

IT'S 11 PM—DO YOU KNOW WHERE YOUR LIQUIDITY IS? THE MEAN-VARIANCE-LIQUIDITY FRONTIER*

Andrew W. Lo^a, Constantin Petrov^b, and Martin Wierzbicki^c

We introduce liquidity into the standard mean-variance portfolio optimization framework by defining several measures of liquidity and then constructing three-dimensional mean-variance-liquidity frontiers in three ways: liquidity filtering, liquidity constraints, and a mean-variance-liquidity objective function. We show that portfolios close to each other on the traditional mean-variance efficient frontier can differ substantially in their liquidity characteristics. In a simple empirical example, the liquidity exposure of mean-variance efficient portfolios changes dramatically from month to month, and even simple forms of liquidity optimization can yield significant benefits in reducing a portfolio's liquidity-risk exposure without sacrificing a great deal of expected return per unit risk.



1 Introduction

Liquidity has long been recognized as one of the most significant drivers of financial innovation, and the collapse of several high-profile hedge funds such as Askin Capital Management in 1994 and

Long Term Capital Management in 1998 has only intensified the financial industry's focus on the role of liquidity in the investment management process. Many studies—in both academic journals and more applied forums—have made considerable progress in defining liquidity, measuring the cost of immediacy and price impact, deriving optimal portfolio rules in the presence of transactions costs, investigating the relationship between liquidity and arbitrage, and estimating liquidity risk premia in the context of various partial and general equilibrium asset-pricing models.¹ However, relatively little attention has been paid to the more practical problem of integrating liquidity directly into the portfolio construction process.²

In this paper, we attempt to remedy this state of affairs by modeling liquidity using simple measures

*Research support from the MIT Laboratory for Financial Engineering and Undergraduate Research Opportunities Program is gratefully acknowledged. We thank Simon Lee for excellent research assistance, and Gifford Fong, Svetlana Sussman, and a referee for helpful comments.

^aHarris & Harris Group Professor, MIT Sloan School of Management, and Chief Scientific Officer, AlphaSimplex Group, LLC, 50 Memorial Drive, Cambridge, MA 02142–1347 (corresponding author).

^bResearch Analyst, Fidelity Management and Research Co., 82 Devonshire Street, Boston, MA 02109.

^c838 Green Street, San Francisco, CA 94133.

such as trading volume and percentage bid/offer spreads, and then introducing these measures into the standard mean-variance portfolio optimization process to yield optimal mean-variance-liquidity portfolios. We begin by proposing several measures of the liquidity ℓ_i of an individual security, from which we define the liquidity ℓ_p of a portfolio $\omega_p \equiv [\omega_{p1}\omega_{p2} \dots \omega_{pn}]'$ as the weighted average $\sum_i \ell_i \omega_{pi}$ of the individual securities' liquidities. Using these liquidity measures, we can construct three types of "liquidity-optimized" portfolios: (a) a mean-variance-efficient portfolio subject to a liquidity filter that each security in the portfolio have a minimum level of liquidity ℓ_o ; (b) a mean-variance-efficient portfolio subject to a constraint that the portfolio have a minimum level of liquidity ℓ_o ; and (c) a mean-variance-liquidity-efficient portfolio, where the optimization problem has three terms in its objective function: mean, variance, and liquidity. Using three different definitions of liquidity—turnover, percentage bid/offer spread, and a nonlinear function of market capitalization and trade size—we show empirically that liquidity-optimized portfolios have some very attractive properties, and that even simple forms of liquidity optimization can yield significant benefits in terms of reducing a portfolio's liquidity-risk exposure without sacrificing a great deal of expected return per unit risk.

In Section 2, we describe our simple measures of liquidity, and we define our three liquidity-optimized portfolios in Section 3. We provide an empirical example of liquidity-optimized portfolios in Section 4 for a sample of 50 US stocks using monthly, daily, and transactions data from January 2, 1997 to December 31, 2001, and we conclude in Section 5.

2 Liquidity Metrics

The natural starting point of any attempt to integrate liquidity into the portfolio optimization

process is to develop a quantitative measure of liquidity, i.e., a liquidity metric. Liquidity is a multi-faceted concept, involving at least three distinct attributes of the trading process—price, time, and size—hence a liquid security is one that can be traded quickly, with little price impact, and in large quantities. Therefore, we are unlikely to find a single statistic that summarizes all of these attributes. To represent these distinct features, we start with the following five quantities on which our final liquidity metrics will be based:

$$\text{Trading volume} \equiv \text{Total number of shares traded at time } t \quad (1)$$

$$\text{Logarithm of trading volume} \equiv \log(\text{Trading volume}) \quad (2)$$

$$\text{Turnover} \equiv \frac{\text{Trading volume}}{\text{Shares outstanding}} \quad (3)$$

$$\text{Percentage bid/ask spread} \equiv \frac{\text{Ask}-\text{Bid}}{(\text{Ask} + \text{Bid})/2} \quad (4)$$

$$\text{Loeb price impact function} \equiv f(\text{Trade size, Market cap}) \quad (5)$$

where the first three variables measure the amount of trading and the last two measure the cost.³

Perhaps the most common measure of the liquidity of a security is its trading volume. It is almost tautological to say that a security is more liquid if it is traded more frequently and in greater quantities. Both trading volume and turnover capture this aspect of liquidity, and because these two variables are so highly correlated (see Tables 3 and 4), we will use only one of the three measures of trading activity (1)–(3) in our empirical analysis. Given Lo and Wang's (2000) motivation for turnover in the context of modern asset-pricing models such as the Capital Asset Pricing Model and the Arbitrage Pricing Theory, we shall adopt turnover (3) as our measure of trading activity.

Another popular measure of the liquidity of a security is the cost of transacting in it, either as buyer or seller, hence the bid/ask spread is a natural candidate. Smaller bid/ask spreads imply lower costs of trading, whereas larger bid/ask spreads are partly attributable to a liquidity premium demanded by market-makers for making markets in illiquid securities.⁴

Finally, market capitalization—the market value of total outstanding equity—has also been proposed as an important proxy for liquidity. Larger amounts of outstanding equity tend to be traded more frequently, and at a lower cost because there will be a larger market for the stock. Of course, even a large amount of outstanding equity can be distributed among a small number of major shareholders, yielding little liquidity for the stock, but this seems to

be the exception rather than the rule. We adopt the specification proposed by Loeb (1983) in which he provides estimates of the percentage round-trip total trading cost including: (a) the market-maker's spread; (b) the price concession; and (c) the brokerage commission. The total trading cost is an array with nine capitalization categories and nine block sizes (see Table II in Loeb, 1983). This matrix provides a good approximation for liquidity, but to account for the continuous nature of market capitalization and block sizes beyond his original specification, we interpolate and extrapolate Loeb's table using a two-dimensional spline.⁵ Figure 1 contains a graphical representation of our parametrization of Loeb's specification, and our Matlab sourcecode is provided in Appendix A.1. To minimize the impact of *ad hoc* extrapolation procedures such as the one we use to extend Loeb




Figure 1 Loeb's (1983) price impact function which gives the percentage total cost as a function of block size and market capitalization, with spline interpolation and linear extrapolation.

(1983) (see footnote 5), we assumed a fixed block size of \$250,000 in all our calculations involving Loeb's liquidity metric, and for this size, the extrapolation/capping of the trading cost is used rather infrequently.

2.1 Liquidity Metrics for Individual Securities

To construct liquidity metrics, we begin by computing (1)–(5) with daily data and then aggregating the daily measures to yield monthly quantities. Monthly trading volume is defined as the sum of the daily trading volume for all the days within the month, and monthly log-volume is simply the natural logarithm of monthly trading volume. Monthly turnover is defined as the sum of daily turnover for all the days within the month (see Lo and Wang, 2000 for further discussion). The monthly bid/ask spread measure is defined as a mean of the daily bid/ask spreads for all the days within the month. And finally, the average monthly Loeb price impact measure is defined as a mean of the corresponding daily measures for all days within the month.

Having defined monthly counterparts to the daily variables (1)–(5), we renormalize the five monthly measures to yield quantities that are of comparable scale. Let $\tilde{\ell}_{it}$ represent one of our five liquidity variables for security i in month t . Then the corresponding *liquidity metric* ℓ_{it} is defined as:

$$\ell_{it} \equiv \frac{\tilde{\ell}_{it} - \min_{k,\tau} \tilde{\ell}_{k\tau}}{\max_{k,\tau} \tilde{\ell}_{k\tau} - \min_{k,\tau} \tilde{\ell}_{k\tau}} \quad (6)$$

where the maximum and minimum in (6) are computed over all stocks k and all dates in the sample so that each of the five normalized measures—which we now refer to as a liquidity metric to distinguish it from the unnormalized variable—takes on values strictly between 0 and 1. Therefore, if the turnover-based liquidity metric for a given security is 0.50 in a particular month, this implies that the

level of turnover exceeds the minimum turnover by 50% of the difference between the maximum and minimum turnover for all securities and across all months in our sample. Note that for consistency, we use the *reciprocal* of the monthly bid/ask spread measure in defining ℓ_{it} for bid/ask spreads so that larger numerical values imply more liquidity, as do the other four measures.

2.2 Liquidity Metrics for Portfolios

Now consider a portfolio p of securities defined by the vector of portfolio weights $\omega_p \equiv [\omega_{p1} \omega_{p2} \cdots \omega_{pn}]'$ where $\omega_p' \mathbf{1} = 1$ and $\mathbf{1} \equiv [1 \cdots 1]'$. Assume for the moment that this is a long-only portfolio so that $\omega_p \geq 0$. Then a natural definition of the liquidity ℓ_{pt} of this portfolio is simply:

$$\ell_{pt} \equiv \sum_{i=1}^n \omega_{pi} \ell_{it} \quad (7)$$

which is a weighted average of the liquidities of the securities in the portfolio.

For portfolios that allow short positions, (7) is not appropriate because short positions in illiquid securities may cancel out long positions in equally illiquid securities, yielding a very misleading picture of the overall liquidity of the portfolio. To address this concern, we propose the following definition for the liquidity metric of a portfolio with short positions, along the lines of Lo and Wang's (2000) definition of portfolio turnover:

$$\ell_{pt} \equiv \sum_{i=1}^n \frac{|\omega_{pi}|}{\sum_{j=1}^n |\omega_{pj}|} \ell_{it}. \quad (8)$$

In the absence of short positions, (8) reduces to (7), but when short positions are present, their liquidity metrics are given positive weight as with the long positions, and then all the weights are renormalized by the sum of the absolute values of the weights.

2.3 Qualifications

Although the liquidity metrics described in Sections 2.1 and 2.2 are convenient definitions for purposes of mean-variance portfolio optimization, they have a number of limitations that should be kept in mind. First, (7) implicitly assumes that there are no interactions or cross-effects in liquidity among securities, which need not be the case. For example, two securities in the same industry might have similar liquidity metrics individually, but may become somewhat more difficult to trade when combined in a portfolio because they are considered close substitutes by investors. This assumption can be relaxed by specifying a more complex “liquidity matrix” in which ℓ_{it} are the diagonal entries but where interaction terms ℓ_{ijt} are specified in the off-diagonal entries. In that case, the liquidity metric for the portfolio p is simply the quadratic form:

$$\ell_{pt} \equiv \sum_{i=1}^n \sum_{j=1}^n \omega_{pi} \omega_{pj} \ell_{ijt}. \quad (9)$$

The off-diagonal liquidity metrics are likely to involve subtleties of the market microstructure of securities in the portfolio as well as more fundamental economic links among the securities, hence for our current purposes, we assume that they are zero.

Second, because (7) is a function only of the portfolio weights and not of the dollar value of the portfolio, ℓ_{pt} is scale independent. While this also holds true for mean-variance analysis as a whole, the very nature of liquidity is dependent on scale to some degree. Consider the case where IBM comprises 10% of two portfolios p and q . According to (7), the contribution of IBM to the liquidity of the overall portfolio would be the same in these two cases: 10% times the liquidity metric of IBM. However, suppose that the dollar value of portfolio p is \$100,000 and the dollar value of portfolio q is \$100 million—is a \$10,000 position in IBM identical to a \$10 million position in terms of liquidity?

At issue is the fact that, except for Loeb’s measure of price impact, the liquidity metrics defined by the variables (1)–(4) are not functions of trade size, hence are scale-independent. Of course, this is easily remedied by reparametrizing the liquidity metric ℓ_{it} so that it varies with trade size, much like Loeb’s price impact function, but this creates at least three additional challenges: (a) there is little empirical evidence to determine the appropriate functional specification⁶; (b) trade size may not be the only variable that affects liquidity; and (c) making ℓ_{it} a function of trade size complicates the portfolio optimization problem considerably, rendering virtually all of the standard mean-variance results scale-dependent. For these reasons, we shall continue to assume scale-independence for ℓ_{it} throughout this study (even for Loeb’s price impact function, for which we fix the trade size at \$250,000), and leave the more challenging case for future research.

More generally, the liquidity variables (1)–(5) are rather simple proxies for liquidity, and do not represent liquidity premia derived from dynamic equilibrium models of trading behavior.⁷ Therefore, these variables may not be stable through time and over very different market regimes. However, given their role in influencing the price, time, and size of transactions in equity markets, the five liquidity metrics defined by (1)–(5) are likely to be highly correlated with equilibrium liquidity premia under most circumstances and should serve as reasonable local approximations to the liquidity of a portfolio.

Finally, because our liquidity metrics are *ad hoc* and not the by-product of expected utility maximization, they have no objective interpretation and must be calibrated to suit each individual application. Of course, we might simply assert that liquidity is a sufficiently distinct characteristic of a financial security that investors will exhibit specific preferences along this dimension, much like for a security’s mean and variance. However, unlike mean and variance, it is difficult to identify plausible preference rankings

for securities of varying liquidity levels. Moreover, there are approximation theorems that derive mean-variance preferences from expected utility theory (see, e.g., Levy and Markowitz, 1979), and corresponding results for our liquidity metrics have yet to be developed.

Nevertheless, liquidity is now recognized to be such a significant factor in investment management that despite the qualifications described above, there is considerable practical value in incorporating even *ad hoc* measures of liquidity into standard mean-variance portfolio theory. We turn to this challenge in Section 3.

3 Liquidity-Optimized Portfolios

Armed with quantitative liquidity metrics $\{\ell_{it}\}$ for individual securities and portfolios, we can now incorporate liquidity directly into the portfolio construction process. There are at least three methods for doing so: (a) imposing a liquidity “filter” for securities to be included in a portfolio optimization program; (b) constraining the portfolio optimization program to yield a mean-variance efficient portfolio with a minimum level of liquidity; and (c) adding the liquidity metric into the mean-variance objective function directly. We describe each of these methods in more detail in Sections 3.1–3.3, and refer to portfolios obtained from these procedures as “mean-variance-liquidity (MVL) optimal” portfolios.⁸

3.1 Liquidity Filters

In this formulation, the portfolio optimization process is applied only to those securities with liquidity metrics greater than some threshold level ℓ_o . Denote by U the universe of all securities to be considered in the portfolio optimization process, and let U_o denote the subset of securities in U for which

$$\ell_{it} \geq \ell_o:$$

$$U_o \equiv \{i \in U : \ell_{it} \geq \ell_o\}. \quad (10)$$

The standard mean-variance optimization process can now be applied to the securities in U_o to yield mean-variance-efficient liquidity-filtered portfolios:

$$\min_{\{\omega\}} \frac{1}{2} \omega' \Sigma_o \omega \quad \text{subject to} \quad (11a)$$

$$\mu_p = \omega' \mu_o \quad (11b)$$

$$1 = \omega' \iota \quad (11c)$$

where μ_o is the vector of expected returns of securities in U_o , Σ_o is the return covariance matrix of securities in U_o , and as μ_p is varied, the set of ω_p^* that solve (11) yields the ℓ_o -liquidity-filtered mean-variance efficient frontier.

3.2 Liquidity Constraints

An alternative to imposing a liquidity filter is to impose an additional constraint in the mean-variance optimization problem:

$$\min_{\{\omega\}} \frac{1}{2} \omega' \Sigma \omega \quad \text{subject to} \quad (12a)$$

$$\mu_p = \omega' \mu \quad (12b)$$

$$\ell_o = \begin{cases} \omega' \ell_t & \text{if } \omega \geq 0 \\ \sum_{i=1}^n \frac{|\omega_{pi}|}{\sum_{j=1}^n |\omega_{pj}|} \ell_{it} & \text{otherwise} \end{cases} \quad (12c)$$

$$1 = \omega' \iota \quad (12d)$$

where μ is the vector of expected returns of securities in the unconstrained universe U , Σ is the return covariance matrix of securities in U , $\ell_t \equiv [\ell_{1t} \cdots \ell_{nt}]'$ is the vector of liquidity metrics for securities in U , and as μ_p is varied, the set of ω_p^* that solve (12) yields the ℓ_o -liquidity-constrained mean-variance-efficient frontier. Note that the liquidity constraint (12c) is in two parts, depending on whether ω is long-only or long-short. For simplicity, we impose a non-negativity restriction on ω in our empirical example so that the constraint reduces to $\ell_o = \omega' \ell_t$.

3.3 Mean-Variance-Liquidity Objective Function

Perhaps the most direct method of incorporating liquidity into the mean-variance portfolio optimization process is to include the liquidity metric in the objective function:⁹

$$\max_{\{\omega\}} \omega' \mu - \frac{\lambda}{2} \omega' \Sigma \omega + \phi \omega' \ell_t \quad (13a)$$

$$\text{subject to } 1 = \omega' \iota, 0 \leq \omega \quad (13b)$$

where λ is the risk tolerance parameter, ϕ determines the weight placed on liquidity, and we have constrained ω to be non-negative so as to simplify the expression for the liquidity of the portfolio.

4 An Empirical Example

To illustrate the practical relevance of liquidity metrics for investment management, we construct the three types of liquidity-optimized portfolios described in Section 3 using historical data for 50 US stocks selected from the University of Chicago's Center for Research in Securities Prices (CRSP) and the New York Stock Exchange's Trades and Quotes (TAQ) database for the sample period from January 2, 1997 to December 31, 2001. These 50 stocks are listed in Table 1, and were drawn randomly from 10 market capitalization brackets, based on December 31, 1996 closing prices. These stocks were chosen to provide a representative portfolio with sufficiently diverse liquidity characteristics, and Appendix A.2 provides a more detailed description of our sampling procedure.¹⁰

In Section 4.1 we review the basic empirical characteristics of our sample of stocks and define the mean and covariance estimators that are the inputs to the liquidity-optimized portfolios described in Sections 3.1–3.3. Section 4.2 contains results for liquidity-filtered portfolios, Section 4.3 contains corresponding results for liquidity-constrained

portfolios, and Section 4.4 contains results for portfolios obtained by optimizing a mean-variance-liquidity objective function.

