

Trading Costs and Returns for US Equities: Estimating Effective Costs from Daily Data

Joel Hasbrouck

Department of Finance
Stern School of Business
New York University
44 West 4th St. Suite 9-190
New York, NJ 10012-1126
212.998.0310

jhasbrou@stern.nyu.edu

This draft: August 12, 2006

Preliminary draft
Comments welcome

For comments on an earlier draft, I am grateful to Yakov Amihud, Lubos Pastor, Jay Shanken, Bill Schwert, Kumar Venkataraman and seminar participants at the University of Rochester, the NBER Microstructure Research Group, the Federal Reserve Bank of New York, Yale University, the University of Maryland, the University of Utah, Emory University and Southern Methodist University. All errors are my own responsibility.

The latest version of this paper and a SAS dataset containing the long-run Gibbs sampler estimates are available on my web site at www.stern.nyu.edu/~jhasbrou.

Trading Costs and Returns for US Equities: Estimating Effective Costs from Daily Data

Abstract

The effective cost of trading is usually estimated from transaction-level trade and quote data. This study proposes a Gibbs estimate that is based on daily closing prices. In a broad sample of US firms over a period when both high-frequency TAQ and daily CRSP data are available (1993-2005), an annual Gibbs estimate based on daily data achieves a correlation of 0.965 with the TAQ value. The approach is extended to a panel specification in which effective costs for individual stocks are driven by a latent common factor. In the comparison sample, the estimated series for the common factor based on daily data achieves a correlation of 0.447 with the corresponding TAQ value at a daily frequency (0.670 at a monthly frequency). The firm-specific factor loadings estimated from daily data are also positively correlated with the loadings estimated from transactions data.

The Gibbs estimates are employed in asset pricing specifications over a longer historical sample (1927-2005). The results offer only weak support for the view that effective cost (as a characteristic) affects expected stock returns, except when interacted with a January seasonal dummy variable. An asset's return covariance with the common factor of effective cost is not found to be a determinant of expected returns. The difference between these results and those of analyses based on other liquidity proxies indirectly suggests the importance of trading volume. The latter quantity is used in most daily liquidity proxies, but does not enter the effective cost estimates constructed here.

JEL classification codes: C15, G12, G20

1. Introduction and summary of results

Investigations into the role of liquidity and transaction costs in asset pricing must generally confront the fact that while many asset pricing tests make use of US equity returns from 1926 onwards, the high-frequency data used to estimate trading costs are usually not available prior to 1983. Accordingly, most studies either limit the sample to the 1983-present period of common coverage or use the longer historical sample with liquidity proxies estimated from daily data. The present paper falls into the latter group. It proposes a new approach to estimate of the effective cost of trading and the common variation in this cost. These estimates are then used in conventional asset pricing specifications with a view to ascertaining the role of trading cost as a characteristic and as a risk factor in explaining expected returns.¹

For the purchase of a security, the effective cost is the execution price less the midpoint of the prevailing bid and ask quotes (and vice versa for a sale). Although it does not resolve the trade-related temporary and permanent (price impact) components of the price change, it is simple to compute (from detailed trade and quote records), easy to interpret, and is widely used as a measure of market quality.² Since 2000 the SEC has required US market centers to report

¹ Recent asset pricing analyses that cover a sample where high frequency data are available include Brennan and Subrahmanyam (1996), Easley, Hvidkjaer and O'Hara (2002), Sadka (2004), Korajczyk and Sadka (2006). Analyses that use proxies based on daily data include Amihud (2002), Pastor and Stambaugh (2003), Acharya and Pedersen (2005), and Spiegel and Wang (2005). Closing daily or annual bid-ask quotes are sometimes available over samples longer than those of the high-frequency data. Studies that use closing spreads include Stoll and Whalley (1983), Amihud and Mendelson (1986), Eleswarapu and Reinganum (1993), Eleswarapu and Reinganum (1993), Reinganum (1990). Chalmers and Kadlec (1998). Easley and O'Hara (2002) survey the area.

² Lee (1993) is representative of early work. Stoll (2006) is a recent example.

summary statistics of their effective costs based on the orders they actually receive and execute (Reg NMS Rule 605, formerly designated Rule 11ac1-5, often called “dash five”).

To estimate the effective cost from daily closing prices, the study starts with the Roll (1984) model of price dynamics. Hasbrouck (2004) suggests a Bayesian Gibbs approach to this model, and applies it to futures transaction data. The present study generalizes the model and applies it to daily CRSP US equity data. The paper develops two models. The basic market-adjusted model generates annual estimates of effective cost at the firm level. The latent common factor version allows for time variation in effective cost with commonality across firms. The latter model generates estimates of the common factor at a daily frequency and, for each firm, an annual estimate of the loading on this factor.

All of the CRSP/Gibbs estimates are compared to corresponding values of effective cost level and variation computed using high-frequency (TAQ) data. This comparison sample spans 1993-2005, and comprises roughly 300 firms per year (approximately 3,900 firm-years). In the comparison sample, the CRSP/Gibbs estimate of average effective cost achieves a correlation of 0.965 with the TAQ value. The CRSP/Gibbs estimate of the common effective cost factor is highly correlated with the cross-firm average of the TAQ values (0.447 at the daily frequency; 0.585, weekly; 0.670, monthly). The CRSP/Gibbs estimates of the firm-level loadings on this factor are moderately correlated (0.328) with the corresponding TAQ-based estimates. Overall, subject to some qualifications discussed in the body of the paper, these findings suggest that the daily Gibbs estimates are strong proxies for the high-frequency measures. The estimates are extended to span the full CRSP daily sample (1926 to present).

The paper then turns to applications of these proxies in asset pricing specifications. The earlier papers in this area view liquidity as a characteristic that drives a wedge between the returns an investor might realize net of trading costs and the gross returns that form the basis for most asset pricing tests (Amihud and Mendelson (1986)). Later analyses emphasize the effects of liquidity variation, both in time and cross-section. Pastor and Stambaugh (2003) note that a trading cost that covaries positively with an asset’s (gross) return increases the risk in the net

return. This effect is magnified if a low (gross) return increases the likelihood of a forced liquidation, and consequent involuntary realization of the trading cost. This argues in favor of treating market-wide liquidity as a risk-factor. Pastor and Stambaugh find that the covariance between an asset's return and the common liquidity factor is priced.

Acharya and Pedersen (2005) analyze an overlapping generations model populated with myopic (one-period) agents. In their model the effects of liquidity on expected returns appear in diverse terms including a characteristic effect (as in the earlier papers), covariance between gross return and the liquidity factor (as in Pastor and Stambaugh), and a term in the covariance between an asset's liquidity and the common liquidity factor. Intuitively, the last term captures the extent to which the trading cost can be diversified in a portfolio. They find support for these effects using the Amihud (2002) illiquidity measure. Korajczyk and Sadka (2006) also find liquidity to be a risk factor, using high-frequency measures constructed over 1983-2000.

Relative to these studies the results for expected returns and effective costs obtained in the present paper, however, are less supportive of liquidity and liquidity-risk effects. In diverse samples across listing venues and time, effective cost as a characteristic tends to be a positive (but not uniformly significant) determinant of expected returns. The relationship is nevertheless marked by strong seasonality. The impact of effective cost in January is uniformly large and significantly positive. This confirms in a larger and broader sample, the seasonality results reported by Eleswarapu and Reinganum (1993).

The analysis finds no support for the role of market-wide effective cost as a risk factor. In estimations of return-generating processes, the addition of this component (or its innovation) to a specification that includes the three Fama-French factors results in a negligible improvement in explanatory power. In expected return specifications, the coefficients of effective-cost betas are close to zero. When the beta of asset's effective cost relative to the market-wide effective cost is added as a characteristic (as in the Acharya-Pedersen framework), the estimated effect on expected return tends to be positive but insignificant.

The present results on the role of liquidity variation are therefore somewhat weaker than the findings of other studies. While a full discussion of the sources of this divergence is deferred until later in the paper, one consideration in particular stands out. The effective cost estimate used in the present paper is a narrow measure of trading cost and is based solely on closing prices. The daily-based liquidity proxies used in all of the other studies incorporate trading volume. The overall results therefore may simply attest to the importance of trading volume, and a definition of liquidity that is broad enough to encompass the dimensions of trading activity that volume reflects. Trading costs per se may be unimportant, while liquidity (in the larger sense of economic distortions attributed to the trading process) may still be relevant.

The paper is organized as follows. Section 2 describes the specification and computational procedures used the basic market-adjusted model, in which the effective cost is assumed static. Section 3 discusses the latent common-factor model, which allows for dynamic common variation in effective cost. Data sources and sample construction are discussed in Section 4. Results for the basic market-adjusted and latent common factor models are presented in Sections 5 and 6. Section 7 applies the effective cost estimates in asset pricing specifications. A discussion of the results in Section 8 concludes the paper.

2. Bayesian estimation of effective cost

a. The Roll model

Roll (1984) suggests a simple model of security prices in a market with transaction costs. The specifications estimated in this paper are variants of the Roll model, but the basic version is useful for describing the estimation procedure. The price dynamics may be stated as:

$$\begin{aligned} m_t &= m_{t-1} + u_t \\ p_t &= m_t + c q_t \end{aligned} \quad (1)$$

where m_t is the log quote midpoint prevailing prior to the t th trade (“efficient price”), p_t is the log trade price. The q_t are direction indicators, which take on the values +1 (for a buy) or –1 (for a sale) with equal probability. The disturbance u_t reflects public information and is assumed to be

uncorrelated with q_t . Roll motivates c as one-half the posted bid-ask spread, but since the model applies to transaction prices, it is natural to view c as the effective cost. The model has essentially the same form under time aggregation. In particular, although the model is sometimes estimated with transaction data (e.g., Schultz (2000)), it was originally applied to daily data, with q_t being interpreted as the direction variable for the last trade of the day.

The Roll model implies

$$\Delta p_t = m_t + c q_t - (m_{t-1} + c q_{t-1}) = c \Delta q_t + u_t, \quad (2)$$

from which it follows that $c = \sqrt{-Cov(\Delta p_t, \Delta p_{t-1})}$, where $Cov(\Delta p_t, \Delta p_{t-1})$ is the first-order autocovariance of the price changes. The usual estimate of c is the sample analog of this, termed here the “moment estimate” because it uses a sample moment (the sample autocovariance) in lieu of the population value, and to distinguish it from the Gibbs estimate.

The moment estimate is feasible, however, only if the first-order sample autocovariance is negative. In samples of daily frequency this is often not the case. In annual samples of daily returns, Roll found positive autocovariances in roughly half the cases. Harris (1990) discusses this and other aspects of this estimator. His results show that positive autocovariances are more likely for low values of the spread. Accordingly, one simple remedy is to assign an a priori value of zero. Another problem arises when there is no trade on a particular day, in which case CRSP reports the midpoint of the closing bid and ask. If these days are retained in the sample, the estimated cost will generally be biased downwards, because the midpoint realizations do not include the cost. If these days are dropped from the sample, there may be heteroscedasticity since the efficient price innovations may span multiple days.

b. Bayesian estimation using the Gibbs sampler

Hasbrouck (2004) advocates a Bayesian approach. Completing the Bayesian specification requires specification of the distribution of u_t . I assume here that $u_t \stackrel{d}{\sim} i.i.d.N(0, \sigma_u^2)$. The prior distributions for parameters are discussed below.

The data sample is denoted $p \equiv \{p_1, p_2, \dots, p_T\}$. The unknowns comprise both the model parameters $\{c, \sigma_u^2\}$ and the latent data, the trade direction indicators $q \equiv \{q_1, \dots, q_T\}$. (Knowing p and q suffices to determine $m \equiv \{m_1, \dots, m_T\}$.) The parameter posterior density $f(c, \sigma_u | p)$ is not obtained in closed-form, but is instead characterized by a random draws (from which means and other summary statistics may be computed). The random draws are generated using a Gibbs sampler whereby each unknown is drawn in turn from its full conditional (posterior) distribution. First, c and q are initialized to arbitrary feasible values. Then c is drawn from $f(c | \sigma_u^2, q, p)$; σ_u^2 is drawn from $f(\sigma_u^2 | c, q, p)$; q_1 is drawn from $f(q_1 | c, \sigma_u^2, q_2, q_3, \dots, q_T, p)$, and so on.

The draws are described in more detail below, but one central feature of the model warrants emphasis. In the expression for Δp_t given by equation (2), if the Δq_t are known (or taken as given), the equation describes a simple linear regression, with coefficient c . The linear regression perspective is a dominant theme of the present analysis. It simplifies implementation using standard results from Bayesian statistics, and suggests ways in which the model may be generalized.

