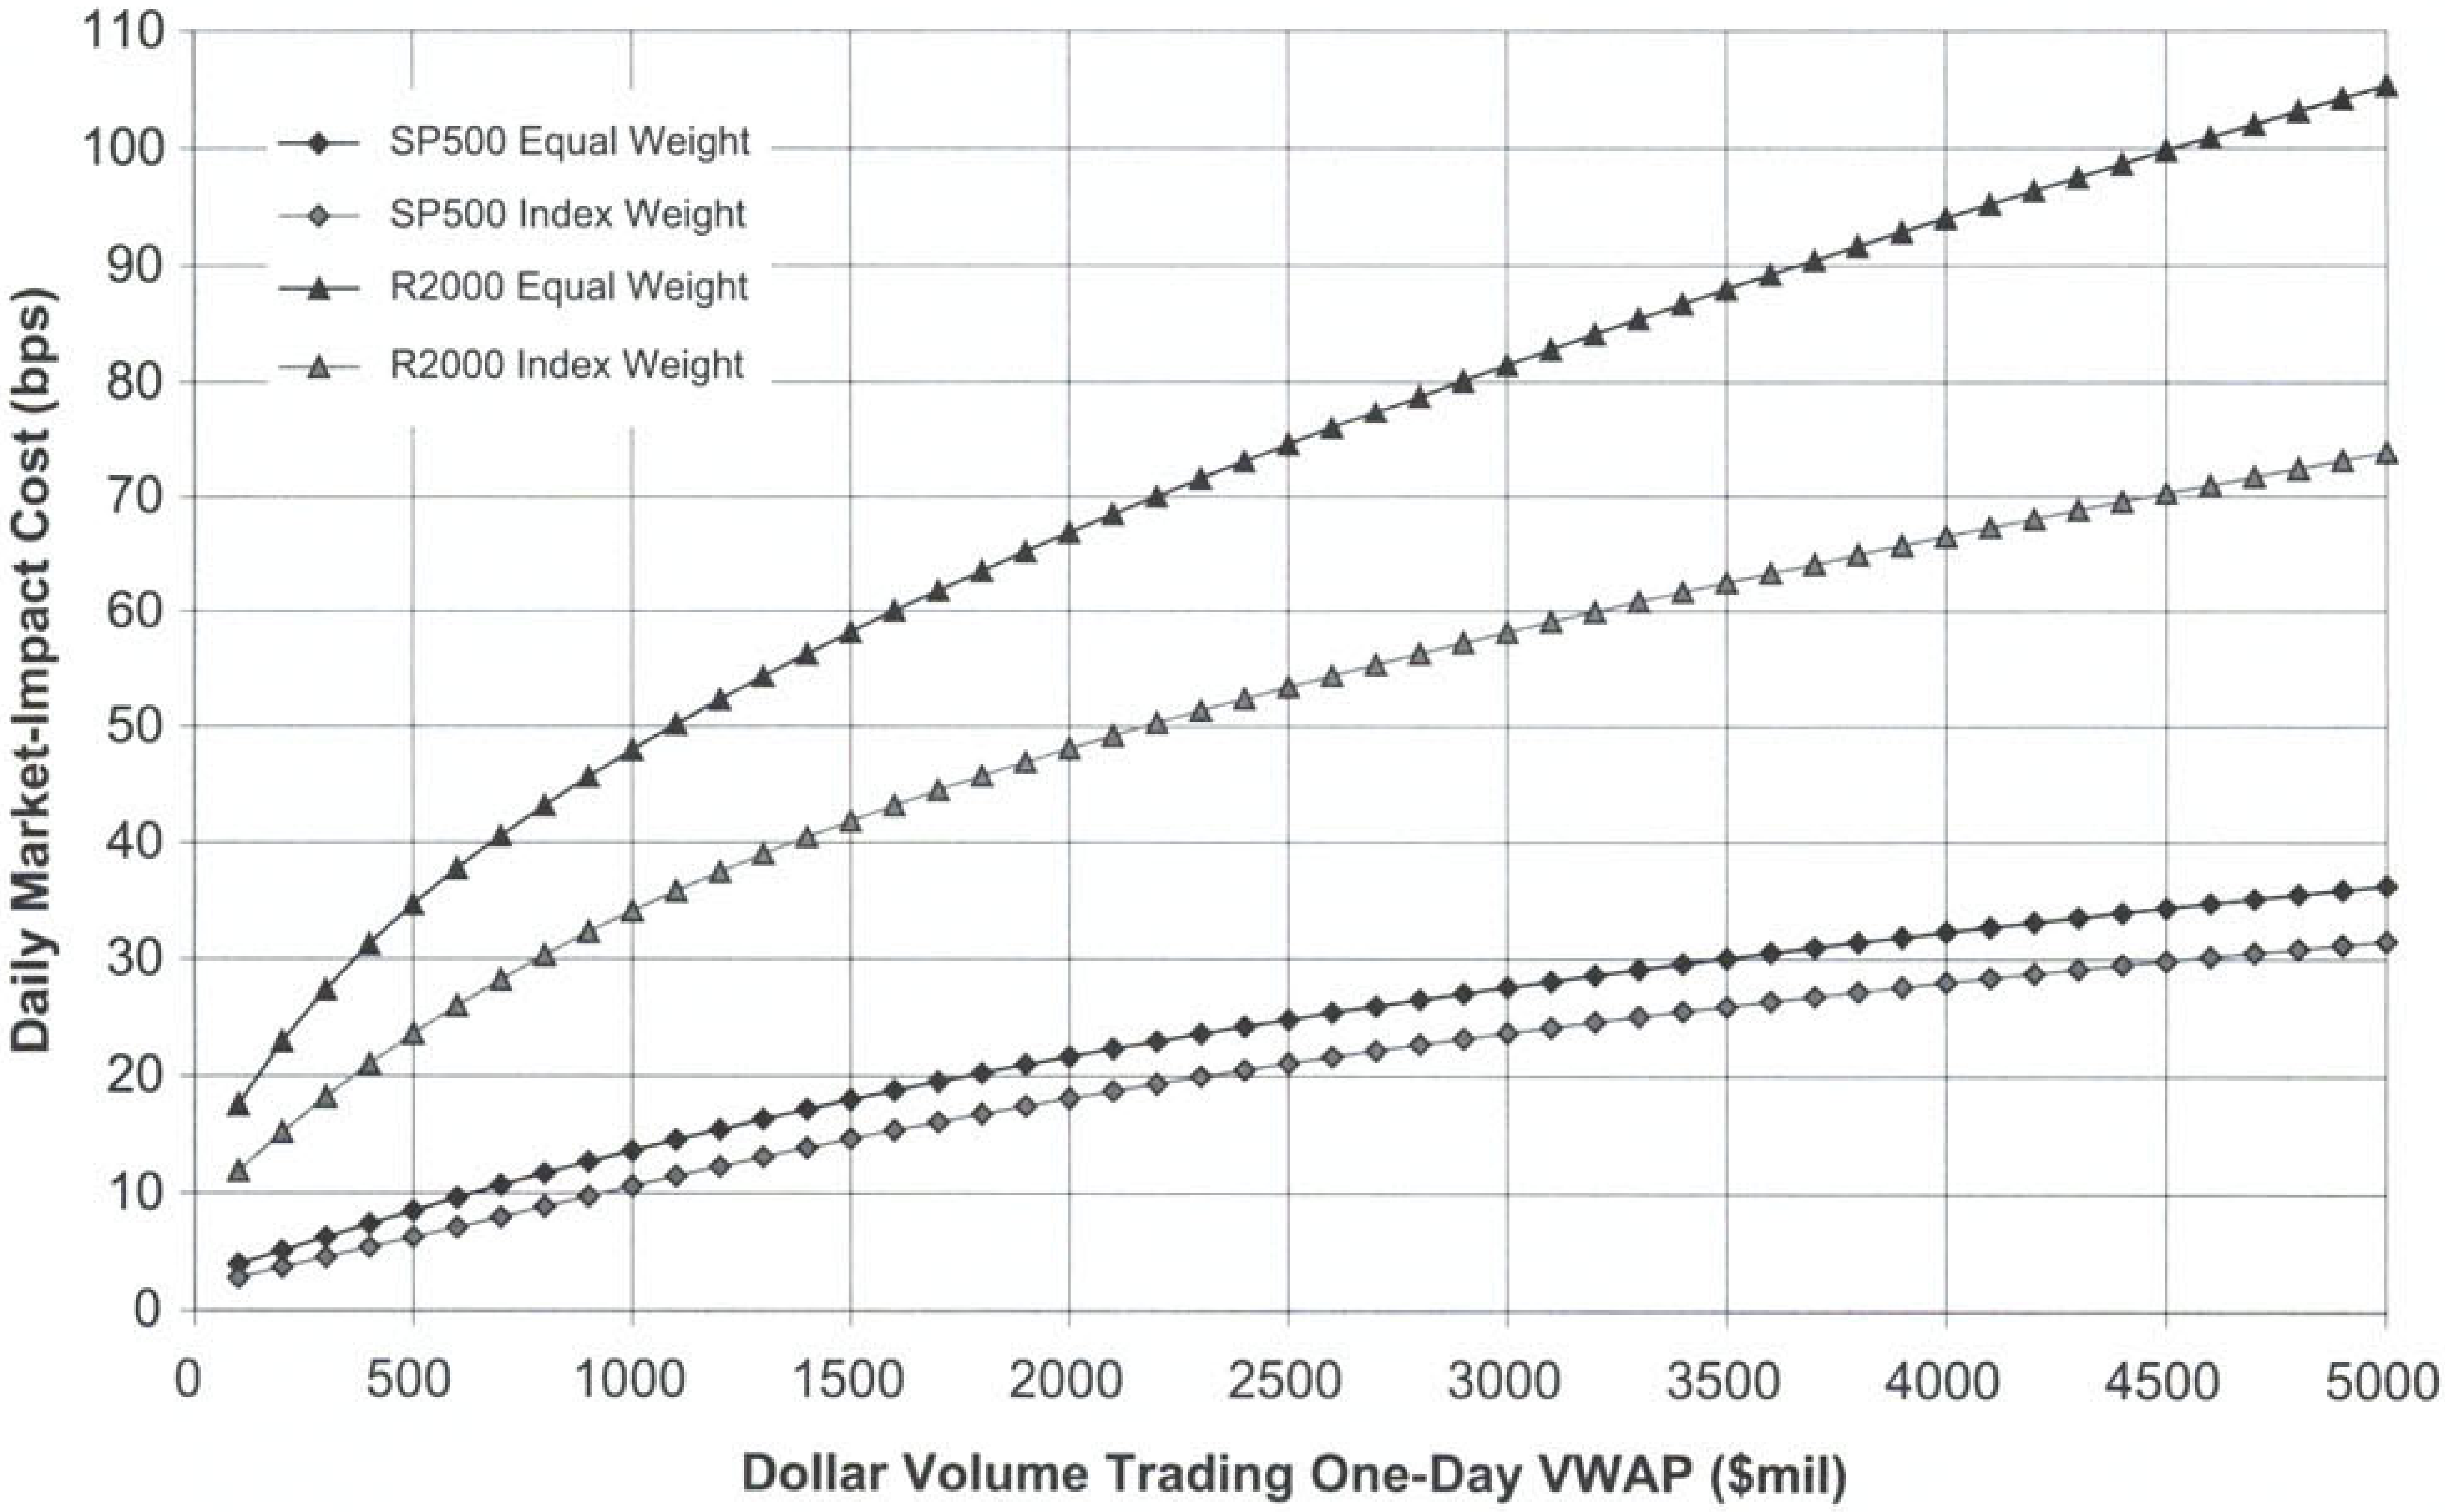

Note: Cost estimates as of August 31, 2006.

EXHIBIT 2
Dollar Amounts Traded (in millions) for Average Percent MDV Breakpoints

Avg. % of MDV	S&P 500 Equal Weighted	S&P 500 Index Weighted	R2000 Equal Weighted	R2000 Index Weighted
0.5%	106	260	12	21
1.0%	213	520	23	41
2.0%	425	1,040	47	82
5.0%	1,063	2,601	117	206
10.0%	2,126	5,201	235	412

Note: Volume estimates as of August 31, 2006.