4.1 Data Summary

Table 2 reports summary statistics for the daily prices, returns, turnovers, volume, bid/ask spreads and Loeb measures for the 50 stocks listed in Table 1. Table 2 shows that the average price generally increases with market capitalization, and the minimum and maximum average prices of \$1.72 and \$72.72 correspond to stocks in the first and tenth brackets, respectively. Average daily returns were generally positive, with the exception of a small negative return for GOT. The lower-bracket stocks exhibit very high historical average returns and volatilities, while the top-bracket stocks displayed the opposite characteristics. For example, the average daily returns and volatilities of the stocks in the first and tenth brackets were 0.27% and 7.13%, and 0.06% and 2.4%, respectively.

The relation between daily turnover and market capitalization is less clear due to the fact that turnover is volume normalized by shares outstanding. In general, the mid-tier stocks exhibited the highest turnover, up to 2.13% a day, whereas the daily turnover of bottom-tier and top-tier stocks were only 0.3%–0.4%. However, a clearer pattern emerges from the raw volume numbers. From the first to the fifth bracket, average daily trading volume is typically less than 100 million shares, but a discrete shift occurs starting in the fifth bracket, where daily volume jumps to 300 million shares or more and generally remains at these higher levels for the higher market-cap brackets.

The opposite pattern is observed with the distribution of the percentage bid/ask spread. For small-cap stocks, the average bid/ask spread varies between 1% and 8%. High bid/ask spreads are observed

Table 1 50 US stocks selected randomly within 10 market capitalization brackets, based on December 31, 1996 closing prices. For comparison, market capitalizations based on December 31, 2002 closing prices are also reported.

Ticker	Name	1996 Market Cap (\$MM)	2001 Market Cap (\$MM)	Market Cap Bracket
MANA	MANATRON INC	4.30	13.37	1
SPIR	SPIRE CORP	6.80	21.48	1
WTRS	WATERS INSTRUMENTS INC	7.13	12.57	1
CTE	CARDIOTECH INTERNATIONAL INC	9.02	15.38	1
NCEB	NORTH COAST ENERGY INC	9.09	51.86	1
ALDV	ALLIED DEVICES CORP	12.11	4.85	2
RVEE	HOLIDAY R V SUPERSTORES INC	12.32	10.36	2
DAKT	DAKTRONICS INC	16.76	153.23	2
ANIK	ANIKA RESEARCH INC	18.49	9.93	2
GMCRC	GREEN MOUNTAIN COFFEE INC	20.93	183.21	2
EQTY	EQUITY OIL CO	39.05	22.84	3
STMI	S T M WIRELESS INC	40.94	10.14	3
LTUS	GARDEN FRESH RESTAURANT CORP	42.07	37.60	3
DISK	IMAGE ENTERTAINMENT INC	45.18	37.99	3
ISKO	ISCO INC	48.17	56.91	3
DWCH	DATAWATCH CORP	52.33	3.33	4
LASE	LASERSIGHT INC	53.85	16.39	4
KVHI	K V H INDUSTRIES INC	54.20	65.00	4
GOT	GOTTSCHALKS INC	54.98	32.87	4
MIMS	M I M CORP	60.21	382.31	4
URS	U R S CORP NEW	77.43	490.28	5
AEOS	AMERICAN EAGLE OUTFITTERS INC	77.99	1,881.07	5
DSPG	D S P GROUP INC	81.09	623.53	5
QDEL	QUIDEL CORP	98.21	218.47	5
EFCX	ELECTRIC FUEL CORP	99.30	42.60	5
AEIS	ADVANCED ENERGY INDUSTRIES INC	114.32	847.71	6
ADVS	ADVENT SOFTWARE INC	223.07	1,689.06	6
MOND	ROBERT MONDAVI CORP THE	269.15	348.92	6
NABI	N A B I	302.87	392.91	6
LAMR	LAMAR ADVERTISING CO	427.07	3,496.69	6
HNCS	H N C SOFTWARE INC	597.69	727.84	7
ART	APTARGROUP INC	632.42	1,255.76	7
GGC	GEORGIC GULF CORP	928.16	586.73	7
CMVT	COMVERSE TECHNOLOGY INC	935.52	4,163.84	7
AHG	APRIA HEALTHCARE GROUP INC	959.10	1,361.88	7
BEC	BECKMAN INSTRUMENTS INC NEW	1,113.07	2,699.91	8
ATG	A G L RESOURCES INC	1,173.01	1,270.66	8
ACXM	ACXIOM CORP	1,229.33	1518.49	8
EAT	BRINKER INTERNATIONAL INC	1,236.62	2,922.79	8
XRAY	DENTSPLY INTERNATIONAL INC NEW	1,277.75	2,605.33	8
BCR	BARD C R INC	1,596.78	3,296.27	9
HIB	HIBERNIA CORP	1,621.76	2,829.27	9

Table 1 (*Continued*)

Ticker	Name	1996 Market Cap (\$MM)	2001 Market Cap (\$MM)	Market Cap Bracket
CTL	CENTURY TELEPHONE ENTRPRS INC	1,846.76	4,628.18	9
NI	N I P S C O INDUSTRIES INC	2,399.93	4,768.67	9
LIZ	LIZ CLAIBORNE INC	2,759.06	2,617.30	9
ATML	ATMEL CORP	3,271.16	3,431.05	10
EMN	EASTMAN CHEMICAL CO	4,292.65	3,008.64	10
CLX	CLOROX CO	5,181.06	9,198.42	10
AEP	AMERICAN ELECTRIC POWER INC	7,708.26	14,026.89	10
GIS	GENERAL MILLS INC	9,970.67	18,947.03	10

between the first and fifth brackets, but starting with the fifth bracket, the spread falls rapidly to values as low as 0.19%. For mid- and top-tier stocks, differences in bid/ask spreads are very small. Loeb's (1983) liquidity metric exhibits the same general patterns—for small-cap stocks, the metric is as high as 28%, but by the fourth bracket, the metric stabilizes between 3% and 1.3%. The standard deviation of this metric for the top-tier stocks is close to zero.

Table 3 contains correlation matrices for the average price, market capitalization, return, turnover, volume and Loeb's metric using daily data from January 2, 1997 to December 31, 2001. The three subpanels correspond to correlation matrices for the combined portfolio of 50 stocks, the large-cap subportfolio (the 26th to 50th stocks in Table 1), and the small-cap subportfolio (the 1st to 25th stocks in Table 1), respectively.¹¹ Some of the correlations in Table 3 are unusually high by construction and need not concern us. For example, since turnover is defined as the ratio of volume to shares outstanding, where the latter is generally a slowly varying function through time, the correlation between the volume and the turnover is higher than 90% in each of the three correlation matrices in Table 3. The same is true for the high negative correlation between Loeb's metric and market capitalization.

The correlations between market capitalization, price, and turnover are more significant, confirming the general trends observed in Table 2. In each subportfolio, higher market capitalization corresponds to higher average prices, and higher turnover and volume. The correlations are the strongest in the small-cap subportfolio where the gradients of all the underlying variables are the highest. For example, the correlations between the market capitalization and the turnover in the combined, large- and small-cap subportfolios are 11.94%, 4.01% and 19.87%, respectively. At 90%, the correlation between the market capitalization and average price in the small-cap subportfolio is particularly strong. The relationship between turnover, volume, and Loeb's metric is particularly important because each metric represents an alternate measure of liquidity. With the notable exception of the correlation between Loeb's metric and turnover in the large-cap subportfolio, all correlations have the correct signs and are statistically significant at a 5% level. For example, for the combined portfolio, the turnover-Loeb and volume-Loeb correlations are -8.83% and -12.40%, respectively. The corresponding correlations for the small-cap subportfolio are -14.91% and -17.37%, respectively. The weak correlation between turnover and Loeb's metric for the large-cap subportfolio can be explained by

Table 2 Summary statistics for daily prices, returns, turnover, volume bid/ask spreads, and Loeb measures for 50 US stocks selected randomly within 10 market capitalization brackets. Statistics for prices, returns, turnover, volume, and the Loeb measure are based on daily data from January 2, 1997 to December 31, 2001. Statistics for bid/ask spreads are based on tick data from January 3, 2000 to December 31, 2001. Trading volume is measured in units of millions of shares per day, and the Loeb measure is computed for a fixed block size of \$250,000.

Stock	Average Price (\$)	Return (%)		Turnover (%)		Volume		Bid/Ask (%)		Loeb (%)	
		Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
MANA	4.43	0.26	6.33	0.29	0.57	9	18	4.98	2.91	24.17	5.56
SPIR	6.44	0.39	8.28	0.53	1.90	19	63	5.15	3.66	19.29	5.06
WTRS	5.38	0.20	5.69	0.31	1.41	5	21	5.13	2.38	28.11	2.44
CTE	1.73	0.29	8.27	0.44	1.43	28	95	6.63	3.57	26.67	4.62
NCEB	2.30	0.23	7.07	0.13	0.49	13	32	8.29	8.07	18.57	7.18
ALDV	2.17	0.15	6.85	0.32	0.68	15	32	5.50	2.90	25.84	4.14
RVEE	3.04	0.11	5.36	0.20	0.36	15	27	5.57	3.96	19.44	3.29
DAKT	10.92	0.28	4.74	0.43	0.69	36	79	1.63	1.03	9.90	6.64
ANIK	5.68	0.09	5.88	0.73	1.91	62	177	5.30	3.01	14.09	7.48
GMCR	13.56	0.28	4.65	0.56	0.86	25	45	1.48	1.07	13.76	6.72
EQTY	2.35	0.11	5.58	0.27	0.39	34	50	3.97	3.25	16.93	3.83
STMI	6.25	0.17	8.05	0.82	2.52	57	177	3.74	2.25	14.19	6.42
LTUS	12.62	0.02	3.25	0.51	0.83	26	43	2.21	1.28	8.26	3.73
DISK	4.60	0.13	5.97	0.57	1.01	84	141	2.50	1.56	8.07	3.59
ISKO	6.74	0.14	5.06	0.10	0.18	5	10	5.34	2.79	14.19	3.48
DWCH	1.92	0.18	9.96	1.09	4.97	100	457	6.32	5.02	23.22	6.27
LASE	5.46	0.05	7.00	1.07	1.44	158	228	2.72	1.80	9.01	4.99
KVHI	5.08	0.19	6.73	0.34	0.95	25	69	2.85	1.70	14.15	6.47
GOT	6.57	-0.01	2.88	0.14	0.32	16	40	2.12	1.19	6.41	3.40
MIMS	5.14	0.33	6.95	1.02	1.80	186	361	3.20	2.75	8.24	4.51
URS	17.75	0.13	2.84	0.27	0.29	41	47	0.81	0.50	3.22	0.30
AEOS	34.97	0.35	4.54	2.13	1.96	931	1,225	0.31	0.17	2.16	0.97
DSPG	29.62	0.26	4.96	2.04	2.43	300	328	0.52	0.24	2.97	0.55
QDEL	4.35	0.17	5.17	0.50	0.70	121	169	1.94	1.16	4.87	2.23
EFCX	5.01	0.18	8.11	1.35	4.10	236	648	1.40	0.71	8.53	5.49
AEIS	27.22	0.28	5.47	1.02	1.51	280	450	0.45	0.22	2.74	0.58
ADVS	46.29	0.24	4.80	1.02	1.08	208	331	0.52	0.36	2.49	0.81
MOND	38.98	0.04	2.69	0.97	1.28	79	104	0.51	0.28	3.15	0.01
NABI	5.40	0.19	6.09	0.66	0.91	236	321	1.57	0.88	3.30	0.62
LAMR	37.91	0.13	3.30	0.68	0.84	340	425	0.34	0.18	1.90	0.77
HNCS	37.57	0.25	5.46	1.55	1.76	404	448	0.54	0.31	2.63	0.54
ATR	35.80	0.09	2.38	0.24	0.23	74	71	0.51	0.28	2.62	0.24
GGC	21.11	0.01	2.74	0.47	0.54	148	169	0.49	0.23	2.94	0.18
CMVT	68.65	0.15	4.45	2.01	1.78	2,065	3,095	0.15	0.07	1.61	0.52
AHG	16.16	0.10	3.99	0.61	0.72	318	376	0.55	0.36	2.66	0.52
BEC	49.48	0.09	1.93	0.48	0.34	170	130	0.25	0.11	1.56	0.34
ATG	19.76	0.04	1.46	0.21	0.17	117	93	0.47	0.23	2.39	0.24
ACXM	22.29	0.07	4.27	0.93	1.17	719	1,005	0.35	0.14	1.89	0.63
EAT	23.90	0.12	2.64	0.60	0.60	451	493	0.33	0.15	1.66	0.54
XRAY	34.75	0.08	2.09	0.52	0.54	256	275	0.37	0.19	1.55	0.30
BCR	43.38	0.10	2.11	0.60	0.81	314	418	0.22	0.11	1.30	0.00

Table 2 (Continued)

Stock	Average Price (\$)	Return (%)		Turnover (%)		Volume		Bid/Ask (%)		Loeb (%)	
		Mean	SD	Mean	SD	Mean	SD	Mean	SD	Mean	SD
HIB	14.91	0.06	2.23	0.30	0.24	449	381	0.51	0.30	1.30	0.02
CTL	39.95	0.09	2.11	0.37	0.34	428	459	0.31	0.13	1.30	0.00
NI	29.05	0.04	1.61	0.39	0.54	503	692	0.33	0.18	1.30	0.00
LIZ	42.34	0.05	2.36	0.64	0.48	385	282	0.23	0.12	1.30	0.00
ATML	22.94	0.11	4.86	2.13	1.42	4,177	3,514	0.28	0.15	1.41	0.35
EMN	50.82	0.01	2.05	0.40	0.28	313	220	0.23	0.12	1.30	0.00
CLX	72.72	0.07	2.37	0.42	0.34	708	690	0.23	0.11	1.30	0.00
AEP	42.15	0.04	1.43	0.26	0.18	595	454	0.19	0.10	1.30	0.00
GIS	58.42	0.06	1.29	0.35	0.27	791	781	0.22	0.11	1.30	0.00

Table 3 Correlation matrices (in percent) for average price, market capitalization, average return, turnover, volume, and the Loeb measure for the combined sample of 50 randomly selected securities (five from each of 10 market capitalization brackets), and large- and small-capitalization subportfolios (the 25 largest and 25 smallest market capitalization securities, respectively, of the 50), using daily data from January 2, 1997 to December 31, 2001. The Loeb measure is computed for a fixed block size of \$250,000.

	Price	Market Cap	Return	Turnover	Volume	Loeb
<i>Combined Sample</i>						
Price	100.0	79.1	6.0	10.5	6.4	-63.1
Market Cap	79.1	100.0	4.8	11.9	19.0	-70.4
Return	6.0	4.8	100.0	7.4	6.2	-4.1
Turnover	10.5	11.9	7.4	100.0	95.0	-8.8
Volume	6.4	19.0	6.2	95.0	100.0	-12.4
Loeb	-63.1	-70.4	-4.1	-8.8	-12.4	100.0
<i>Large Capitalization Stocks</i>						
Price	100.0	67.5	5.5	0.1	-6.8	-43.3
Market Cap	67.5	100.0	4.1	4.0	14.3	-52.4
Return	5.5	4.1	100.0	-0.4	-1.6	-2.9
Turnover	0.1	4.0	-0.4	100.0	92.9	-2.8
Volume	-6.8	14.3	-1.6	92.9	100.0	-7.4
Loeb	-43.3	-52.4	-2.9	-2.8	-7.4	100.0
<i>Small Capitalization Stocks</i>						
Price	100.0	90.7	6.5	20.8	19.6	-82.9
Market Cap	90.7	100.0	5.5	19.9	23.7	-88.4
Return	6.5	5.5	100.0	15.3	13.9	-5.3
Turnover	20.8	19.9	15.3	100.0	97.1	-14.9
Volume	19.6	23.7	13.9	97.1	100.0	-17.4
Loeb	-82.9	-88.4	-5.3	14.9	-17.4	100.0

the lack of variation in Loeb's metric at higher capitalization levels, a feature evident in Table 2. High positive return-volume and return-turnover correlations in the small-cap subportfolio—13.92% and 15.29%, respectively—are also noteworthy, and is not observed in the large-cap subportfolio.

Table 4 is similar to Table 3 except for the addition of another liquidity metric, the percentage bid/ask spread. Because our source of bid/ask spread data was available only starting on January 3, 2000, all the correlations were re-estimated with the more recent two-year sample from January 3, 2000 to December 31, 2001.¹² The patterns in Table 4 are similar to those in Table 3. Market capitalization is positively correlated with average price, turnover, and volume, and is negatively correlated with Loeb's metric and the bid/ask spread. For the combined portfolio, the turnover-Loeb and the volume-Loeb correlations as well as the turnover-bid/ask and the volume-bid/ask correlations are of the order of -10% , that is, they have the correct sign and are statistically significant. For the large-cap subportfolio, turnover-Loeb, volume-Loeb, turnover-bid/ask, and volume-bid/ask correlations are all statistically insignificant. For the combined portfolio, and large- and small-cap subportfolios, the bid/ask-Loeb correlations are strong and equal to 27.48%, 14.68% and 40.29%, respectively.

Tables 3 and 4 confirm that the correlations between the various liquidity measures—turnover, volume, Loeb's metric, and the bid/ask spread—are generally consistent with each other, yet are not all perfectly correlated, hence each measure seems to capture certain aspects of liquidity not reflected in the others. The single exception is volume and turnover, which are extremely highly correlated, so we eliminate volume and log-volume from consideration and confine our attention to the following three liquidity measures in our empirical analysis: turnover, bid/ask spreads, and Loeb's metric.

To compute mean-variance-liquidity frontiers, we require estimates of the expected return μ and covariance matrix Σ of the 50 stocks in our sample. Using daily returns data from January 2, 1997 to December 31, 2001, we compute the following standard estimators:

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T \mathbf{R}_t \quad (14a)$$

$$\hat{\Sigma} = \frac{1}{T-1} \sum_{t=1}^T (\mathbf{R}_t - \hat{\mu})(\mathbf{R}_t - \hat{\mu})' \quad (14b)$$

where $\mathbf{R}_t \equiv [R_{1t} \cdots R_{50t}]'$ is the vector of date- t returns of the 50 stocks in our sample. We convert these estimates to a monthly frequency by multiplying by 21, the number of trading days per month. Liquidity-optimized portfolios may then be constructed with these estimates and any one of the liquidity metrics defined in Section 2.