Simulating the coefficient(s) in a linear regression coefficients

The standard Bayesian normal regression model is $y = Xb + e$ where y is a column vector of n observations of the dependent variable, X is an $(n \times k)$ matrix of fixed regressors, β is a vector of coefficients and the residuals are zero-mean multivariate normal $e \sim N(0, \Omega_e)$. Given Ω_e and a normal prior on β , $\beta \sim N(\mu_\beta, \Omega_\beta)$, the posterior is $\beta \sim N(\mu_\beta^*, \Omega_\beta^*)$ where $\mu_\beta^* = (X' \Omega_e^{-1} X + \Omega_\beta^{-1})^{-1} (X' \Omega_e^{-1} y + \Omega_\beta^{-1} \mu_\beta)$ and $\Omega_\beta^* = (X' \Omega_e^{-1} X + \Omega_\beta^{-1})^{-1}$. Carlin and Louis (2000), Lancaster (2004), and Geweke (2005) are contemporary textbook treatments.

In the present applications it is often necessary to impose inequality restrictions on the β . Typically, one or more coefficients is restricted to the positive domain. It is straightforward to show that when the β prior is restricted as $\underline{\beta} < \beta < \bar{\beta}$, the posterior has the same parameters as in the unrestricted case, but truncated to the same interval as the prior (see, for example, Geweke (2005), section 5.3.1). Hajivassiliou, McFadden and Ruud (1996) discuss computationally efficient procedures for making random draws from truncated multivariate normal distributions.

Simulating the error covariance matrix

The usual conjugate prior for Ω_e in the general normal case is the inverted Wishart (see Carlin and Taylor, p. 146). All of the applications in the present analysis, however, involve diagonal (but not necessarily scalar) Ω_e , i.e., where the diagonals are of the form $\sigma_i^2 I$ for $i = 1, \dots, m$ (finite). In this simpler case, it is convenient to handle the σ_i^2 independently. Assume that a set of regression errors e_i for $i = 1, \dots, n$ is i.i.d. $N(0, \sigma^2)$. If the parameter prior is $\sigma^2 \sim IG(\alpha, \beta)$ where IG denotes the inverted gamma distribution, then the posterior is $\sigma^2 \sim IG(\alpha^*, \beta^*)$ where $\alpha^* = \alpha + n/2$ and $\beta^* = \beta + \sum e_i^2/2$, (Kim and Nelson (2000), p. 176).

Simulating the trade direction indicators

The remaining step in the sampler involves drawing $q \equiv \{q_1, \dots, q_T\}$ when c and σ_u^2 are known. The procedure is sequential. The first draw is $q_1 | q_2, \dots, q_T$, the second draw is $q_2 | q_1, q_3, \dots, q_T$, the third draw is $q_3 | q_1, q_2, q_4, \dots, q_T$, etc., where the “|” notation denotes the conditional draw. The full set of conditioning information includes the price changes $\Delta p \equiv \{\Delta p_2, \dots, \Delta p_T\}$ and the parameters c and σ_u^2 .

The first realization of u_i to enter the observed prices is u_2 . This may be written as a function of q_1 as $u_2(q_1) = \Delta p_2 - cq_2 + cq_1$ (given q_2 , etc.). A priori, $u_2 \sim N(0, \sigma_u^2)$ and $q_1 = \pm 1$ with equal probability. The posterior odds ratio of a buy vs. a sell is

$$\frac{\Pr(q_1 = +1 | q_2, \dots)}{\Pr(q_1 = -1 | q_2, \dots)} = \frac{f(u_2(q_1 = +1))}{f(u_2(q_1 = -1))}$$

Where f is the normal density function with mean zero and variance σ_u^2 . The right hand side of this is easily computed, and q_1 is drawn using the implied (Bernoulli) probability.

To draw q_2 , note that given everything else, we may write $u_2(q_2) = \Delta p_2 - cq_2 + cq_1$ and $u_3(q_2) = \Delta p_3 - cq_3 + cq_2$. Given the assumed serial independence of the u_i , the posterior odds ratio is

$$\frac{\Pr(q_2 = +1 | q_1, q_3, \dots)}{\Pr(q_2 = -1 | q_1, q_3, \dots)} = \frac{f(u_2(q_2 = +1)) f(u_3(q_2 = +1))}{f(u_2(q_2 = -1)) f(u_3(q_2 = -1))}$$

Again, we compute the right-hand side and make the draw. In this fashion, we progress through the remaining q_t . For all draws of q_t (except the first and last) the posterior odds ratio involves only the adjacent disturbances u_t and u_{t+1} . The posterior odds ratio for the last draw is

$$\frac{\Pr(q_T = +1 | q_1, \dots, q_{T-1})}{\Pr(q_T = -1 | q_1, \dots, q_{T-1})} = \frac{f(u_T(q_T = +1))}{f(u_T(q_T = -1))}$$

In some samples, for a subset of times, the trade directions may be known. These q_t may simply be left at their known values. A related situation arises from the CRSP convention of reporting quote midpoints on days with no trades. For these days we fix $q_t = 0$, implying that $p_t = m_t$, i.e. that the quote midpoint is observed without error. This may be formally justified by positing a more general model that admits the possibility of no trade. If the no-trade probability is denoted π , for example, the general model would allow q_t to take on values 0, +1, and -1 with probabilities π , $(1-\pi)/2$, and $(1-\pi)/2$. Assuming that the no-trade days are known, that buys and sells are equally likely given a trade occurrence, and that we do not wish to estimate or characterize π , however, the more general model is observationally equivalent to the simpler one.

Another sort of observational equivalence is slightly more troublesome. It is natural to assume that trading costs are (at least on average) non-negative, i.e., $c > 0$. This is an economic assumption, however. From a statistical viewpoint, the model is observationally equivalent to one in which $c < 0$ and all trade directions have the opposite signs (“buys” have $q_t = -1$, etc.). Simulated posteriors for c are in consequence bimodal, symmetric about the origin. To rule out this “mirror” situation, it is convenient to impose the restriction $c > 0$ on the prior.

Bayesian analyses sometimes use improper priors, often with the purpose of establishing an explicit connection to classical frequentist approaches. For example, letting Ω_β^{-1} approach zero (in some norm) in the Bayesian regression model discussed above leads to posterior estimates that converge to the usual frequentist ones. The present situation does not, however, admit this device. The regressors in equation (2) are the Δq_t , which are simulated. If the q_t drawn in one iteration (sweep) of the sampler all happen to have the same sign, then all of the $\Delta q_t = 0$, and the computed regression is uninformative (for this sweep). In this case, a draw must be made

from the prior distribution. Although this is an infrequent occurrence, it effectively rules out a prior for c that is proper but extremely diffuse.

c. The basic market-adjusted model and sampler specification

The models estimated in this paper generalize on the basic Roll model in various respects. The discussion now turns to the first of two models actually estimated. It is straightforward to add other regressors to equation (2). The motivation for doing so is that, intuitively, the procedure tries to allocate transaction price changes between “true” (efficient price) returns and transient trading costs. Anything that helps explain either component will sharpen the resolution. Return factors are obvious candidates for supplemental regressors. The basic market-adjusted model is:

$$\Delta p_t = c\Delta q_t + \beta r_{mt} + u_t, \quad (3)$$

where r_{mt} is the excess market return on day t . It is assumed that the market return is independent of Δq_t . This would be the case if the trade direction indicators for the component securities are mutually independent, implying a diversification of bid-ask bounce. Note that although the present goal is improved estimation of c , it is likely that estimation of β will also be enhanced.

In the present applications (all involving US equity data), the prior for c is the normal density with mean parameter equal to zero and variance parameter equal to 0.05^2 restricted to nonnegative values, denoted $N^+(\mu = 0, \sigma^2 = 0.05^2)$. The μ and σ^2 appearing here are only formal parameters: the actual mean and variance of the distribution will differ due to the truncation. The prior for β is $N(\mu = 1, \sigma^2 = 1)$; that for σ_u^2 is inverted gamma, $IG(\alpha = 1 \times 10^{-12}, \beta = 1 \times 10^{-12})$.

The sampler then follows the following program:

- Step 0 (initializations). Although the limiting behavior of the sampler is invariant to starting values, “reasonable” initial guesses may hasten convergence. The q_t that do not correspond to midpoint reports are set to the sign of the most recent price change, with q_1 set (arbitrarily) to +1. σ_u^2 is initially set to 0.0004 (roughly corresponding to a 30%

annual idiosyncratic volatility). No initial values are required for c and β , as they are drawn first.

- Step 1. Based on the most recently simulated values for σ_u^2 and the q_t , compute the posterior for the regression coefficients (c and β) and make a new draw.
- Step 2. Given c, β and the q_t , compute the implied u_t , update the posterior for σ_u^2 and make a new draw.
- Step 3. Given c, β and σ_u^2 , make draws for q_1, q_2, \dots, q_T . Go to step 1.

To ease the computational burden, each sampler is run for only 1,000 sweeps. Although this value is small by the standards of most MC analyses, it appears to be sufficient in the present case. Experimentation with up to 10,000 sweeps did not materially affect the mean parameter estimates. Of the 1,000 draws for each parameter, the first 200 are discarded to “burn in” the sampler, i.e., remove the effect of starting values. The average of the remaining 800 draws (in principle the posterior mean) is used as a point estimate of the parameter in subsequent analysis

d. An illustration

The essential properties of the estimator may be illustrated by considering two simulated price paths. The paths correspond to situations typical of US equities. Both paths are of length 250 (roughly a year of daily observations). The standard deviation for the efficient price innovation is $\sigma_u = 0.02$ (i.e., about two percent, corresponding to an annual standard deviation of about thirty-two percent). For simplicity, $\beta = 0$. One simulated series of u_t and one simulated series of q_t are used for both paths. The price paths are identical except for the scaling of the effective cost: c is either set to 0.01 or 0.10, and $\beta = 0$. For each path the Gibbs sampler is run for 10,000 sweeps, with the first 2,000 discarded. The remaining 8,000 draws are used to characterize the posteriors.

Figure 1 illustrates the simulated 90% confidence regions for the parameter posteriors. Panel A depicts the posterior when $c = 0.01$; Panel B, when $c = 0.10$. To facilitate comparisons, the horizontal axes (σ_u) are identical. The vertical axes (c) are shifted, but have the same scale.

The results are striking. In Panel A ($c = 0.01$), the joint confidence region is large and negatively sloped. In Panel B ($c = 0.10$), the confidence region is circular, centered around the population values, and compact.

To develop the intuition for this result, recall that the Gibbs procedure generates conditional random draws for the trade direction indicators. These draws characterize the posteriors for the trade direction indicators, and the sharpness of these posteriors corresponds very closely to what one might guess on the basis of looking at the price paths. When c is large relative to the efficient price increments, the price path appears distinctly “spikey” (with many reversals), as a consequence of the large bid-ask bounce. It is easy to confidently identify buys and sells, and the parameter posterior is concentrated. When c is small, however, the reversals are less distinct. It is less certain whether a given trade is a buy or sell. The allocation of the price change between the transient (bid-ask) component and the permanent change in the security value is less clear. This naturally leads to greater uncertainty (less concentration) and the negative correlation (downward slope) implied by the posterior in Panel A.

This illustration has implications for studies of US equities. The posted half-spread in a large, actively traded issue might be roughly one penny on a share price of \$50, implying $c = 0.0002$. No approach using daily trade data is likely to achieve a precise estimate of such a magnitude. The posted half-spread for a thinly traded issue might be twenty-five cents on a five-dollar stock, implying $c = 0.05$. This is likely to be estimated much more precisely.

3. Estimating variation in effective cost

The estimation of a liquidity measure is rarely an end in itself. One usually seeks to explain liquidity variation in the cross-section (across firms) or over time, or to relate this variation to other quantities of economic interest. This section describes a general approach to modeling liquidity variation and the specific model used to characterize latent commonality.

To assess variation in any liquidity estimate, the simplest strategy is to partition the sample across firms and/or time periods, form estimates over the subsamples, and use the subsample values in further analysis. Cross-sectional variation in liquidity for US equities, for

example, might be analyzed by computing an estimate for each firm using a year of daily data and then regressing these estimates against capitalization, etc. Time variation is often characterized using estimates formed over shorter intervals, typically one month (as in, e.g., Pastor and Stambaugh (2003) or Acharya and Pedersen (2005)).

This two-step approach is most attractive when the subsamples are large enough that the estimation errors in the liquidity measures are small relative to the between-subsample variation of interest. The analysis of the estimates in the second step may involve additional procedures to minimize the effects of these errors, such as forming portfolio averages. The first estimation, though, is simply the procedure discussed in the last section, and needs no further elaboration.

An alternative approach is to model the liquidity variation directly within the price change specification. This general technique is widely used in other financial econometrics contexts. In asset pricing applications, for example, time variation in betas and risk-premia is commonly modeled by placing parametric functions, typically linear projections on conditioning variables, directly in the return specifications (see Jagannathan, Skoulakis and Wang (2006) and references therein). The remainder of this section develops this one-step approach for modeling effective costs.

a. The case of observed liquidity determinants

The approach follows from the interpretation of the price-change equation as a linear regression specification. By using linear projections to model cost variation, estimation can proceed by repeated applications of the Bayesian regression model. The price change equation in all cases may be written as

$$\Delta p_{it} = c_{it}q_{it} - c_{i,t-1}q_{i,t-1} + \beta_i r_{mt} + u_{it} \text{ for } i = 1, \dots, N \text{ firms and } t = 2, \dots, T \quad (4)$$

Here, c_{it} denotes the cost for firm i at time t . It is assumed that the q_{it} and u_{it} are independent across firms. All commonality in efficient price movements is driven by the market factor.