EXHIBIT 3
S&P 500 Trade Lists: Top 50% Breadth vs. Top 50% Weight

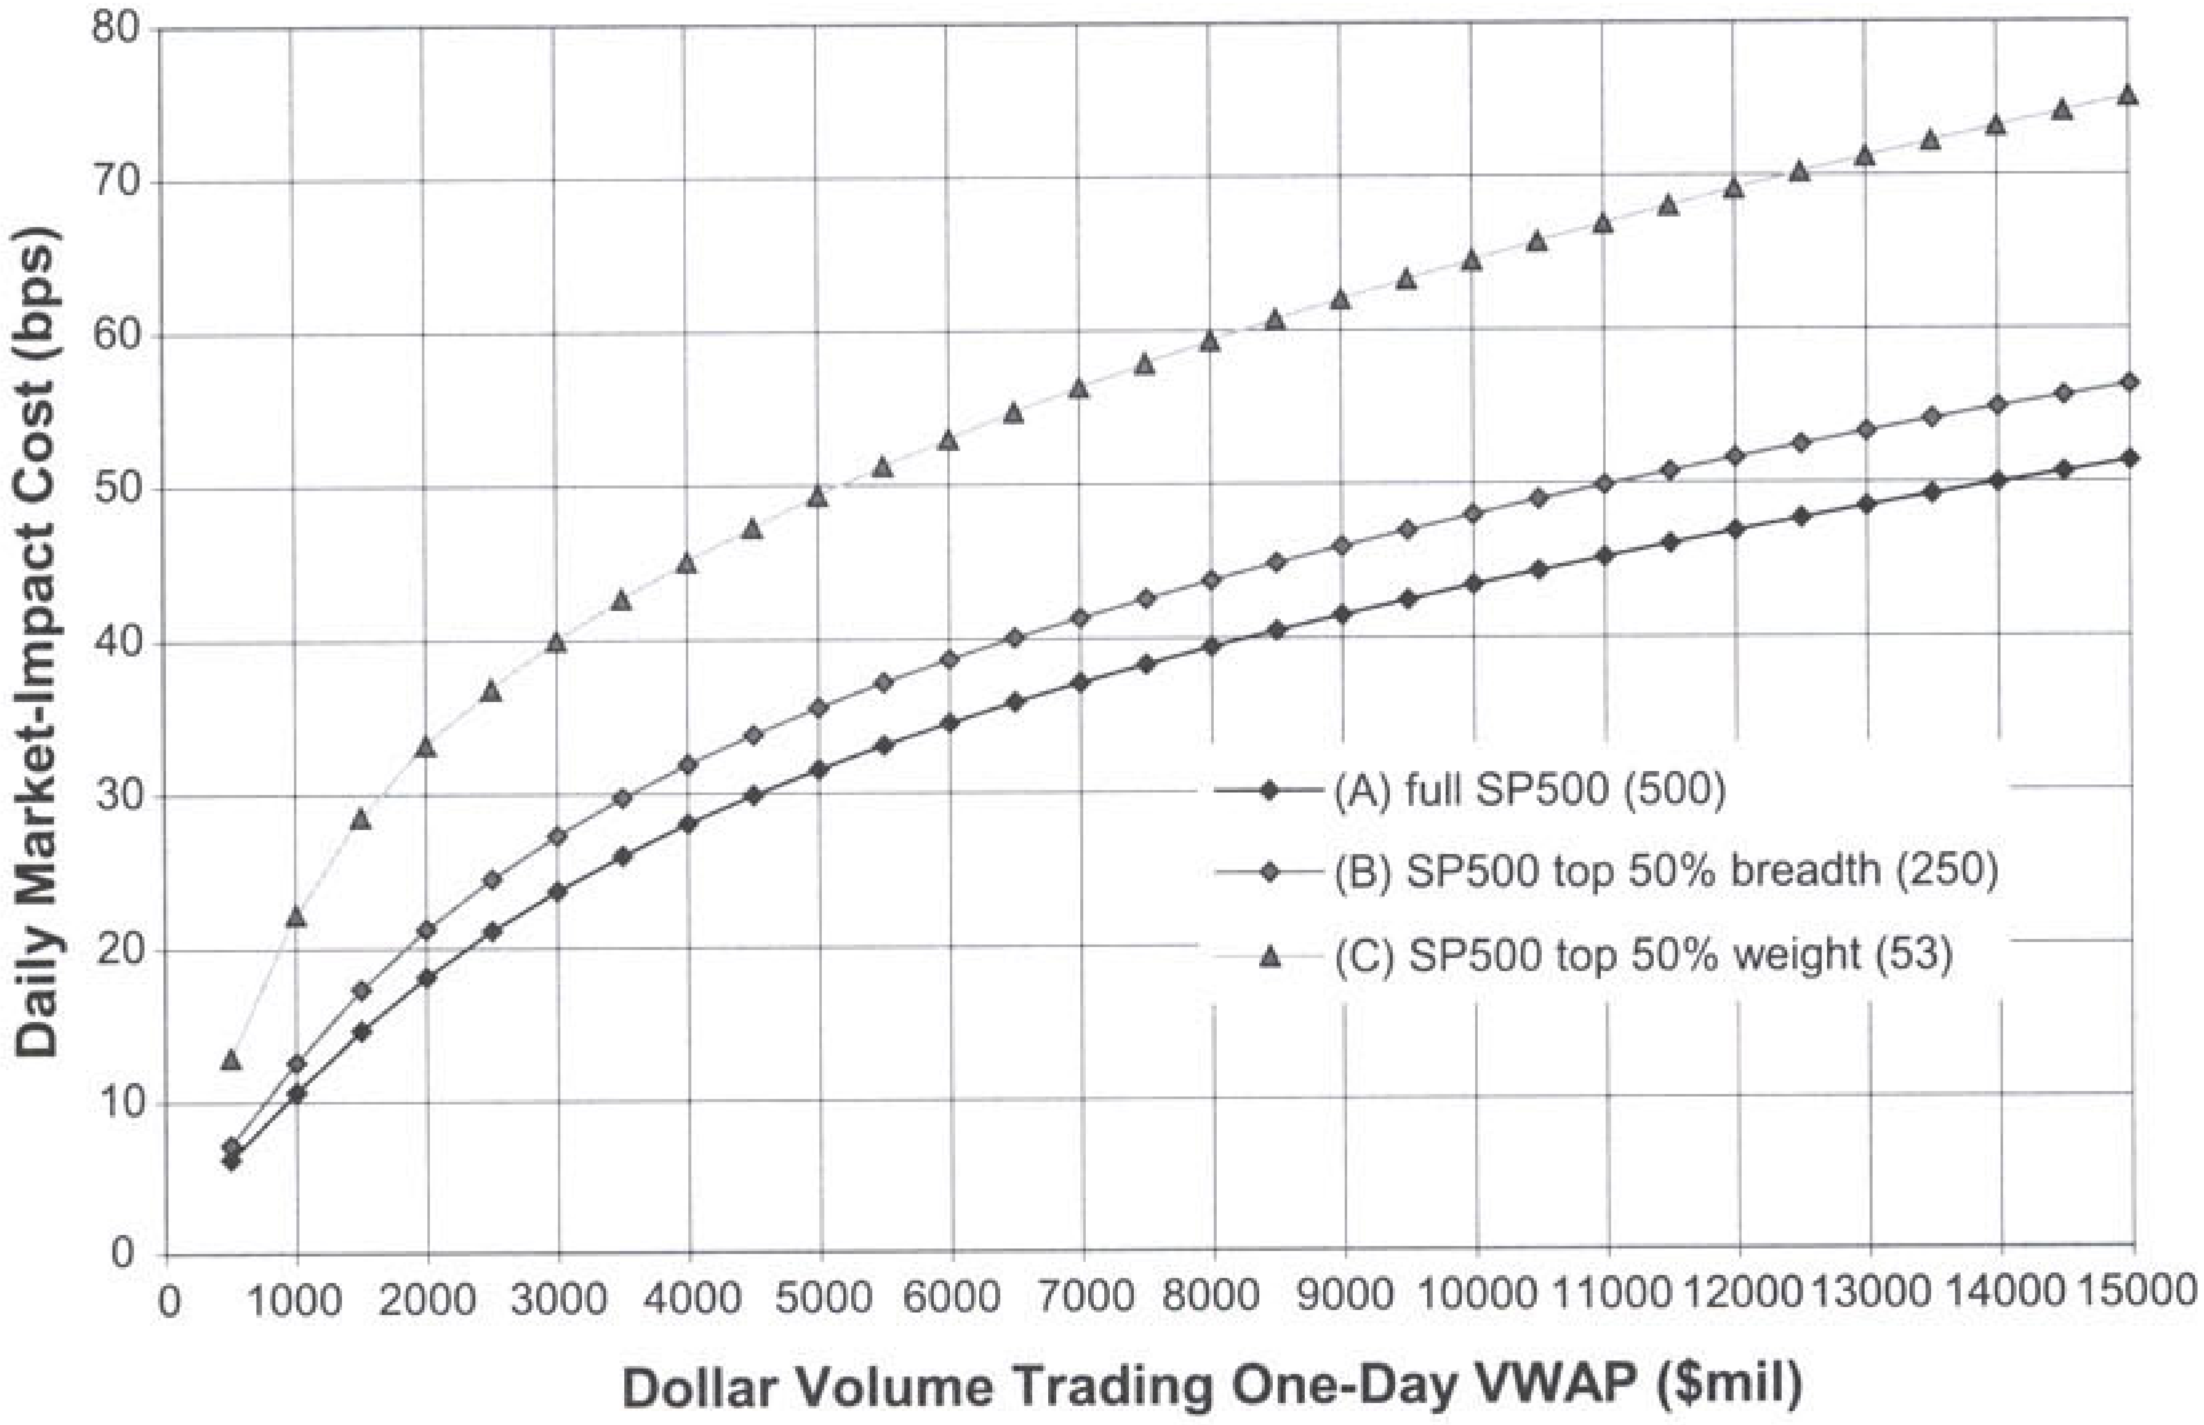

Note: Cost estimates as of August 31, 2006.

The relationship between a portfolio's degree of concentration and its liquidity as portrayed in Exhibit 1 deserves a bit more discussion, as it is directly related to the investment strategy that tracks the index. Since the equal-weighted R2000 trade list in Exhibit 1 reaches into more illiquid names than does the index-weighted trade list (relative to a cash benchmark), it costs the investor more in terms of market impact. Exhibit 3 shows the reverse effect and depicts three distinct S&P 500–based equal-weighted trade lists: (A) the full-index 500-name trade list, (B) the top-weighted 50% of the index (250 names), and (C) the top 50% in index weight, which contains only 53 names. It is clear that less concentration is resulting in less market impact in this case, as (A) dominates (B), which dominates (C). Constructing a portfolio with minimal tracking error to its benchmark could mean gravitating to names with the highest concentration of benchmark weight, as in (C), at the expense of impact-reducing breadth. Exhibit 3 not only contrasts with