To underscore the fact that liquidity can vary considerably from one month to the next, in Sections 4.2–4.4 we will construct liquidity-optimized portfolios for the months listed in Table 5, which include the start and end of our sample as controls, as well as months that contain significant liquidity events such as the default of Russian government debt in August 1998 and the terrorist attacks of September 11, 2001.

4.2 The Liquidity-Filtered Frontier

Given estimates $\hat{\mu}$ and $\hat{\Sigma}$ of the mean and covariance matrix of the 50 stocks in our sample, we can readily extract the filtered counterparts μ_o and Σ_o with which to construct the liquidity-filtered mean-variance frontier according to Section 3.1. For expositional convenience, we focus only on one of the three liquidity metrics—turnover—in this section, and will consider the other two liquidity metrics in Section 4.3.¹³

Table 4 Correlation matrices (in percent) for average price, market capitalization, average return, turnover, volume, the Loeb measure, and bid/ask spreads for the combined sample of 50 randomly selected securities (five from each of 10 market capitalization brackets), and large- and small-capitalization subportfolios (the 25 largest and 25 smallest market capitalization securities, respectively, of the 50), using daily data from January 3, 2000 to December 31, 2001. The Loeb measure is computed for a fixed block size of \$250,000, and bid/ask spreads are daily averages based on intradaily tick data.

	Price	Market Cap	Return	Turnover	Volume	Loeb	Bid/Ask
<i>Combined Sample</i>							
Price	100.0	87.9	7.9	14.3	10.2	-60.6	-31.0
Market Cap	87.9	100.0	7.0	11.0	12.6	-66.7	37.9
Return	7.9	7.0	100.0	6.6	6.0	-5.0	-0.4
Turnover	14.3	11.0	6.6	100.0	97.7	-8.6	-8.6
Volume	10.2	12.6	6.0	97.7	100.0	-9.2	-10.6
Loeb	-60.6	-66.7	-5.0	-8.6	-9.2	100.0	27.5
Bid/ask	-31.0	-37.9	-0.4	-8.6	-10.6	27.5	100.0
<i>Large Capitalization Stocks</i>							
Price	100.0	84.4	7.2	2.0	-4.1	-39.6	-26.0
Market Cap	84.4	100.0	6.6	-1.1	0.9	-44.4	-34.8
Return	7.2	6.6	100.0	0.1	-0.5	-3.5	-1.0
Turnover	2.0	-1.1	0.1	100.0	96.8	0.7	-5.4
Volume	-4.1	0.9	-0.5	96.8	100.0	0.6	-8.1
Loeb	-39.6	-44.4	-3.5	0.7	0.6	100.0	14.7
Bid/ask	-26.0	-34.8	-1.0	-5.4	-8.1	14.7	100.0
<i>Small Capitalization Stocks</i>							
Price	100.0	91.4	8.7	26.7	24.5	-81.6	-36.0
Market Cap	91.4	100.0	7.4	23.0	24.2	-89.1	-41.0
Return	8.7	7.4	100.0	13.1	12.6	-6.6	0.2
Turnover	26.7	23.0	13.1	100.0	98.6	-18.0	-11.8
Volume	24.5	24.2	12.6	98.6	100.0	-19.0	-13.2
Loeb	-81.6	-89.1	-6.6	-18.0	-19.0	100.0	40.3
Bid/ask	-36.0	-41.0	0.2	-11.8	-13.2	40.3	100.0

In Table 6 we report the means and standard deviations of two benchmark portfolios—the global minimum-variance portfolio, and the tangency portfolio—and the Sharpe ratio of the tangency portfolio for various levels of the liquidity filter for each of the months listed in Table 5.¹⁴ For each set of portfolios of a given month, the first row—with “Liquidity Metric” set to 0.00—corresponds to portfolios with no liquidity filters imposed, hence these refer to the usual mean-variance benchmark portfolios. Subsequent rows of a given month

correspond to portfolios with increasingly stricter liquidity filters imposed at fixed increments until the liquidity filter yields too few securities to construct a meaningful efficient frontier (four securities or less).

Consider the first group of rows in Table 6, for December 1996, the start of our sample period. Without any liquidity filtering, the tangency portfolio has an expected monthly return of 4.13% and a monthly return standard deviation of 5.72%,

Table 5 Significant months during the sample period from December 1996 to December 2001 for which liquidity-optimized portfolios are constructed.

Date	Event
December 1996	Beginning of sample
August 1998	Russian default/LTCM
October 1998	Fall of 1998
March 2000	First peak of S&P 500
July 2000	Second peak of S&P 500
April 2001	First bottom of S&P 500
September 2001	9/11 terrorist, attacks, second bottom of S&P 500
December 2001	End of sample

implying a monthly Sharpe ratio of 0.65.¹⁵ However, with a liquidity filter of 2.29 imposed—only stocks with liquidity metrics greater than or equal to 2.29 are included in the portfolio—the tangency portfolio changes to one with an expected return of 4.23%, a standard deviation of 8.20%, and a Sharpe ratio of 0.46. Although the expected return increases, the standard deviation increases more than proportionally so as to yield a Sharpe ratio that is only 71% of the unfiltered portfolio's Sharpe ratio. As the liquidity filter threshold ℓ_o in (10) is increased, the Sharpe ratio of the tangency portfolio will continue to decrease since it represents the best risk/reward trade-off available for a given set of securities, and portfolios with lower values of ℓ_o include all the securities of portfolios with higher values of ℓ_o but not vice-versa. For the month of December 1996, a liquidity filter of 9.15 yields a Sharpe ratio for the tangency portfolio of 0.39, almost half the value of the unfiltered portfolio's Sharpe ratio.

However, the trade-off between liquidity and the risk/reward profile of the efficient frontier is quite different during March 2000, the height of the bull market when the first peak of the S&P 500

is attained. For the same level of liquidity, 2.29, the Sharpe ratio of the tangency portfolio is 0.64, virtually identical to that of the unfiltered portfolio.¹⁶ In contrast to December 1996, liquidity seems to be less problematic in March 2000, with little or no material impact of liquidity filtering on the Sharpe ratio. In fact, even in the extreme case of a filter of 9.15, the resulting Sharpe ratio is 0.50 in March 2000, which is higher than the Sharpe ratio of the December 1996 filtered tangency portfolio with a filter of 2.29. In fact, a filter level of 22.86 is required in March 2000 to yield a Sharpe of 0.40, which is approximately the risk/reward profile of the portfolio with the most extreme liquidity filter in December 1996, a filter of 9.15.

The results in Table 6 are more readily appreciated via graphical representation since we have now expanded the focus from two dimensions (mean and variance) to three (mean, variance, and liquidity). Figures 2 and 3 display liquidity-filtered mean-variance-liquidity (MVL) efficient frontiers for each of the dates in Table 5. At the “ground level” of each of the three-dimensional coordinate cubes in Figures 2 and 3, we have the familiar expected-return and standard-deviation axes. The liquidity threshold ℓ_o of (10) is measured along the vertical axis. In the plane of ground level, the liquidity level is zero hence the efficient frontier is the standard Markowitz mean-variance efficient frontier, and this frontier will be identical across all the months in our sample since the estimated mean $\hat{\mu}$ and covariance matrix $\hat{\Sigma}$ are based on the entire sample of daily data from January 2, 1997 to December 31, 2001 and do not vary over time. However, as the liquidity metric is used to filter the set of securities to be included in constructing the mean-variance-efficient frontier, the risk/reward profile of the frontier will change, as depicted by the color of the surface. By construction, the liquidity of a filtered portfolio is always greater than or equal to the liquidity threshold ℓ_o , and since the normalization of all liquidity metrics is performed

Table 6 Monthly means and standard deviations of tangency and minimum-variance portfolios of liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, March 2000, July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

Date	Liquidity Metric	Tangency		Min Var		Sharpe
		Mean	SD	Mean	SD	
1996-12	0.00	4.13	5.72	1.53	3.37	0.65
1996-12	2.29	4.23	8.20	1.49	4.91	0.46
1996-12	4.57	5.72	13.04	2.49	8.58	0.40
1996-12	6.86	6.32	15.10	2.51	9.71	0.39
1996-12	9.15	6.41	15.36	5.29	14.14	0.39
1998-08	0.00	4.13	5.72	1.53	3.37	0.65
1998-08	2.29	4.22	6.94	1.60	4.29	0.55
1998-08	4.57	5.96	13.69	1.84	7.69	0.40
1998-08	6.86	6.36	15.28	2.47	9.61	0.39
1998-08	9.15	6.36	16.21	4.06	12.77	0.37
1998-10	0.00	4.13	5.72	1.53	3.37	0.65
1998-10	2.29	3.53	6.52	1.48	3.86	0.48
1998-10	4.57	4.13	8.59	1.79	5.38	0.43
1998-10	6.86	6.07	13.96	2.42	9.27	0.40
1998-10	9.15	6.07	13.96	2.80	9.60	0.40
1998-10	11.43	6.18	14.75	2.70	9.68	0.39
2000-03	0.00	4.13	5.72	1.53	3.37	0.65
2000-03	2.29	4.25	6.02	1.60	3.57	0.64
2000-03	4.57	4.31	6.90	1.69	4.20	0.56
2000-03	6.86	4.98	8.86	2.44	6.36	0.51
2000-03	9.15	5.71	10.63	4.25	9.41	0.50
2000-03	11.43	5.69	10.61	4.53	9.58	0.50
2000-03	13.72	6.01	11.54	5.07	10.72	0.48
2000-03	16.00	6.09	12.60	4.92	11.41	0.45
2000-03	18.29	6.11	12.64	5.13	11.69	0.45
2000-03	20.58	6.14	14.44	4.86	12.80	0.40
2000-03	22.86	6.14	14.44	4.86	12.80	0.40
2000-03	25.15	4.32	16.00	3.68	14.62	0.24
2000-07	0.00	4.13	5.72	1.53	3.37	0.65
2000-07	2.29	3.88	6.43	1.53	3.79	0.54
2000-07	4.57	4.98	10.55	2.33	7.50	0.43
2000-07	6.86	5.94	13.17	3.90	11.14	0.42
2000-07	9.15	6.34	15.69	4.85	13.80	0.38
2001-04	0.00	4.13	5.72	1.53	3.37	0.65
2001-04	2.29	4.40	6.88	1.52	3.90	0.58
2001-04	4.57	6.45	11.67	2.47	7.70	0.52
2001-04	6.86	6.36	13.61	2.73	9.54	0.44

Table 6 (Continued)

Date	Liquidity Metric	Tangency		Min Var		Sharpe
		Mean	SD	Mean	SD	
2001-04	9.15	6.46	15.33	4.19	12.84	0.39
2001-04	11.43	6.68	16.79	3.97	13.26	0.37
2001-09	0.00	4.13	5.72	1.53	3.37	0.65
2001-09	2.29	4.27	6.84	1.70	4.19	0.56
2001-09	4.57	5.34	11.02	1.54	6.39	0.44
2001-09	6.86	5.30	11.09	2.32	7.22	0.44
2001-09	9.15	6.50	14.49	5.55	13.41	0.41
2001-12	0.00	4.13	5.72	1.53	3.37	0.65
2001-12	2.29	3.84	6.46	1.47	3.75	0.53
2001-12	4.57	5.04	10.17	1.50	5.94	0.45
2001-12	6.86	5.37	11.53	2.26	8.30	0.43
2001-12	9.15	6.50	14.51	4.43	12.40	0.42

cross-sectionally as well as through time, the color and the height of the frontiers at different dates have the same meaning and can be compared to one another.

Figures 2 and 3 show that as the liquidity filter is imposed, the frontier located at ground level rises steeply—implying relatively little impact on the risk/reward trade-off—until the liquidity threshold reaches the level of the least liquid stock in the portfolio. When the threshold ℓ_o is incremented further, some of the illiquid stocks fail to satisfy the liquidity filter and are eliminated from the filtered portfolio. As the number of stocks in the portfolio is reduced in this fashion, the MVL frontier becomes less efficient and the frontier surface shifts inward, in the north-east direction.¹⁷ For sufficiently high liquidity thresholds, too few securities satisfy the filter and it becomes impossible to compute a non-degenerate MVL frontier, hence the graph ends beyond these levels.¹⁸

The evolution of the MVL-efficient frontier is highly dependent on the underlying trends in the liquidity distribution. During our 5-year sample

period, the average monthly turnover of our randomly selected portfolio of 50 stocks grew steadily from 0.56% in 1997 to 0.90% in 2000, along with the level of the market. In 2001, the market declined dramatically and the average turnover decreased to 0.70%. The higher moments of turnover—the standard deviation, skewness, and kurtosis—followed similar but somewhat more dramatic trends. At 0.17% and 0.16% in 1997 and 1998, respectively, the standard deviation of turnover was almost unchanged as the market rallied. In 2000, when average turnover peaked at 0.90%, the standard deviation of turnover also peaked at 0.42%, i.e., the distribution of turnover expanded. At the same time, extremely high skewness and kurtosis during the boom years of 1999 and 2000 indicated that a small number of stocks enjoyed very active trading. As markets declined in 2001, the moments of the distribution of turnover returned to their 1997 levels.

These patterns are borne out by the graphs in Figures 2 and 3. The upper left subplot in Figure 2 shows the MVL-efficient frontier calculated using turnover in December 1996. At this point in time,

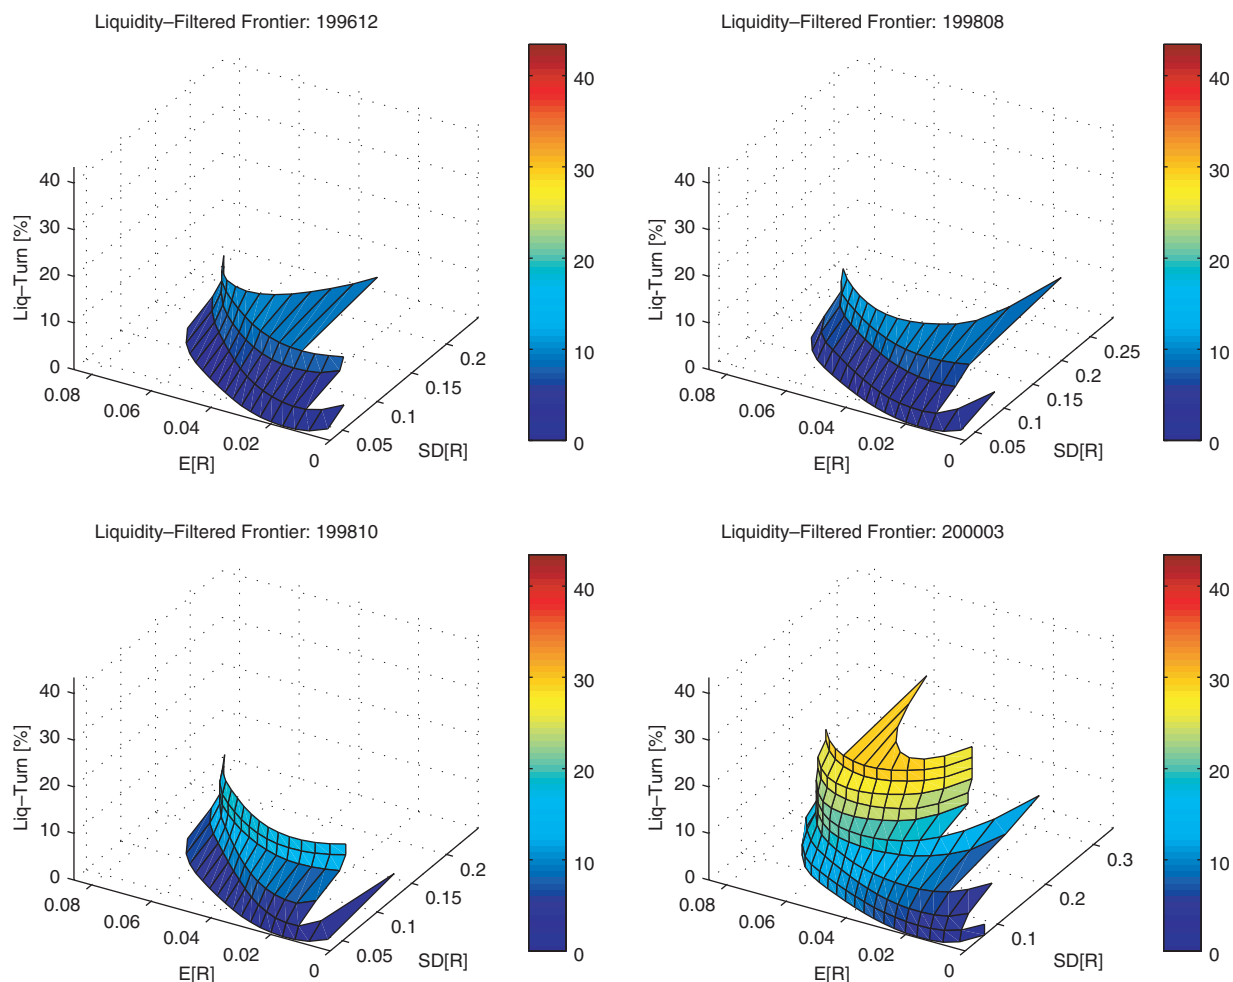


Figure 2 Liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, and March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month. Color strips to the right of each figure provide the correspondence between liquidity levels and the spectrum.

the turnover distribution was quite compressed by historical standards and its mean was relatively low. When the liquidity filter is raised, the frontier shifts to the northeast and its risk/return profile deteriorates. Similar patterns are observed in the upper right and lower left subplots in Figure 2, corresponding to August 1998 and October 1998, respectively. Although the levels of the S&P 500 in both months were similar, the liquidity conditions were apparently more favorable in October 1998,

which is depicted by a brighter color and steeper MVL surface in the latter case. In March 2000 (lower right subplot of Figure 2), the market reached its peak. During that time, the mean and standard deviation of turnover were both very high, making the liquidity filter almost irrelevant up to a very high liquidity threshold. However, during the bear market of late 2000 and 2001 (Figure 3), liquidity deteriorated considerably and the MVL-efficient frontier flattens out to levels comparable with 1996.