To modify the Gibbs sampler developed for the basic market-adjusted model, note that at the point where we need to simulate the q_{it} , the values of c_{it} will be known (taken as given).

Thus, these draws may be accomplished with a straightforward modification of the procedure described in section 2.b: it suffices to replace all terms involving cq_t with $c_{it}q_{it}$.

We now turn to the specification of c_{it} . Let $c_{it} = Z_{it}\gamma_i$ where Z_{it} is a set of known conditioning variables and γ_i is a firm-specific coefficient vector. Candidate variables might include forecast volatility, market capitalization, earnings surprises, and/or dummy variables for splits, changes in regulation, etc. With this functional form, the price change may be written

$$\Delta p_{it} = (q_{it}Z_{it} - q_{i,t-1}Z_{i,t-1})\gamma_i + \beta_i r_{mt} + u_{it}, \quad (5)$$

Thus, given all other variables and parameters, γ_i and β_i are regression coefficients. The draws may be accomplished by applying the Bayesian normal regression model. Since the u_{it} are assumed independent across firms, the computation may be performed separately for each firm.

It will generally be necessary, however, to impose some restrictions in order to insure the non-negativity of the c_{it} . One might, for example, transform the conditioning variables so that they are non-negative, and impose $\gamma_i > 0$ in the coefficient prior.

b. The latent common factor (LCF) model

Bayesian Gibbs sampling approaches have been applied to multivariate models involving latent factors (Geweke and Zhou (1996) present a treatment of the APT, for example.) Thus, with sufficient additional structure, it is not even essential that the conditioning variables be observable. A major goal of the present study is characterization of common variation in effective cost. This is accomplished with the model:

$$c_{it} = \gamma_{0i} + \gamma_{1i}z_t \quad (6)$$

where z_t is an unobserved factor common to the effective costs of all firms. Putting this into the price change equation and rearranging yields:

$$\Delta p_{it} - (q_{it} - q_{i,t-1})\gamma_{0i} - \beta_i r_{mt} = (\gamma_{1i}q_{it})z_t - (\gamma_{1i}q_{i,t-1})z_{t-1} + u_{it} \quad (7)$$

Written in this form, the z_t are coefficients in a panel regression involving N firms and $T-1$ price changes:

$$y = X \begin{bmatrix} z_1 \\ z_2 \\ \vdots \\ z_T \end{bmatrix} + U \quad (8)$$

where

$$y_{N(T-1) \times 1} = \begin{bmatrix} \vdots \\ \Delta p_{i2} - (q_{i2} - q_{i1})\gamma_{0i} - \beta_i r_{m2} \\ \vdots \\ \Delta p_{iT} - (q_{iT} - q_{i,T-1})\gamma_{0i} - \beta_i r_{mT} \\ \vdots \end{bmatrix} \quad X_{N(T-1) \times T} = \gamma_{li} \begin{bmatrix} \vdots \\ -q_{i1} & q_{i2} & 0 \\ 0 & -q_{i2} & q_{i3} & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & -q_{i,T-1} & q_{i,T} \\ \vdots \end{bmatrix}$$

$$U_{N(T-1) \times 1} = \begin{bmatrix} \vdots \\ u_{i2} \\ \vdots \\ u_{iT} \\ \vdots \end{bmatrix} \quad Cov(U)_{N(T-1) \times N(T-1)} = \begin{bmatrix} \sigma_{u1}^2 I_{T-1} & & \\ & \ddots & \\ & & \sigma_{uN}^2 I_{T-1} \end{bmatrix}$$

It is common in factor analysis to hypothesize that the unobserved factor is a standard normal variate. Normalizing the mean and variance to zero and unity fixes a scaling of the factor loadings (γ_{0j} and γ_{1j}) that would otherwise be indeterminate. As the present application also requires non-negativity, the prior is that the z_t are identically and independently distributed as $N^+(0,1)$ variates.

As noted, the specification is essentially a panel regression, the form of which fits within the Bayesian regression framework summarized earlier. This panel regression is included as an additional step in a sweep of the sampler. The prior for γ_{0i} is $N^+(\mu = 0, \sigma^2 = 0.05^2)$; and for γ_{1i} , $N^+(\mu = 0, \sigma^2 = 0.02^2)$. In starting up the sampler, the z_t are initialized to random draws from $N^+(0,1)$.

At first glance the LCF model might seem to make impossible demands on the data. The price change specification, for example, contains terms such as $\gamma_{li} z_t q_{it}$, that are the product of three unobserved quantities. The best evidence in support of the procedure is the comparison of

daily-based and high-frequency estimates presented in Section 6.a. It is useful at this point, though, to note certain structural features of the model that facilitate identification.

Firstly, the model attributes all variation and commonality in effective costs to z_t . There is no idiosyncratic variation in the effective costs, nor is there commonality in the trade directions. Return commonality, of course, is still allowed via the βr_{mt} term in the specification. Secondly, the distributional assumptions are strong, even by the usual standards of Bayesian analysis. The present analysis not only assumes pervasive normality, but also nonnegativity. This assumption, when invoked for all the determinants of c_{it} , ensures the nonnegativity of c_{it} itself. This in turn helps resolve the reversal components of price changes, implicitly identifying q_{it} . The common factor z_t is essentially identified by the cross-sectional features of the price changes (and the normalization). In the panel regression, a realization of z_t enters $2N$ price changes (the price change at times t and $t+1$ for each of the N securities). Alternatively, viewed from the perspective of inference, $2N$ price changes are contributing to the estimation of each z_t .

c. Extensions

The procedure used here involves little more than repeated application of the standard Bayesian normal regression model. The approach can be applied to any liquidity measure that is obtained as a regression coefficient. The Pastor-Stambaugh gamma measure is such a quantity. The Amihud illiquidity measure is defined as the average of a ratio, but one might construct a similar quantity as the coefficient in a regression of $|r_t|$ against dollar volume. The Amivest liquidity ratio might be modified likewise. (Note, however, that in the present application, the distributions of key latent variables, q_{it} and z_t in particular, are concentrated. Sample distributions of trading volumes are very diffuse. This may create convergence problems in modeling their regression coefficients.)

4. Data and implementation

a. Sample construction

Most of the Gibbs estimates in the paper are computed in annual samples of daily data. These data are taken from the 1926-2005 CRSP daily dataset, restricted to ordinary common shares (CRSP share code 10 or 11) that had a valid price for the last trading day of the year, and had no changes of listing venue or large splits within the last three months of the year. For purposes of assessing the performance of the Gibbs estimates, the analysis uses TAQ data produced by the NYSE for the period covering 1993-2005. The asset pricing tests also using the Fama-French return factors (downloaded from Ken French's web site).

In consideration of computational limits described more fully below, the full latent common factor model is estimated in each year (1926-2005) only for a random sample of 150 firms (300, after 1985) that possessed a full data record for that year (and had no splits or changes in listing venue during the year). For day t , the average (across draws) of z_t is taken as a point estimate of the effective cost factor on that day. These estimates are then used as fixed regressors in estimating the LCF model for the remaining CRSP firms. Broad CRSP coverage of Nasdaq stocks starts in 1985. In this and subsequent years, the sample of firms used to estimate the full latent common factor model consists of 150 listed (NYSE/Amex firms) and 150 Nasdaq firms, randomly selected from a sample stratified by market capitalization. Prior to 1985, the sample is limited to 150 listed firms.

The 300 firms/year in 1993-2005 are also used as the basis for the comparison sample. Liquidity measures for these firms were estimated from the TAQ dataset. These 3,900 CRSP firm-years were matched to TAQ subject to the criteria of: agreement of ticker symbol; uniqueness of ticker symbol; the correlation (over the year) between the TAQ and CRSP closing prices had to be above 0.9; and, on fewer than 2% of the days did TAQ report trades when CRSP did not (or vice versa). Subject to these criteria, 3,777 firms were matched between TAQ and CRSP. Summary statistics for the comparison sample.

Gibbs estimates (indeed, all Markov chain Monte Carlo procedures) tend to be computationally intensive. For a sample of N firms over T days, each sweep for the basic market-adjusted model requires N ordinary least squares time-series regression over the firm's return series (N regressions of size T). Each sweep of the latent-common factor model, however, also requires a generalized least-squares panel regression with NT observations. Additional effort for the latent common factor model also arises from unbalanced data. The basic model can be estimated separately for each firm. If the data record for a given firm only covers a portion of what is generally available for other firms, the price-change regression is simply computed using a shorter sample. Computation of the panel GLS regression, however, requires construction of large matrices that are correctly aligned with respect to firm and time. The computational time and programming overhead necessary to accommodate firms with incomplete records was substantial. These considerations motivated the use of restricted samples described above.

b. TAQ liquidity measures

In the comparison sample, the effective cost for firm i on day t is computed as a trade-weighted average for all trades relative to the prevailing quote midpoint. Similar results were obtained using unweighted averages.³ In principle the effective cost measures the cost of an order executed as a single trade. When the order is executed in multiple trades, the price impact of a trade also contributes to the execution cost. For each firm in the comparison sample, a representative price impact coefficient is estimated as the λ_i coefficient in:

$$\Delta p_{it} = \lambda_i \left(\text{Signed} \sqrt{\text{Dollar Volume}} \right)_{it} + \varepsilon_{it}. \quad (9)$$

³ The prevailing quote is assumed to be the most recent quote posted two seconds or more prior to the trade. This is within the “1 to 2 seconds” rule that Piwowar and Wei (2006) find optimal for their 1999 sample, but it is likely that reporting conventions have changed over the sample used here.

The specification was estimated using price changes and signed volume aggregated over five-minute intervals. A separate estimate was computed for each month. Reported summary statistics are based on the average of the monthly values. Variants of specification (9) were used, with qualitatively similar results.

c. CRSP liquidity measures

The study considers various alternative daily liquidity proxies. The simplest is the moment estimate of the effective cost based on the traditional Roll model, that is $\sqrt{-Cov(\Delta p_{i,t}, \Delta p_{i,t-1})}$. When the autocovariance is positive, the moment estimate is set to zero. (This occurs for roughly one-third of the firm-years in the comparison sample.) The statistics reported in the paper use only those days on which trading occurred, but similar results are obtained when all prices (including non-trade days) are used.

In addition, the analysis includes the proportion of days with no price changes relative to the previous close (Lesmond, Ogden and Trzcinka (1999)) and the Amihud (2002) illiquidity measure $(I = |return|/|Dollar\ volume|)$. The study does not include Pastor and Stambaugh (2003) gamma measure because the authors caution against its use as a liquidity measure for individual stocks, noting the large sampling error in the individual estimates (p. 679).

5. Results for the basic market-adjusted model

a. Comparison sample

Table 1 presents summary statistics for the TAQ and CRSP liquidity variables. Since the effective costs are logarithmic, the means correspond to effective costs of about one percent. Proportion of zero returns is restricted to the unit interval by construction. At its median value, the TAQ-based price impact coefficient λ implies that a \$10,000 buy order would move the log price by $\sqrt{10,000} \times 7 \times 10^{-6} = 0.0007$, i.e., seven basis points. The median value for the illiquidity ratio suggests that \$10,000 of daily volume would move the price by $10,000 \times 0.07 \times 10^{-6} = 0.0007$, as well. The summary statistics of both the CRSP moment and

Gibbs estimates of effective costs are close to the TAQ values. All liquidity measures exhibit extreme values; the coefficients of skewness and kurtosis are large, particularly for the illiquidity measure.

The discussion now focuses more closely on effective costs. Figure 2 presents annual box-and-whisker plots the TAQ and CRSP/Gibbs estimates. There are several notable features of the TAQ values. First, the distributions do not appear stationary. Although the fifth percentile (indicated by the lower limit of the whisker) is relatively stable, the ninety-fifth percentile (upper limit of the whisker) has declined from about 0.05 in 1993 to 0.02 in 2005. The median (marked by the horizontal line in the box) has declined from roughly 0.01 in 1993 to 0.004 in 2005. This decline may reflect changes in trading technology and regulation, but it may also arise from changes in the composition of the sample.

The second important feature is that cross-sectional variation generally appears to be much larger than the aggregate time series variation. The smallest range between the fifth and ninety-fifth percentiles is about 0.01 (in 2005), and for most the sample the range is at least 0.02. This dominates the roughly 0.006 decline in the median over the period. This suggests that tests of liquidity effects are likely to be more powerful if they are based on cross-sectional variation.