EXHIBIT 4

Average Market Impact of Trading $10 Million in R3000 Stocks

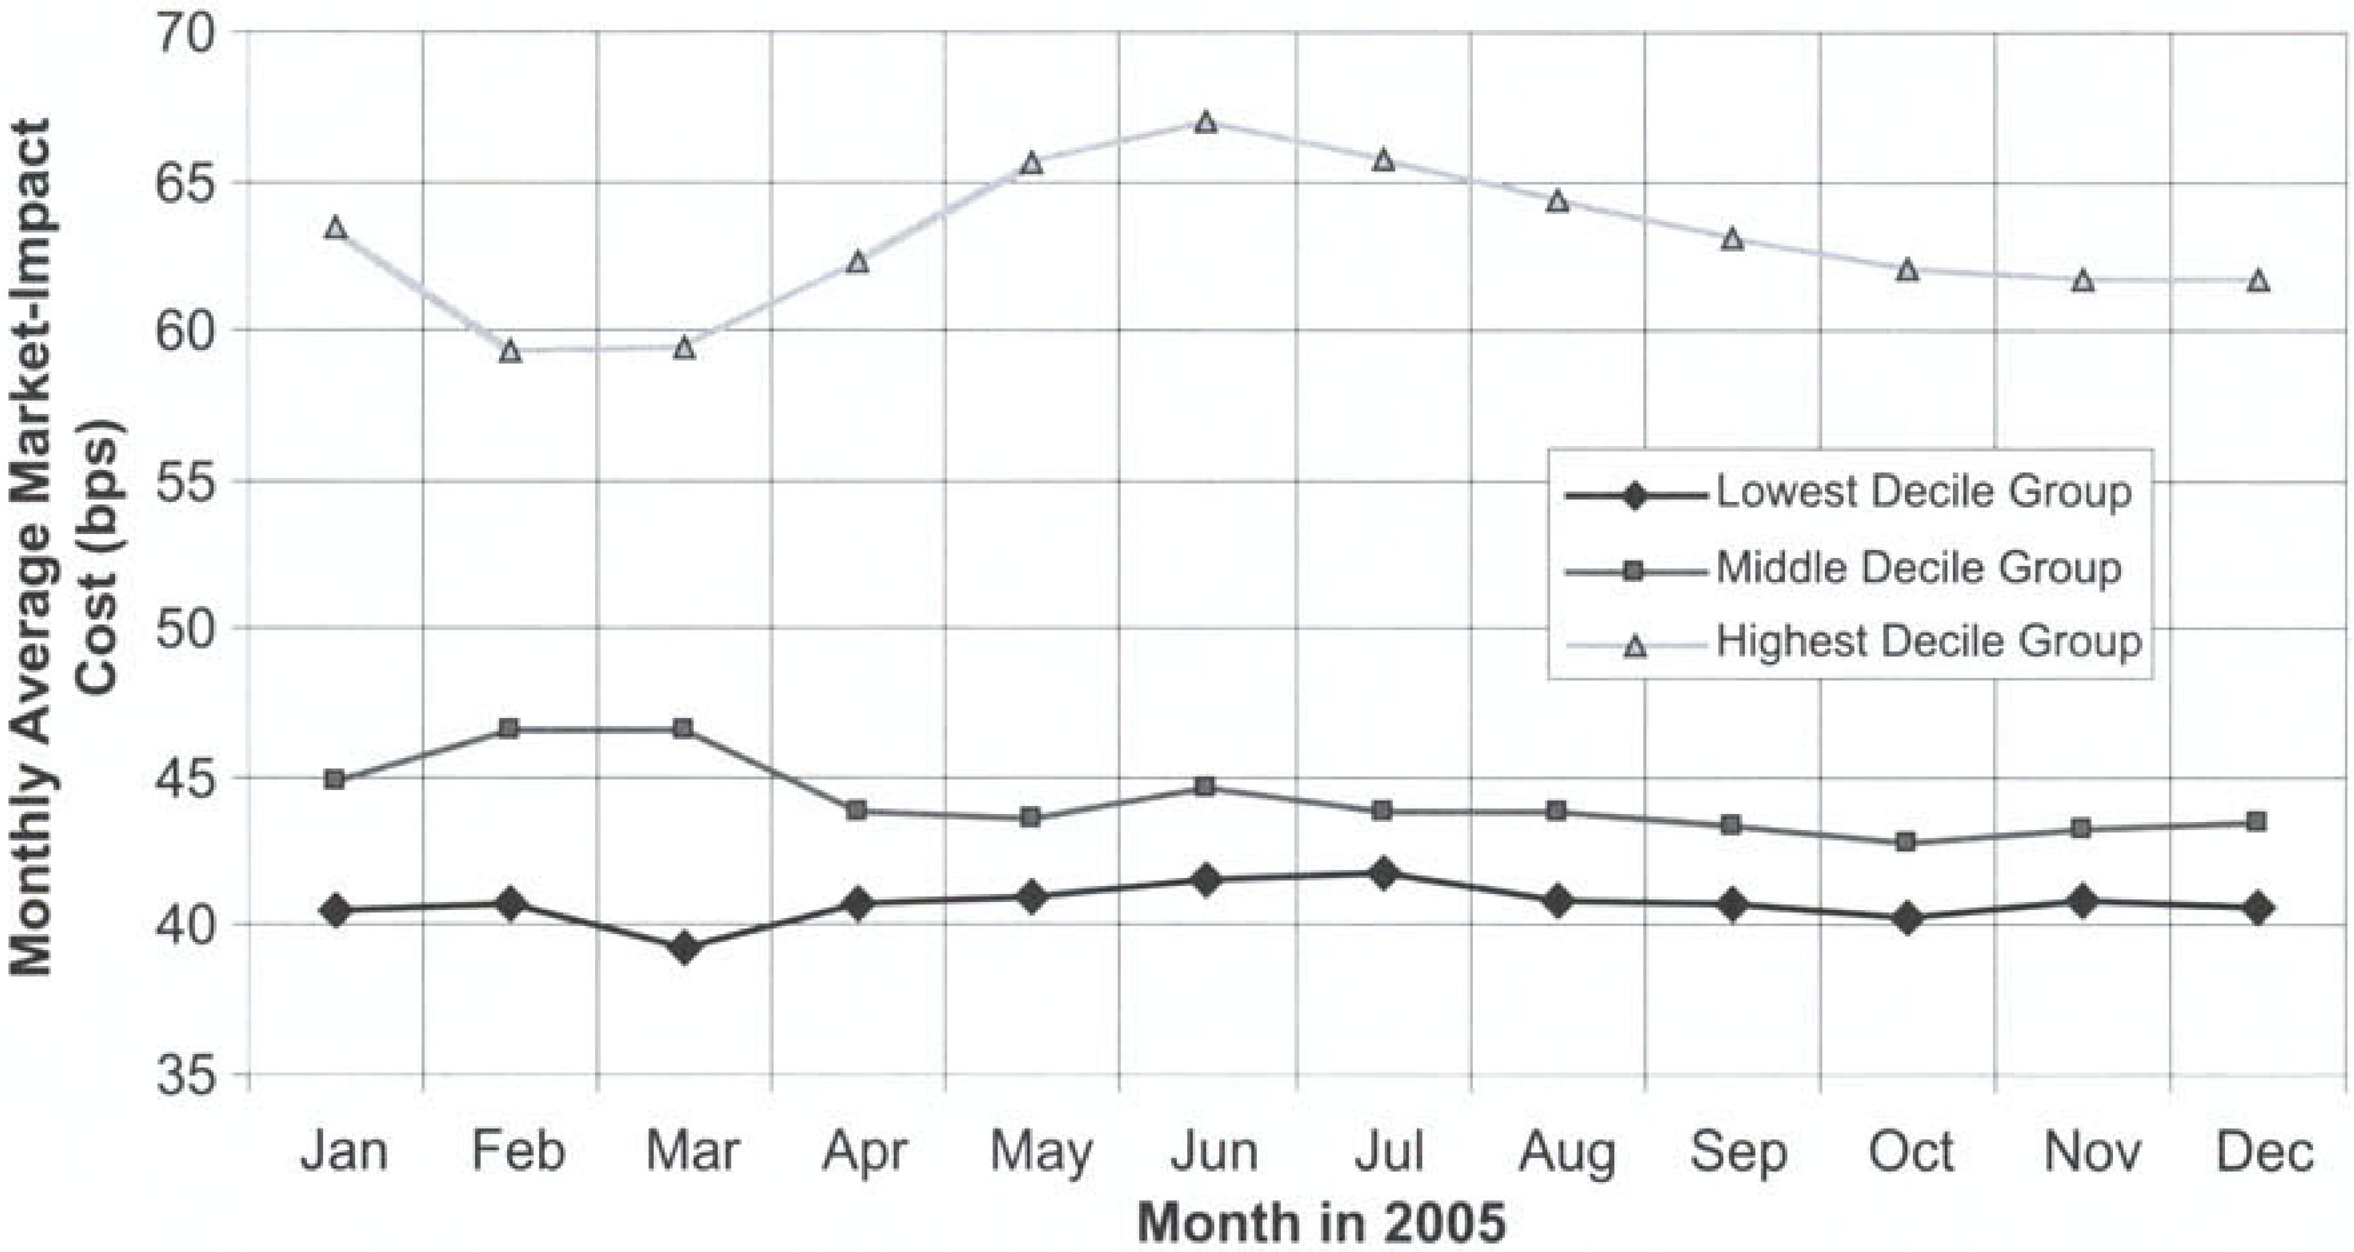

the scenario depicted in Exhibit 1, but also illustrates potential pitfalls of not including precise market-impact costs into common portfolio construction practices. Of course, a more informed analysis of a portfolio concentration's effect on its overall liquidity involves a thorough understanding of how trading costs are distributed in a cross section of securities. For a more in-depth study of cross-sectional distribution of liquidity measures (price impact is one such measure), see Chordia, Shivakumar, and Subrahmanyam [2001] and their references.

The optimal amount of assets under management often cannot be determined for each fund in isolation, because many funds follow the same investment strategies. A simple way to see this is by plotting the average market impact costs of trading $10 million in stocks that are Russell 3000 (R3000) constituents. We separate the R3000 list into 10 deciles according to the realized absolute returns in each month of 2005 and then compute each decile's average market-impact cost.[7] Exhibit 4 depicts the results. The evidence has two implications for studies on the relation between investment strategies and trading costs and, in particular, for studies on fund capacity. First, alphas and costs are related. The decile having the most extreme returns is also the most expensive to trade. Second, average market-impact costs are time varying, which points to systematic liquidity effects. While our evidence is not comprehensive, it offers conclusions similar to those of Chan et al. [2006]. Those authors use autocorrelations in hedge fund returns as opposed to actual transaction costs to proxy for illiquidity, but in the end they also report "considerable swings" in illiquidity over time. Whether our results are related to momentum or to some other statistical arbitrage strategy is not important. Our main point is more general—anyone who runs an investment strategy should be prepared to pay higher than expected costs due to liquidity shortages in stocks that are massively bought or sold (we use absolute returns as a proxy for such events). The substantial losses many quant funds suffered in August 2007 seem to corroborate this conjecture. The concept of systematic liquidity has recently attracted a lot of interest; more information is provided in Chollete, Naes, and Skjeltorp [2007], Domowitz, Hansch, and Wang [2005], and Kamara, Lou, and Sadka [2007].

MERGING ALPHAS AND COSTS: SIMULATION

In the previous section, we laid the groundwork for a more informed discussion of fund capacity by highlighting the varied market-impact consequences of trading different lists with different turnover levels. The focus was to quantify the dependency between market-impact costs and overall portfolio liquidity. We showed that liquidity could be affected by changing the breadth of the portfolio or by changing its investable universe. We also noted that liquidity effects differ depending on the investment strategies that are followed, implying that capacity is a strategy-specific phenomenon. Ultimately, however, liquidity alone cannot explain the notion of a fund's optimal operating capacity—it must be combined with expected return, or alpha. In this section, we perform Monte Carlo simulations to illustrate how to jointly model alpha and liquidity.

We start by assuming that a hypothetical long–short portfolio manager is acting on an imperfect alpha signal ($IC = 5\%$) in order to maximize the annual return subject to transaction costs. We assume that the investable universe is the S&P 500 and that the portfolio is held from January to December 2005.[8] We also assume that the portfolio manager is mandated not to exceed a monthly turnover of T (we vary T from 0% to 150%).[9] Finally, we assume that the portfolio manager is running a fund of size S (we consider four values for S: \$2 billion, \$5 billion, \$10 billion, and \$20 billion).

On the first day of every month, starting January 1, 2005, the portfolio manager estimates the alpha of every stock in the S&P 500. We simulate this alpha as

$$\alpha \sim N(IC * r, \sigma^2 * (1 - IC^2)) \tag{1}$$

where r is the actual (realized) stock return in the following month,[10] σ is a forward-looking estimate of stock volatility from the ITG proprietary U.S. monthly risk model,[11] and IC is a proxy for the manager's stock-picking ability. A detailed description of the simulation exercise is presented in the appendix.