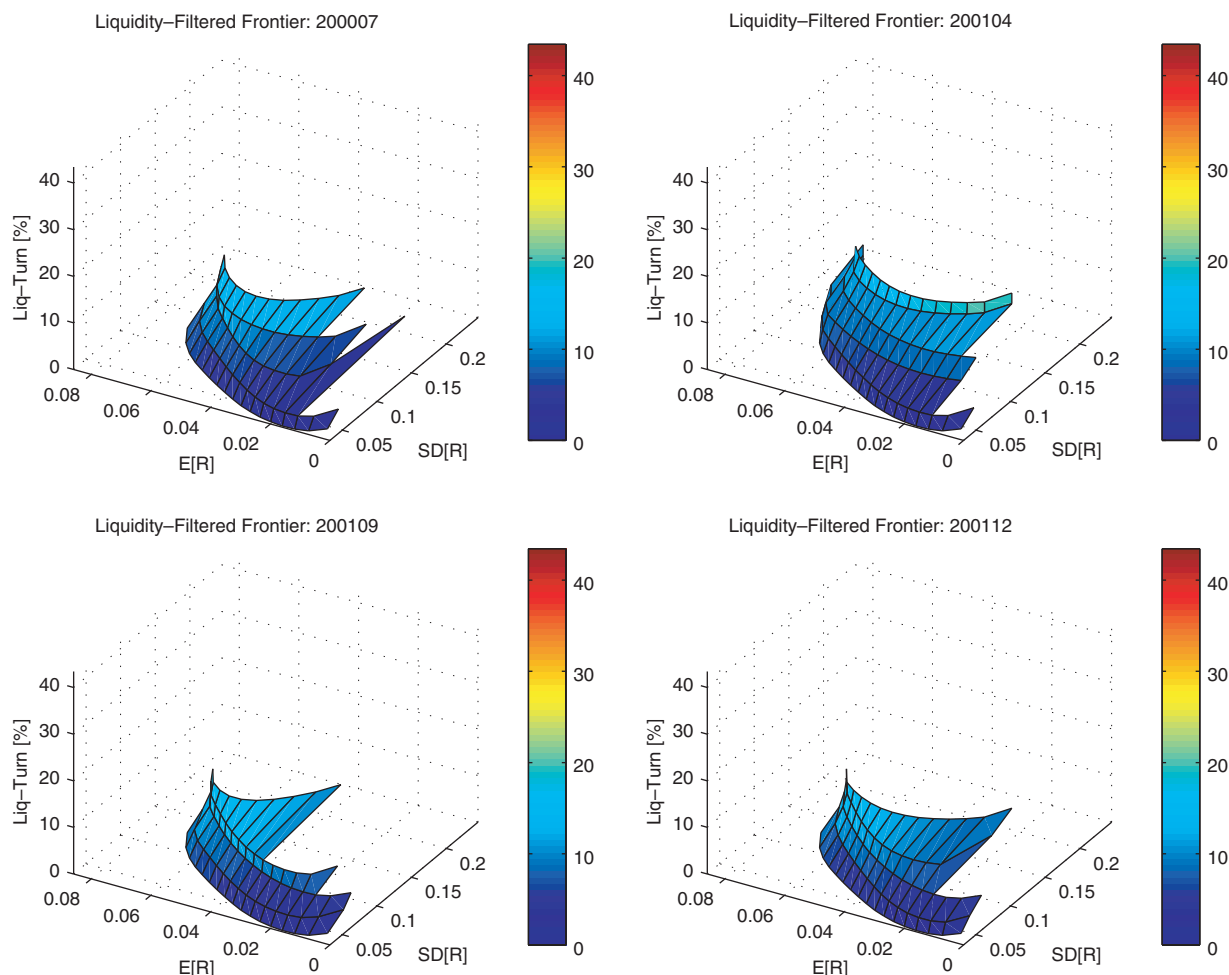


Figure 3 Liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances are estimated with daily returns data from January 2, 1997 to December 31, 2001. Color strips to the right of each figure provide the correspondence between liquidity levels and the spectrum.

An alternative to describing the evolution of the MVL surface is to select a small number of characteristic points on this surface and to plot the trajectories of these points in mean-standard deviation-liquidity space through time. For any mean-variance-efficient frontier, the most relevant point is, of course, the tangency portfolio. In Figures 4 and 5, the *trajectories* of the tangency portfolio are plotted for various levels of the liquidity filter and over time. Each point along the trajectory corresponds to the tangency portfolio of

the efficient frontier for a given liquidity threshold ℓ_ρ . The numerical value of the threshold (in percent) is displayed next to the tangency point, and the position of each point is projected onto the ground-level plane for visual clarity. In addition, two sets of lines are drawn on the ground-level plane: a straight line connecting the riskless portfolio to each tangency portfolio (whose slope is the Sharpe ratio of the tangency portfolio), and curved lines which are MVL frontiers for various levels of the liquidity filter. For each figure, the trajectory

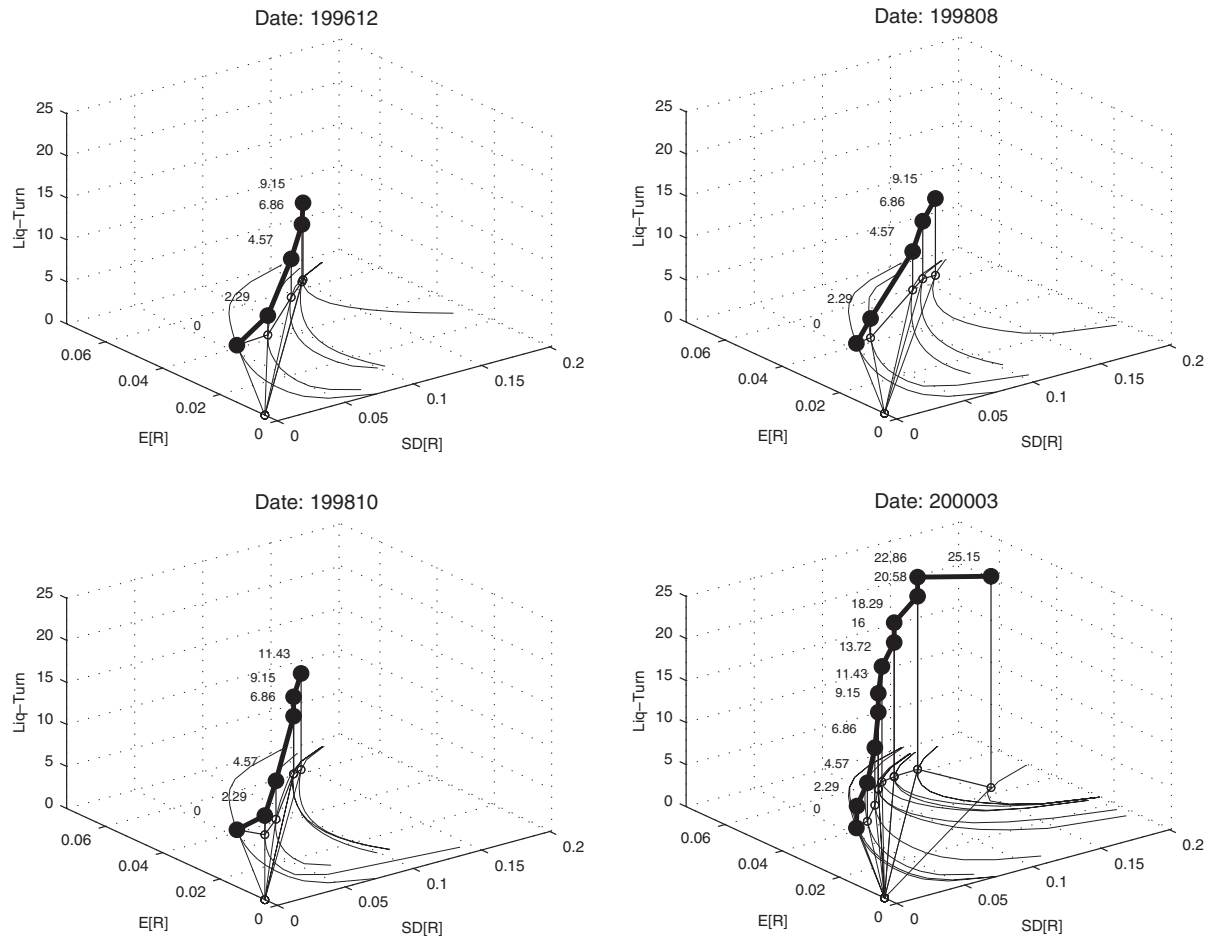


Figure 4 Trajectories of the tangency portfolio for liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, and March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

of the tangency point starts at the same location on the ground-level plane. In the absence of any liquidity effects, the trajectory of the tangency portfolio would be vertical and its projection onto the ground-level plane would coincide with its starting point, but because the liquidity filter does have an impact in filtering out certain securities, as the threshold increases, the trajectory of the tangency portfolio moves eastward and away from the viewer. The ground-level projection of the tangency trajectory moves initially in the east/northeast direction but always yielding less desirable Sharpe ratios. In

some cases, as the liquidity threshold increases, the ground-level projection of the tangency portfolio turns southeast, yielding tangency portfolios with higher volatility and lower expected return, but with higher levels of liquidity (see, for example, the lower right subplot, for March 2000, in Figure 4). At some point, when it becomes impossible for any of the 50 randomly selected securities to satisfy the liquidity filter, the trajectory terminates. The dynamics of the trajectory of the tangency portfolio is a qualitative alternative to assessing the impact of liquidity on the characteristics of a mean-variance optimal portfolio.

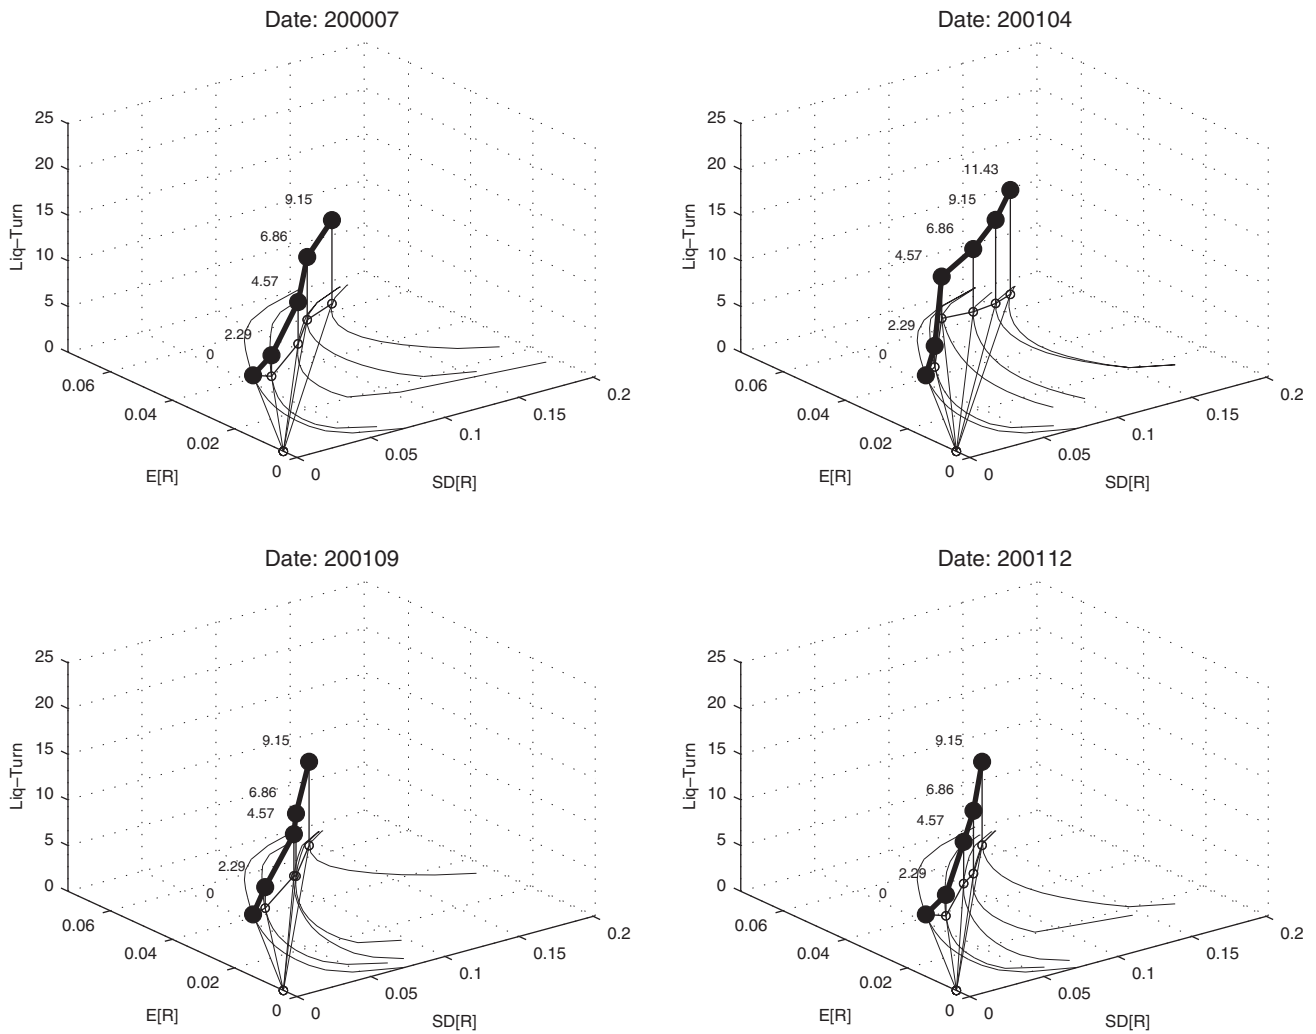


Figure 5 Trajectories of the tangency portfolio for liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

The graphs in Figures 4 and 5 show that for successively higher liquidity filters, the risk/reward profile of the efficient frontier—as measured by the tangency portfolio—worsens, but at different rates for different months. Figure 6 depicts the time variation of this trade-off more explicitly by graphing the trajectories of Sharpe ratios as a function of the liquidity filter for each of the months in Table 5. This two-dimensional representation of a three-dimensional object is a

simple way to highlight the trade-off between liquidity and investment performance. When the level of liquidity is high (March 2000), the Sharpe ratio declines rather slowly in response to rising levels of liquidity filtering, but when liquidity conditions are poor (September 2001), the Sharpe ratio falls precipitously as the liquidity threshold is increased. For liquidity-filtered portfolios, the decline in performance takes the form of discrete jumps because the liquidity threshold changes the




Figure 6 Sharpe ratio trajectories of tangency portfolios of liquidity-filtered MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric, as a function of the liquidity filter, for the months of December 1996, August 1998, October 1998, March 2000, July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month. Thicker lines represent trajectories from more recent months.

composition of the portfolio by filtering out illiquid stocks. We shall see in Section 4.3 that imposing liquidity constraints can smooth out these jumps.

4.3 The Liquidity-Constrained Frontier

The liquidity-filtered portfolios described in Section 4.2 illustrate the potential value of

incorporating simple notions of liquidity into the portfolio construction process, but a more direct approach is to impose liquidity constraints directly into the optimization problem as described in Section 3.2. Table 7 summarizes the characteristics of liquidity-constrained portfolios for the same 50 stocks considered in Section 4.2 using the same liquidity metric, monthly normalized turnover.

Table 7 Monthly means and standard deviations of tangency and minimum-variance portfolios of liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks, (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, March 2000, July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

Date	Liquidity Threshold	Tangency		Min Var		Sharpe	Date	Liquidity Threshold	Tangency		Min Var		Sharpe
		Mean	SD	Mean	SD				Mean	SD			
1996-12	0.00	4.13	5.72	1.53	3.37	0.65	2000-03	34.29	6.01	14.74	4.84	13.44	0.38
1996-12	2.29	4.13	5.72	1.53	3.39	0.65	2000-03	36.58	6.03	16.08	4.84	14.66	0.35
1996-12	4.57	4.99	7.36	1.69	4.15	0.62	2000-03	38.87	6.03	17.61	4.86	16.08	0.32
1996-12	6.86	5.71	9.53	1.98	5.69	0.55	2000-03	41.15	6.00	19.33	4.85	17.70	0.29
1996-12	9.15	5.78	11.18	2.26	7.66	0.48	2000-03	43.44	5.83	20.85	4.76	19.45	0.26
1996-12	11.43	5.65	13.03	2.61	9.88	0.40	2000-07	0.00	4.13	5.72	1.53	3.37	0.65
1996-12	13.72	5.28	14.86	2.83	12.39	0.33	2000-07	2.29	4.13	5.72	1.53	3.37	0.65
1998-08	0.00	4.13	5.72	1.53	3.37	0.65	2000-07	4.57	4.12	5.70	1.73	3.62	0.65
1998-08	2.29	4.13	5.72	1.53	3.38	0.65	2000-07	6.86	4.96	7.23	1.97	4.42	0.63
1998-08	4.57	4.81	6.93	1.76	4.09	0.63	2000-07	9.15	5.92	9.38	2.33	5.61	0.59
1998-08	6.86	5.90	9.44	2.14	5.57	0.58	2000-07	11.43	6.14	10.61	2.70	7.06	0.54
1998-08	9.15	6.11	10.97	2.60	7.56	0.52	2000-07	13.72	6.17	11.78	3.09	8.67	0.49
1998-08	11.43	6.12	12.69	3.16	9.84	0.45	2000-07	16.00	6.24	13.25	3.50	10.37	0.44
1998-08	13.72	6.13	14.95	3.81	12.38	0.38	2000-07	18.29	6.36	15.08	3.91	12.15	0.39
1998-10	0.00	4.13	5.72	1.53	3.37	0.65	2000-07	20.58	6.51	17.26	4.32	14.00	0.35
1998-10	2.29	4.13	5.72	1.53	3.37	0.65	2001-04	0.00	4.13	5.72	1.53	3.37	0.65
1998-10	4.57	4.13	5.72	1.55	3.42	0.65	2001-04	2.29	4.13	5.72	1.53	3.37	0.65
1998-10	6.86	4.46	6.33	1.66	3.75	0.64	2001-04	4.57	4.16	5.77	1.63	3.66	0.65
1998-10	9.15	4.98	7.42	1.76	4.33	0.61	2001-04	6.86	5.33	7.95	1.69	4.45	0.61
1998-10	11.43	5.52	8.69	1.90	5.09	0.59	2001-04	9.15	5.90	9.53	1.94	5.59	0.57
1998-10	13.72	5.62	9.38	2.02	5.98	0.55	2001-04	11.43	5.92	10.45	2.09	6.95	0.53
1998-10	16.00	5.66	10.10	2.25	6.98	0.52	2001-04	13.72	5.80	11.48	2.31	8.48	0.47
1998-10	18.29	5.63	10.85	2.45	8.03	0.48	2001-04	16.00	5.55	12.63	2.55	10.10	0.40
1998-10	20.58	5.56	11.67	2.65	9.13	0.44	2001-04	18.29	5.28	14.19	2.78	11.80	0.34
1998-10	22.86	5.51	12.62	2.84	10.27	0.40	2001-09	0.00	4.13	5.72	1.53	3.37	0.65
1998-10	25.15	5.37	13.51	3.02	11.46	0.37	2001-09	2.29	4.13	5.72	1.53	3.37	0.65
1998-10	27.44	4.96	13.97	3.17	12.70	0.32	2001-09	4.57	4.13	5.72	1.79	3.65	0.65
2000-03	0.00	4.13	5.72	1.53	3.37	0.65	2001-09	6.86	4.63	6.57	2.10	4.42	0.64
2000-03	2.29	4.13	5.72	1.53	3.37	0.65	2001-09	9.15	5.49	8.23	2.50	5.52	0.61
2000-03	4.57	4.13	5.72	1.53	3.37	0.65	2001-09	11.43	6.05	9.65	2.92	6.86	0.58
2000-03	6.86	4.13	5.72	1.73	3.48	0.65	2001-09	13.48	6.34	10.87	3.40	8.36	0.54
2000-03	9.15	4.12	5.70	1.97	3.82	0.65	2001-09	16.00	6.44	11.99	4.04	10.01	0.50
2000-03	11.43	4.54	6.41	2.24	4.33	0.64	2001-09	18.29	6.55	13.48	4.75	11.83	0.45
2000-03	13.72	5.06	7.38	2.52	4.98	0.63	2001-12	0.00	4.13	5.72	1.53	3.37	0.65
2000-03	16.00	5.61	8.47	2.79	5.73	0.61	2001-12	2.29	4.13	5.72	1.53	3.37	0.65
2000-03	18.29	5.77	9.04	3.06	6.55	0.59	2001-12	4.57	4.11	5.70	1.67	3.64	0.65
2000-03	20.58	5.87	9.64	3.33	7.43	0.57	2001-12	6.86	4.96	7.19	1.91	4.52	0.63
2000-03	22.86	5.93	10.26	3.60	8.35	0.54	2001-12	9.15	5.88	9.14	2.33	5.81	0.59
2000-03	25.15	5.96	10.95	3.87	9.31	0.51	2001-12	11.43	6.35	10.68	2.87	7.35	0.55
2000-03	27.44	5.98	11.74	4.14	10.29	0.47	2001-12	13.72	6.55	12.02	3.47	9.06	0.51
2000-03	29.72	6.00	12.64	4.42	11.31	0.44	2001-12	16.00	6.69	13.49	4.24	10.97	0.46
2000-03	32.01	6.01	13.62	4.67	12.36	0.41	2001-12	18.29	6.80	15.13	5.07	13.11	0.42