The general features of the CRSP/Gibbs distributions closely match those of the TAQ. To more rigorously assess the quality of the CRSP/Gibbs estimates and other liquidity proxies, Table 2 presents the correlation coefficients. The standard (Pearson) correlation between the TAQ and CRSP/Gibbs estimate of effective cost is 0.965.⁴ The Spearman correlation, a more appropriate measure if the proxy is being used to rank liquidity, is 0.872. Because liquidity proxies are often used in specifications with explanatory variables that are themselves likely to

⁴ This and other reported correlations are computed as a single estimate, pooled over years and firms. The values are very similar, though, to the averages of annual cross-sectional correlations. Over the 13-year sample, the lowest estimated correlation between the CRSP/Gibbs estimate and the TAQ value was 0.903 (in 2005).

be correlated with liquidity, the table also presents partial correlations that control for logarithm of end-of-year share price and logarithm of market capitalization (Pearson: 0.943; Spearman: 0.678). Not only are the CRSP/Gibbs estimates strong proxies in the sense of correlation, but they are also good point estimates of the TAQ measures. A regression of the latter against the former would ideally have unit slope and zero intercept. In the comparison sample, the estimated regression is $c_i^{TAQ} = 0.001 + 0.935 c_i^{CRSP/Gibbs} + e_i$. Finally, by any of the four types of correlation considered here, the conventional moment estimate of effective cost is dominated by the CRSP/Gibbs estimator.

The table also reports correlations for the alternative TAQ and CRSP liquidity measures. The two TAQ-based liquidity measures (effective cost and price impact coefficient) are moderately positively correlated (0.513, Pearson). This is qualitatively similar to the findings of Korajczyk and Sadka (2006). Among the daily proxies, the Amihud illiquidity measure is most strongly correlated with the TAQ-based price impact coefficient, with the CRSP/Gibbs effective cost estimate being second.

b. Historical estimates of effective cost, 1926-2005

The basic market-adjusted model is estimated annually for all ordinary common shares in the CRSP daily data base. Figure 3 graphs effective costs, separately for NYSE/Amex (listed) and Nasdaq, averaged over market capitalization quartiles.

Effective costs for NYSE/Amex issues (upper graph) exhibit considerable variation over time. The highest values are found immediately after the 1929 crash and during the Depression. It is likely that this reflects historic lows for per-share prices coupled with a tick size that remained at one-eighth of a dollar, which together imply an elevated proportional cost. Subsequent peaks in effective cost generally also coincide with local minima of per share prices. After the Depression, however, average effective costs don't rise above one percent for the three highest capitalization quartiles. The largest variation is confined to the bottom capitalization quartile.

The Nasdaq estimates (lower graph) begin in 1985. As for the listed sample, the largest variation arises in the lowest capitalization quartile. The temporal variation, however, may also reflect changes in sample composition. In the early 1990s, Nasdaq listed firms that were especially young and volatile (Fama and French (2004); Fink, Fink, Grullon and Weston (2006)).

6. Results for the latent common factor (LCF) model

a. Comparison sample

Analysis of the comparison sample is aimed at investigating the correlation of CRSP/Gibbs estimates with TAQ values. In the LCF model interest centers on the estimated latent liquidity factor z_t and the factor loadings (γ_{0i} and γ_{1i}) in equ. (6). We first consider the factor itself. To facilitate comparison, note that equ. (6) averaged over all firms in the year yields $\bar{c}_t = \bar{\gamma}_0 + \bar{\gamma}_1 z_t$, where the bars indicate cross-firm averages. In principle, therefore, z_t should be perfectly correlated with the cross-firm average effective cost.

The factor z_t is estimated at a daily frequency, and the TAQ average effective cost may be computed at a daily frequency as well. It is helpful, though, to begin with a graphic presentation of the weekly averages (Figure 4). To remove the long-term time trend (previously discussed in connection with Figure 2), and enhance comparability, both series are standardized within each year to have zero mean and unit variance. Vertical lines in the figure demarcate years. The TAQ average is plotted on the top graph; the CRSP/Gibbs factor z_t on the bottom. The skew in the CRSP/Gibbs series (pronounced peaks, absence of valleys) is simply a consequence of the nonnegativity requirement imposed on the factor. Nevertheless many of the peaks in the two series correspond.

More formally, the (Pearson) correlation between the two weekly series is 0.585; the Spearman correlation is 0.594. Choice of averaging period affects the correlations. The correlation at the daily frequency (the highest available) drop modestly, to 0.447 (and 0.450 for the Spearman). The asset pricing tests presented later, however, are conducted at a monthly

frequency. The correlation between z_t and the TAQ measure averaged over monthly intervals is 0.670 (and 0.689 for the Spearman).

We now turn to the estimated factor loadings. For each firm/year in the comparison sample, we estimate the regression:

$$c_{i,t} = \gamma_{i0}^{TAQ} + \gamma_{i1}^{TAQ} \bar{c}_t^{TAQ} + e_{i,t} \quad (10)$$

where \bar{c}_t^{TAQ} is the cross-firm average effective cost. This regression is simply the analog of the linear specification (6) used in the daily CRSP/Gibbs analysis. In principle, $\gamma_{i,1}^{TAQ}$ should be identical to $\gamma_{i,1}$. In the comparison sample, the estimated correlation is 0.328, Pearson (0.365, Spearman). While this is lower than most of the proxy correlations reported, it should be noted that in the asset pricing tests these proxies are averaged within portfolios, which presumably enhances the precision.

In summary, the analysis of the comparison sample establishes a good case for the validity of the LCF CRSP/Gibbs estimates as proxies for the corresponding TAQ values.

b. Results in the full CRSP sample, 1926-2005

Figure 5 graphs the effective cost common factor z_t over the 1926-2005 period. For visual clarity, the figure plots monthly averages. Many of the peaks sensibly correspond to contemporaneous news events, of which several are identified. The small drop in the average level of z_t post-1985 coincides with the inclusion of Nasdaq firms in the CRSP data (and in the panel sample used to estimate the factor).

When the liquidity factor is viewed as a risk factor in modeling stock returns, it is sometimes more appropriate to focus on the innovation in the series (i.e., the new information). The innovations are constructed as AR(1) residuals. (This specification was chosen by minimizing the Bayesian Information Criterion across ARMA specifications through the fifth order. Due to the Nasdaq inclusion, separate estimations were made for the pre- and post-1985 periods.)

7. Asset pricing results

This section presents empirical analyses aimed at determining whether the level and covariation of effective cost is a priced characteristic and whether the common component of effective cost is a priced risk factor.

a. Specifications

The empirical analysis follows the GMM approach summarized in Cochrane (2005) (pp. 241-243), modified to allow for characteristics and applied to portfolios constructed according to various rankings. The specification for expected returns is

$$ER_t = \beta\lambda + Z_t\delta \quad (11)$$

where R_t is a vector of excess returns relative to the risk-free rate for N assets; λ is a K -vector of factor risk premia; β is a matrix of factor loadings; Z_t is an $N \times M$ matrix of characteristics; and δ is an M -vector of coefficients for the characteristics. The factor loadings are the projection coefficients in the K -factor return generating process:

$$R_t = a + \beta f_t + u_t \quad (12)$$

where a is a constant vector; f_t is a vector of factor realizations; and, u_t is a vector of idiosyncratic zero-mean disturbances. The equilibrium conditions that follow from the usual economic arguments imply $\delta = 0$ and $a = \beta(\lambda - Ef_t)$. If all factors are excess returns on traded portfolios (a condition that is sometimes, but not always, met in the present analyses), the second conclusion reduces to $a = 0$.

The parameter estimates are equivalent to those obtained from a two-pass procedure in which estimates of β are obtained via ordinary least squares (OLS) time-series regression of (12), and then used on the right-hand side in an OLS estimation of (11). In practice (as described in Cochrane) these two steps are combined into a single GMM estimation. By doing this, the

standard errors of the λ and δ estimates are corrected for the estimation error in the β values (as well as heteroscedasticity).⁵

The results reported are a representative among a large set of potential specifications. Three sets of factors are considered. The first set consists solely of the Fama-French excess market return $r_{mt} - r_{ft}$ factor. The second set adds the Fama-French smb_t and hml_t factors. The third set consists of the three FF factors and the innovation in the liquidity common factor \tilde{z}_t .

Three specifications for the set of characteristics Z_t are considered:

- (basic) the level estimate of effective cost from the basic market-adjusted model: c_{it} where c_{it} is the (portfolio average) of the cost estimates over the prior year.
- (common factor) intercept and slope estimates from the LCF model: the portfolio average of γ_{0i} and γ_{1i} estimated over the prior year.
- (seasonal basic) a January dummy variable, both by itself and interacted with the level estimate of the effective cost: d_t^{Jan} , $d_t^{Jan} c_{it}$, and $(1 - d_t^{Jan}) c_{it}$

As the characteristics are not de-meant, Z_t also includes a constant term.

⁵ More precisely, the moment conditions used in estimation are:

$$E \begin{bmatrix} R_t - (a + \beta f_t) \\ f_t' \otimes (R_t - (a + \beta f_t)) \\ \beta' (R_t - \beta \lambda - Z_t \delta) \\ Z_t' (R_t - \beta \lambda - Z_t \delta) \end{bmatrix} = 0$$

These suffice to identify estimates of a , β , λ , and δ that equal those from the two-pass OLS procedure. The first two (vector) conditions are the $N(K+1)$ normal equations that identify the estimates of a and β ; the second two conditions are the $K+M$ normal equations that identify the estimates of λ and δ . Cochrane shows that under the assumption of normality, the GMM standard errors are asymptotically equivalent to those constructed with the Shanken (1992) correction.

The basic specification can be motivated as a straightforward test of whether effective cost is a priced characteristic. The common factor specification extends this test to encompass liquidity covariation. The seasonal basic specification examines the prominence of January seasonality.

b. Portfolio formation

Portfolios are formed annually based on information available at the end the prior year: market capitalization at the close of the prior year; and, CRSP/Gibbs estimates of the basic market-adjusted and latent common factor models estimated over the prior year. Results are reported for two sets of portfolios. Twenty-five effective cost/beta portfolios are formed by independent quintile rankings on effective cost and beta estimated using the basic market adjusted model. Note that although the Gibbs estimate of beta is used for constructing the rankings, the beta used in the expected return specification (11) is the estimate from the return-generating process (12). This makes the results more comparable to those of other studies, and ensures that differences in results are primarily due to differences in liquidity measures.

Twenty-five effective cost intercept/loading portfolios are formed by independent quintile rankings on γ_{0i} and γ_{1i} estimated using the latent common factor model. This second set maximizes variation across the portfolios in γ_{0i} and γ_{1i} , and so can reasonably be expected to illuminate the effects of stochastic liquidity variation and covariation.

Separate portfolio sets are formed for NYSE, Amex and Nasdaq listings. Although securities from all listing venues should in principle be priced according to the same model, data limitations (noted above) precluded forming a single set of portfolios with approximately constant characteristics over the full sample.

c. Properties of the factors

Table 3 presents summary statistics for the factors discussed above and related series over the three sample periods. All three Fama-French factors have positive average returns in all

sample periods. The risk-free rate is the most persistent series. Moderate positive autocorrelation is also exhibited by the common liquidity factor, but not its innovation series.

Table 4 presents the correlations between these series. Most importantly, the effective cost factor is not highly correlated with any of the three Fama-French factors. It is, however, moderately positively correlated with $|r_{mt} - r_{ft}|$. This is what might be expected from the positive association between spreads (and effective costs) and volatility. The effective cost factor innovation is slightly negatively correlated with the market return and size factors.

d. Results for the effective cost/beta portfolios

To characterize their general features, Table 5 reports means for firm counts and other variables for the odd-numbered effective cost/beta portfolios. Note that the effective cost in the highest quintile is sharply higher, relative to the lower quintiles. This is consistent with the positive skewness of effective costs noted in connection with Table 1. Also, sorting on effective cost leads to a similar ranking in the intercept and loading coefficient estimates (γ_{0i} and γ_{1i}).

Table 6 reports estimates of the expected return specifications. Results for NYSE (1927-2005), Amex (1962-2005) and Nasdaq (1985-2005) samples are given in Panels A, B and C, respectively. For brevity, Table 6 does not report the estimates of the return generating process (cf. equation (12)). One feature of these estimates, however, is noteworthy. Specification (1) employs excess market return as the sole factor; specification (2) adds the Fama-French size and book-to-market factors; specification (3) also includes the innovation in the latent common factor of effective cost. In the NYSE (1927-2005) sample, across all twenty-five portfolios, the average adjusted R^2 for the return-generating model is 0.762 when only excess market return is used. Adding the two Fama-French factors increases the average to 0.870. With the further addition of the effective cost common factor innovation, this increases to 0.872. Thus, the incremental explanatory power of the effective cost factor is weak.

This weakness is consistent with the general insignificance of the estimated factor risk premia for \tilde{z}_t in specifications (3) and (6). Only in specification (3) for the NYSE is this coefficient large, and in that case it has the wrong sign.

