THE JOINT PRICING OF VOLATILITY AND LIQUIDITY*

Federico M. Bandi,[†] Claudia E. Moise,[‡] and Jeffrey R. Russell[§]

November 30, 2008

Abstract

We evaluate the joint pricing of market volatility and market liquidity. Combining recent advances in continuous-time econometrics with no-arbitrage arguments, we extract novel proxies for market volatility and market illiquidity from a single time-series of high-frequency SPIDERS transaction prices. When considered *individually*, illiquidity and volatility shocks are shown to be strongly negatively correlated with market returns as well as with the returns on the size, book-to-market, and momentum portfolios. Shocks to illiquidity and shocks to volatility are found to be individually negatively priced. In *joint* specifications, shocks to volatility drive out shocks to illiquidity leading to a drastic reduction in the statistical significance of the illiquidity factor loadings. While innovations in illiquidity and innovations in volatility may still be jointly negatively priced, as is shown for our data, the factor loadings associated with volatility shocks provide a more accurate assessment of risk. When interpreting shocks to illiquidity and shocks to volatility as proxies for a more fundamental distress factor, this result is suggestive of the superior robustness of the latter.

JEL Classification: G12
Keywords: Market volatility, market illiquidity, cross-section of stock returns, ETFs

---

*We thank Torben Andersen, Mark Grinblatt, John Heaton, Renè Garcia, Abraham Lioui, Florencio López-de-Silanes, George Tauchen and seminar participants at Bank of Japan, Case Western Reserve, Chicago Booth, Chicago FED, Duke Economics, Edhec Business School, the Federal Reserve Board, Kent State, McGill, George Mason, Texas A&M, the Latin American Meetings of the Econometric Society in Mexico City, and the Society for Financial Econometrics Inaugural Conference for helpful comments. We are grateful to Akiko Fujimoto for providing data on some aggregate liquidity proxies. We are also grateful to Gene Fama and Ken French for providing their factors and return data.

[†]Chicago Booth, 5807 South Woodlawn Avenue Chicago, IL 60637, tel: (773) 834-4352, e-mail: federico.bandi@chicagogsb.edu

[‡]Weatherhead School of Management, Case Western Reserve University, Cleveland, OH 44106, tel: (216)-368-4778, fax: (216)-368-6249, e-mail: claudia.moise@case.edu

[§]Chicago Booth, 5807 South Woodlawn Avenue Chicago, IL 60637, tel: (773) 834-0720, e-mail: jeffrey.russell@chicagogsb.edu

1 Introduction

Two successful strands of the recent asset pricing literature have emphasized the importance of market volatility (e.g., Ang et al., 2006, Adrian and Rosenberg, 2006, and Moise, 2006) and market liquidity (e.g., Acharya and Pedersen, 2005, Pástor and Stambaugh, 2003, and Sadka, 2003) as systematic risk factors priced in the cross-section of stock returns. The intuition is straightforward. Aggregate illiquidity and aggregate volatility are high in less favorable states of the world. Assets whose returns are more positively correlated with their innovations will provide a hedge, thereby requiring relatively lower expected returns.[1]

Admittedly, the joint pricing of market volatility and market liquidity has hardly been investigated. Even though they are arguably the result of different economic phenomena (market volatility being related with changes in fundamental asset values and market illiquidity being the outcome of aggregate trading frictions affecting fundamental asset values *if* priced), volatility and illiquidity are positively correlated (in terms of innovations) and relatively higher in negative states of the world. Should one take the view that volatility and illiquidity may be fundamental pricing factors, as sometimes done in the literature, then it is meaningful to ask whether their individual pricing ability is preserved, when jointly considered. Should one believe that they may be proxies for a more fundamental factor (or factors) varying with the state of the economy, an argument we find more convincing, then again it would seem relevant to ask whether their individual explanatory power is subsumed in a model which allows for the other proxy to be present.

We make two contributions to the literature. The first contribution is methodological. By combining recent advances in the econometrics of high-frequency data with classical no-arbitrage arguments, we extract novel proxies for market volatility and market illiquidity from a single time-series of high-frequency SPIDERS transaction prices. The joint evaluation of the cross-sectional pricing implications of (innovations in) market volatility and (innovations in) market illiquidity represents our second contribution.