The simulated alphas are normalized so that positive alphas sum to one and negative alphas sum to negative one. They are then netted against preexisting position weights in the portfolio to create buy and sell rankings. We pick the top-ranked 200 stocks for the buy list and the bottom-ranked 200 stocks for the sell list. We then rescale each side of the trade list by another constant so that the entire list achieves dollar neutrality and the prescribed turnover limit, T. After determining the desired vector of weights, we start trading. We do not assume that the prescribed turnover is accomplished in one day, but rather that it is spread evenly throughout the month (about 21 trading days). We assume the VWAP trading strategy for each day. Accordingly, the return achieved each month is divided evenly between that of the legacy and target portfolios. While this procedure in practice assumes an alpha decay of zero, it does a fair job of accounting for implementation delay with the commensurate effects on both target-return capture and the risks of holding the legacy portfolio. Finally, we adjust the vector of realized returns for transaction costs consisting of fixed per-share commission cost and market-impact cost. We assume a fixed commission rate of \$0.01 per share throughout. The market-impact cost is derived from ITG Cost Curves data assuming the daily VWAP trading strategy previously described. The Cost Curves data contain expected ACE model forecasts of implementation shortfall cost along with the standard deviation of the cost for varying order sizes. Exhibit 5 provides a graphical illustration of the procedure.

Subtracting the estimated transaction cost marks the end of a single simulation draw.[12] We compute the realized net return and then repeat the procedure 1,000 times for each turnover level T and fund size S, and then present the average results in Exhibit 6. We do not take market-impact costs into consideration during the scoring