In contrast to the liquidity-filtered portfolios of Table 6, the results in Table 7 show that the performance of liquidity-constrained portfolios is considerably more attractive, with generally higher Sharpe ratios for the same liquidity thresholds and smoother transitions as the threshold is increased. For example, for the month of December 1996, an increase in the liquidity threshold from 0.00 to 2.29 yields a drop in the Sharpe ratio from 0.65 to 0.46 for the liquidity-filtered portfolios in Table 6, but Table 7 shows no decline in the Sharpe ratio for the liquidity-constrained portfolios. In fact, for every month in Table 5, imposing a liquidity constraint of 2.29 has virtually no impact on the Sharpe ratio, and in some months, e.g., March 2000, the threshold can be increased well beyond 2.29 without any loss in performance for the tangency portfolio.

The intuition for these improvements lies in the fact that in contrast to liquidity filtering—which eliminates securities that fall below the liquidity threshold—liquidity-constrained portfolios generally contain all 50 securities and the portfolio weights are adjusted accordingly so as to achieve the desired liquidity threshold. Rather than simply dropping securities that fall below the liquidity threshold, the liquidity-constrained portfolios underweight them and overweight the more liquid securities, yielding Sharpe ratios that are larger than those of liquidity-filtered portfolios for the same liquidity threshold, and smoother functions of the liquidity threshold.

The intuition for the advantages of liquidity constraints over liquidity filtering is not tied to the turnover liquidity metric, but carries over to the other two metrics as well. Table 8 summarizes the characteristics of liquidity-constrained portfolios for all three liquidity metrics—turnover, Loeb, and bid/ask spread—during March 2000 and December 2001. For all three metrics, and during both months, it is clear that the Sharpe ratios of the tangency portfolio are generally unaffected by the

first few levels of liquidity constraints, in contrast to the behavior of the liquidity-filtered portfolios of Table 6.¹⁹ However, Table 8 does show that the three metrics behave somewhat differently as market conditions change. During the height of the market in March 2000, the turnover and Loeb metrics yield a larger number of feasible liquidity-constrained efficient portfolios than the bid/ask metric, but in the midst of the bear market in December 2001, it is the Loeb and bid/ask metrics that yield more feasible efficient portfolios. While this may seem to suggest that the Loeb metric is the most robust of the three, the comparison is not completely fair since we have fixed the block size for the Loeb metric at \$250,000, and the price impact of such a transaction is likely to be quite different between March 2000 and December 2001.²⁰ The three liquidity metrics capture distinct—albeit overlapping—aspects of liquidity, and which metric is most useful depends intimately on the nature of the application at hand.

A graphical representation of the turnover-constrained MVL frontier renders an even clearer illustration of the difference between liquidity-filtered and liquidity-constrained portfolios. Figures 7 and 8 contain the liquidity-constrained counterparts to Figures 2 and 3. In the upper left subplot of Figure 7, which contains the MVL frontier for December 1996, the period when the distribution of average turnover was at its historically low mean and standard deviation, the sail-like surface is rather flat and covers relatively little surface area. The infeasibility of the constrained portfolio optimization problem at higher liquidity thresholds is responsible for the tattered edges of the surface starting at the fourth liquidity level (note that the size of the liquidity increments is identical across all months and all the axes have the same scale). At the highest levels of liquidity, only the most liquid segments of the MVL frontier appear in Figure 7. Because of the generally positive correlation between liquidity and market capitalization,

Table 8 Monthly means and standard deviations of tangency and minimum-variance portfolios of liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), for three liquidity metrics—turnover, Loeb's (1983) price impact measure, and bid/ask spread—for March 2000 and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

Liquidity Threshold	Tangency		Min Var		Sharpe	Liquidity Threshold	Tangency		Min Var		Sharpe
	Mean	SD	Mean	SD			Mean	SD	Mean	SD	
March 2000						84.17	3.29	5.09	1.45	3.48	0.56
<i>Turnover-Constrained Portfolios</i>						89.12	3.22	5.28	1.42	3.58	0.53
0.00	4.13	5.72	1.53	3.37	0.65	94.08	3.18	5.63	1.39	3.71	0.49
2.29	4.13	5.72	1.53	3.37	0.65	<i>Bid/Ask-Constrained Portfolios</i>					
4.57	4.13	5.72	1.53	3.37	0.65	0.00	4.13	5.72	1.53	3.37	0.65
6.86	4.13	5.72	1.73	3.48	0.65	2.46	4.13	5.72	1.53	3.37	0.65
9.15	4.12	5.70	1.97	3.82	0.65	4.91	4.13	5.72	1.53	3.37	0.65
11.43	4.54	6.41	2.24	4.33	0.64	7.37	4.13	5.72	1.53	3.37	0.65
13.72	5.06	7.38	2.52	4.98	0.63	9.82	3.94	5.45	1.54	3.37	0.64
16.00	5.61	8.47	2.79	5.73	0.61	12.28	3.60	5.09	1.54	3.37	0.62
18.29	5.77	9.04	3.06	6.55	0.59	14.73	3.29	5.01	1.45	3.47	0.57
20.58	5.87	9.64	3.33	7.43	0.57	17.19	3.10	5.45	1.35	3.75	0.49
22.86	5.93	10.26	3.60	8.35	0.54	19.65	3.24	7.06	1.36	4.16	0.40
25.15	5.96	10.95	3.87	9.31	0.51	22.10	3.98	11.23	1.24	5.20	0.32
27.44	5.98	11.74	4.14	10.29	0.47	December 2001					
29.72	6.00	12.64	4.42	11.31	0.44	<i>Turnover-Constrained Portfolios</i>					
32.01	6.01	13.62	4.67	12.36	0.41	0.00	4.13	5.72	1.53	3.37	0.65
34.29	6.01	14.74	4.84	13.44	0.38	2.29	4.13	5.72	1.53	3.37	0.65
36.58	6.03	16.08	4.84	14.66	0.35	4.57	4.11	5.70	1.67	3.64	0.65
38.87	6.03	17.61	4.86	16.08	0.32	6.86	4.96	7.19	1.91	4.52	0.63
41.15	6.00	19.33	4.85	17.70	0.29	9.15	5.88	9.14	2.33	5.81	0.59
43.44	5.83	20.85	4.76	19.45	0.26	11.43	6.35	10.68	2.87	7.35	0.55
<i>Loeb-Constrained Portfolios</i>						13.72	6.55	12.02	3.47	9.06	0.51
0.00	4.13	5.72	1.53	3.37	0.65	16.00	6.69	13.49	4.24	10.97	0.46
4.95	4.13	5.72	1.53	3.37	0.65	18.29	6.80	15.13	5.07	13.11	0.42
9.90	4.13	5.72	1.53	3.37	0.65	<i>Loeb-Constrained Portfolios</i>					
14.85	4.13	5.72	1.53	3.37	0.65	0.00	4.13	5.72	1.53	3.37	0.65
19.81	4.13	5.72	1.53	3.37	0.65	4.95	4.13	5.72	1.53	3.37	0.65
24.76	4.13	5.72	1.53	3.37	0.65	9.90	4.13	5.72	1.53	3.37	0.65
29.71	4.13	5.72	1.53	3.37	0.65	14.85	4.13	5.72	1.53	3.37	0.65
34.66	4.13	5.72	1.53	3.37	0.65	19.81	4.13	5.72	1.53	3.37	0.65
39.61	4.13	5.72	1.53	3.37	0.65	24.76	4.13	5.72	1.53	3.37	0.65
44.56	4.13	5.72	1.53	3.37	0.65	29.71	4.13	5.72	1.53	3.37	0.65
49.51	4.15	5.75	1.53	3.37	0.65	34.66	4.13	5.72	1.53	3.37	0.65
54.46	4.06	5.62	1.54	3.37	0.65	39.61	4.13	5.72	1.53	3.37	0.65
59.42	3.88	5.36	1.54	3.37	0.64	44.56	4.13	5.72	1.53	3.37	0.65
64.37	3.73	5.18	1.54	3.37	0.64	49.51	4.13	5.72	1.53	3.37	0.65
69.32	3.60	5.06	1.54	3.37	0.63	54.46	4.12	5.71	1.54	3.37	0.65
74.27	3.49	5.01	1.53	3.38	0.61	59.42	4.00	5.53	1.54	3.37	0.65
79.22	3.38	4.99	1.49	3.42	0.59						

Table 8 (*Continued*)

Liquidity Threshold	Tangency		Min Var		Sharpe	Liquidity Threshold	Tangency		Min Var		Sharpe	
	Mean	SD	Mean	SD			Mean	SD	Mean	SD		
64.37	3.81	5.27	1.53	3.37	0.64	14.73	4.13	5.72	1.53	3.37	0.65	
69.32	3.64	5.08	1.54	3.37	0.63	17.19	4.13	5.72	1.53	3.37	0.65	
74.27	3.52	5.00	1.54	3.38	0.62	19.65	4.12	5.71	1.54	3.37	0.65	
79.22	3.38	4.95	1.52	3.41	0.59	22.10	4.13	5.73	1.54	3.37	0.65	
84.17	3.27	5.00	1.46	3.48	0.57	24.56	4.17	5.78	1.547	3.37	0.65	
89.12	3.17	5.16	1.42	3.57	0.53	27.01	4.08	5.64	1.54	3.37	0.65	
94.08	3.07	5.44	1.39	3.70	0.48	29.47	3.97	5.48	1.54	3.37	0.65	
	<i>Bid/Ask-Constrained Portfolios</i>						31.92	3.84	5.30	1.54	3.37	0.64
0.00	4.13	5.72	1.53	3.37	0.65	34.38	3.72	5.16	1.54	3.37	0.64	
2.46	4.13	5.72	1.53	3.37	0.65	36.84	3.60	5.01	1.54	3.37	0.63	
4.91	4.13	5.72	1.53	3.37	0.65	39.29	3.49	4.91	1.54	3.37	0.62	
7.37	4.13	5.72	1.53	3.37	0.65	41.75	3.38	4.83	1.53	3.37	0.61	
9.82	4.13	5.72	1.53	3.37	0.65	44.20	3.29	4.79	1.51	3.38	0.60	
12.28	4.13	5.72	1.53	3.37	0.65	46.66	3.19	4.77	1.46	3.40	0.58	

and the fact that the large-cap stocks in our sample have modest expected returns and volatilities as compared to the smaller-cap stocks, at higher liquidity threshold levels portfolios on the MVL frontier consist mostly of defensive large-cap equities.

In the upper right sub-plot of Figure 7 (August 1998), liquidity conditions have improved—the MVL frontier rises up from the ground-level plane almost vertically, and up to the third liquidity threshold, the shape of the frontier remains almost unaffected by the liquidity constraint. In the lower left sub-plot of Figure 7 we observe a dramatic increase in liquidity—the MVL frontier is twice as tall as the December 1996 frontier, and the level of liquidity at which the surface starts bending to the right is significantly higher than in the previous figures. In the lower right subplot of Figure 7, corresponding to the first peak in the S&P 500 (March 2000), the MVL frontier is at its tallest and it is apparent that the liquidity constraint is irrelevant up to a very high liquidity threshold.

Figure 8 tells a very different story. The shape and height of the MVL frontier change dramatically

starting with the upper left subplot for July 2000 (the second peak of the S&P 500) and moving clockwise to April 2001 (the first bottom of the S&P 500), September 2001 (the terrorist attacks on 9/11) and December 2001 (the last month of the simulation). In the face of the bear market of 2000–2001, liquidity conditions have clearly deteriorated, and Figure 8 provides a detailed roadmap of the dynamics of this trend.

The dynamics of liquidity-constrained MVL frontiers can also be seen through the trajectories of the tangency portfolio, contained in Figures 9–11. As with the liquidity-filtered trajectories in Figures 4–6, the trajectories in Figures 9 and 10 originate at the same point on the ground-level plane because the lowest-level frontier is unaffected by the liquidity constraint, and the trajectories remain vertical until the first liquidity threshold, at which point they begin to move initially in the northeast direction and, in some cases, eventually turning towards the southeast direction, until they reach a sufficiently high liquidity threshold where the tangency portfolios no longer exist.

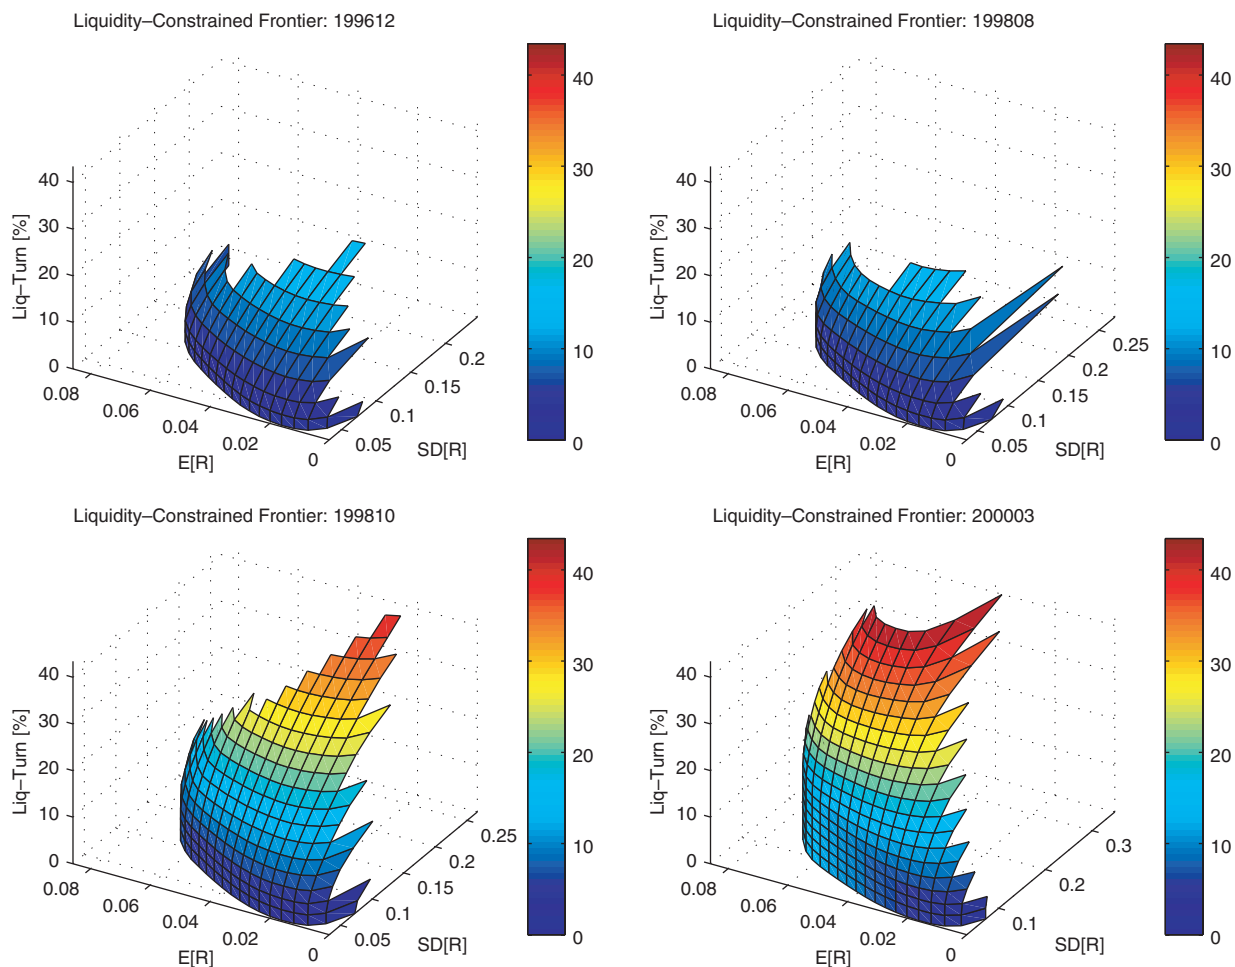


Figure 7 Liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, and March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month. Color strips to the right of each figure provide the correspondence between liquidity levels and the spectrum.