Specification (4) includes as a characteristic the average effective cost from the basic market-adjusted model. Its coefficient is positive in all samples, but statistically significant only for the Amex. Specifications (5) and (6) include the intercepts and loadings from the latent common factor model. The coefficients are positive, but (again, with the exception of the Amex sample) of marginal significance.⁶

Specification (7) examines the seasonality of the effective cost result. The January dummy d_t^{Jan} is included to pick up seasonality unrelated to effective cost. The interacted variables $d_t^{Jan}c_{it}$ and $(1 - d_t^{Jan})c_{it}$ are of more interest. In all three samples the coefficient of $d_t^{Jan}c_{it}$ is significantly positive. This implies that effective cost plays a particularly large role in January.

It is difficult, however, to account for the magnitude of the coefficients. Unlike some liquidity proxies, the effective cost can be directly interpreted in the context of simple trading strategies. An agent executing a round-trip purchase and sale of a stock in principle pays twice the effective cost. Thus, even under the extreme assumption that the marginal agent is pursuing such a strategy (selling at December's closing bid and buying at January's closing ask), the coefficient of effective cost should be at most two. In the NYSE and Amex samples, the estimated coefficients exceed four.

e. Results for the liquidity intercept/loading portfolios

Forming portfolios on the basis of stocks' γ_{0i} and γ_{1i} estimates should in principle make it more likely to detect the effects of liquidity risk on expected returns. Table 7 reports variable means for the odd-numbered portfolios. The ordering of the portfolio averages for γ_{0i} and γ_{1i} are similar to those of the portfolios formed on effective cost and beta, but the ranges are larger.

⁶ Spiegel and Wang (2005) also find a weak liquidity effect using the Gibbs estimates for effective cost developed in an earlier draft of this paper. They furthermore find that in explaining returns, effective cost is dominated by idiosyncratic volatility.

Estimates of the expected return specifications are given in Table 8. The results are similar to those found for the effective cost/beta portfolios. The coefficients of loading on the innovation in the effective cost common factor are small. The γ_{0i} and γ_{1i} coefficient estimates are positive. They are generally of marginal significance (with the exception, this time, of the Nasdaq estimates). The seasonality pattern for the effective cost level is similar to that found for the effective cost/beta portfolios.

8. Discussion and conclusion

The results presented in the last section suggest that the unexpected stochastic variation in aggregate effective cost is not strongly related to stock returns, that a firm's sensitivity to this factor (as a characteristic) has weak explanatory power for expected returns, and that the level of effective cost is related to expected returns mainly through a seasonal component. The seasonality of liquidity effects is noted in Eleswarapu and Reinganum (1993). The present analysis confirms the presence of this phenomenon in a longer and broader sample.

The equivocal findings regarding the importance of effective cost variation and risk, however, contrast with the stronger conclusions found by Pastor and Stambaugh (2003), Acharya and Pedersen (2005), and Korajczyk and Sadka (2006) using different liquidity measures. There are various possible explanations for this. First, the CRSP/Gibbs estimates of effective cost may not be sufficiently precise proxies for the values actually used by agents in making their decisions. This seems unlikely, however, since the analysis of the comparison sample establishes strong correlation between the CRSP/Gibbs estimates and those formed directly from the trade and quote data. Second, the asset-pricing specifications used here may lack the power necessary to detect stochastic liquidity effects. The papers mentioned above span a range of approaches comparable to that found in other asset-pricing contexts, but the present paper employs a general method used by at least one (Acharya and Pedersen).

A third possibility is that effective cost *per se* may not be the relevant trading cost measure used by investors. As noted earlier it does not explicitly measure the price impact effects that come into play when the trading strategy involves splitting an order over time.

Although effective cost and price impact are conceptually distinct, however, they are in practice correlated. From Table 1, the correlation between effective cost and price impact (both estimated from TAQ) is 0.513, suggesting that effective cost is a partial proxy for price impact in the cross-section. Korajczyk and Sadka (2006) find high canonical correlations between the common factors extracted from effective costs and those extracted from price impact, suggesting that the proxy relationship also picks up time series variation. This provides a basis for the assertion that results estimated using effective cost have relevance for other liquidity measures.

The effective cost used to measure liquidity in the present study is unique, however, in one important respect. Alone among the daily-based liquidity proxies commonly used in asset pricing studies (the Pastor-Stambaugh gamma, the Amivest liquidity ratio and the Amihud illiquidity ratio), the effective cost estimate does not incorporate volume. This can be viewed as a limitation, since many microstructure-based measures (such as the price impact) involve a size-related component. On the other hand, most of these measures involve signed order flow, instead of the unsigned volume used in the daily proxies. The microstructure measures also generally assume that order flow is exogenous to price and liquidity dynamics. In fact, volume endogeneity with price dynamics arises from portfolio rebalancing, momentum trading, hedging and other price-driven strategies. The feedback from trading costs to order placement strategy causes volume to depend on liquidity variation.

Thus, although effective cost is a narrow measure of trading cost, measures derived from volume may reflect factors that extend beyond the usual notion of liquidity as immediacy. That these measures have power for explaining expected returns may indicate the importance of defining liquidity broadly enough to encompass the full range of costs and distortions associated with the trading process. Such definitions and interpretations, however, are not invariable straightforward. Chordia, Subrahmanyam and Anshuman (2001) find strong explanatory power in summary measures of trading activity such as the level and volatility of turnover. Surprisingly they find that turnover volatility is negatively related to expected returns. This is contrary to the

notion that turnover volatility might be acting as proxy for liquidity risk. Further exploration of alternative definitions and measures of liquidity may yet offer clarification.

9. References

- Acharya, Viral V. and Lasse Heje Pedersen, 2005, Asset pricing with liquidity risk. *Journal of Financial Economics* 77: 375-410.
- Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets* 5(1): 31-56.
- Amihud, Yakov and Haim Mendelson, 1986, Asset pricing and the bid-ask spread. *Journal of Financial Economics* 17(2): 223-249.
- Brennan, Michael J. and Avanidhar Subrahmanyam, 1996, Market microstructure and asset pricing: on the compensation for illiquidity in stock returns. *Journal of Financial Economics* 41(3): 441-464.
- Carlin, Bradley P. and Thomas A. Louis, 2000, *Bayes and Empirical Bayes Methods for Data Analysis*. London, Chapman and Hall.
- Chalmers, John M. R. and Gregory B. Kadlec, 1998, An empirical examination of the amortized spread. *Journal of Financial Economics* 48(2): 159-188.
- Chordia, Tarun, Avanidhar Subrahmanyam and V. Ravi Anshuman, 2001, Trading activity and expected stock returns. *Journal of Financial Economics* 59: 3-32.
- Cochrane, John H., 2005, *Asset Pricing*. Princeton, Princeton University Press.
- Easley, David, Soeren Hvidkjaer and Maureen O'Hara, 2002, Is information risk a determinant of asset returns? *Journal of Finance* 57(5): 2185-2221.
- Easley, David and Maureen O'Hara, 2002, Microstructure and asset pricing. *Handbook of Financial Economics*. G. M. Constantinides, M. Harris and R. M. Stulz. New York, Elsevier.
- Eleswarapu, Venkat R. and Marc R. Reinganum, 1993, The seasonal behavior of the liquidity premium in asset pricing. *Journal of Financial Economics* 34: 373-386.
- Fama, Eugene F. and Kenneth R. French, 2004, New lists: fundamentals and survival rates. *Journal of Financial Economics* 72: 229-269.
- Fink, Jason, Kristin Fink, Gustavo Grullon and James P. Weston, 2006, Firm age and fluctuations in idiosyncratic risk, Jones School, Rice University.
- Geweke, John, 2005, *Contemporary Bayesian Statistics and Econometrics*. New York, John Wiley and Sons.
- Geweke, John and Guofu Zhou, 1996, Measuring the pricing error of the arbitrage pricing theory. *Review of Financial Studies* 9: 557-587.

- Hajivassiliou, Vassilis, Daniel McFadden and Paul Ruud, 1996, Simulation of multivariate normal rectangle probabilities and their derivatives - Theoretical and computational results. *Journal of Econometrics* 72(1-2): 85-134.
- Harris, Lawrence E., 1990, Statistical properties of the Roll serial covariance bid/ask spread estimator. *Journal of Finance* 45(2): 579-590.
- Hasbrouck, Joel, 2004, Liquidity in the futures pits: Inferring market dynamics from incomplete data. *Journal of Financial and Quantitative Analysis* 39(2).
- Jagannathan, Ravi, Georgios Skoulakis and Zhenyu Wang, 2006, The analysis of the cross-section of security returns. *Handbook of Financial Econometrics*. L. Hansen and Y. Ait-Sahalia, Elsevier North-Holland.
- Kim, Chang-Jin and Charles R. Nelson, 2000, *State-space models with regime switching*. Cambridge, Massachusetts, MIT Press.
- Korajczyk, Robert A. and Ronnie Sadka, 2006, Commonality across alternative measures of liquidity, Kellogg School, Northwestern University.
- Lancaster, Tony, 2004, *An Introduction to Modern Bayesian Econometrics*. Malden (MA), Blackwell Publishing.
- Lee, Charles M. C., 1993, Market integration and price execution for NYSE-listed securities. *Journal of Finance* 48(3): 1009-1038.
- Lesmond, David A., Joseph P. Ogden and Charles A. Trzcinka, 1999, A new estimate of transactions costs. *Review of Financial Studies* 12(5): 1113-1141.
- Pastor, Lubos and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns. *Journal of Political Economy* 111(3): 642-685.
- Piwowar, Michael S. and Li Wei, 2006, The sensitivity of effective spread estimates to trade-quote matching algorithms. *International Journal of Electronic Markets* 16(2): 112-129.
- Reinganum, Marc R., 1990, Market microstructure and asset pricing: an empirical Investigation of NYSE and NASDAQ securities. *Journal of Financial Economics* 28: 127-147.
- Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance* 39(4): 1127-1139.
- Sadka, Ronnie, 2004, Liquidity risk and asset pricing. University of Washington.
- Schultz, Paul H., 2000, Regulatory and legal pressures and the costs of Nasdaq trading. *Review of Financial Studies* 13(4): 917-957.
- Shanken, Jay, 1992, On the estimation of beta pricing models. *Review of Financial Studies* 5(1-34).

Spiegel, Matthew and Xiatong Wang, 2005, Cross-sectional variation in stock returns: liquidity and idiosyncratic risk, Yale University.

Stoll, Hans R., 2006, Electronic trading in stock markets. *Journal of Economic Perspectives* 20: 153-174.

Stoll, Hans R. and Robert E. Whalley, 1983, Transaction cost and the small firm effect. *Journal of Financial Economics* 12: 57-79.

Table 1. Summary statistics for the comparison sample, 1993-2005

The comparison sample consists of approximately 150 Nasdaq firms and 150 NYSE/Amex firms selected in a capitalization-stratified random draw in each of the years 1993-2005. Values in the table are based on annual estimates for the 3,777 firms that could be matched between CRSP and TAQ. Effective cost is the difference between the log transaction price and the prevailing log quote midpoint. For each firm, the TAQ estimate is the annual average of this value over all trades, trade-weighted. The CRSP moment estimate is $\sqrt{-Cov(\Delta p_t, \Delta p_{t-1})}$ where Δp_t is the log price change and the covariance is estimated over all trading days in the year. The estimate is set to zero if the covariance is positive. The CRSP Gibbs values are estimates from the basic market-adjusted model; Proportion of zero returns is the fraction of trading days that had a zero price change from the previous day. The Amihud (2002) illiquidity measure is $I = \frac{|return|}{|Dollar\ volume|}$, averaged over all days with non-zero volume. The price impact coefficient is λ in the regression $\Delta p_t = \lambda \left(Signed \sqrt{Dollar\ Volume} \right)_t + \varepsilon_t$, estimated annually using log price changes and signed dollar volumes aggregated over five-minute intervals.

<i>Estimate</i>	<i>Source</i>	<i>Mean</i>	<i>Median</i>	<i>Std. Dev.</i>	<i>Skewness</i>	<i>Kurtosis</i>
<i>Effective cost</i>	TAQ	0.0106	0.0054	0.0146	4.61	54.7
<i>Effective cost</i>	CRSP Gibbs	0.0112	0.0061	0.0141	4.97	62.8
<i>Effective cost</i>	CRSP Moment	0.0106	0.0056	0.0152	4.35	52.1
<i>Proportion of zero returns</i>	CRSP	0.1363	0.1071	0.1171	1.02	0.9
<i>Price impact ($\lambda \times 10^6$)</i>	TAQ	28.1500	7.4098	70.6173	7.84	101.2
<i>Amihud Illiquidity ratio ($I \times 10^6$)</i>	CRSP	3.6592	0.0709	20.0366	16.56	395.8
<i>Market capitalization (\$ Million)</i>	CRSP	2,587.7190	196.9200	14,407.3199	18.55	502.9
<i>Price (end of year, \$/share)</i>	CRSP	20.8442	14.5000	29.4357	11.38	229.8

Table 2. Correlations between liquidity proxies for the comparison sample

The comparison sample consists of approximately 150 Nasdaq firms and 150 NYSE/Amex firms selected in a capitalization-stratified random draw in each of the years 1993-2005. Values in the table are based on annual estimates for the 3,777 firms that could be matched between CRSP and TAQ. Effective cost is the difference between the log transaction price and the prevailing log quote midpoint. For each firm, the TAQ estimate is the annual average of this value over all trades, trade-weighted. The CRSP moment estimate is $\sqrt{-Cov(\Delta p_t, \Delta p_{t-1})}$ where Δp_t is the log price change and the covariance is estimated over all trading days in the year. The estimate is set to zero if the covariance is positive. The CRSP Gibbs values are estimates from the basic market-adjusted model; Proportion of zero returns is the fraction of trading days that had a zero price change from the previous day. The Amihud (2002) illiquidity measure is $I = |return|/|Dollar\ volume|$, averaged over all days with non-zero volume. The price impact coefficient is λ in the regression $\Delta p_t = \lambda(\text{Signed}\sqrt{Dollar\ Volume})_t + \varepsilon_t$, estimated annually using log price changes and signed dollar volumes aggregated over five-minute intervals. Partial correlations are adjusted for log(end-of-year price) and log(market capitalization).