EXHIBIT 5
Portfolio Selection and Rebalance Procedure

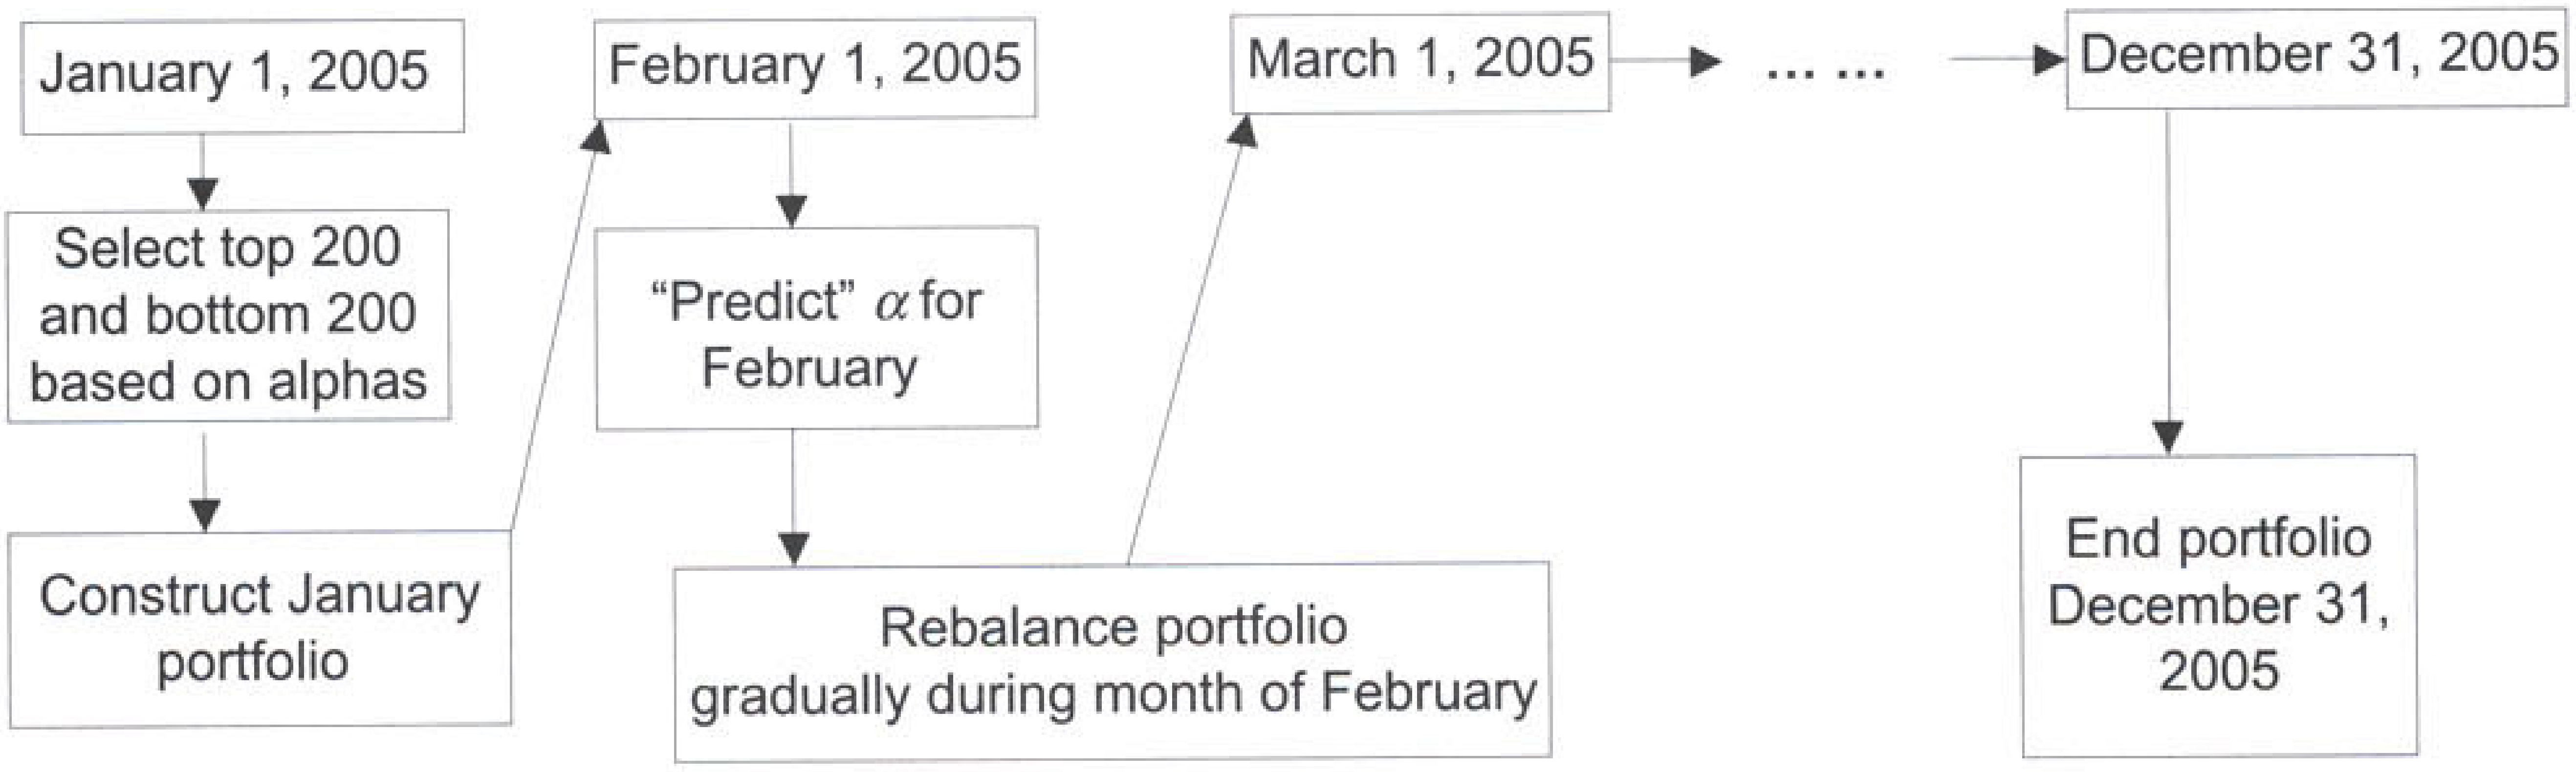

EXHIBIT 6
Optimal Turnover Rate for Different *AUM* Levels

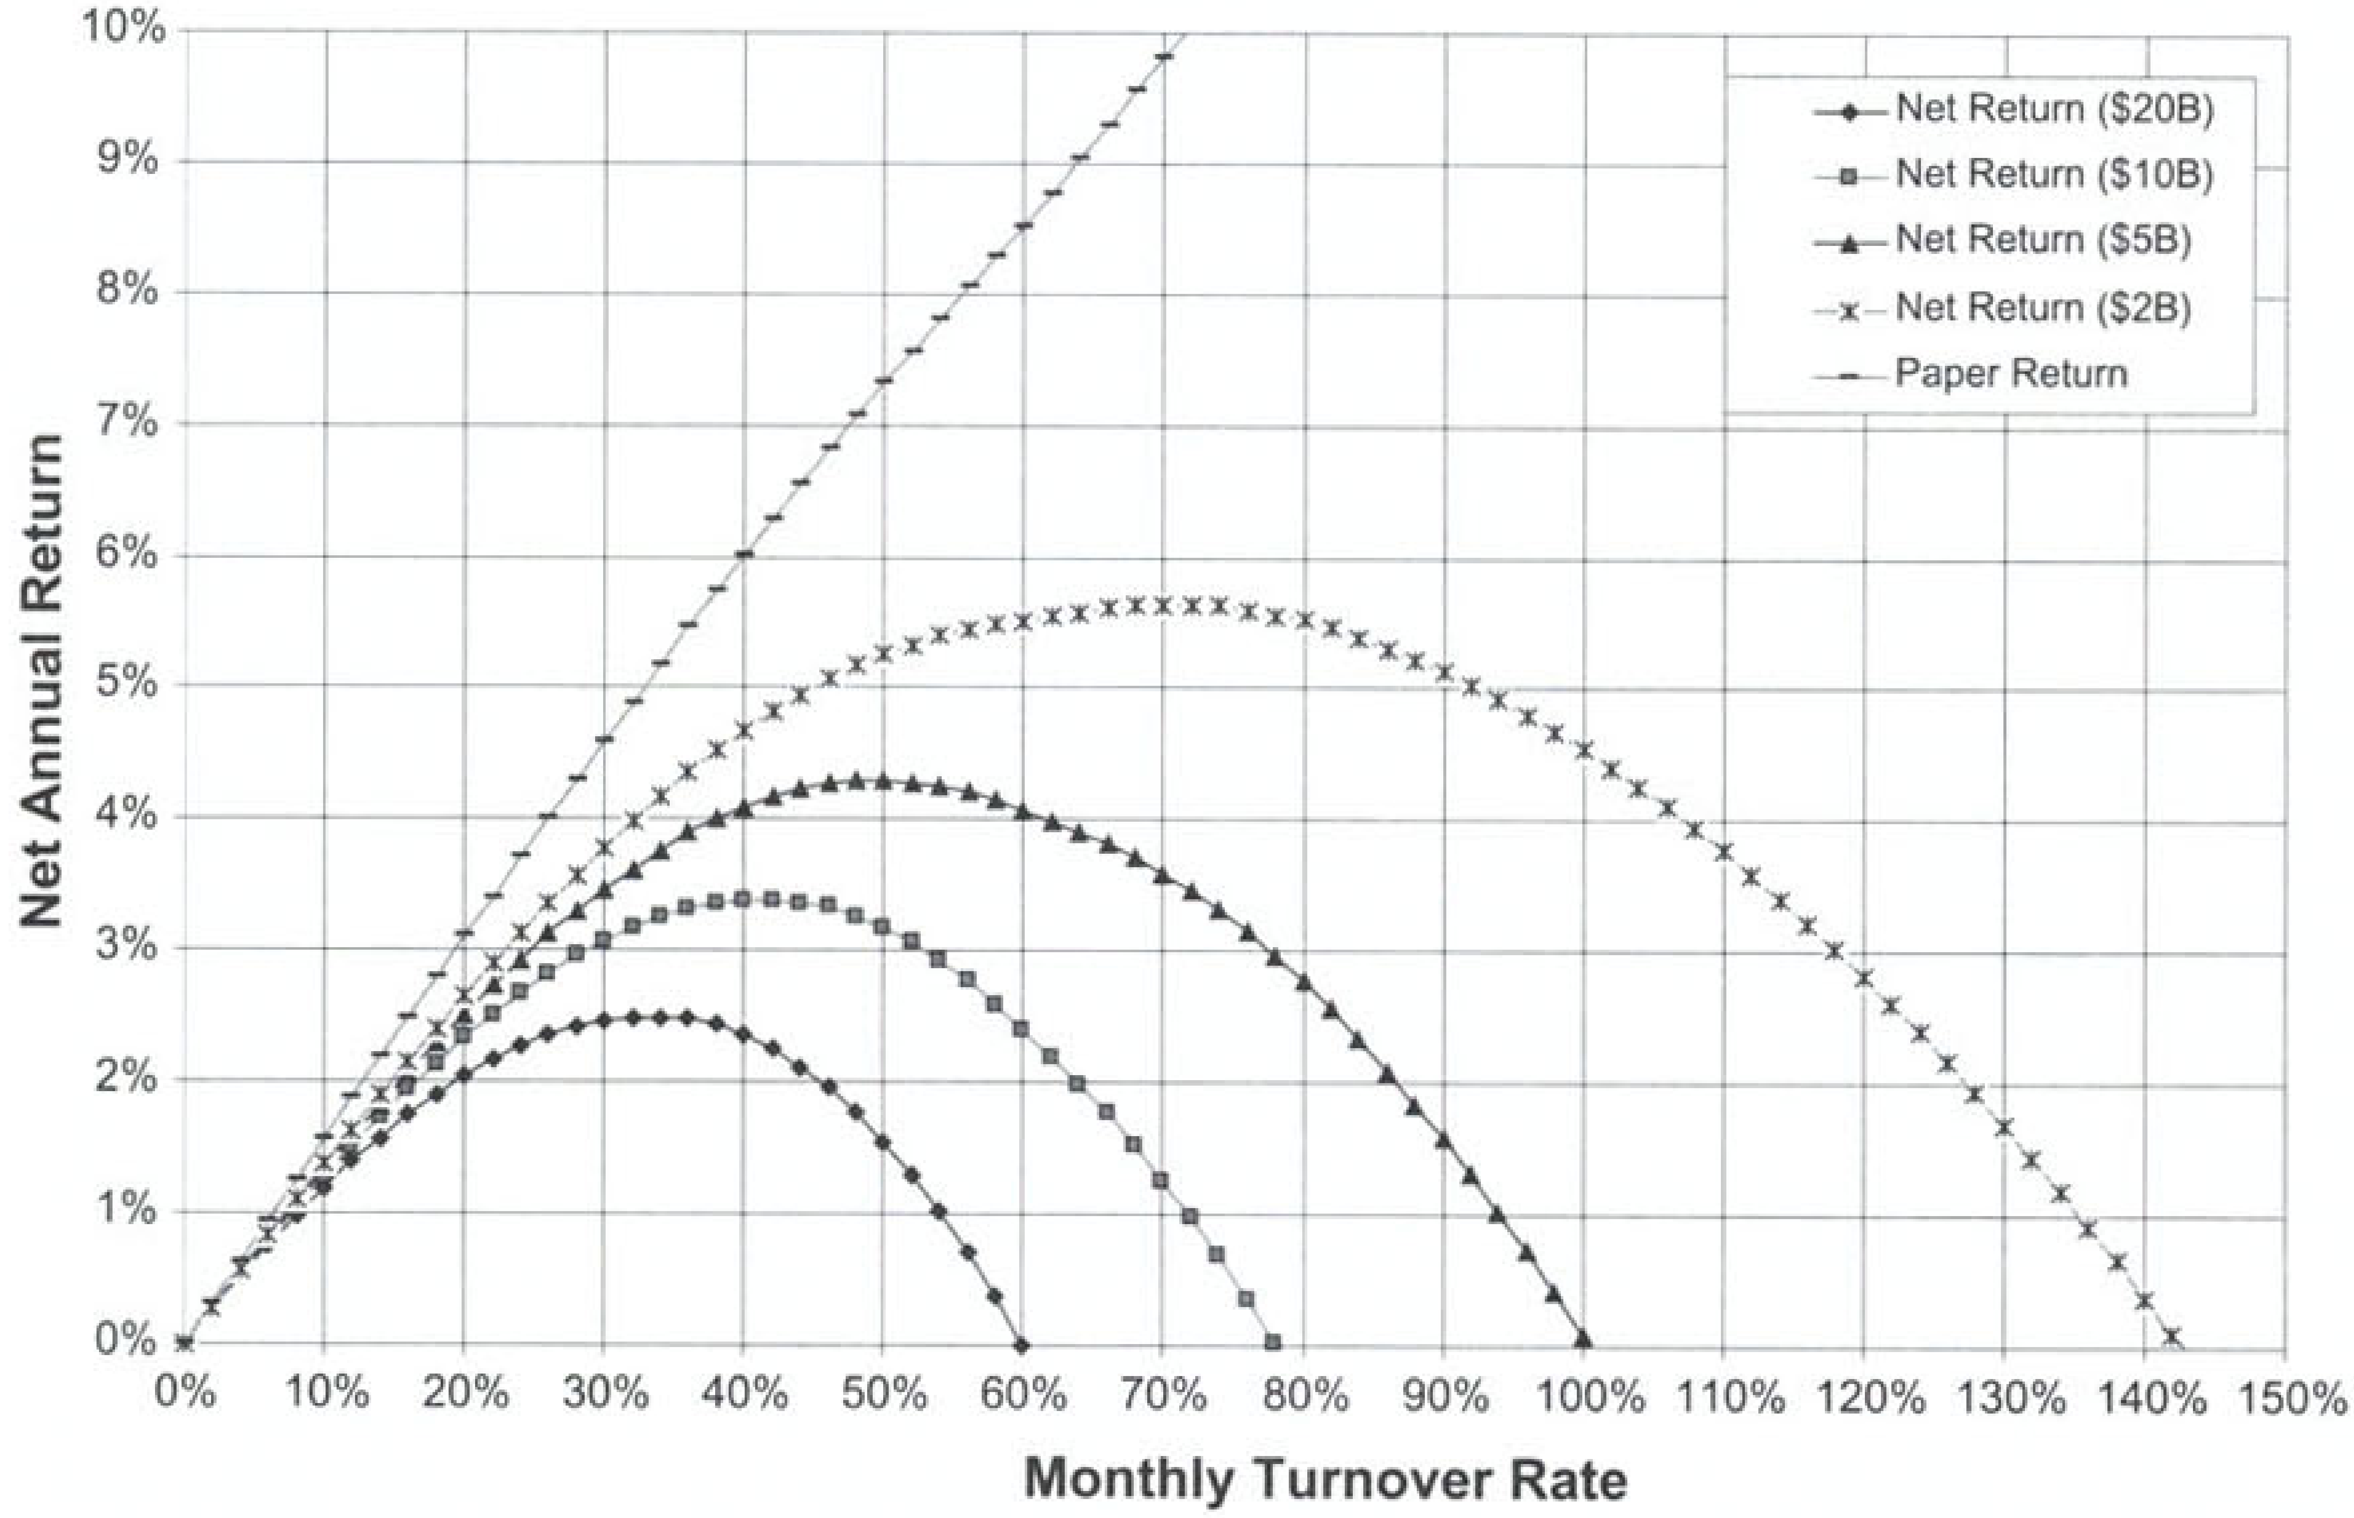

and rebalancing procedure itself. Market-impact cost estimates are computed on the back end of the process in order to illustrate the consequences of the "blind pursuit" of alpha at different *AUM* and turnover levels.

Several results are apparent. First, because of transaction costs, the paper return dominates the returns of the portfolios at all *AUM* and turnover levels (except zero). Second, where turnover is relatively low (i.e., below 15%) all four *AUM* levels show an almost linear increase in net return. Finally, as turnover increases past 20%, net returns for the larger *AUM* portfolios peak and drop much more quickly. For the $10 billion portfolio, net return maximizes at 3.40% with 40% monthly turnover, while the $20 billion portfolio maximizes at 2.50% net return with 34% monthly turnover. The $2 billion portfolio achieves a net return closer to 6%, and we can assume that even lower *AUM* levels would continue to approach the paper return (the paper return includes neither fixed commission costs nor market-impact costs).

Our simulations in Exhibit 6 lead us to the concept of an "efficient frontier" for turnover as Exhibit 7 illustrates. Let us consider what happens as our simulated manager undertakes a move from $2 billion to $5 billion under management.

If our manager holds turnover constant, he moves from point (A) straight down to point (B). In order to achieve the highest return under this scenario, the manager will have to decrease turnover to move to point (C). Point (B) represents not just an inefficient portfolio but an inefficient *strategy* that is paying too much in market-impact costs from excessive turnover relative to skill level represented by *IC*. This concept of efficient turnover holds not only in cases when *AUM* is changing and capacity is under investigation. If at $2 billion our manager started at any point to the left of (A) (e.g., with a mandated monthly turnover of 20%), he would have to allow higher turnover in order to maximize returns at point (A).

The analysis illustrated in Exhibit 7 is consistent with maximizing the individual fund investor's utility, which happens when a fund's net return reaches its maximum. Alternatively, we can run the analysis from the macro perspective (i.e., maximizing the total wealth the fund creates). In this case, we should model the amount of value-added instead of net return.[13] Exhibit 8 shows the dynamics of the value-added as a function of start-of-the-period *AUM*. To demonstrate the importance of choosing the turnover level carefully, we depict two scenarios. First, we assume that the fund manager picks optimal