Because SPIDERS represent ownership of a trust invested in the S&P500 index, changes in SPIDERS fundamental value reflect changes in the index's fundamental value. In addition, since SPIDERS can be redeemed for the underlying portfolio of S&P500 stocks (or created in exchange for the underlying portfolio of assets), deviations of SPIDERS transaction prices from fundamental values signal pervasive market frictions rendering arbitrages harder to implement. We provide a method to separate the volatility of SPIDERS *unobserved* fundamental values (used as a proxy

[1] The cross-sectional relation between expected stock returns and idiosyncratic, rather than systematic, liquidity has been investigated in numerous papers, including Amihud and Mendelson (1986), Brennan and Subrahmanyam (1996), Datar et al. (1998), and Elaswarapu (1997). Amihud (2002), Jones (2002), and Fujimoto (2003), among others, study the time-series properties of market excess returns and market liquidity. The related effect of idiosyncratic asymmetric information in cross-sectional asset pricing is discussed in Easley et al. (2002). Amihud et al. (2006) and Cochrane (2005) provide discussions of the current state of the literature on liquidity and asset prices.

for market volatility) from the volatility of the difference between SPIDERS transaction prices and their fundamental values (used as a proxy for market illiquidity). The method solely requires the computation of averages of high-frequency SPIDERS transaction prices sampled at different, "optimally-selected," frequencies.

When *individually* considered, we show that shocks to illiquidity and shocks to volatility are strongly negatively correlated with market returns, as well as with the returns on the size, book-to-market, and momentum portfolios. Consider the size decile portfolios, for instance. For *both* illiquidity and volatility, these individual correlations (and the related factor loadings) increase monotonically (while remaining negative) when going from small cap stocks to large cap stocks. Hence, large cap stocks are less (negatively) correlated with both illiquidity and volatility then small cap stocks, thereby requiring lower expected returns, as empirically found in practice. Not surprisingly, the individual prices of volatility and illiquidity risk are found to be negative in the cross-section of monthly 25 size- and value-sorted Fama-French portfolios.

When *jointly* considered, shocks to volatility largely drive out the statistical significance of shocks to illiquidity thereby leading to inaccurately estimated factor loadings in the latter case. Ignoring the precision of the first-stage factor loadings, as often done in practise, we find that both of our derived proxies are again negatively priced. This is due to factor loadings (with respect to volatility and illiquidity, jointly) which are negative but largely increase in the size dimension (when going from small cap stocks to large cap stocks) while generally decreasing in the book-to-market dimension (when going from growth stocks to value stocks). For our sample, the performance of a 3-factor model with market returns, (innovations in) market variance, and (innovations in) market illiquidity is shown to be similar to the performance of the classical Fama-French 3-factor model, when pricing the Fama-French portfolios.

In sum, when risk is proxied by innovations in market illiquidity and market volatility *individually* or *jointly*, we find statistically significant (negative) *prices of risk* associated with both measures. The *quantities of risk* (the factor loadings) associated with the two proxies are precisely estimated in individual models. In joint specifications, however, only the volatility factor loadings are estimated accurately thereby leading to a fundamental lack of robustness of illiquidity pricing (even in the presence of statistically-significant illiquidity risk prices) in models allowing for market volatility. Conversely, when interpreting volatility and illiquidity as proxies for a more fundamental distress factor, this result points to an important robustness of market volatility. We stress that such a robustness is not a by-product of the way market illiquidity and market volatility are measured in this paper. The result would hold under more classical measures in the literature, as is shown below.

The remainder of the paper is structured as follows. In Section 2 we discuss the defining features of SPIDERS. Section 3 proposes a formation mechanism for SPIDERS transaction prices.