Figure 11 summarizes the trajectories of Figures 9 and 10 by plotting the Sharpe ratio as a function of the liquidity threshold for each of the months in Table 5. In contrast to the liquidity-filtered trajectories of Figure 6, the liquidity-constrained trajectories of Figure 11 are all concave, and each trajectory is comprised of three distinct segments. The first segment—beginning at the left boundary of the graph—is parallel to the liquidity axis, indicating that liquidity constraints have no effect

on the tangency portfolio's Sharpe ratio. The second segment is decreasing and concave, implying Sharpe ratios that decline at increasingly faster rates as the liquidity threshold is increased. The third segment is decreasing but linear, implying Sharpe ratios that decline with increasing liquidity thresholds, but at a constant rate.

Intuitively, an optimal MVL portfolio—one that balances all three characteristics in some

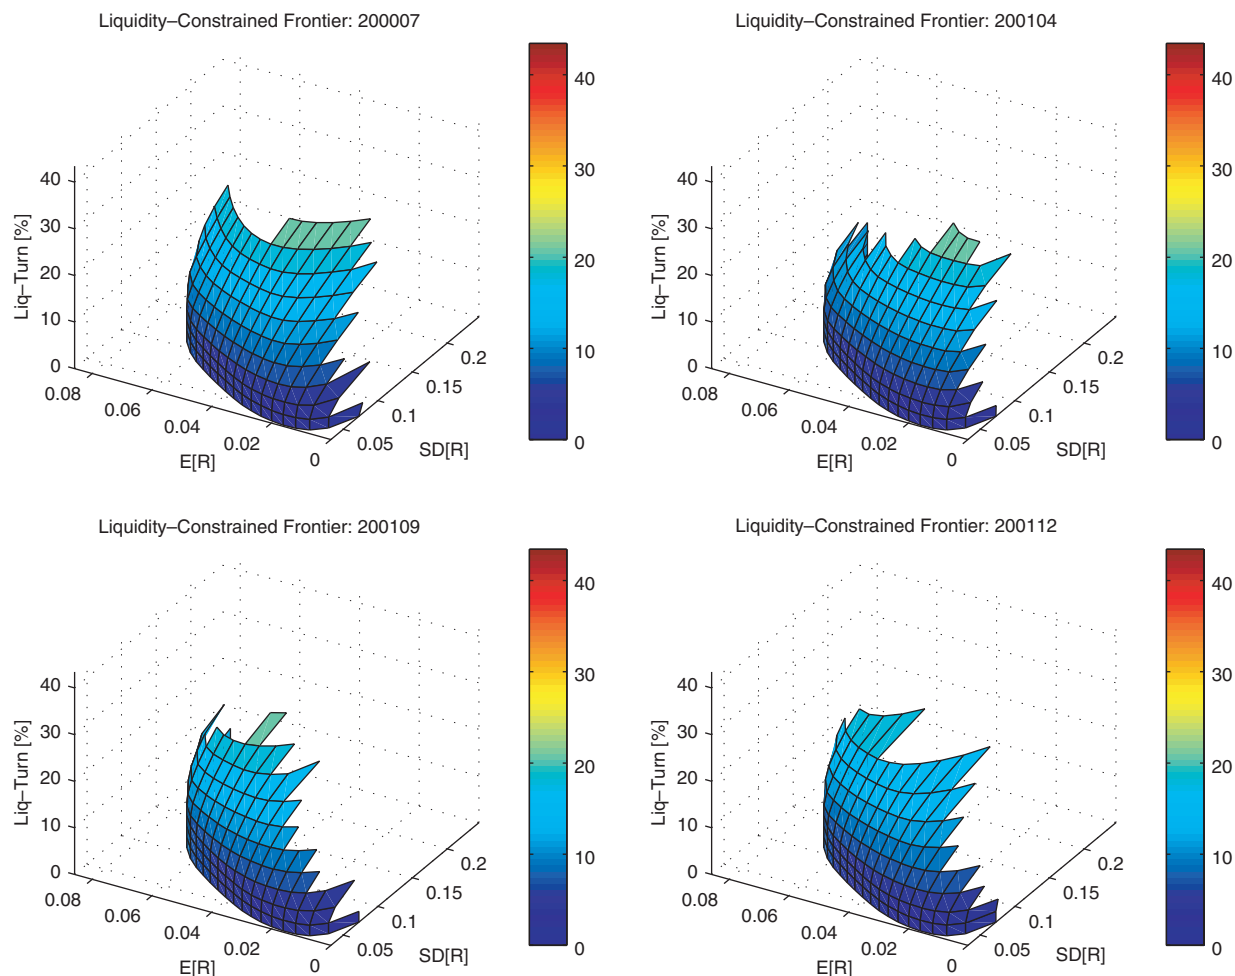


Figure 8 Liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances are estimated with daily returns data from January 2, 1997 to December 31, 2001. Color strips to the right of each figure provide the correspondence between liquidity levels and the spectrum.

fashion—should be located somewhere along the second segments of the Sharpe ratio curves in Figure 11. It is along these segments that marginal increases in the liquidity threshold yield increasingly higher costs in terms of poorer Sharpe ratios, hence there should be some liquidity threshold along this segment that balances an investor's preference for liquidity and the risk/reward profile of the tangency portfolio. Of course, turning this heuristic argument into a formal procedure for construction MVL-optimal portfolios requires the

specification of preferences for mean, variance, and liquidity, which is precisely the approach developed in Section 3.3 and implemented in Section 4.4.

4.4 The Mean-Variance-Liquidity Frontier

Although the most direct method for incorporating liquidity into the portfolio construction process is to specify an objective function that includes liquidity as in Section 3.3, this assumes that investors

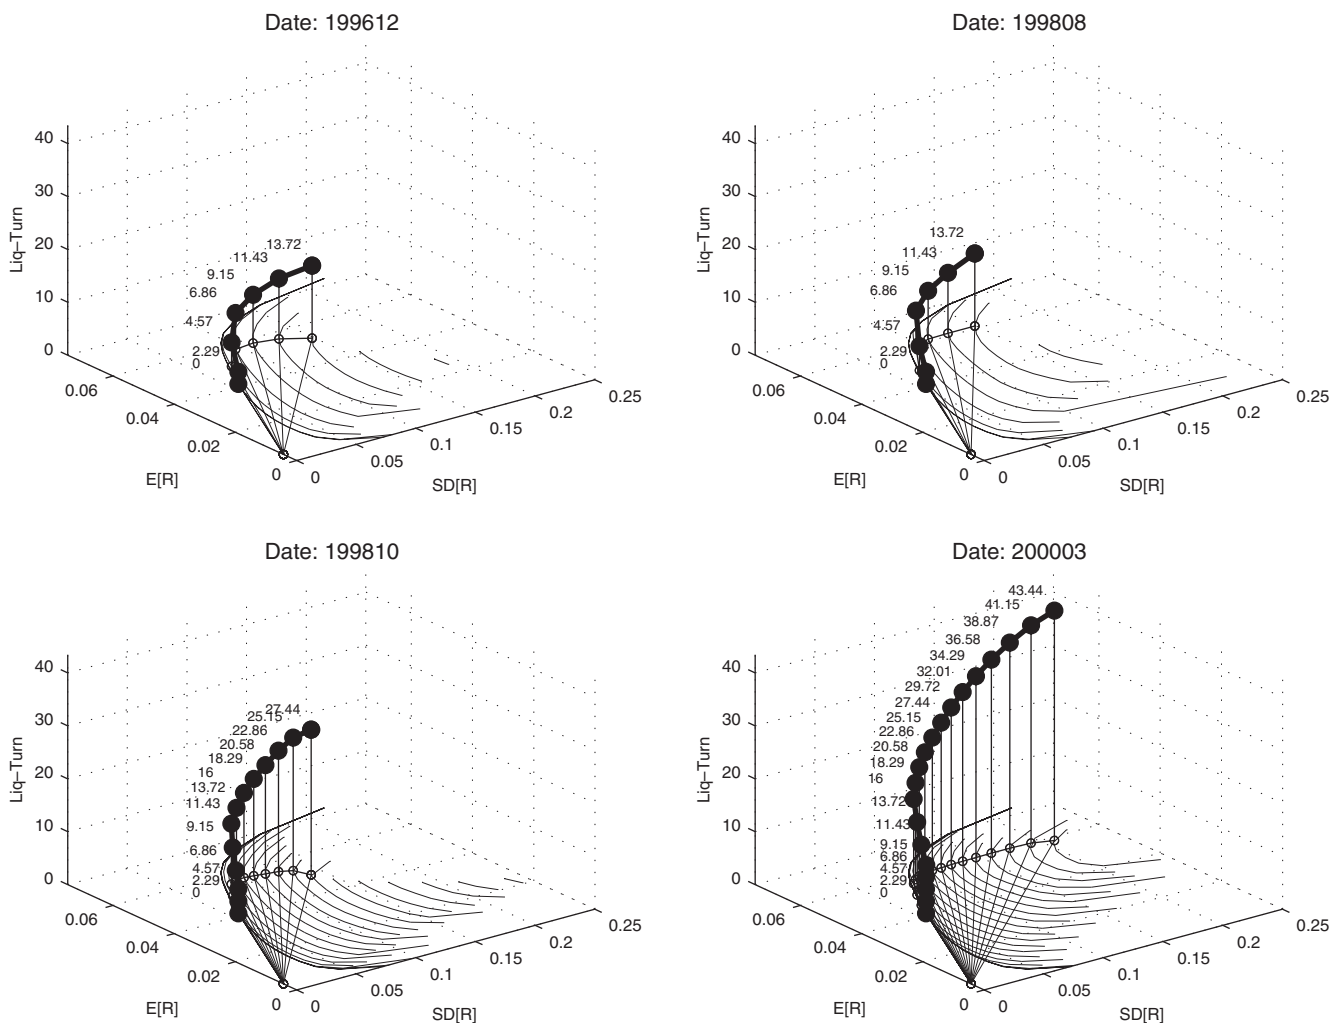


Figure 9 Trajectories of the tangency portfolio for liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of December 1996, August 1998, October 1998, and March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

are able to articulate their preferences for liquidity. This may not be true given that liquidity has only recently become an explicit factor in the investment process of many individual and institutional investors. However, by providing various calibrations of the MVL objective function (13) and their empirical implications for our sample of 50 stocks, we hope to develop a more formal understanding of liquidity preferences in the mean-variance context.

Recall from (13) of Section 3.3 that the MVL objective function is given by:

$$\begin{aligned} \max_{\{\omega\}} \quad & \omega' \mu - \frac{\lambda}{2} \omega' \Sigma \omega + \phi \omega' \ell_t \\ \text{subject to} \quad & 1 = \omega' \iota, \quad 0 \leq \omega \end{aligned}$$

where ϕ represents the weight placed on liquidity. Figure 12 contains four graphs—the expected return, standard deviation, liquidity, and Sharpe ratio of the optimal portfolio—each as a function

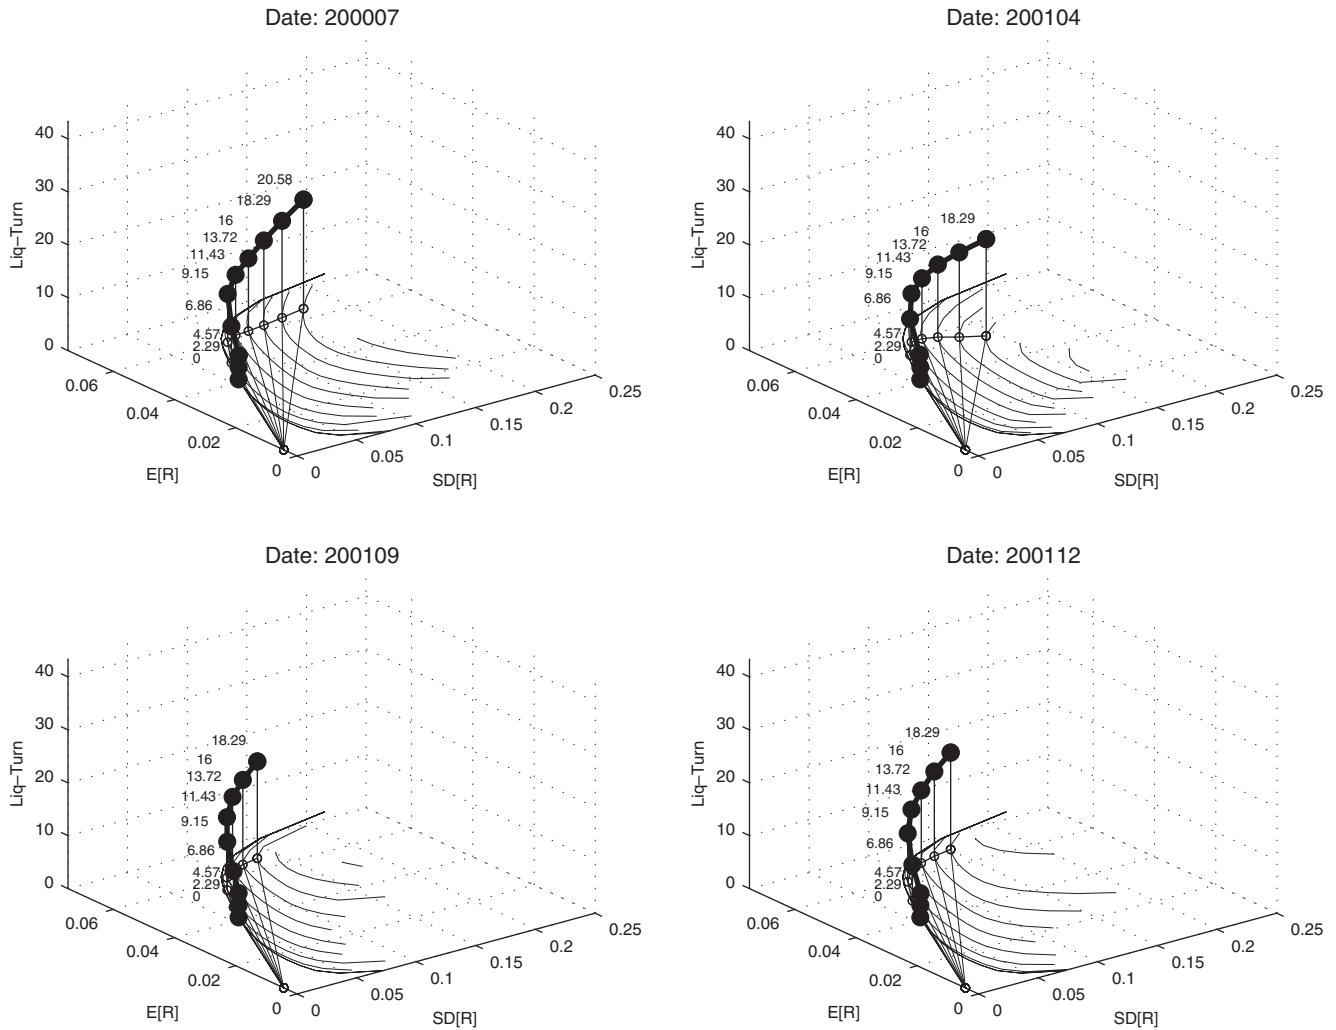


Figure 10 Trajectories of the tangency portfolio for liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric for the months of July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month.

of the risk aversion parameter λ , and for various values of the liquidity parameter ϕ where the liquidity metric used is monthly normalized turnover. When $\phi \equiv 0$, (13) reduces to the standard Markowitz–Tobin mean-variance portfolio optimization problem. As the risk aversion parameter λ increases along the horizontal axis in Figure 12, both the expected return and the standard deviation of the optimal portfolio decline as the investor places increasingly higher penalties on the portfolio’s risk. Up to $\lambda =$

10, the standard deviation declines faster than the expected return, leading to a rising Sharpe ratio curve. After reaching its peak at $\lambda = 10$, the Sharpe ratio begins to decline.