EXHIBIT 7
Where Capacity Meets the Road: The Turnover Efficient Frontier

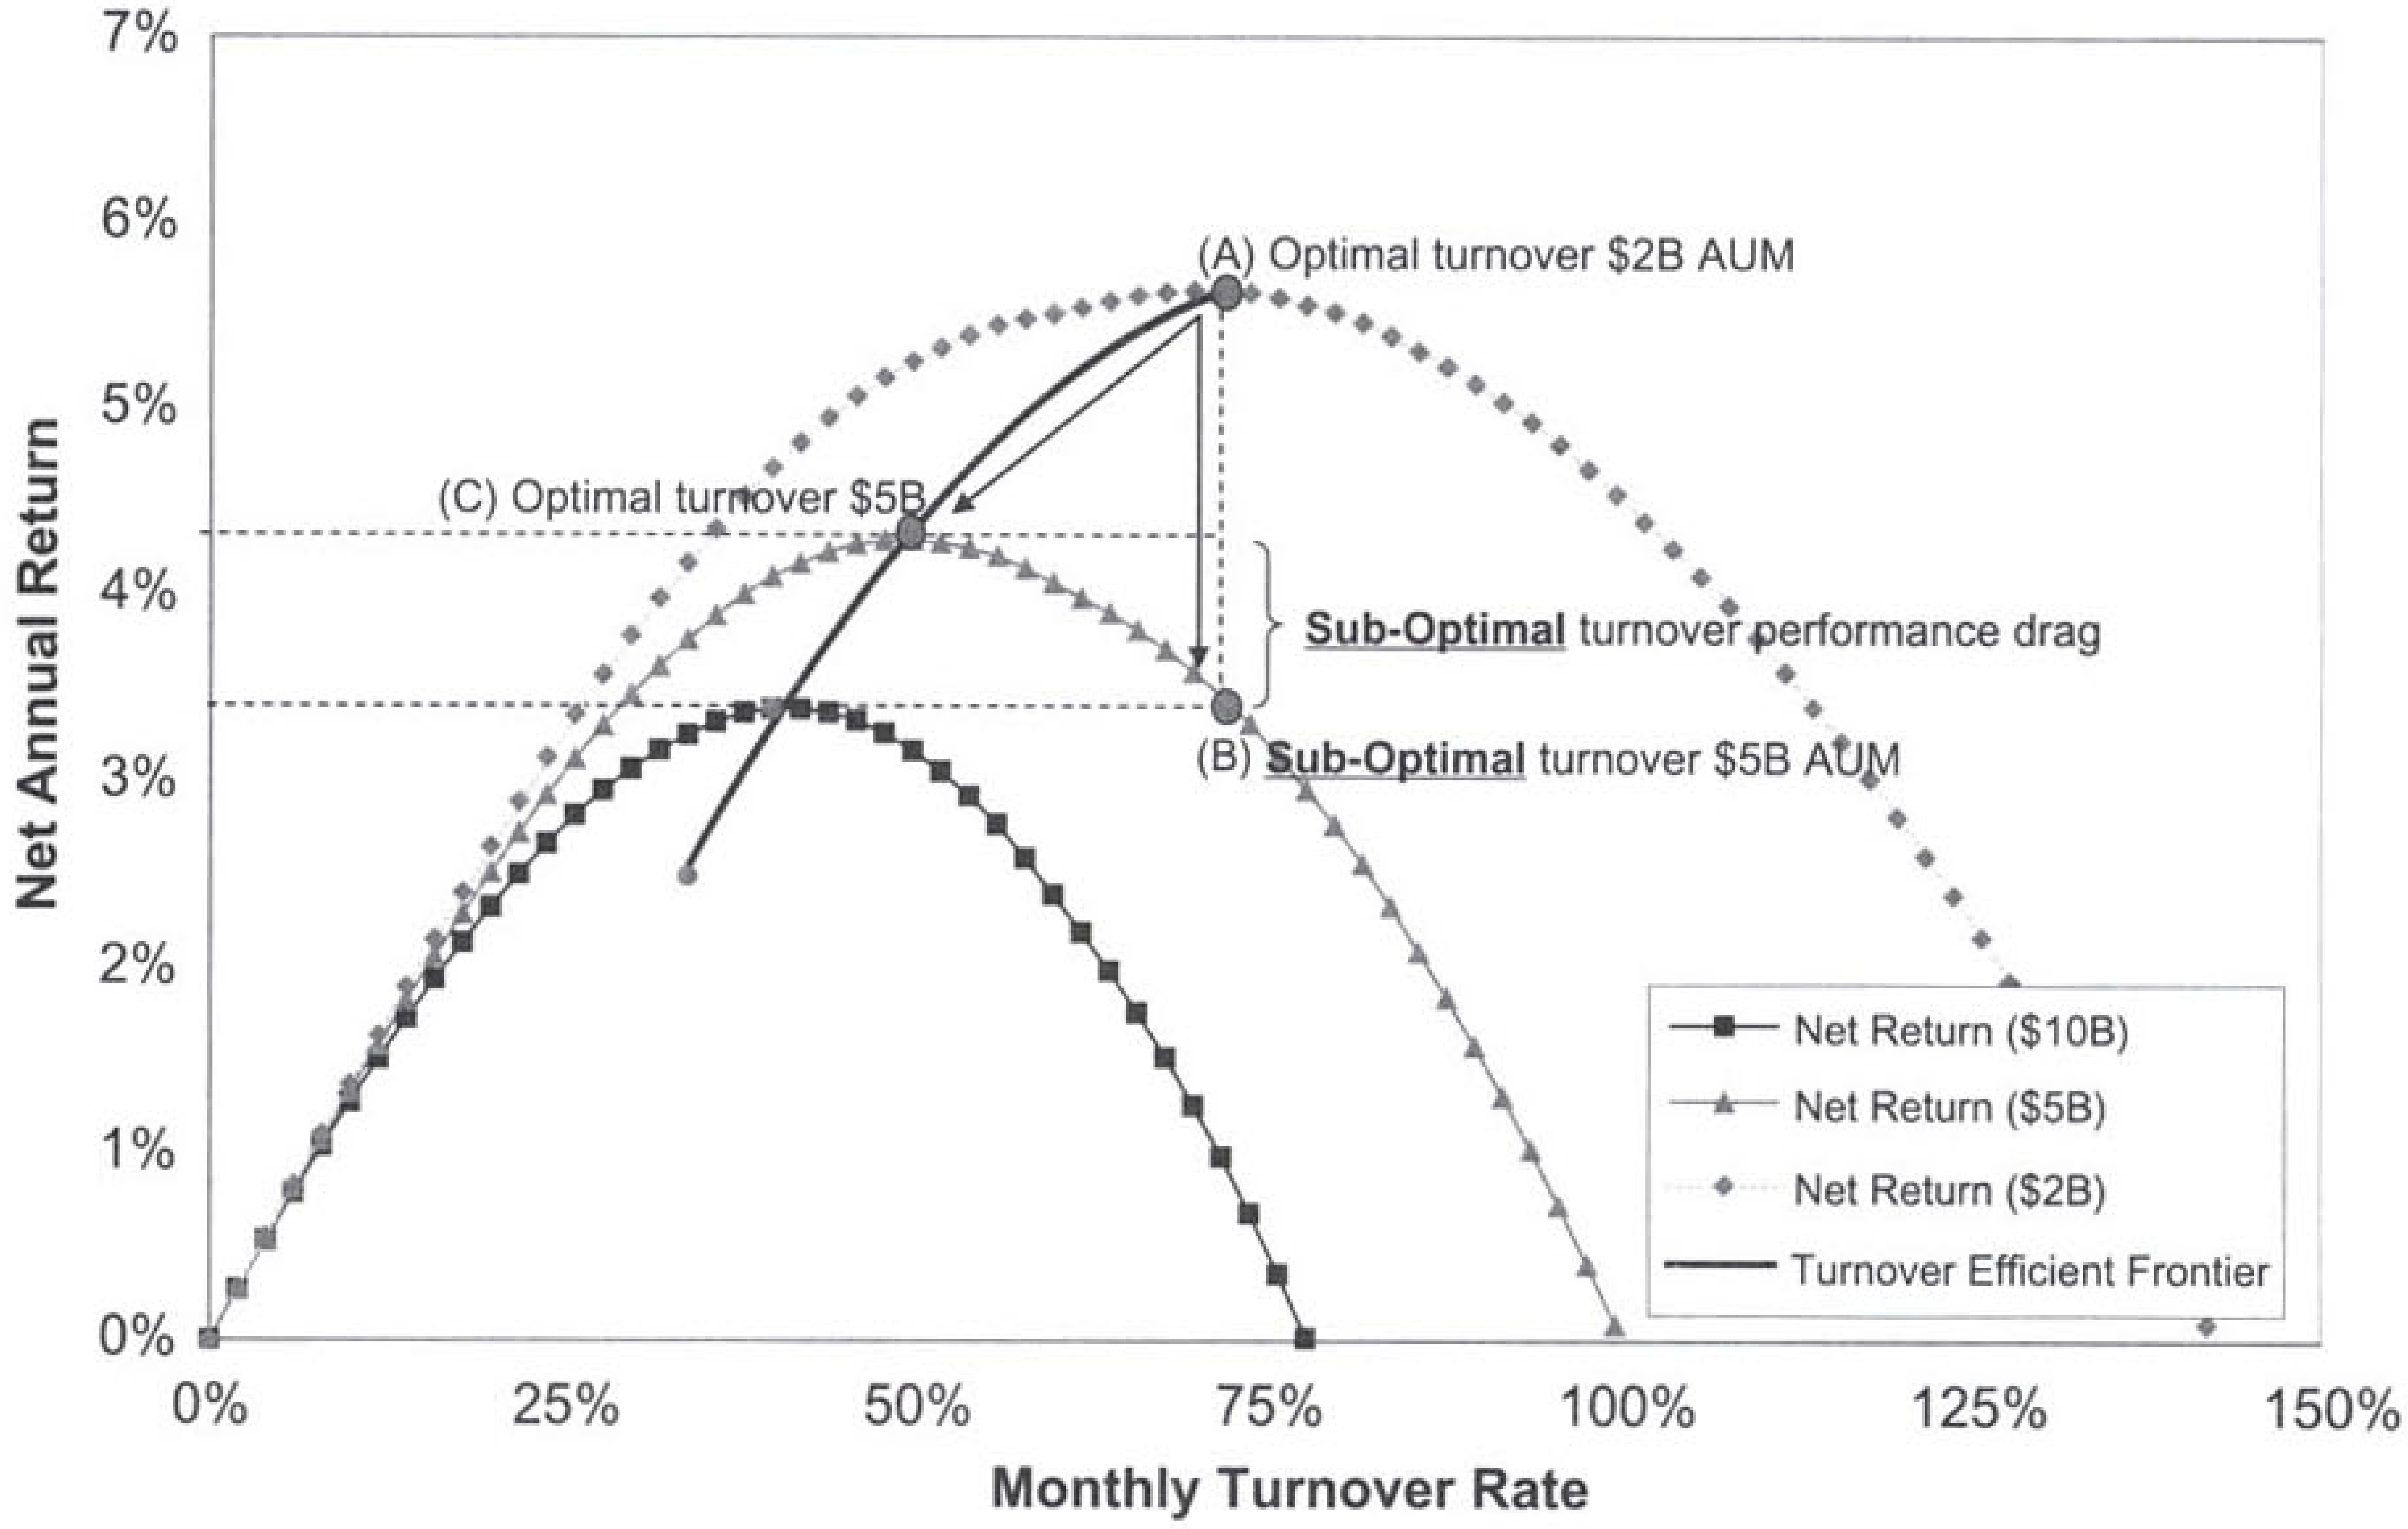

EXHIBIT 8
Value-Added for Different *AUM* Levels

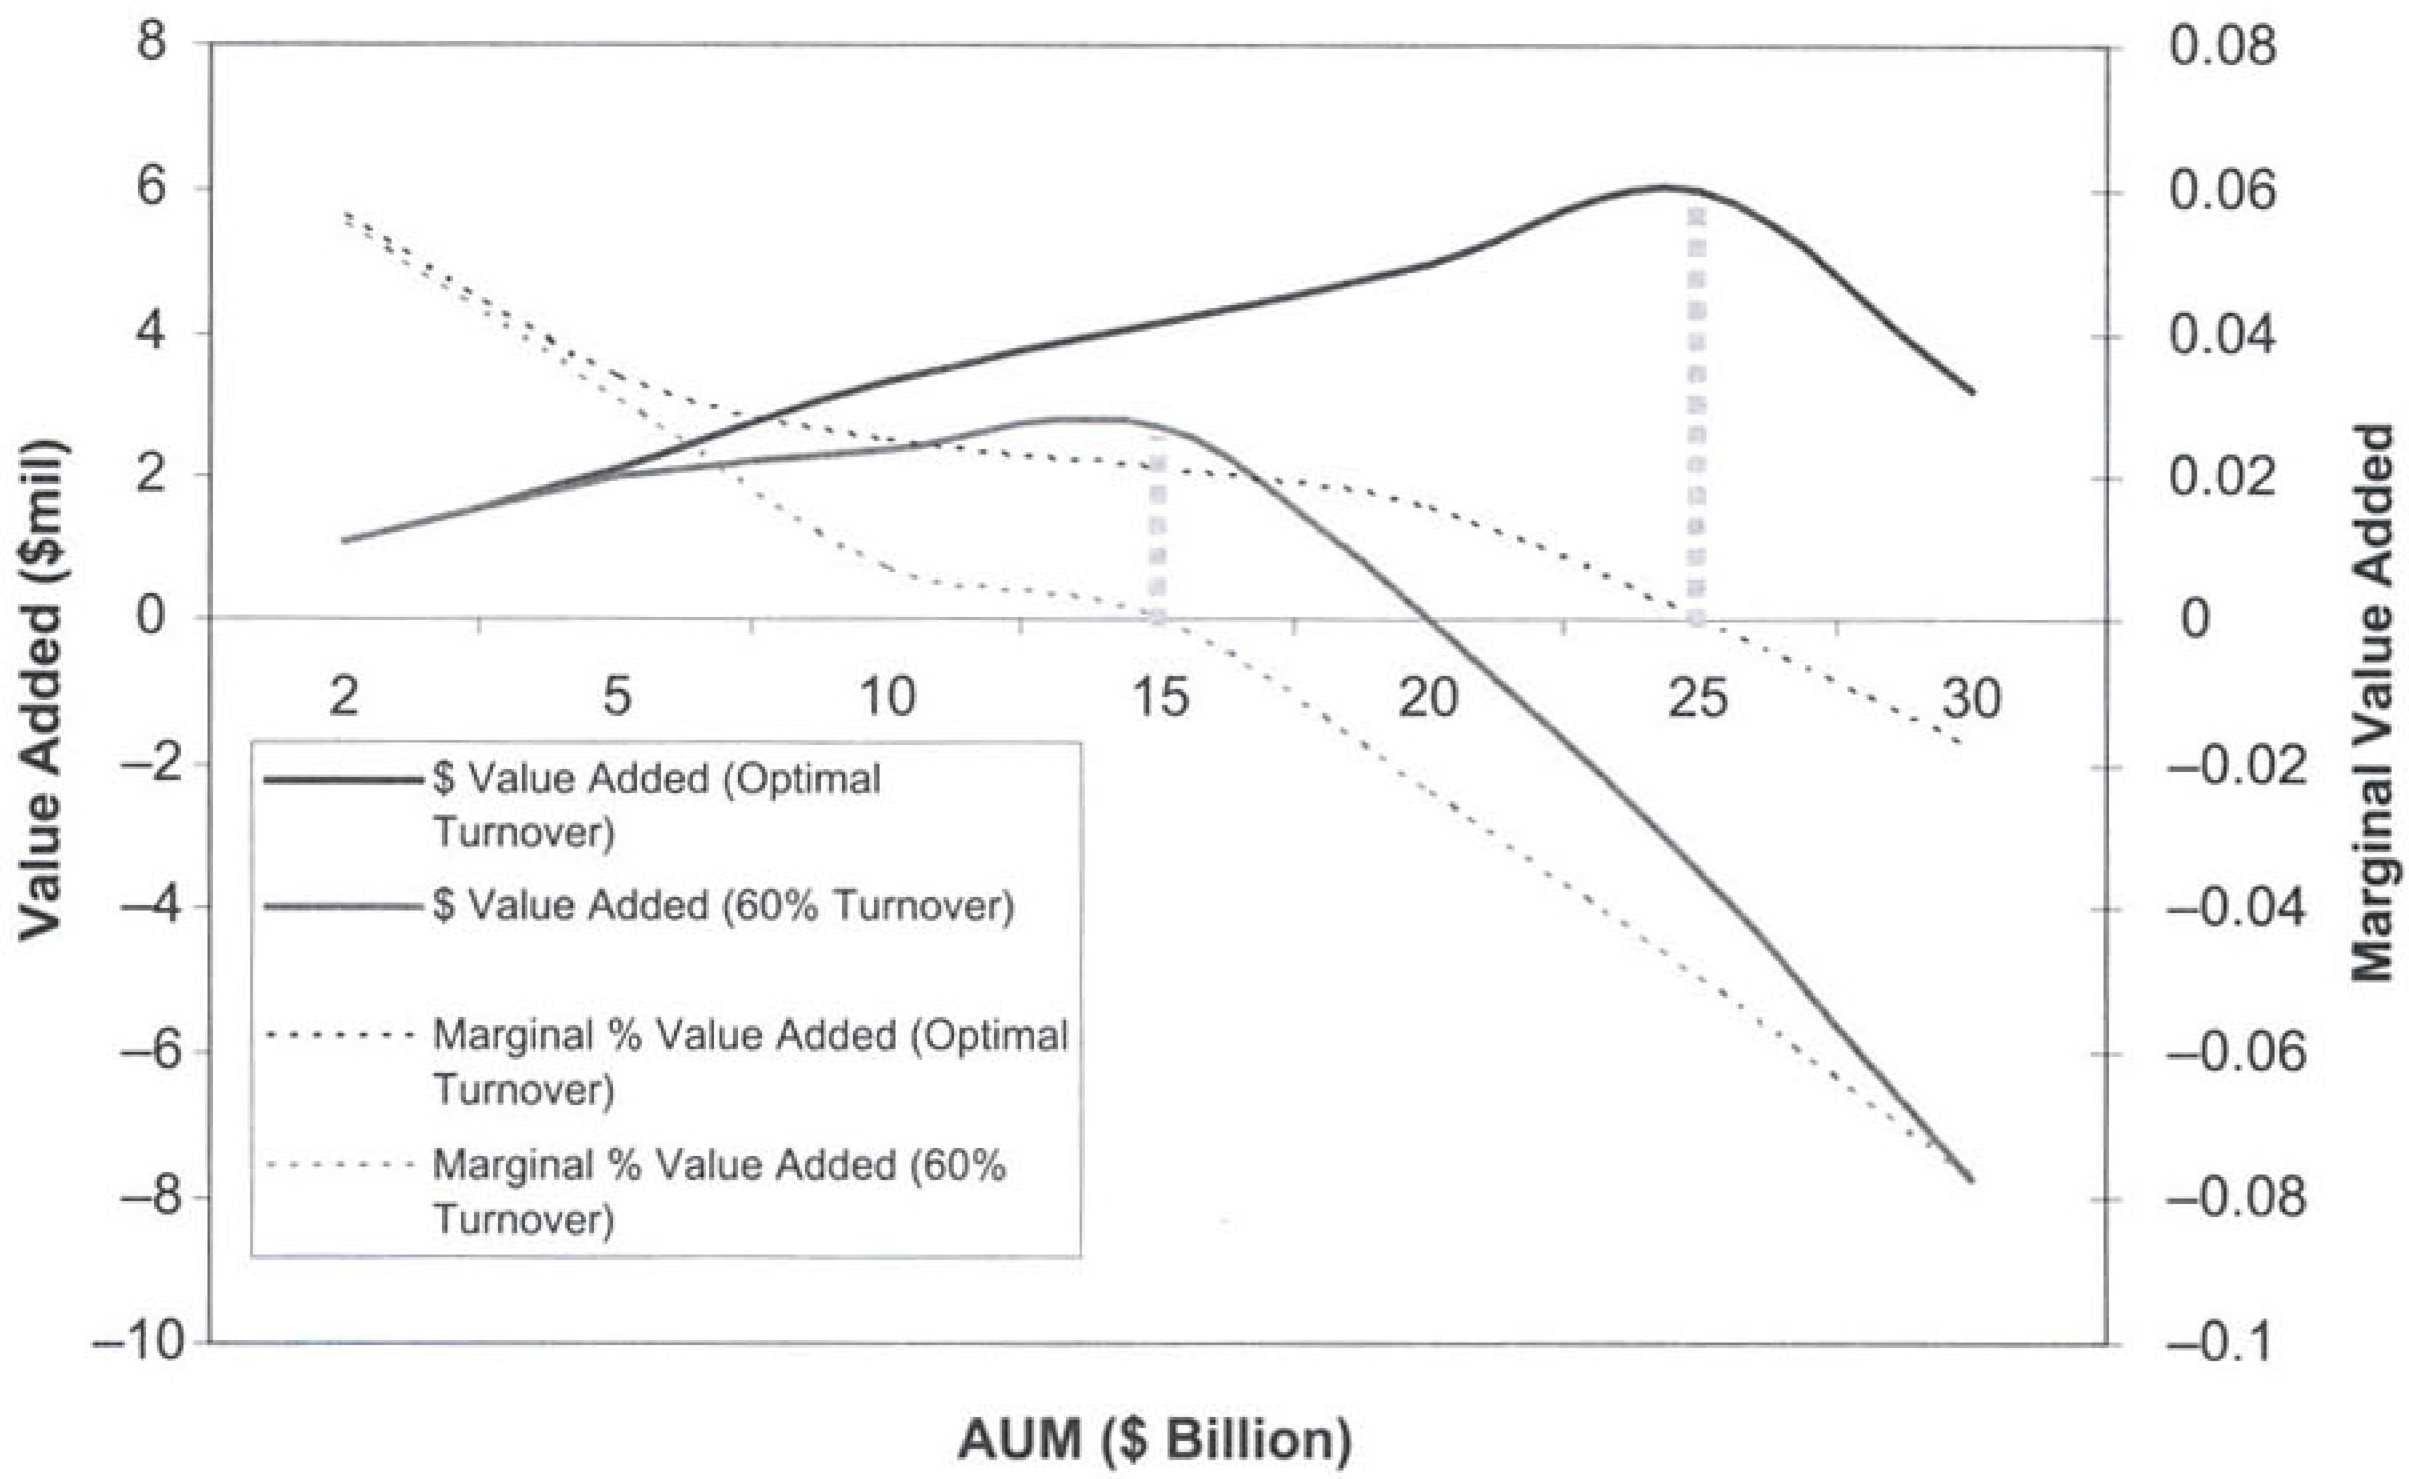

turnover at each *AUM* level, and second, we plot value-added under the assumption that the fund manager keeps the monthly turnover at 60%.

The ability to choose the optimal turnover at each level (which corresponds to moving along the efficient frontier of Exhibit 7) provides the fund manager with the opportunity to grow the fund to about \$25 billion, compared to maintaining 60% turnover at each *AUM* level, which would result in the marginal value-added turning negative at about \$15 billion.

Formalizing the preceding discussion leads to a simple variation of the classical Markowitz optimization problem,

$$
\begin{aligned}
&\max_{\omega_t, AUM} \omega_t'\mu * AUM - \tau * f(AUM, \omega_t, \omega_{t-1}^*) \\
&s.t. \\
&\omega_t'\Sigma\omega_t \le \sigma_0^2 \\
&A\omega_t = b \\
&C\omega_t \ge d
\end{aligned}
\tag{2}
$$

where μ is the $n \times 1$ vector of expected returns, τ is the coefficient of aversion to transaction cost losses, Σ is the $n \times n$ symmetric covariance matrix of asset returns, f is the price impact function, and *AUM* is the amount of assets under management. The vector of weights ω_t contains n decision variables, while the vector of optimal weights at the end of the previous period is denoted ω_{t-1}^*. An upper limit on the portfolio variance is denoted by σ_0^2. We introduce an $m \times n$ matrix, A, and an $m \times 1$ vector, b, to stand for possible linear equality constraints, and a $p \times n$ matrix, C, together with a $p \times 1$ vector, d, to stand for possible linear inequality constraints.

While in the classical Markowitz formulation the optimization is performed over the vector of decision variables ω_t, the problem defined in expression (2) adds *AUM* to the decision set. The objective function is reformulated so that, in expression (2), expected value-added rather than portfolio return is maximized, which is consistent with the discussion relating to Exhibit 8.

Revisiting the "fundamental law of active management" in the context of the net alpha–turnover trade-off sheds light on an interesting question. Specifically, where within the constructs of the fundamental law can we attribute the *increased return* that results from *decreased turnover* in moving from point (B) to point (C)? In its original form, the law states that the expected net active return, $E(R_A)$, can be enhanced by accepting higher active risk, increasing breadth, or improving forecasting ability,