This mechanism justifies our identification procedure for market volatility and market illiquidity. Section 4 expands on the logic of our illiquidity measure. Section 5 evaluates the empirical properties of both proxies. Particular emphasis is placed on the negative (roughly monotonic) relation between size, book-to-market, and momentum portfolio returns and volatility/illiquidity shocks. The cross-sectional pricing of illiquidity and volatility is discussed in Section 6. Section 7 evaluates robustness. Section 8 concludes. Technical details, figures, and tables are in the Appendix.

2 SPIDERS

Standard & Poor's depository receipts (SPDR or SPIDERS) represent shares in a trust which owns stocks in the same proportion as that found in the S&P500 index. They trade like a stock (with the ticker symbol SPY on the Amex) at approximately one-tenth of the level of the S&P500 index, and are used by large institutions and traders either as bets on the overall direction of the market or as a means of passive management.

SPIDERS are exchange traded funds (ETFs).[2] They can be redeemed for the underlying portfolio of assets *at the end of the trading day*. Equivalently, investors have the right to obtain newly issued SPIDERS shares from the fund company in exchange for a basket of securities that mirrors the SPIDERS' portfolio. This implies that SPIDERS, like other ETFs, must trade at a value that is near net asset value (NAV). If they traded above their NAV, arbitrageurs would purchase the basket of underlying securities for a lower price and force the fund company to issue new shares. Conversely, if they traded below their NAV, arbitrageurs would buy shares and redeem them for the underlying portfolio of securities (see Cherkes et al., 2006, for further discussions and comparisons between ETFs and closed-end funds). Similarly, SPIDERS' values will not deviate much from NAVs[3] *during the day* either, since the future convergence of prices would open up the possibility for simple, immediate investment opportunities. Assume trading prices are higher than NAVs. An arbitrageur could sell SPIDERS short,[4] buy the underlying basket of security, wait for price convergence, and unwind the position for an initial profit.

As Elton et al. (2002) and Engle and Sarkar (2002) point out, the process of share deletion/creation acts as an extremely effective mechanism in keeping prices close to NAV and assuring that potential differences disappear quickly. Conversely, since arbitrages require acquisition of the

[2] A growing academic literature focuses on ETFs. Among other issues, the existing work studies the dynamics of price deviations from net asset value (Engle and Sarkar, 2002), compares the return from holding ETFs (specifically, SPIDERS) to the return from holding the underlying index (Elton et al., 2002), analyzes the tax implications of ETFs (Poterba and Shoven, 2002), investigates price discovery (Hasbrouck, 2002) and competition (Boehmer and Boehmer, 2002) in the ETF market.

[3] The NAV is computed at market close. During the day, an estimated value of the portfolio called Indicative Optimized Portfolio Value (IOPV) is posted. The IOPV is provided every 15 seconds using the most recent transaction price of each component of the portfolio.

[4] Contrary to individual stocks, SPIDERS can be short-sold on a down tick.

underlying basket of securities, the extent of deviations from NAV should signal pervasive market frictions rendering arbitrages harder to implement.

Importantly, rather than focusing on deviations of trade prices from NAV or IOPV, we measure deviations from unobserved *fundamental values*. Given the basket nature of SPIDERS, the fundamental value of a SPIDER share should coincide with the fundamental value of the index. Fundamental values are only approximated by the NAVs at close and by the IOPVs during the day. Measuring deviations from fundamental values rather than from NAV is important as emphasized, for example, by Engle and Sarkar (2002). The NAV is evaluated at the closing transaction price of each of the assets. However, each closing transaction price could be higher or lower than the individual fundamental value. In addition, the closing transaction could occur earlier in the day, especially for less frequently-traded stocks.[5] Similarly, the IOPVs may be stale in that they are posted at equispaced intervals of 15 seconds, are evaluated at trade prices and, hence, likely deviate from fundamental values. More generally, it is well-known from classical market microstructure theory that transaction prices (and, as said, NAVs and IOPVs are computed at transaction values) differ from fundamental values. The size of these deviations will, again, depend on liquidity.[6] Hence, aggregate illiquidity may affect both the size of the deviations between SPIDERS prices and NAVs (given the previous arbitrage arguments) and the size of the deviations between the transaction prices of the underlying securities (which lead to the NAVs and the IOPVs) and unobserved fundamental values. Lower aggregate liquidity may therefore be expected to lead to larger overall deviations. In light of these arguments, the next section will discuss a procedure to identify the size of the deviations between SPIDERS transaction prices and the unobserved fundamental values of the underlying portfolio of security (and, of course, SPIDERS shares).