Once liquidity is allowed to enter the objective function, i.e., $\phi > 0$, the dynamics of the optimal portfolio become more complex. For expositional convenience, we focus our comments exclusively on the Sharpe ratio of the optimal portfolio. The




Figure 11 Sharpe-ratio trajectories of tangency portfolios of liquidity-constrained MVL-efficient frontiers for 50 randomly selected stocks (five from each of 10 market capitalization brackets), based on a monthly normalized turnover liquidity metric, as a function of the liquidity threshold, for the months of December 1996, August 1998, October 1998, March 2000, July 2000, April 2001, September 2001, and December 2001. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001 and do not vary from month to month. Thicker lines represent trajectories from more recent months.

interaction between the penalty for risk and the payoff for liquidity in (13) depends on the interaction between the cross-sectional distributions of liquidity and volatility in our sample. Typically, a security's liquidity metric and volatility are both correlated with market capitalization, e.g., large-cap stocks usually exhibit lower volatility and higher liquidity than smaller-cap counterparts. In this case, when a

MVL objective function is optimized, the risk and liquidity components act in the same direction—an increment in either λ or ϕ , apart from differences in scale, has the same qualitative impact on the optimal portfolio's characteristics. On the other hand, if the correlations between the liquidity metric and volatility are weak, then the interactions between the second and third terms in the objective function

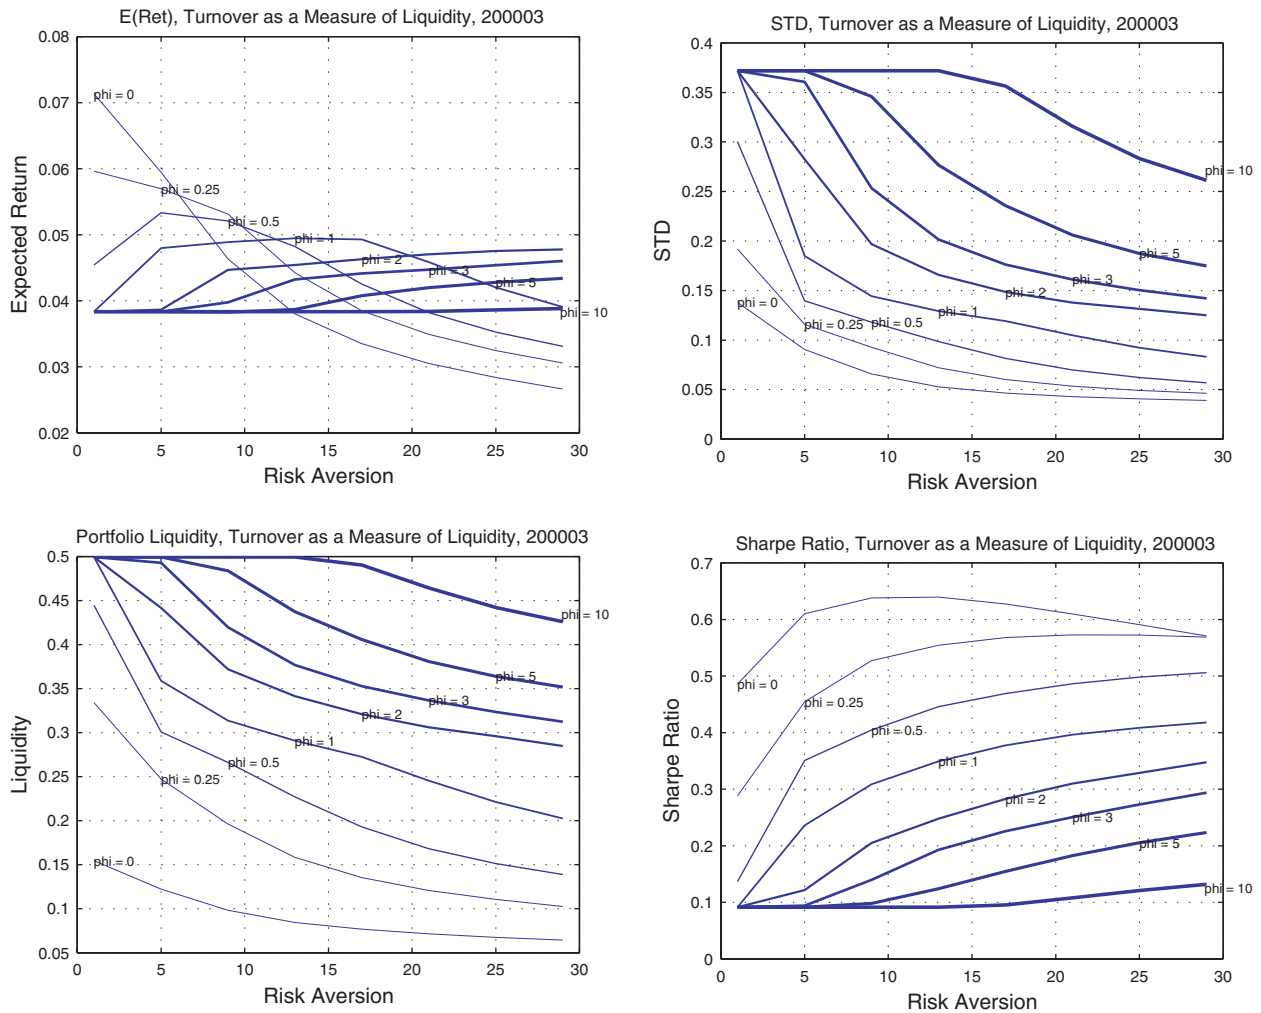


Figure 12 Properties of optimal MVL portfolios using a monthly normalized turnover liquidity metric for 50 randomly selected stocks (five from each of 10 market capitalization brackets), for the month of March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001, and “phi” denotes the liquidity parameter where a value of 0.00 implies that liquidity is not included in the portfolio optimization problem.

(13) are more complicated. Figure 13 plots daily cross-sectional correlations between raw turnover and rolling 20-day return standard deviations for the sample of 50 stocks, and with the notable exception of the year 2000, the correlation between liquidity and volatility is weak, hence there are indeed three distinct components in optimizing (13): expected return, risk, and liquidity. This is confirmed in Figure 12 for cases where $\phi > 0$. The addition of liquidity in the mean-variance objective

function results in lower Sharpe ratios for all values of λ and ϕ , and risk aversion and liquidity act as countervailing forces in the objective function.

It should be emphasized that the specific interactions between λ and ϕ are quite sensitive to the liquidity metric used. For example, Figure 14 displays the same relations as in Figure 12 but using the bid/ask spread as the liquidity metric instead of turnover. A comparison of the two figures shows




Figure 13 Daily cross-sectional correlations between raw turnover and rolling 20-day standard deviations for 50 randomly selected stocks (five from each of 10 market capitalization brackets) from January 2, 1997 to December 31, 2001.

some significant differences in the dynamics of the Sharpe ratio for the MVL-optimal portfolio. With the bid/ask liquidity metric, the tightening of both risk aversion and liquidity thresholds shifts the optimal portfolio qualitatively in the same direction—towards larger-cap, less risky stocks. An increase in the liquidity preference parameter ϕ accelerates the migration of portfolio toward larger-cap stocks driven by an increase in the risk aversion parameter, λ . This is due to the fact that during our sample period, the cross-sectional correlation between individual bid/ask spreads and volatilities is quite high and positive, and since liquidity is inversely proportional to bid/ask spread, the

correlation between liquidity and volatility in this case is strong and negative.

5 Conclusions

In this paper, we have proposed several measures of liquidity for individual securities and their portfolios, and have developed the implications of these measures for standard mean-variance portfolio analysis. In particular, there are at least three ways to integrate liquidity into the investment process once a liquidity metric is defined—through liquidity filtering, liquidity constraints, and a

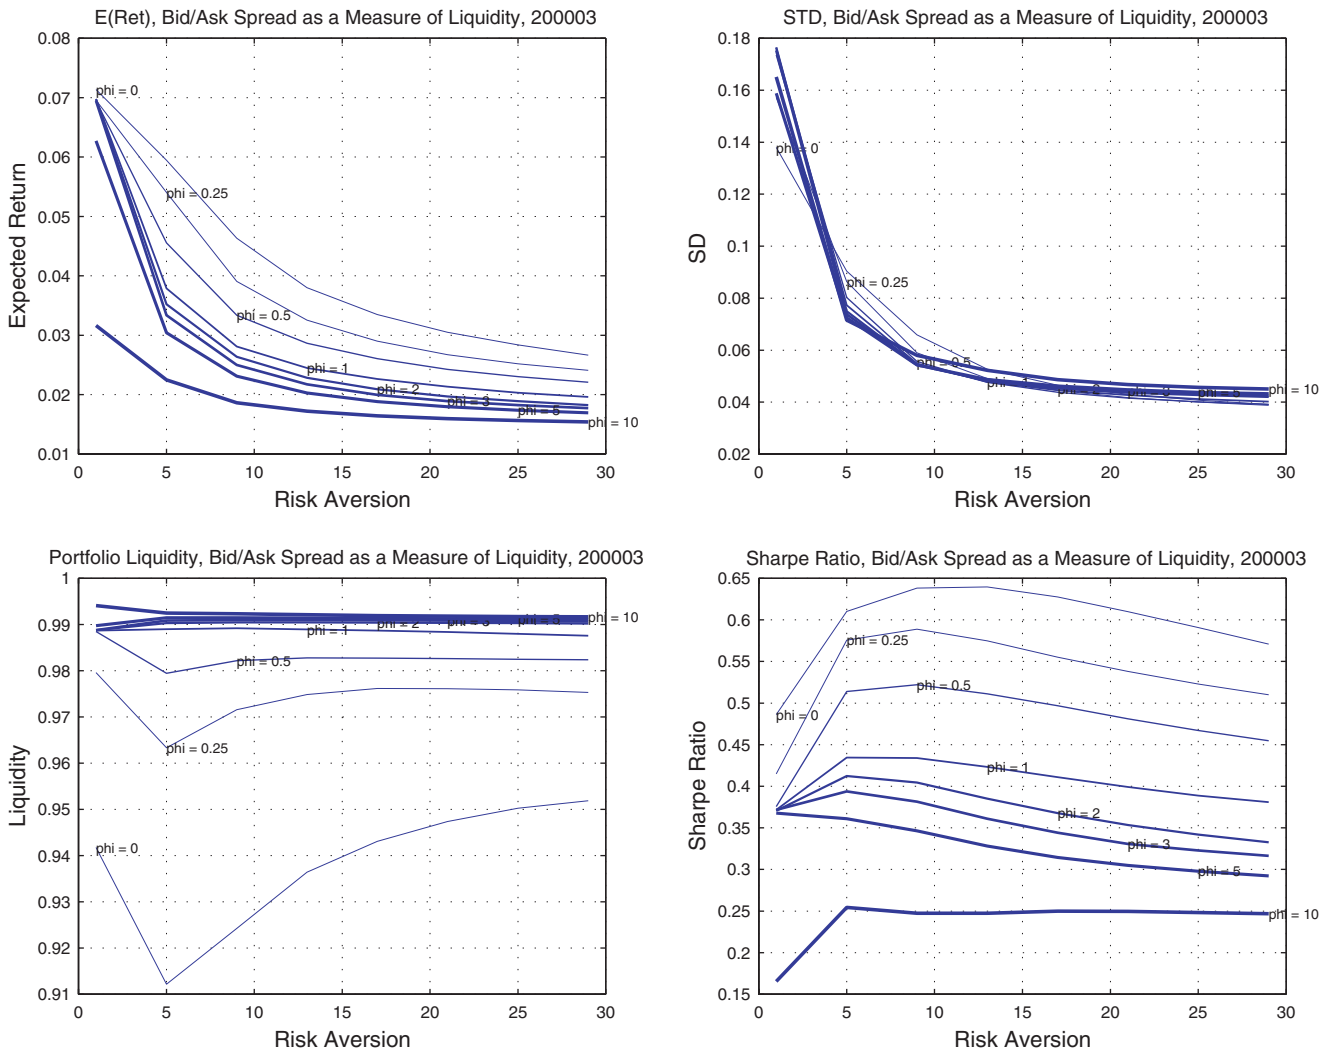


Figure 14 Properties of optimal MVL portfolios using a monthly normalized bid/ask spread liquidity metric for 50 randomly selected stocks (five from each of 10 market capitalization brackets), for the month of March 2000. Expected returns and covariances of the 50 individual securities are estimated with daily returns data from January 2, 1997 to December 31, 2001, and “phi” denotes the liquidity parameter where a value of 0.00 implies that liquidity is not included in the portfolio optimization problem.

direct optimization of a mean-variance-liquidity objective function. In a simple empirical example of 50 randomly selected securities, we have shown that in many cases, even the simplest liquidity-based portfolio optimization procedures can yield mean-variance-efficient portfolios that are considerably more liquid than their standard counterparts. More importantly, because liquidity varies so much over time, the mean-variance-liquidity

landscape is a highly dynamic surface that contains pitfalls and opportunities, and which should be managed carefully and purposefully.

Because the integration of liquidity directly into portfolio management processes has not yet become standard practice, many aspects of our analysis can be improved upon and extended. Our liquidity metrics are clearly simplistic and not based on any

equilibrium considerations, and our definition of portfolio liquidity as the weighted average of individual securities' liquidity measures may not be the best definition in all contexts. Better methods of measuring liquidity will obviously lead to better MVL portfolios.²¹ The dynamics of liquidity should also be modeled explicitly, in which case static mean-variance optimization may no longer be appropriate but should be replaced by dynamic optimization methods such as stochastic dynamic programming. Preferences for liquidity must be investigated in more detail—do such preferences exist, and if so, are they stable and how should they best be parametrized? Finally, we have ignored estimation error in the portfolio construction process, and just as sampling variation affects mean and covariance matrix estimators, liquidity estimators will also be subject to sampling variation and this may have significant impact on the empirical properties of MVL portfolios.²²

We believe we have only begun to explore the many practical implications of liquidity for investment management, and our framework adds an important new dimension—literally as well as figuratively—to the toolkit of quantitative portfolio managers. In particular, with three dimensions to consider, portfolio management can no longer operate within a purely numerical paradigm, and three- and four-dimensional visualization techniques will become increasingly central to industrial applications of portfolio optimization. We plan to explore these issues in ongoing and future research, and hope to have provided sufficient “proof-of-concept” in this paper for the benefits of incorporating liquidity into the investment process.

Appendix A

In this appendix we provide Matlab sourcecode for our extension of Loeb's (1983) price impact

function in Section A.1, and the details of our sample selection procedure in Section A.2.

A.1 Matlab Loeb Function `tloeb`

```
function tloeb
% the default value for the Loeb (1983)
spread/price cost b = 50;
% cap range
xi = [ 0.01 10 25 50 75 100 500 1000 1500 3000 ];
% block size range, in $1,000's
yi = [0.01 5 25 250 500 1000 2500 5000 10000 20000]
% original Loeb (1983) measure of liquidity
% (Table II)
zi = [
17.3 17.3 27.3 43.8 NaN NaN NaN NaN NaN NaN ;
8.9 8.9 12.0 23.8 33.4 NaN NaN NaN NaN NaN ;
5.0 5.0 7.6 18.8 25.9 30.0 NaN NaN NaN NaN ;
4.3 4.3 5.8 9.6 16.9 25.4 31.5 NaN NaN NaN ;
2.8 2.8 3.9 5.9 8.1 11.5 15.7 25.7 NaN NaN ;
1.8 1.8 2.1 3.2 4.4 5.6 7.9 11.0 16.2 NaN ;
1.9 1.9 2.0 3.1 4.0 5.6 7.7 10.4 14.3 20.0 ;
1.9 1.9 1.9 2.7 3.3 4.6 6.2 8.9 13.6 18.1 ;
1.1 1.1 1.2 1.3 1.7 2.1 2.8 4.1 5.9 8.0 ;
1.1 1.1 1.2 1.3 1.7 2.1 2.8 4.1 5.9 8.0 ];
nx = size(xi,2); ny = size(yi,2);
% array of indices of last non-NaN points in zi
matrix along mcap dimension nonnan = [ 4 4 5 6 7
8 9 ];
% deal with NaN's in zi matrix
% loop over rows
for i = 1: size(xi,2) -3
% last non-nan point
f = nonnan(i);
for j=f+1:1:ny
% Loeb cost based on simple linear extra-
% polation starting from the end points
zi(i,j) = zi(i,f)+(zi(i,f)-zi(f-1))*(yi(j)-
yi(f))/(yi(f) - y(f-1));
% cap the cost zi by b=50% if cost >50%;
if zi(i,j) > 50; zi(i,j)=b;
end;
% If trade size > 20% of market cap (not
% T. Loeb's original 5% ), zi is still NaN
if (yi(j)/1000) >0.2*xi(i); zi(i,j) = NaN;
end;
```

```

end
end
zi
% produce arrays acceptable by MATLAB for 3D
% graphics
for i=1:ny
    for j=1:nx
        x(i,j) = (xi(j));
        y(i,j) = (yi(i));
        z(i,j) = zi(j,i);
    end
end
end
% determine max-min for interpolation
maxx = max(xi); minx = min(xi); maxy = max(yi);
miny = min(yi);
% the number of nodes in each direction
N=40; dx = (maxx - minx)/N;
dy = (maxy - miny)/N;
% interpolated arrays
for i=1:N
    for j=1:N
        x1(i,j)=xi(1)+dx*j;
        y1(i,j)=yi(1)+dy*i;
    end
end
end
% plot extended Loeb function
mesh((x1), (y1), interp2(x, y, z, x1, y1,
'linear')) view(30,50); colormap(jet); grid on;
xlabel('Cap [$1,000,000]', 'FontSize', 8);
ylabel('Block [$1000]', 'FontSize', 8)
zlabel('Spread/Price Cost [%]');
% title('Loeb (1983) Total Spread/ Price Cost');
print -depsc p:\\msl\\tloeb.eps

```

A.2 Sampling Procedure

The process by which we selected our sample of 50 stocks and constructed our dataset for the empirical example consisted of the following five steps:

1. Using CRSP, we selected all ordinary common stocks having CRSP share code, SHRCD, equal to 11 or 10 for December 1996, the last pre-sample month, and for December 2001, the last in-sample month. ADRs, SBIs, units, certificates, closed-end funds and REITs were excluded. From these two sets of stocks, one for

December 1996 and one for December 2001, we selected a common subset.

2. From this common subset we selected stocks with valid daily returns which have never been delisted during the in-sample period. For each stock, we calculated the initial market capitalization as of the last trading day, December 31, 1996, of the pre-sample period.
3. We split the final subset of stocks into 10 capitalization categories, in millions US dollars (see Loeb, 1983):

0.1 10 25 50 75 100 500 1,000

1,500 3,000 \geq 3,000

The filtering is concluded by random selection of five stocks from each capitalization category.

4. For each stock in our randomly selected portfolio, we downloaded the data items listed in Table A.1 from the daily CRSP database, and calculated the daily market capitalization, in thousands of dollars, by multiplying the absolute value of price, $|PRC|$, by number of shares outstanding, SHROUT, and daily turnover, TURN, by dividing the daily trading volume, VOL, by the current number of shares outstanding, SHROUT.
5. For each randomly selected stock, using the CRSP TICKER symbol as the key, we downloaded from the TAQ database the tick-by-tick BID and ASK prices, calculated tick-by-tick bid/ask spreads, averaged the spreads for each day, and combined them with the remaining CRSP data set. The TAQ data, which are used exclusively for bid/ask spread calculations, start in January 2000, while the CRSP data start in January 1997. Missing daily bid/ask spreads in the 2000–2001 period (this problem is particularly acute for small cap stocks) were backfilled with valid ex-post values. For example, if a valid bid/ask spread at t_1 is $s(t_1)$, and the bid/ask spreads at t_2 and t_3 are missing because there

Table A.1 Data items extracted from CRSP Daily Master File.

Variable	Definition
CUSIP	CUSIP identifier
PERMNO	CRSP permanent number
PERMCO	CRSP permanent company number
TICKER	Exchange ticker symbol
COMNAM	Company name
SHRCD	Share code
SICCD	Standard industrial classification code
DATE	Trading date
BIDLO	Bid or low price
ASKHI	Ask or high price
PRC	Actual close (positive number) or the average between BIDLO and ASKHI (negative number)
VOL	Trading volume, units of one share
RET	Daily total return, including dividends
SHROUT	Number of shares outstanding, in thousands

were no quotes in the TAQ database, then we assign $s(t_2) = s(t_3) = s(t_1)$.