	<i>Eff. cost (TAQ)</i>	<i>Eff. cost, Gibbst</i>	<i>Eff. cost, Moment</i>	<i>Prop. zero returns</i>	<i>Price Impact (TAQ)</i>	<i>Illiquidity</i>
Pearson correlation						
<i>Eff. cost (TAQ)</i>	1.000	0.965	0.878	0.611	0.513	0.612
<i>Eff. cost, Gibbs</i>	0.965	1.000	0.917	0.579	0.450	0.589
<i>Eff. cost, Moment</i>	0.878	0.917	1.000	0.451	0.378	0.504
<i>Prop. zero returns</i>	0.611	0.579	0.451	1.000	0.311	0.252
<i>Price impact (TAQ)</i>	0.513	0.450	0.378	0.311	1.000	0.668
<i>Illiquidity</i>	0.612	0.589	0.504	0.252	0.668	1.000
Spearman correlation						
<i>Eff. cost (TAQ)</i>	1.000	0.872	0.636	0.770	0.735	0.937
<i>Eff. cost, Gibbs</i>	0.872	1.000	0.791	0.620	0.577	0.778
<i>Eff. cost, Moment</i>	0.636	0.791	1.000	0.417	0.363	0.592
<i>Prop. zero returns</i>	0.770	0.620	0.417	1.000	0.510	0.704
<i>Price impact (TAQ)</i>	0.735	0.577	0.363	0.510	1.000	0.824
<i>Illiquidity</i>	0.937	0.778	0.592	0.704	0.824	1.000
Pearson partial correlation						
<i>Eff. cost (TAQ)</i>	1.000	0.943	0.805	0.366	0.268	0.567
<i>Eff. cost, Gibbs</i>	0.943	1.000	0.866	0.359	0.189	0.517
<i>Eff. cost, Moment</i>	0.805	0.866	1.000	0.193	0.107	0.397
<i>Prop. zero returns</i>	0.366	0.359	0.193	1.000	0.068	0.103
<i>Price impact (TAQ)</i>	0.268	0.189	0.107	0.068	1.000	0.610
<i>Illiquidity</i>	0.567	0.517	0.397	0.103	0.610	1.000
Spearman partial correlation						
<i>Eff. cost (TAQ)</i>	1.000	0.678	0.382	0.564	0.024	0.631
<i>Eff. cost, Gibbs</i>	0.678	1.000	0.682	0.285	-0.123	0.361
<i>Eff. cost, Moment</i>	0.382	0.682	1.000	0.101	-0.182	0.288
<i>Prop. zero returns</i>	0.564	0.285	0.101	1.000	-0.021	0.341
<i>Price impact (TAQ)</i>	0.024	-0.123	-0.182	-0.021	1.000	0.375
<i>Illiquidity</i>	0.631	0.361	0.288	0.341	0.375	1.000

Table 3. Summary statistics for return factors and related series

r_{ft} is the one-month Treasury bill rate (Ibbotson and Associates); r_{mt} is the CRSP valued-weighted average NYSE/Amex/Nasdaq return; smb_t and hml_t are the Fama-French size and value/growth factors (from Kenneth French's website). z_t is the average monthly effective cost common factor, estimated from the latent common factor model; \tilde{z}_t is the corresponding innovations series estimated as the AR(1) residuals.

		<i>Mean</i>	<i>Std.Dev.</i>	<i>First-order autocorrelation</i>
1926-2005	r_{ft}	0.0030	0.0026	0.9727
	$r_{mt} - r_{ft}$	0.0065	0.0547	0.1070
	$ r_{mt} - r_{ft} $	0.0391	0.0388	0.2249
	smb_t	0.0024	0.0336	0.0761
	hml_t	0.0041	0.0359	0.1772
	z_t	0.4861	0.2763	0.5522
	\tilde{z}_t	0.0004	0.2287	0.0285
1962-2005	r_{ft}	0.0046	0.0023	0.9524
	$r_{mt} - r_{ft}$	0.0045	0.0445	0.0600
	$ r_{mt} - r_{ft} $	0.0345	0.0283	0.0682
	smb_t	0.0023	0.0320	0.0655
	hml_t	0.0046	0.0289	0.1301
	z_t	0.4365	0.2586	0.5653
	\tilde{z}_t	-0.0107	0.2125	-0.0016
1985-2005	r_{ft}	0.0039	0.0017	0.9494
	$r_{mt} - r_{ft}$	0.0069	0.0443	0.0423
	$ r_{mt} - r_{ft} $	0.0348	0.0282	0.1228
	smb_t	0.0006	0.0344	-0.0305
	hml_t	0.0033	0.0319	0.0938
	z_t	0.3855	0.2802	0.5763
	\tilde{z}_t	-0.0040	0.2300	-0.0000

Table 4. Correlations for return factors

r_{ft} is the one-month Treasury bill rate (Ibbotson and Associates); r_{mt} is the CRSP valued-weighted average NYSE/Amex/Nasdaq return; smb_t and hml_t are the Fama-French size and value/growth factors (from Kenneth French's website). z_t is the average monthly effective cost common factor, estimated from the latent common factor model; \tilde{z}_t is the corresponding innovations series estimated as the AR(1) residuals.

	r_{ft}	$r_{mt} - r_{ft}$	$ r_{mt} - r_{ft} $	smb_t	hml_t	z_t	\tilde{z}_t
r_{ft}	1.000	-0.069	-0.067	-0.059	0.013	-0.109	-0.031
$r_{mt} - r_{ft}$	-0.069	1.000	0.069	0.326	0.216	-0.063	-0.197
$ r_{mt} - r_{ft} $	-0.067	0.069	1.000	0.045	0.193	0.401	0.369
smb_t	-0.059	0.326	0.045	1.000	0.094	-0.021	-0.150
hml_t	0.013	0.216	0.193	0.094	1.000	0.100	0.029
z_t	-0.109	-0.063	0.401	-0.021	0.100	1.000	0.828
\tilde{z}_t	-0.031	-0.197	0.369	-0.150	0.029	0.828	1.000

Table 5. Summary statistics for portfolios constructed on effective cost and beta rankings

Twenty-five portfolios are constructed as the intersection of independent quintile rankings on the Gibbs estimates of effective cost (c_{it}) and beta (β_{it}), estimated over the prior year using the basic market-adjusted model. Table reports mean values for odd-numbered portfolios over the sample. γ_{0i} and γ_{1i} are the Gibbs estimates of the intercept and loading for the latent common factor model.

	c_{it} Rank	β_{it} Rank	No. firms	c_{it}	β_{it}	γ_{0i}	γ_{1i}
<i>NYSE 1927-2005</i>	1	1	42.8	0.0021	0.346	0.0020	0.0028
		3	48.4	0.0021	0.901	0.0020	0.0035
		5	30.8	0.0023	1.626	0.0021	0.0045
	3	1	41.6	0.0056	0.323	0.0047	0.0043
		3	42.2	0.0056	0.901	0.0043	0.0057
		5	46.3	0.0056	1.731	0.0042	0.0071
	5	1	45.9	0.0223	0.283	0.0193	0.0085
		3	39.7	0.0194	0.899	0.0155	0.0097
		5	47.9	0.0174	1.725	0.0127	0.0119
<i>Amex 1962-2005</i>	1	1	26.5	0.0032	0.129	0.0034	0.0034
		3	25.7	0.0032	0.625	0.0034	0.0049
		5	23.6	0.0035	1.462	0.0038	0.0064
	3	1	25.1	0.0101	0.103	0.0090	0.0057
		3	24.2	0.0102	0.628	0.0088	0.0078
		5	27.9	0.0101	1.455	0.0086	0.0107
	5	1	26.0	0.0379	0.039	0.0350	0.0110
		3	28.3	0.0351	0.631	0.0317	0.0117
		5	21.0	0.0366	1.376	0.0319	0.0153
<i>Nasdaq 1985-2005</i>	1	1	28.7	0.0051	0.092	0.0049	0.0049
		3	119.7	0.0044	0.611	0.0042	0.0061
		5	177.6	0.0046	1.561	0.0048	0.0087
	3	1	108.8	0.0146	0.049	0.0135	0.0062
		3	108.7	0.0144	0.606	0.0127	0.0092
		5	102.1	0.0141	1.505	0.0121	0.0128
	5	1	177.0	0.0459	-0.021	0.0431	0.0124
		3	102.9	0.0480	0.597	0.0443	0.0156
		5	43.3	0.0482	1.410	0.0435	0.0187

Table 6. Expected return estimates for portfolios constructed on effective cost and beta rankings

Table reports estimates for β and δ in the specification $ER_t = \beta\lambda + Z_t\delta$ where R_t is a vector of excess returns (relative to the risk-free rate); λ is a K -vector of factor risk premia; β is a matrix of factor loadings; Z_t is an $N \times M$ matrix of characteristics; and δ is an M -vector of coefficients for the characteristics. The factor loadings are the projection coefficients in the K -factor return generating process: $R_t = a + \beta f_t + u_t$, estimated via OLS. The factors considered are the excess market return ($r_{mt} - r_{ft}$), the Fama-French size and book-to-market factors (smb_t and hml_t), and the innovation in the common liquidity factor (\tilde{z}_t). The characteristics considered are an intercept, the Gibbs estimate of effective cost using the basic market-adjusted model, the liquidity intercept and loading (γ_{0i} and γ_{1i}) estimated from the latent common factor model, a January monthly dummy variable (d_t^{Jan}), and the January dummy interacted with the Gibbs estimate of the effective cost ($d_t^{Jan}c_{it}$ and $(1 - d_t^{Jan})c_{it}$). Coefficient estimates are ordinary least squares. T-statistics are corrected for joint estimation of factor loadings, heteroscedasticity and autocorrelation. In each panel, the model is estimated for twenty-five portfolios constructed as the intersection of independent quintile rankings on the Gibbs estimates of effective cost (c_{it}) and beta (β_{it}), estimated over the prior year using the basic market-adjusted model.

A. NYSE, 1927-2005.

		(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00757 (3.67)	0.00565 (1.65)	0.00302 (1.05)	0.00448 (0.78)	0.00412 (0.62)	0.00482 (0.74)	0.00448 (0.76)
	smb_t		0.00637 (0.67)	0.00079 (0.14)	-0.00695 (-1.19)	-0.00673 (-0.83)	-0.00696 (-0.84)	-0.00695 (-1.17)
	hml_t		-0.00084 (-0.05)	0.01381 (1.50)	-0.00384 (-0.36)	-0.00436 (-0.33)	-0.00287 (-0.21)	-0.00384 (-0.35)
	\tilde{z}_t			-0.17580 (-2.99)			-0.03750 (-0.40)	
	<i>Charac- teristics</i>	<i>Intercept</i>				0.00186 (0.44)	0.00075 (0.16)	-0.00043 (-0.07)
	c_{it}				1.11512 (1.47)			
	γ_{0i}					1.21937 (1.53)	1.21771 (1.53)	
	γ_{1i}					0.44924 (0.35)	0.45227 (0.35)	
	d_t^{Jan}							0.01026 (1.07)
	$d_t^{Jan}c_{it}$							4.18978 (3.05)
	$(1 - d_t^{Jan})c_{it}$							0.83560 (1.03)

Table 6. Expected return estimates for portfolios constructed on effective cost and beta rankings (continued)