$$E(R_A) \approx IC \cdot \sqrt{N} \cdot \sigma_A \tag{3}$$

where *IC* is the information coefficient, or the cross-sectional correlation between alphas (forecasted active returns) and realized active returns; *N* is the breadth, or the number of independent investment decisions; and σ_A is the active risk, or tracking error. Clarke, de Silva, and Thorley [2002] extended expression (3) to include the transfer coefficient *TC*,

$$E(R_A) = IC \cdot TC \cdot \sqrt{N} \cdot \sigma_A \tag{4}$$

where *TC* is the cross-sectional correlation between risk-adjusted alphas and active weights. As Clarke, de Silva, and Thorley [2002] stated, one "can think of the [*TC*] as an additional adjustment to breadth, *N*, that reflects the reduction in independent bets because of constraints." In other words, any portfolio constraints that impinge on the full realization of alpha forecasts within the portfolio will cause *TC* to be less than one.

We believe that the fundamental law of active management could be invoked to interpret the optimal capacity of the fund. However, in order to do so, one has to reformulate *IC* in terms of net (not gross) alpha. If *IC* is reformulated as the correlation between predicted and realized *net* alphas, then increasing *AUM* would be accompanied by decreasing *IC* (due to transaction costs) and *TC* (due to imposed constraints on weights aimed at controlling transaction costs), or both. Exhibit 9 provides the values of *TC* and net *IC* computed for 0% turnover, optimal turnover, and maximum turnover for which net alpha is still non-negative for every level of *AUM* considered.[14]

Notably, the dynamics of both *IC* and *TC* across different turnover levels agree with our prediction: *TC* is rising with increased turnover (lessened constraints on the portfolio weights) and net *IC* is falling (more drag on portfolio performance due to turnover costs). Predictably, net *IC* is always smaller than average "raw" *IC*, which we assumed to be equal to 0.05. Without computing the volatility of the portfolio for each *AUM*-turnover level (σ_A from expression (4)), we cannot empirically match

EXHIBIT 9

TC and IC for Different *AUM* and Turnover Levels

		0% Turnover	Optimal Turnover	Max Turnover with Non-Negative Net Alpha
$2B	*IC*	0.0276	0.0271	0.0260
	TC	0.6442	0.7422	0.7496
$5B	*IC*	0.0279	0.0276	0.0256
	TC	0.6296	0.7424	0.7545
$10B	*IC*	0.0272	0.0267	0.0264
	TC	0.5880	0.6690	0.6955
$20B	*IC*	0.0263	0.0239	0.0224
	TC	0.6008	0.6603	0.6813
$30B	*IC*	0.0270	0.0249	0.0232
	TC	0.6064	0.6805	0.6951

the observed portfolio returns with the predicted ones, but the overall trend is in line with our expectations.