Notes

¹ See, for example, Acharya and Pedersen (2002), Aiyagari and Gertler (1991), Atkinson and Wilmott (1995), Amihud and Mendelson (1986b), Bertsimas and Lo (1998), Boyle and Vorst (1992), Chordia, Roll and Subrahmanyam (2000, 2001a,b, 2002), Chordia, Subrahmanyam, and Anshuman (2001), Cohen *et al.* (1981), Constantinides (1986), Davis and Norman (1991), Dumas and Luciano (1991), Epps (1976), Garman and Ohlson (1981), Gromb and Vayanos (2002), Grossman and Laroque (1990), Grossman and Vila (1992), Heaton and Lucas (1994, 1995), Hodges and Neuberger (1989), Holmstrom and Tirole (2001), Huang (2002), Litzenberger and Rolfo (1984), Leland (1985), Liu and Longstaff (2000), Lo, Mamaysky, and Wang (2001), Magill and Constantinides (1976), Morton and Pliska (1995), Pastor and Stambaugh (2002), Sadka (2003), Shleifer and Vishny (1997), Tuckman and Vila (1992), Vayanos (1998), Vayanos and Vila (1999), and Willard and Dybvig (1999).

² Of course, many studies have considered the practical significance of trading costs or “slippage” in investment management, e.g., Arnott and Wagner (1990), Bertsimas and Lo (1998), Bodurtha and Quinn (1990), Brinson, Hood,

and Beebower (1986, 1991), Chan and Lakonishok (1993, 1995), Collins and Fabozzi (1991), Cuneo and Wagner (1975), Gammill and Pérold (1989), Hasbrouck and Schwartz (1988), Keim and Madhavan (1997), Leinweber (1993, 1994), Loeb (1983), Pérold (1988), Schwartz and Whitcomb (1988), Stoll (1993), Treynor (1981), Wagner and Banks (1992), Wagner and Edwards (1993), and the papers in Sherrerd (1993). None of these studies focuses squarely on the quantitative trade-off between expected return, risk, and liquidity. However, Michaud (1989) observes that standard mean-variance portfolio optimization does not take liquidity into account, and proposes liquidity constraints and quadratic penalty functions in a mean-variance framework in Michaud (1998, Chapter 12).

³ The third dimension of liquidity—time to completion of a purchase or sale—is obviously missing from this list, but only because of lack of data. With access to time-stamped orders of a large institutional trading desk, time-based measures of liquidity can easily be constructed as well.

⁴ See, for, example, Amihud and Mendelson (1986a,b), Glosten and Milgrom (1985), Lo, Mamaysky, and Wang (2001), Tinić (1972), and Vayanos (1998).

⁵ Loeb’s original matrix does not allow for a block sizes in excess of 5% of a stock’s total market capitalization which, in our sample, would imply a maximum block size of $5\% \times \$2.84 \text{ MM} = \0.142 MM , a relatively small number. To relax this restriction, we extrapolate the total cost function to allow for block sizes of up to 20% of market capitalization, where the extrapolation is performed linearly by fixing the capitalization level and using the last two available data points along the block-size dimension. The maximum total cost is capped at 50%, an arbitrary large number. For example, for the \$0–10 MM capitalization sector (see Table II in Loeb, 1983) and block sizes of \$5,000, \$25,000 and \$250,000 the total spread/price costs are 17.3%, 27.3% and 43.8%, respectively. The cost at the next block size of \$500,000 is computed as:

$$\begin{aligned} & \min [50\%, 43.8\% + (\$500,000 - \$250,000) \\ & \quad \times (43.8\% - 27.3\%) / (\$50,000 - \$25,000)] = 50\%. \end{aligned}$$

⁶ However, see Bertsimas and Lo (1998), Chan and Lakonishok (1993, 1995), Hausman, Lo, and MacKinlay (1992), Kraus and Stoll (1972), Lillo, Farmer, and Mantegna (2003), and Loeb (1983) for various approximations in a number of contexts.

⁷ This literature is vast, and overlaps with the literature on financial asset-pricing models with transactions costs. Some of the more relevant examples include Amihud

and Mendelson (1986b), Bagehot (1971), Constantinides (1986), Demsetz (1968), Gromb and Vayanos (2002), Lo, Mamaysky and Wang (2001), Tiniç (1972), Vayanos (1998), and Vayanos and Vila (1999). For a more complete list of citations, see the references contained in Lo, Mamaysky and Wang (2001).

⁸ For expositional convenience, all of our tables and graphs use standard deviations in place of variances as risk measures. Nevertheless, we shall continue to refer to graphs of efficient frontiers as “mean-variance-liquidity efficient frontiers” despite the fact that standard deviation is the x -axis, not variance. We follow this convention because the objective function on which our efficient frontiers are based are mean-variance objective functions, and because “mean-standard deviation-liquidity” is simply too cumbersome a phrase to use more than once.

⁹ See, for example, Michaud (1998, Chapter 12).

¹⁰ For comparison, Table 1 also reports market capitalizations based on December 31, 2001 prices. From December 31, 1996 to December 31, 2001, the average portfolio market capitalization increased twofold, with mid-tier market-capitalization stocks—those in the 5th, 6th and 7th brackets—experiencing the biggest gains. The market capitalization of the top-tier stocks increased less dramatically. By the end of the sample, the original capitalization-based ranking was generally well preserved—the correlation between the 1996 and 2001 year-end market capitalizations was over 95%.

¹¹ Since 1,256 observations were used to calculate the correlation coefficients, the 95% confidence interval under the null hypothesis of zero correlation is $[-5.6\%, 5.6\%]$.

¹² For this 2-year sample, the 95% confidence interval under the null hypothesis of zero correlation is $[-8.9\%, 8.9\%]$.

¹³ Results for the Loeb and bid/ask metrics are qualitatively identical to those for turnover, hence we omit them to conserve space. However, they are available upon request.

¹⁴ Throughout this study, we assume a fixed value of 0.4308% per month for the riskless return R_f .

¹⁵ These values may seem rather high, especially in the context of current market conditions. There are two explanations: (a) our sample period includes the tail end of the remarkable bull market of the 1990s, and contains some fairly spectacular high-flyers such as North Coast Energy (571% 5-year return from 1996 to 2001), Daktronics (914% 5-year return), and Green Mountain Coffee (875% 5-year return); (b) we are using a relatively small sample of 50 stocks, which is considerably less well-diversified than other well-known portfolios such as the S&P 500 or the Russell 2000, and the lack of diversification will tend to yield higher expected returns (especially given

the small-cap component in our portfolio) and higher standard deviations.

¹⁶ Recall that the only difference between the December 1996 and March 2000 portfolio inputs is the liquidity metrics for each stock; the estimated means and covariance matrix are the same for both months, i.e., the values obtained by applying (14) to the entire sample of daily returns from January 2, 1997 to December 31, 2001.

¹⁷ Within each liquidity plane (planes that are parallel to ground level), portfolios to the north have higher expected return, and portfolios to the east have higher standard deviation.

¹⁸ We refrain from computing MVL frontiers when the number of securities falls below 5.

¹⁹ Recall that each of the liquidity metrics has been normalized to take on values strictly between 0 and 1, hence liquidity thresholds are comparable across metrics and are denominated in units of percent of the range of the original liquidity measure.

²⁰ In fact, this observation suggests that the Loeb function—as well as any other realistic measure of price impact—varies with market conditions, and such dependencies should be incorporated directly into the specification of the price impact function, i.e., through the inclusion of “state variables” that capture the salient features of the market environment at the time of the transactions. See Bertsimas and Lo (1998) and Bertsimas, Hummel, and Lo (2000) for examples of such specifications.

²¹ See, for example, Chordia, Roll, and Subrahmanyam (2000, 2001, 2002), Getmansky, Lo, and Makarov (2003), Glosten and Harris (1988), Lillo, Farmer, and Mantegna (2003), Lo, Mamaysky, and Wang (2001), Pastor and Stambaugh (2002), and Sadka (2003) for alternate measures of liquidity.

²² See, for example, Jobson and Korkie (1980, 1981), Klein and Bawa (1976, 1977), and Michaud (1998).

References

- Acharya, V., and Pedersen, L. (2002). “Asset Pricing with Liquidity Risk.” Unpublished working paper, London Business School.
- Aiyagari, R., and Gertler, M. (1991). “Asset Returns with Transaction Costs and Uninsured Individual Risk.” *Journal of Monetary Economics*, 27, 311–331.
- Amihud, Y., and Mendelson, H. (1986a). “Asset Pricing and the Bid-Asked Spread.” *Journal of Financial Economics* 17, 223–249.
- Amihud, Y., and Mendelson, H. (1986b). “Liquidity And Stock Returns.” *Financial Analysts Journal* 42, 43–48.

- Arnott, R., and Wagner, W. (1990). "The Measurement And Control of Trading Costs." *Financial Analysts Journal* 46, 73–80.
- Atkinson, C., and Wilmott, P. (1995). "Portfolio Management with Transaction Costs: An Asymptotic Analysis of the Morton and Pliska Model." *Mathematical Finance*, 357–367.
- Bagehot, W. (a.k.a. Jack Treynor), (1971). "The Only Game in Town." *Financial Analysts Journal* 22, 12–14.
- Bertsimas, D., and Lo, A. (1998). "Optimal Control of Execution Costs." *Journal of Financial Markets* 1, 1–50.
- Bertsimas, D., Hummel, P., and Lo, A. (2000). "Optimal Control of Execution Costs for Portfolios." *Computing in Science & Engineering* 1, 40–53.
- Bodurtha, S., and Quinn, T. (1990). "Does Patient Program Trading Really Pay?" *Financial Analysts Journal* 46, 35–42.
- Brinson, G., Hood, R., and Beebower, G. (1986). "Determinants of Portfolio Performance." *Financial Analysts Journal* 42, 39–44.
- Brinson, G., Singer, B., and Beebower, G. (1991). "Determinants of Portfolio Performance II: An Update." *Financial Analysts Journal* 47, 40–48.
- Chan, L., and Lakonishok, J. (1993). "Institutional Trades and Intra-Day Stock Price Behavior." *Journal of Financial Economics* 33, 173–199.
- Chan, L., and Lakonishok, J. (1995). "The Behavior of Stock Prices Around Institutional Trades." *Journal of Finance* 50, 1147–74.
- Chordia, T., Roll, R., and Subrahmanyam, A. (2000). "Commonality in Liquidity." *Journal of Financial Economics* 56, 3–28.
- Chordia, T., Roll, R., and Subrahmanyam, A. (2001). "Market Liquidity and Trading Activity Source." *Journal of Finance* 56, 501–530.
- Chordia, T., Roll, R., and Subrahmanyam, A. (2002). "Order Imbalance, Liquidity, and Market Returns." *Journal of Financial Economics* 65, 111–130.
- Chordia, T., Subrahmanyam, A., and Anshuman, V. (2001). "Trading Activity and Expected Stock Returns." *Journal of Financial Economics* 59, 3–32.
- Cohen, K., Maier, S., Schwartz, R., and Whitcomb, D. (1981). "Transaction Costs, Order Placement Strategy and Existence of the Bid-Ask Spread." *Journal of Political Economy* 89, 287–305.
- Collins, B., and Fabozzi, F. (1991). "A Methodology for Measuring Transaction Costs." *Financial Analysts Journal* 47, 27–36.
- Constantinides, G. (1986). "Capital Market Equilibrium with Transaction Costs." *Journal of Political Economy* 94, 842–862.
- Cunéo, L., and Wagner, W. (1975). "Reducing the Cost of Stock Trading." *Financial Analysts Journal* 26, 35–44.
- Davis, M., and Norman, A. (1990). "Portfolio Selection with Transactions Costs." *Mathematics of Operations Research* 15, 676–713.
- Demsetz, H. (1968). "The Cost of Transacting." *Quarterly Journal of Economics* 82, 35–53.
- Dumas, B., and Luciano, E. (1991). "An Exact Solution to a Dynamic Portfolio Choice Problem under Transactions Costs." *Journal of Finance* 46, 577–595.
- Epps, T. (1976). "The Demand For Brokers' Services: The Relation Between Security Trading Volume And Transaction Cost." *Bell Journal of Economics* 7, 163–196.
- Gammill, J., and Perold, A. (1989). "The Changing Character Of Stock Market Liquidity." *Journal of Portfolio Management* 15, 13–18.
- Garman, M., and Ohlson, J. (1981). "Valuation Of Risky Assets In Arbitrage-Free Economies With Transactions Costs." *Journal of Financial Economics* 9, 271–280.
- Getmansky, M., Lo, A., and Makarov, I. (2003). "Econometric Models of Serial Correlation and Illiquidity in Hedge Fund Returns." unpublished working paper, MIT Sloan School of Management.
- Glosten, L., and Harris, L. (1988). "Estimating the Components of the Bid/Ask Spread." *Journal of Financial Economics* 21, 123–142.
- Glosten, L., and Milgrom, P. (1985). "Bid, Ask, and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders." *Journal of Financial Economics* 13, 71–100.
- Gromb, D., and Vayanos, D. (2002). "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs." *Journal of Financial Economics* 66, 361–407.
- Grossman, S., and Laroque, G. (1990). "Asset Pricing and Optimal Portfolio Choice in the Presence of Illiquid Durable Consumption Goods." *Econometrica* 58, 25–52.
- Hasbrouck, J., and Schwartz, R. (1988). "Liquidity and Execution Costs in Equity Markets." *Journal of Portfolio Management* 14, 10–16.
- Hausman, J., Lo, A., and MacKinlay, C. (1992). "An Ordered Probit Analysis of Transaction Stock Prices." *Journal of Financial Economics* 31, 319–379.
- Heaton, J., and Lucas, D. (1996). "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing." *Journal of Political Economy* 104, 443–487.
- Holmstrom, B., and Tirole, J. (2001). "LAPM: A Liquidity-Based Asset Pricing Model." *Journal of Finance* 57, 1837–1867.
- Huang, M. (2002). "Liquidity Shocks and Equilibrium Liquidity Premia." to appear in *Journal of Economic Theory*.

- Jobson, J., and Korkie, R. (1980). "Estimation for Markowitz Efficient Portfolios." *Journal of the American Statistical Association* 75, 544–554.
- Jobson, J., and Korkie, R. (1981). "Performance Hypothesis Testing with the Sharpe and Treynor Measures." *Journal of Finance* 36, 889–908.
- Keim, D., and Madhavan, A. (1997). "Transactions Costs and Investment Style: An Inter-Exchange Analysis of Institutional Equity Trades." *Journal of Financial Economics* 46, 265–292.
- Klein, R., and Bawa, V. (1976). "The Effect of Estimation Risk on Optimal Portfolio Choice." *Journal of Financial Economics* 3, 215–231.
- Klein, R., and Bawa, V. (1977). "The Effect of Limited Information and Estimation Risk on Optimal Portfolio Diversification." *Journal of Financial Economics* 5, 89–111.
- Kraus, A., and Stoll, H. (1972). "Price Impacts of Block Trading on the New York Stock Exchange." *Journal of Finance* 27, 569–588.
- Leinweber, D. (1993). "Using Information From Trading in Trading and Portfolio Management." In K. Sherrerd, ed.: *Execution Techniques, True Trading Costs, and the Microstructure of Markets*. Charlottesville, VA: Association for Investment Management and Research.
- Leinweber, D. (1994). "Careful Structuring Reins In Transaction Costs." *Pensions and Investments* July 25, 19.
- Lillo, F., Farmer, D., and Mantegna, R. (2003). "Master Curve for Price-Impact Function." *Nature* 421, 129–130.
- Liu, J., and Longstaff, F. (2000). "Losing Money on Arbitrages: Optimal Dynamic Portfolio Choice in Markets with Arbitrage Opportunities." Unpublished working paper, Anderson Graduate School of Management, UCLA.
- Lo, A., Mamaysky, H., and Wang, J. (2001). "Asset Prices and Trading Volume Under Fixed Transactions Costs." NBER Working Paper No. W8311.
- Lo, A., and Wang, J. (2000). "Trading Volume: Definitions, Data Analysis, and Implications of Portfolio Theory." *Review of Financial Studies* 13, 257–300.
- Loeb, T. (1983). "Trading Cost: The Critical Link Between Investment Information and Results." *Financial Analysts Journal* 39, 39–44.
- Michaud, R. (1989). "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* 45, 31–42.
- Michaud, R. (1998). *Efficient Asset Management: A Practical Guide to Stock Portfolio Optimization and Asset Allocation*. Boston, MA: Harvard Business School Press.
- Morton, A., and Pliska, S. (1995). "Optimal Portfolio Management with Fixed Transaction Costs." *Mathematical Finance* 5, 337–356.
- Pastor, L., and Stambaugh, R. (2002). "Liquidity Risk and Expected Stock Returns." To appear in *Journal of Political Economy*.
- Pérol, A. (1988). "The Implementation Shortfall: Paper Versus Reality." *Journal of Portfolio Management* 14, 4–9.
- Sadka, R. (2003). "Momentum, Liquidity Risk, and Limits to Arbitrage." Unpublished working paper, Kellogg Graduate School of Management, Northwestern University.
- Schwartz, R., and Whitcomb, D. (1988). "Transaction Costs and Institutional Investor Trading Strategies." *Monograph Series in Finance and Economics 1988–2/3*. New York: Salomon Brothers Center for the Study of Financial Institutions, New York University.
- Sherrerd, K., ed. (1993). *Execution Techniques, True Trading Costs, and the Microstructure of Markets*. Charlottesville, VA: Association for Investment Management and Research.
- Shleifer, A., and Vishny, R. (1997). "The Limits of Arbitrage." *Journal of Finance* 52, 35–55.
- Stoll, H. (1993). *Equity Trading Costs*. Charlottesville, VA: Association for Investment Management and Research.
- Tinic, S., (1972). "The Economics of Liquidity Services." *Quarterly Journal of Economics* 86, 79–93.
- Tuckman, B., and Vila, J. (1992). "Arbitrage With Holding Costs: A Utility-Based Approach." *Journal of Finance* 47, 1283–1302.
- Vayanos, D. (1998). "Transaction Costs and Asset Prices: A Dynamic Equilibrium Model." *Review of Financial Studies* 11, 1–58.
- Vayanos, D., and Vila, J. (1999). "Equilibrium Interest Rate and Liquidity Premium With Transaction Costs." *Econometric Theory* 13, 509–539.
- Wagner, W. (1993). "Defining and Measuring Trading Costs." In K. Sherrerd, ed.: *Execution Techniques, True Trading Costs, and the Microstructure of Markets*. Charlottesville, VA: Association for Investment Management and Research.
- Wagner, W., and Banks, M. (1992). "Increasing Portfolio Effectiveness Via Transaction Cost Management." *Journal of Portfolio Management* 19, 6–11.
- Wagner, W., and Edwards, M. (1993). "Best Execution." *Financial Analysts Journal* 49, 65–71.
- Willard, G., and Dybvig, P. (1999). "Empty Promises and Arbitrage." *Review of Financial Studies* 12, 807–834.