We employ high-frequency transaction prices on SPIDERS obtained from the Trade and Quote (TAQ) database in CRSP for the period February 1993 - March 2005. We use the entire consolidated market. The cross-sectional asset pricing tests employ monthly return data on the 25 size- and value-sorted Fama-French portfolios over the same period.

We now formalize our assumed high-frequency SPIDERS price formation mechanism and our methodology to separate the volatility of SPIDERS fundamental values from the volatility of SPIDERS deviations from fundamental values, i.e., the factor proxies.

[5] In addition, SPIDERS continue to trade 15 minutes after the NYSE closes. This is another source of error for the posted NAV.

[6] At the stock level and, by aggregation, at the index level, these deviations might also depend on asymmetric information (see, e.g., the discussion in Stoll, 2000, and the references therein).

3 Extracting the proxies

Building on the intuition laid out above, we express the logarithmic SPIDERS transaction price prevailing at the end of a trading day t of length h as

$$\widetilde{p}_{th} = p_{th} + \psi_{th} \qquad t = 1, 2, ..., T, \tag{1}$$

where p is the *unobservable* fundamental value and ψ is an equally *unobservable* price deviation. As said, higher market liquidity should lead to smaller price deviations ψ.

Now divide each trading day into M (equispaced, for notational simplicity) sub-periods. The j-th intra-daily continuously-compounded return between day $t-1$ and day t is defined as

$$\widetilde{r}_{j,t} = \widetilde{p}_{(t-1)h+j\delta} - \widetilde{p}_{(t-1)h+(j-1)\delta} \qquad j = 1, 2, ..., M, \tag{2}$$

where $\delta = h/M$ is the interval over which the intra-daily returns are computed. Thus, similarly to the observed price process, the observed return process comprises a fundamental return component, $r_{j,t} = p_{(t-1)h+j\delta} - p_{(t-1)h+(j-1)\delta}$, as well as a deviation component, $\nu_{j,t} = \psi_{(t-1)h+j\delta} - \psi_{(t-1)h+(j-1)\delta}$, i.e.,

$$\widetilde{r}_{j,t} = r_{j,t} + \nu_{j,t} \qquad t = 1, 2, ..., T, \quad j = 1, 2, ..., M. \tag{3}$$

Our objects of interest are the variance of the unobserved daily fundamental returns $r_t = p_{th} - p_{(t-1)h}$ and the variance of the unobserved intra-daily price deviations $\nu_{j,t}$. The former will proxy for daily market-wide variance, the latter will proxy for daily aggregate illiquidity.

To this extent, for each day in our sample, we compute $\frac{\sum_{j=1}^{M_t} \widetilde{r}_{j,t}^2}{M_t}$ and $\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$, where $\frac{\delta}{M_t}$ is the highest frequency at which intra-daily returns are observed and $\frac{\delta}{M_t^*}$ is an appropriately-chosen optimal frequency. Note that M_t and M_t^* have a subscript t to make their dependence on time fully apparent. Under empirically-reasonable assumptions on the unobservable components r and ν, it can be shown that $\frac{\sum_{j=1}^{M_t} \widetilde{r}_{j,t}^2}{M_t}$ estimates consistently the variance of the intra-daily deviations (i.e., $\mathbf{E}(\nu^2)$) as $M_t \to \infty$ (i.e., for a large number of intra-daily observations). The quantity $\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$ (realized variance) will, in general, not estimate the variance of the unobserved fundamental returns consistently for any M_t^* choice. Appropriate selection of M_t^* (as described in the Appendix) will, however, lead to optimization of the estimator's mean-squared deviations from the object of interest (i.e., fundamental return variance). The chosen number of intra-daily returns M_t^* will, of course, be larger, the smaller the size of the price deviations ν.