In practice, we can incorporate both nominal alpha forecasts and market-impact estimates by using an optimizer that properly scales both terms simultaneously in the objective function. Several commercial portfolio optimizers, such as BARRA and ITG Opt®, include the penalty on the *expected* transaction cost in the objective function. Because the transaction costs that appear in the optimization problem are only estimates, having a well-specified market-impact cost function is crucial.

Many active managers are interested in the specific levels of MDV traded in addition to turnover since they do not want to trade beyond a certain fraction of MDV for mandated and/or risk-control reasons (e.g., 10% or 20% of MDV). Translating turnover levels into the average percent MDV traded for the various *AUM* levels demonstrates that the average percent MDV traded increases linearly as either turnover or *AUM* increases, due to the fixed-breadth trade lists and the one-day VWAP trading strategy that we used. We found that for most combinations of turnover and *AUM* that we studied, the corresponding percent MDV traded was below 25%, a range that most portfolio managers would consider reasonable.

To summarize, we made the following assumptions about the portfolio formation and rebalancing process in our simulations: 1) *IC* is fixed for the whole period, 2) the rebalancing frequency is fixed at one month, and 3) the investable universe is the S&P 500. The first two assumptions are innocuous; relaxing them would complicate our modeling without shedding additional light on our main focus. Relaxing the third assumption, however, could provide an opportunity for further research along similar lines. Unreported evidence (see Endnote 8) indicates that the use of less-liquid, small-capitalization stocks from the Russell 2000 would result in lower optimal-turnover rates for given *AUM* and *IC* levels. Perhaps researchers could achieve additional insight by focusing more explicitly on liquidity groups instead of capitalization groups, which typically define the major indices. Such an inquiry, of course, fits within the more general effort of refining stock selection procedures either to increase the optimal-turnover level per the level of *AUM* (along with higher returns) or to increase a fund's threshold capacity for a given level of return.

CONCLUSION

We studied the dependence between an equity strategy's performance and its size using stock-specific market impact costs and alphas and demonstrated the benefits of using a bottom-up approach for calculating a fund's trading costs by linking contemporaneous market-impact costs and absolute returns. The evidence of an inverse relationship between the size of the fund and its net return due to rising market-impact costs is compelling. For instance, given a monthly *IC* of 5%, growing a theoretical active fund's assets from $5 billion to $20 billion decreases annual net return by 2% and decreases optimal monthly turnover from 50% to 34%. Similarly, keeping turnover at the optimal level allows the fund to grow to a larger *AUM* before the value-added of a corresponding strategy peaks. We demonstrated this by using the new concept of an efficient turnover frontier. A typical circumstance in which liquidity becomes particularly important is in changing the breadth of a portfolio or in portfolio transitioning. For example, efforts to increase assets while pursuing smaller tracking error could be frustrating, especially when suitable derivative products are not available. Looking for common factors that allow for more systematic approaches to determining capacity limits across all funds and asset classes is a promising area for future research. This research can enhance the development of new structured products and/or institutional mandates.

APPENDIX

Simulating α

Our goal is to simulate a reasonable forecast of next month's alpha. We assume that a money manager will perform this task every month but, in reality, the forecasting horizon may range from several hours to several years. We also assume that the precision with which forecasted alpha will be close to realized (and observed) return, r, is directly proportional to the manager's forecasting ability, IC, and inversely proportional to a stock's volatility, σ.

Suppose that r and α follow a bivariate normal distribution,

$$\begin{pmatrix} r \\ \alpha \end{pmatrix} \sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma^2, & IC*\sigma^2 \\ IC*\sigma^2, & \sigma^2 \end{pmatrix}\right)$$

Using the conditional distribution of bivariate normal distribution, we obtain

$$(\alpha | r = \bar{r}) \sim N(IC*\bar{r}, \sigma^2*(1-IC^2))$$

This is equivalent to

$$\alpha = IC*r_i + \sigma*\varepsilon*\sqrt{1-IC^2} \quad \text{with } e \sim N(0,1) \qquad \text{(A-1)}$$

Equation (A-1) provides the basis for our simulations.

ENDNOTES

We are grateful to Milan Borkovec, Michael Chigirinskiy, Ian Domowitz, Hans Heidle, Scott Kartinen, Nicolo Menez, and an anonymous referee for many helpful comments that substantially improved this article.

[1]Median daily volume (MDV) is a stock's 21-day median daily share volume.

[2]The ITG Agency Cost Estimator (ACE) is a mathematical model that provides a pre-trade forecast of the price impact cost of an order. The ACE measure of execution costs includes both the bid–ask spread and the price impact of the trade. Explicit cost components (e.g., commissions) can be easily added to the ACE estimate to obtain a total cost of trading. The version of the model used in this article is 2.2.

[3]We use $IC = 0.05$ throughout the article.

[4]Vangelisti [2006] used "a process and alphas very similar to those used to manage a GMO emerging market strategy."

[5]The VWAP of a stock over a period of time is the volume-weighted average price paid per share during that period by all market participants. VWAP strategies are defined as buying and selling a fixed number of shares at an average price that is as close to the VWAP as possible. In practice this amounts to matching the volume pattern of the underlying stock over a specified time horizon (typically one day). Therefore, a trader who follows a VWAP strategy will trade more during periods when the market volume is high and less during periods when the market volume is low. For further discussion of VWAP strategies, see, for instance, Madhavan [2002].

[6]For example, the evidence we are not reporting for the sake of brevity indicates that the ordering of the international markets from least to most expensive in fixed-dollar trade terms coincides with the available floats. McDonald and Richardson [2006] suggested that the index providers' decision to classify South Korea as an "emerging" market is largely for the convenience of institutional investors. With South Korea in their universe, emerging market funds can grow much larger without facing prohibitive turnover and maintenance costs, as South Korean stocks are much more liquid relative to stocks from other emerging markets.

[7]The mean absolute return is 0.46% for the lowest decile and 22.4% for the highest decile. The corresponding mean returns, however, are a mere 0.01% and 0.5%, respectively (all numbers are monthly).

[8]We also ran the same exercise for the R2000 universe and obtained similar results, which are available upon request.

[9]We define monthly turnover as a sum of dollar amounts of sales and purchases in a given month divided by the average *AUM* in that month.

[10]In other words, we assume that the benchmark is cash.

[11]The ITG U.S. Risk Models are equity risk models that assist portfolio managers, researchers, and traders in measuring, analyzing, and managing risk in a variety of market environments and applications. The models are delivered for daily, weekly, and monthly horizons.

[12]We start with cash on January 1 and keep turnover at 100% in each simulation for January; thus, we can achieve a reasonably steady-state portfolio by the end of January. We also drop market-impact and commission costs from January so as not to skew the results with the initial transition from cash. Because of the latter, we ignore January when calculating annual returns (i.e., we calculate only steady-state returns, treating January as a burn-in period).

[13]We thank the referee for this suggestion.

[14]We compute net IC as the cross-sectional correlation coefficient between net realized return and α from Equation (1). We compute TC as the cross-sectional correlation coefficient between α from Equation (1) and the weight computed as described in the paragraph immediately following Equation (1).

REFERENCES

Authers, J. "ETF Industry Is Replacing Clarity with Disquieting Confusion." *Financial Times*, September 30, 2006, p. 12.

Berk, J., and R. Green. "Mutual Fund Flows and Performance in Rational Markets." *Journal of Political Economy*, Vol. 112, No. 6 (2004), pp. 1269–1295.

Bogle, J.C. *Bogle on Mutual Funds: New Perspectives for the Intelligent Investor*. Burr Ridge, IL: Irwin Professional Publishing, 1994.

Chan, N., M. Getmansky, S. Haas, and A. Lo. "Systemic Risk and Hedge Funds." Working Paper No. 4535-05, MIT Sloan School of Management, 2005. Available at http://ssrn.com/abstract=671443.

Chollete, L., R. Naes, and J. Skjeltorp. "What Captures Liquidity Risk? Order-Based Versus Trade-Based Liquidity Measures." Working Paper, 2007. Available at http://ssrn.com/abstract=967598.

Chordia, T., L. Shivakumar, and A. Subrahmanyam. "The Cross-Section of Daily Variation in Liquidity." Working Paper, London Business School, 2001. Available at http://ssrn.com/abstract=263051.

Clarke, R., H. de Silva, and S. Thorley. "Portfolio Constraints and the Fundamental Law of Active Management." *Financial Analysts Journal*, Vol. 58, No. 5 (2002), pp. 48–66.

——. "Performance Attribution and the Fundamental Law." *Financial Analysts Journal*, Vol. 61, No. 5 (2005), pp. 70–83.

Coval, J., and E. Stafford. "Asset Fire Sales (and Purchases) in Equity Markets." Working Paper No. 05-077, Harvard Business School, 2005.

Domowitz, I., O. Hansch, and X. Wang. "Liquidity Commonality and Return Co-movement." *Journal of Financial Markets*, Vol. 8, No. 4 (2005), pp. 351–376.

Grinold, R.C. "The Fundamental Law of Active Management." *Journal of Portfolio Management*, Vol. 15, No. 3 (1989), pp. 30–37.

——. "Alpha Is Volatility Times IC Times Score, or Real Alphas Don't Get Eaten." *Journal of Portfolio Management*, Vol. 20, No. 4 (1994), pp. 9–16.

Grinold, R.C., and R. Kahn. "The Efficiency Gains of Long–Short Investing." *Financial Analysts Journal*, Vol. 56, No. 6 (2000), pp. 40–53.

ITG, Inc. "ITG ACE—Agency Cost Estimator." 2007. Available at http://www.itg.com/news_events/papers/ACE_White_Paper_200705.pdf.

Kahn, R., and S. Shaffer. "The Surprisingly Small Impact of Asset Growth on Expected Alpha." *Journal of Portfolio Management*, Vol. 32, No. 1 (2005), pp. 49–60.

Kamara, A., X. Lou, and R. Sadka. "The Divergence of Liquidity Commonality in the Cross-Section of Stocks." Working Paper, 2007. Available at http://ssrn.com/abstract=943040.

Korajczyk, R., and R. Sadka. "Are Momentum Profits Robust to Trading Costs?" *Journal of Finance*, Vol. 59, No. 3 (2004), pp. 1039–1082.

Madhavan, A. "VWAP Strategies." *Investment Guides*, Spring 2002, pp. 32–39.

Maxey, D. "Mutual Funds Pass $10 Trillion Mark." *Wall Street Journal*, November 30, 2006, p. C11.

McDonald, I., and K. Richardson. "For South Korea, 'Emerging' Label Can Be a Burden." *Wall Street Journal*, July 12, 2006, p. C1.

Perold, A., and R. Salomon. "The Right Amount of Assets under Management." *Financial Analysts Journal*, Vol. 47, No. 3 (1991), pp. 31–39.

Rüffer, R., and L. Stracca. "What Is Global Excess Liquidity and Does It Matter?" Working Paper Series No. 696, European Central Bank, 2006.

Sadka, R. "Momentum and Post-Earnings Announcement Drift Anomalies: The Role of Liquidity Risk." *Journal of Financial Economics*, Vol. 80, No. 2 (2006), pp. 309–349.

Vangelisti, M. "The Capacity of an Equity Strategy." *Journal of Portfolio Management*, Vol. 32, No. 2 (2006), pp. 44–50.

Wermers, R. "Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transactions Costs, and Expenses." *Journal of Finance*, Vol. 55, No. 4 (2000), pp. 1655–1695.

Disclaimer

Any opinions expressed herein reflect the judgment of the individual authors at this date and are subject to change, and they do not necessarily represent the opinions or views of Investment Technology Group, Inc., Realindex Investments or MSCI Inc. The information contained herein has been taken from trade and statistical services and other sources we deem reliable, but we do not represent that such information is accurate or complete, and it should not be relied upon as such. The analyses discussed herein are derived from aggregated ITG client data and are not meant to guarantee future performance or results. This report is for informational purposes and is neither an offer to sell nor a solicitation of an offer to buy any security or other financial instrument. This report does not provide any form of advice (investment, tax, or legal). No part of this article may be reproduced or retransmitted in any manner without permission. 030308-49280.