B. Amex, 1962-2005.

<i>Variable</i>		(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00844 (2.64)	-0.00750 (-1.99)	-0.00882 (-2.10)	0.00917 (0.84)	0.00018 (0.01)	-0.00053 (-0.05)	0.00917 (0.92)
	smb_t		0.00933 (2.49)	0.01224 (2.08)	-0.02910 (-2.69)	-0.02733 (-2.57)	-0.03108 (-2.45)	-0.02910 (-3.05)
	hml_t		0.01144 (3.85)	0.00831 (1.46)	-0.00604 (-0.81)	-0.00190 (-0.26)	0.00099 (0.12)	-0.00604 (-0.84)
	\tilde{z}_t			-0.06599 (-0.76)			0.11158 (1.09)	
	<i>Intercept</i>				0.01594 (2.55)	0.01447 (2.55)	0.01741 (2.61)	0.01440 (2.31)
<i>Charac- teristics</i>	c_{it}				1.57349 (3.16)			
	γ_{0i}					1.30782 (2.84)	1.32085 (2.84)	
	γ_{1i}					1.44763 (1.52)	1.43494 (1.50)	
	d_t^{Jan}							0.01839 (1.68)
	$d_t^{Jan} c_{it}$							5.80286 (6.20)
	$(1 - d_t^{Jan}) c_{it}$							1.18900 (3.03)

Table 6. Expected return estimates for portfolios constructed on effective cost and beta rankings (continued)

C. Nasdaq, 1985-2005

	<i>Variable</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00898 (2.07)	-0.00474 (-0.86)	-0.00375 (-0.61)	0.00601 (0.68)	0.00463 (0.62)	0.00380 (0.42)	0.00601 (0.75)
	smb_t		0.01379 (2.31)	0.01262 (1.83)	0.00095 (0.07)	-0.01268 (-1.11)	-0.01325 (-1.26)	0.00095 (0.08)
	hml_t		0.01099 (3.26)	0.01202 (3.25)	0.00825 (1.94)	0.00897 (1.81)	0.00864 (1.84)	0.00825 (1.98)
	\tilde{z}_t			0.01524 (0.18)			0.03160 (0.31)	
	<i>Intercept</i>				-0.00288 (-0.46)	-0.00455 (-0.65)	-0.00307 (-0.43)	-0.00428 (-0.66)
<i>Charac- teristics</i>	c_{it}				0.28111 (1.36)			
	γ_{0i}					0.11675 (0.52)	0.11556 (0.51)	
	γ_{1i}					1.94321 (1.88)	1.94179 (1.87)	
	d_t^{Jan}							0.01687 (0.92)
	$d_t^{Jan} c_{it}$							2.44436 (5.58)
	$(1 - d_t^{Jan}) c_{it}$							0.08445 (0.45)

Table 7. Summary statistics for portfolios constructed on liquidity factor intercept and loading rankings

Twenty-five portfolios are constructed as the intersection of independent quintile rankings on the Gibbs estimates of γ_{0i} and γ_{1i} estimated over the prior year using the latent common factor model. Effective cost (c_{it}) and beta (β_{it}) are estimated over the prior year using the basic market-adjusted model. Table reports mean values for odd-numbered portfolios over the sample.

	γ_{0i} Rank	γ_{1i} Rank	No. firms	c_{it}	β_{it}	γ_{0i}	γ_{1i}
<i>NYSE 1927-2005</i>	1	1	26.9	0.0015	0.724	0.0015	0.0019
		3	47.9	0.0026	0.933	0.0015	0.0049
		5	47.3	0.0049	1.126	0.0014	0.0127
	3	1	46.6	0.0036	0.710	0.0041	0.0018
		3	43.0	0.0054	0.998	0.0041	0.0049
		5	38.8	0.0090	1.194	0.0041	0.0131
	5	1	51.2	0.0121	0.738	0.0143	0.0017
		3	37.3	0.0168	0.984	0.0163	0.0049
		5	51.4	0.0259	1.058	0.0194	0.0147
<i>Amex 1962-2005</i>	1	1	22.2	0.0024	0.384	0.0028	0.0022
		3	27.2	0.0038	0.670	0.0026	0.0064
		5	25.8	0.0075	0.925	0.0027	0.0185
	3	1	29.4	0.0071	0.448	0.0087	0.0020
		3	24.8	0.0095	0.797	0.0086	0.0064
		5	23.0	0.0158	0.901	0.0088	0.0190
	5	1	20.3	0.0250	0.541	0.0292	0.0022
		3	24.9	0.0337	0.653	0.0338	0.0065
		5	32.0	0.0436	0.704	0.0357	0.0200
<i>Nasdaq 1985-2005</i>	1	1	64.9	0.0034	0.623	0.0039	0.0026
		3	123.8	0.0048	0.906	0.0037	0.0077
		5	114.7	0.0088	1.080	0.0037	0.0225
	3	1	124.4	0.0112	0.499	0.0128	0.0022
		3	106.0	0.0139	0.767	0.0127	0.0076
		5	93.1	0.0201	0.771	0.0128	0.0228
	5	1	112.4	0.0353	0.282	0.0397	0.0022
		3	92.5	0.0444	0.457	0.0446	0.0076
		5	135.2	0.0574	0.499	0.0484	0.0255

Table 8. Expected return estimates for portfolios constructed on liquidity factor intercept and loading rankings.

Table reports estimates for β and δ in the specification $ER_t = \beta\lambda + Z_t\delta$ where R_t is a vector of excess returns (relative to the risk-free rate); λ is a K -vector of factor risk premia; β is a matrix of factor loadings; Z_t is an $N \times M$ matrix of characteristics; and δ is an M -vector of coefficients for the characteristics. The factor loadings are the projection coefficients in the K -factor return generating process: $R_t = a + \beta f_t + u_t$, estimated via OLS. The factors considered are the excess market return ($r_{mt} - r_{ft}$), the Fama-French size and book-to-market factors (smb_t and hml_t), and the innovation in the common liquidity factor (\tilde{z}_t). The characteristics considered are an intercept, the Gibbs estimate of effective cost using the basic market-adjusted model, the liquidity intercept and loading (γ_{0i} and γ_{1i}) estimated from the latent common factor model, a January monthly dummy variable (d_t^{Jan}), and the January dummy interacted with the Gibbs estimate of the effective cost ($d_t^{Jan}c_{it}$ and $(1 - d_t^{Jan})c_{it}$). Coefficient estimates are ordinary least squares. T-statistics are corrected for joint estimation of factor loadings, heteroscedasticity and autocorrelation. In each panel, the model is estimated for twenty-five portfolios constructed as the intersection of independent quintile rankings on the Gibbs estimates of effective cost (c_{it}) and beta (β_{it}), estimated over the prior year using the basic market-adjusted model.

A. NYSE, 1927-2005.

Variable		(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00755 (3.67)	0.00710 (3.12)	0.00764 (3.20)	0.01662 (1.08)	0.01949 (0.88)	0.01941 (0.92)	0.01662 (1.05)
	smb_t		0.00412 (2.71)	0.00381 (2.55)	-0.00693 (-1.06)	-0.00666 (-0.91)	-0.00663 (-0.86)	-0.00693 (-1.03)
	hml_t		-0.00124 (-0.44)	-0.00239 (-0.75)	-0.01184 (-1.04)	-0.01442 (-0.99)	-0.01429 (-1.03)	-0.01184 (-1.00)
	\tilde{z}_t			0.03501 (0.42)			-0.02305 (-0.16)	
	<i>Charac- teristics</i>	<i>Intercept</i>				-0.00780 (-0.63)	-0.01156 (-0.62)	-0.01154 (-0.63)
	c_{it}				1.18642 (1.42)			
	γ_{0i}					1.21408 (1.46)	1.21443 (1.46)	
	γ_{1i}					0.57275 (0.63)	0.57190 (0.62)	
	d_t^{Jan}							0.01279 (1.27)
	$d_t^{Jan}c_{it}$							3.94894 (2.67)
	$(1 - d_t^{Jan})c_{it}$							0.93529 (1.05)

Table 8. Expected return estimates for portfolios constructed on liquidity factor intercept and loading rankings. (continued)

B. Amex, 1962-2005.

	<i>Variable</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00868 (2.68)	-0.00769 (-1.65)	-0.00814 (-1.35)	0.03185 (1.41)	0.00829 (0.36)	0.01045 (0.41)	0.03185 (1.54)
	smb_t		0.00456 (0.94)	0.00530 (0.60)	-0.02795 (-2.51)	-0.02298 (-2.43)	-0.02933 (-2.07)	-0.02795 (-2.83)
	hml_t		0.02199 (2.61)	0.02143 (1.99)	-0.01496 (-1.13)	-0.01450 (-1.06)	-0.01208 (-0.77)	-0.01496 (-1.24)
	\tilde{z}_t			-0.02314 (-0.17)			0.17932 (0.92)	
	<i>Intercept</i>				-0.00323 (-0.22)	0.01040 (0.57)	0.01346 (0.74)	-0.00429 (-0.31)
<i>Characteristics</i>	c_{it}				1.56484 (2.87)			
	γ_{0i}					1.33776 (2.76)	1.36922 (2.82)	
	γ_{1i}					1.08496 (1.73)	1.01960 (1.59)	
	d_t^{Jan}							0.01279 (1.19)
	$d_t^{Jan} c_{it}$							6.20826 (6.59)
	$(1 - d_t^{Jan}) c_{it}$							1.14271 (2.66)

Table 8. Expected return estimates for portfolios constructed on liquidity factor intercept and loading rankings. (continued)

C. Nasdaq, 1985-2005

	<i>Variable</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Factors</i>	$r_{mt} - r_{ft}$	0.00943 (2.16)	-0.00720 (-1.12)	-0.00681 (-0.88)	0.00223 (0.21)	-0.03122 (-1.43)	-0.03159 (-1.27)	0.00223 (0.23)
	smb_t		0.01458 (1.98)	0.01412 (1.56)	-0.00051 (-0.03)	-0.01509 (-0.85)	-0.01524 (-0.83)	-0.00051 (-0.03)
	hml_t		0.02159 (3.48)	0.02265 (3.75)	0.00767 (0.55)	0.00087 (0.05)	0.00053 (0.02)	0.00767 (0.59)
	\tilde{z}_t			0.04805 (0.69)			0.06739 (0.19)	
	<i>Intercept</i>				0.00192 (0.25)	0.04034 (1.64)	0.04095 (1.24)	0.00051 (0.06)
<i>Characteristics</i>	c_{it}				0.29523 (1.10)			
	γ_{0i}					0.19558 (0.73)	0.19568 (0.73)	
	γ_{1i}					1.18753 (2.10)	1.18551 (1.96)	
	d_t^{Jan}							0.01699 (0.95)
	$d_t^{Jan} c_{it}$							2.39502 (5.20)
	$(1 - d_t^{Jan}) c_{it}$							0.10434 (0.43)

Figure 1. Posteriors for simulated price paths

A quote-midpoint series of length 250 (roughly a year's worth of daily data) is simulated using using a volatility $\sigma_u = 0.02$; 250 realizations are also generated for the trade direction indicators (q_t). Using these values, two price series are simulated: one using an effective cost of $c=0.01$, the other with $c=0.10$. For each series, the joint parameter posterior is estimated using 10,000 draws of a Gibbs sampler. The shaded regions indicate the ninety-percent confidence regions. In panels, the horizontal (σ_u) axis and the scale of the vertical (c) axis are identical.

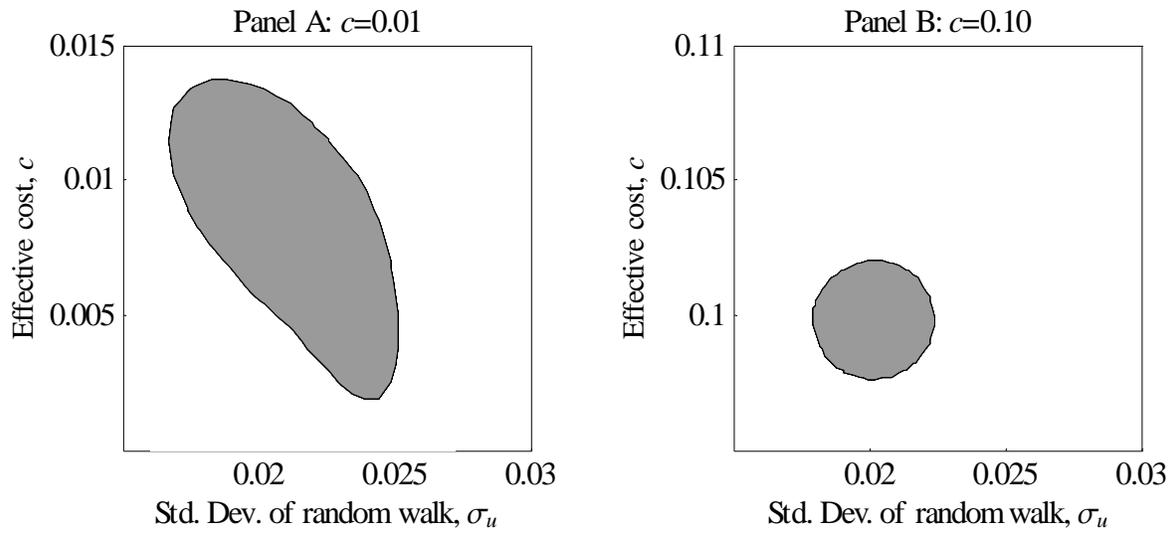


Figure 2. TAQ and CRSP/Gibbs estimates of effective cost in the comparison sample

The comparison sample consists of approximately 150 Nasdaq firms and 150 NYSE/Amex firms selected in a capitalization-stratified random draw in each of the years 1993-2005. For each firm in each year, the effective cost is estimated from TAQ data and from CRSP daily data using the Gibbs procedure. The figure depicts the cross-sectional distributions for these estimates year-by-year, with TAQ estimates on the left and Gibbs estimates on the right. The upper and lower ranges of the box-and-whisker figures demarcate the fifth and ninety-fifth percentiles; the upper and lower edges of the boxes correspond to the twenty-fifth and seventy-fifth percentiles; the line drawn across the box indicates the median.