While we refer the reader to the Appendix for technical details, here we find it important to briefly emphasize the economic intuition underlying the construction of $\frac{\sum_{j=1}^{M_t} \widetilde{r}_{j,t}^2}{M_t}$ and $\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$. Classical work in market microstructure theory implies that meaningful updates to fundamental prices should be less frequent than meaningful changes in transaction prices (see, e.g., O'Hara, 1994). The former depend on the way informed agents form expectations about future cash flows and hence hinge on potentially infrequent updates to the private information set. The latter depend on the trading process. Because the uninformed agents learn from order flow, non-negligible (discrete) changes in transaction prices may occur regardless of the transaction frequency, i.e., even if trades occur very close in time. This implies that the observed intra-daily returns $\widetilde{r}_{j,t}$ are dominated by the deviation component $\nu_{j,t}$ when the trades are frequent. Conversely, they are largely dominated by the fundamental return component $r_{j,t}$ when return sampling is performed at low frequencies. This simple intuition clarifies why sample second moments of observed returns sampled at the highest available frequency $\frac{\delta}{M_t}$, such as $\frac{\sum_{j=1}^{M_t} \widetilde{r}_{j,t}^2}{M_t}$, identify the second moment of the return deviations ν. This is due to fundamental returns that wash out at high frequencies. By the same type of reasoning, $\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$ will give us information about the (integrated, over the trading day) volatility of the fundamental return process if observed returns are sampled at appropriately-chosen, lower frequencies $\frac{\delta}{M_t^*}$. Importantly, the selection of $\frac{\delta}{M_t^*}$ may be conducted "optimally" based, for example, on a mean-squared error criterion. The Appendix imposes statistical assumptions on $r_{j,t}$ and $\nu_{j,t}$ which make this intuition rigorous while fully justifying the adopted estimators. Details are also provided about the construction of the optimal M_t^*. Alternative approaches and potential extensions are discussed.

In sum, for every day in our sample we use (potentially standardized) sums of observable intra-daily returns sampled at optimal (time-varying) frequencies to identify the variance components of the returns' *unobservable* components, r and ν. These daily measures are subsequently aggregated to the monthly level, as we discuss below.

Importantly, since $\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$ is computed over a 6-hour trading period (from 10 a.m. to 4 p.m.), in order to convert it into a genuinely daily measure, we correct it for lack of overnight returns. We do so by multiplying it by the constant factor $\zeta = \frac{\frac{1}{T}\sum_{t=1}^{T} R_t^2}{\frac{1}{T}\sum_{t=1}^{T} \widehat{V}_t}$, where R_t and $\widehat{V}_t = \sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$ are the daily SPIDERS return and the (6-hour) optimally-sampled realized variance measure over day t. This correction ensures that the average of the corrected variance estimates coincides with the variance of the daily returns (see, e.g., Fleming et al., 2001, 2003).

3.1 *Aggregation*

We evaluate pricing at the monthly level. We therefore average the deviation variances across days in a specific month k to obtain monthly measures:

$$\widehat{\mathbf{E}}_k(\nu^2) = \frac{1}{\#Days} \sum_{t=1}^{\#Days} \frac{\sum_{j=1}^{M_t} \widetilde{r}_{j,t,k}^2}{M_t}. \tag{4}$$

As for the fundamental return variance, we sum the daily realized variances across days in a month to, again, obtain the corresponding monthly values:

$$\widehat{V}_k = \sum_{t=1}^{\#Days} \zeta \sum_{j=1}^{M_t^*} \widetilde{r}_{j,t,k}^2. \tag{5}$$

Finally, we compute innovations. In the price deviation case, we have

$$IFV_k = FV_k - FV_{k-1} = \sqrt{\widehat{\mathbf{E}}_k(\nu^2)} - \sqrt{\widehat{\mathbf{E}}_{k-1}(\nu^2)}. \tag{6}$$