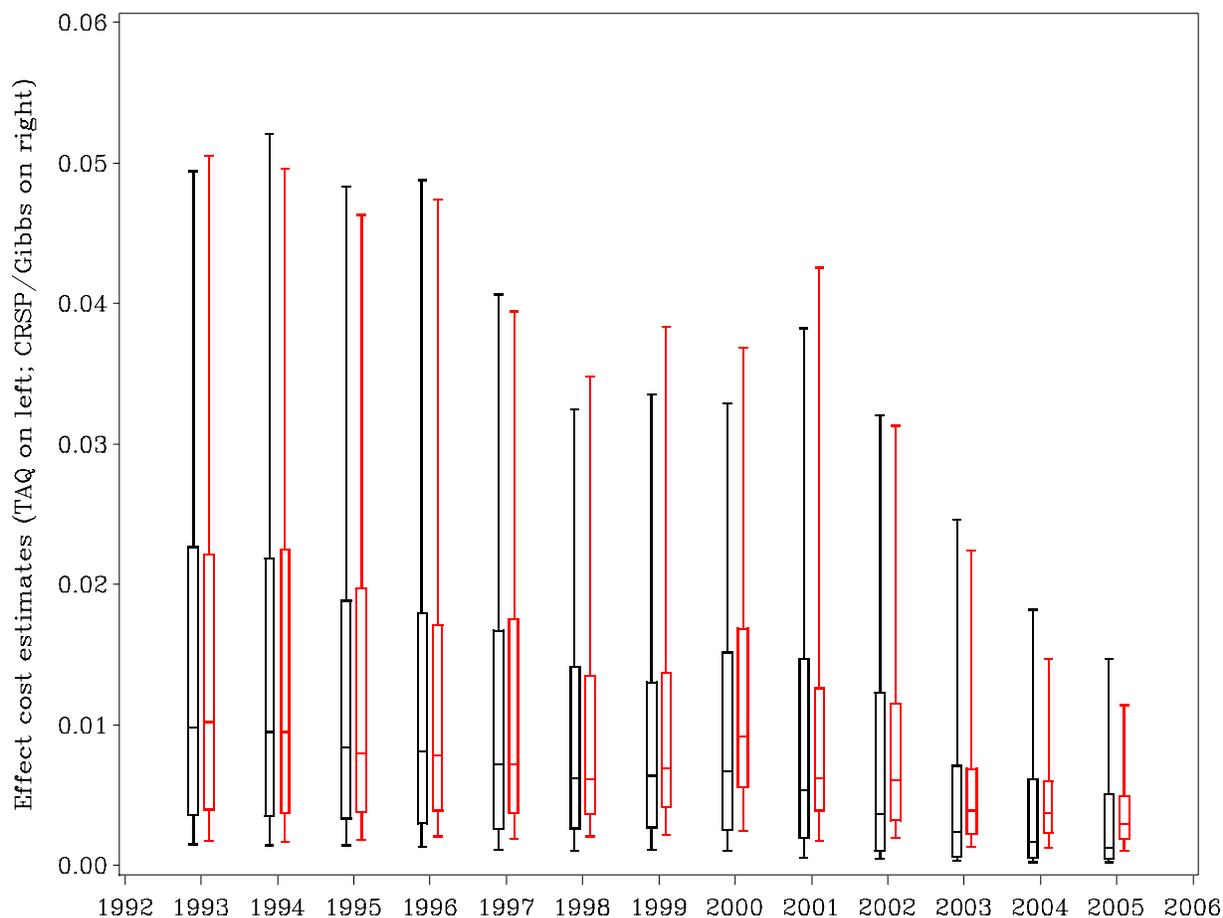


Figure 3. Average effective costs 1926-2005

Average Gibbs effective cost estimates for all ordinary common shares in the CRSP daily database. NYSE/Amex and Nasdaq firms are analyzed separately; subsamples are quartiles based on end-of-year market capitalization. Fama-French NYSE breakpoints are used to construct the samples.

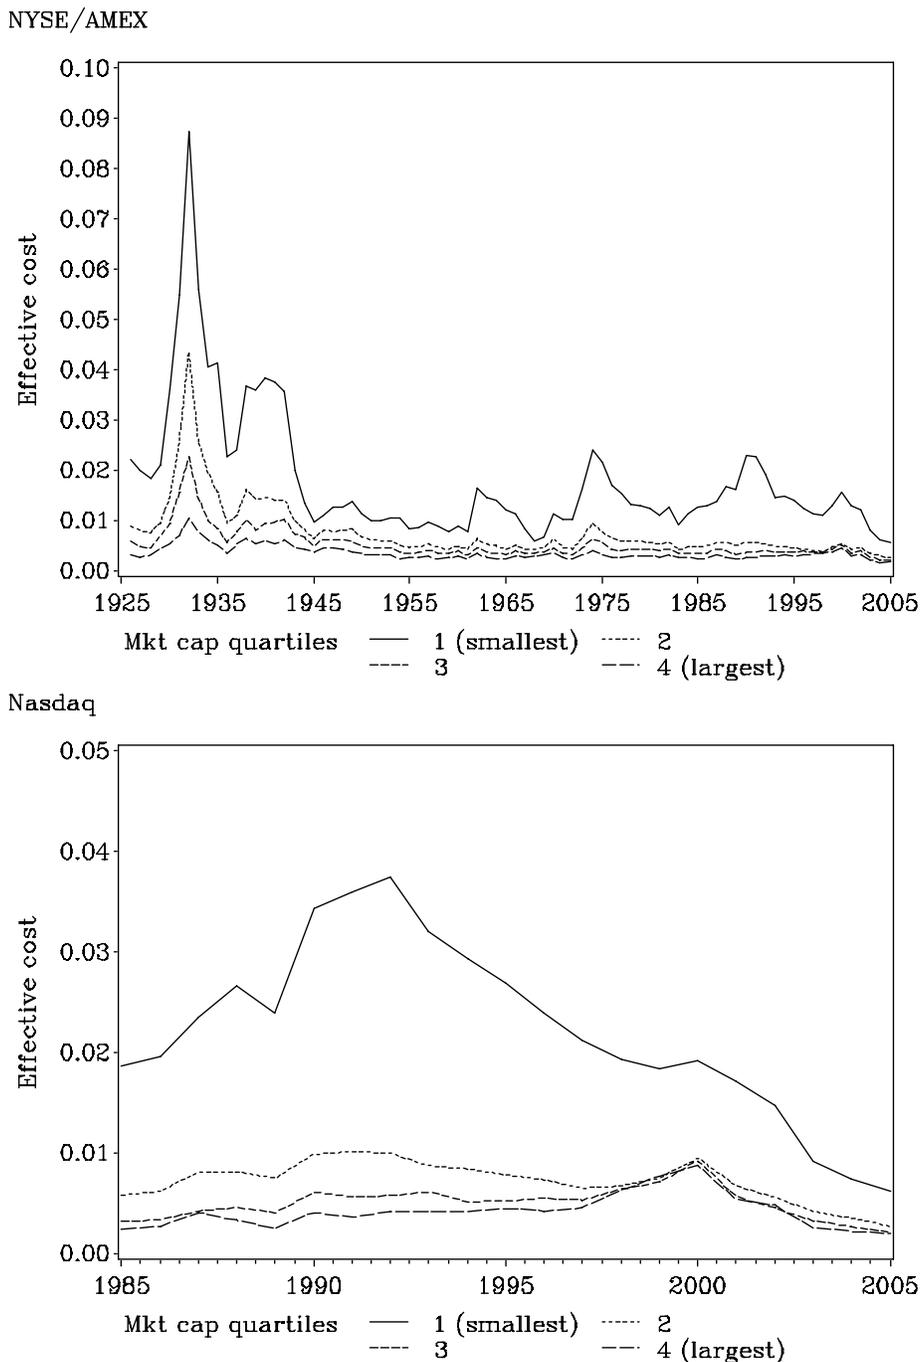


Figure 4. Effective cost commonality in the comparison sample, 1993-2005

The comparison sample consists of approximately 150 Nasdaq firms and 150 NYSE/Amex firms selected in a capitalization-stratified random draw in each of the years 1993-2005. Using TAQ data, the daily effective cost is computed for each firm and averaged weekly (top graph). Using CRSP data, the latent common factor (LCF) model of effective costs is estimated over the panel sample consisting of (approximately) 300 firms in each year. The procedure generates daily estimates for the liquidity factor. The plotted values are weekly averages (bottom graph). Both TAQ and CRSP estimates are standardized (annually) to have zero mean and unit variance.

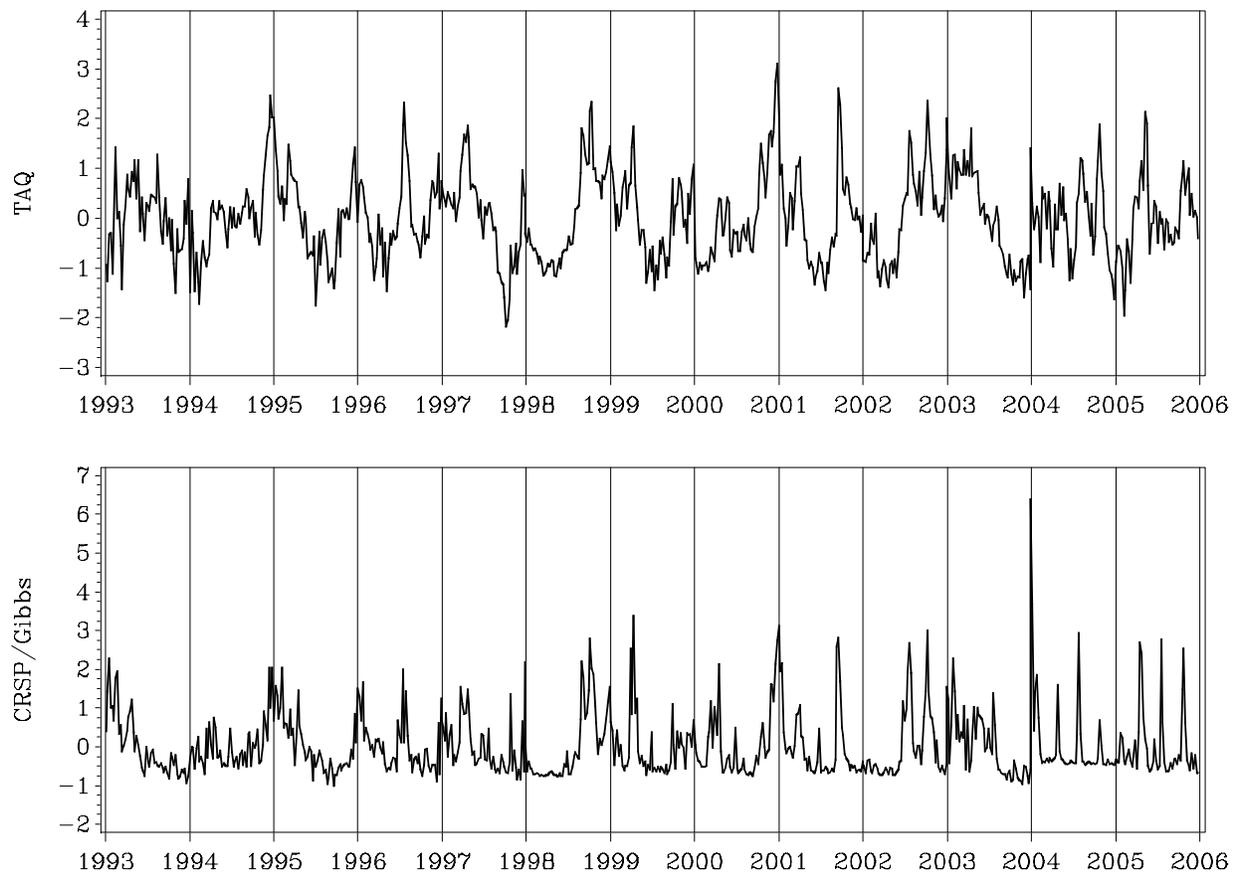


Figure 5. Effective cost common factor, 1926-2005

For each year, 1926-2005, 150 firms are drawn from the CRSP NYSE/Amex firms using capitalization-stratified sampling. In year 1985-2005, this sample is augmented by 150 Nasdaq firms. The latent common factor is estimated for each year over these panels (150 or 300 stocks). The figure depicts monthly averages of the estimated common factor.