In the fundamental return case, we have

$$IV_k = \sqrt{\widehat{V}_k} - \sqrt{\widehat{V}_{k-1}}. \tag{7}$$

To summarize, we hypothesize that (i) innovations in SPIDERS fundamental price variance reflect innovations in market variance and (ii) innovations in SPIDERS deviation variance reflect changes in aggregate illiquidity. As discussed, both hypothesis are justified by the basket nature of SPIDERS. The latter hypothesis also relies on no-arbitrage arguments. Put it differently, even though SPIDERS trade like any other stock, we expect innovations in SPIDERS price deviations to be a less noisy measure of innovations in the overall market liquidity than innovations in individual stocks' price deviations. As generally argued in the industry (see, e.g., Gastineau, 2001, and Spence, 2002), the liquidity properties of an ETF should reflect the liquidity properties of the underlying portfolio of securities.[7] To this extent, in spite of the elusive nature of liquidity and its various facets,[8] the difference between transaction prices and fundamental values constitutes an extremely natural liquidity measure (Bandi and Russell, 2005). This is the difference which is quantified in this work.

4 More on the logic of the liquidity measure

Using the notation from the previous section, the price deviation (with respect to fundamental value) for trading a share of stock s in the S&P500 basket at time j on day t can be written as

$$\widetilde{p}_{s,j} - p_{s,j} = \psi_{s,j} = (1 - 2I_{sell})\psi_{s,j}^*, \tag{8}$$

[7] On Yahoo Finance, for example, we read: "*Some investors appear to believe that the liquidity of an ETF is dependent on the fund's average trading volume, or the number of shares traded per day. However, this is not the case. Rather, a better measure of ETF liquidity is the liquidity of the underlying stocks in the index.*"

[8] It is typical to define liquidity as "the ability to trade large quantities with small price impacts in a short amount of time."

where $\psi^*_{s,j} \geq 0$ is the cost of buying or selling, the subscript j is short for $(t-1)h + j\delta$, and I_{sell} is a sell indicator taking on the value 0 for a buy order and 1 for a sell order.[9] The price deviation for trading the S&P500 portfolio at time j is a (value-weighted) average of price deviations expressed as

$$\sum_{s=1}^{500} w_s \psi_{s,j} = (1 - 2I_{sell}) \sum_{s=1}^{500} w_s \psi^*_{s,j}, \tag{9}$$

where $0 > w_s > 1$ and $\sum_{s=1}^{500} w_s = 1$. By the nature of SPIDERS, as discussed earlier, the price difference for trading a SPIDERS share at time j, i.e., $\psi_{spy,j}$, should approximately be equal to $\sum_{s=1}^{500} w_s \psi_{s,j}$.[10] Hence, the variance of ψ_{spy}, our object of interest, should roughly represent the variance of the portfolio's price deviations (associated with buy or sell orders) from the portfolio's fundamental value. This variance is small if, of course, individual stocks trade near fundamental values. More explicitly,

$$\mathbf{E}(\psi^2_{spy}) \approx \mathbf{E}\left((1 - 2I_{sell})^2 \left(\sum_{s=1}^{500} w_s \psi^*_s \right)^2 \right) = \mathbf{E}\left(\left(\sum_{s=1}^{500} w_s \psi^*_s \right)^2 \right) = \sum_{s,u=1}^{500} w_s w_u \mathbf{E}(|\psi_s|\,|\psi_u|), \tag{10}$$

if the sell indicator and the individual price deviations ψ^* are independent. Should the price deviations be cross-sectionally *independent and identically distributed*, then

$$\mathbf{E}(\psi^2_{spy}) \approx (\mathbf{E}(|\psi_s|))^2, \tag{11}$$

and our assumed measure would capture the common (across stocks) squared expected absolute price deviation from fundamental value. Should the price deviations be cross-sectionally *independent but not identically distributed*, then

$$\mathbf{E}(\psi^2_{spy}) \approx \left(\sum_{s=1}^{500} w_s \mathbf{E}(|\psi_s|) \right)^2, \tag{12}$$

and the measure would be a squared *weighted average* of expected absolute price deviations from fundamental values. Finally, should the ψs be cross-sectionally *dependent* (as likely the case in the presence of aggregate liquidity shocks) *and not identically distributed*, then

$$\mathbf{E}(\psi^2_{spy}) \approx \left(\sum_{s=1}^{500} w_s \mathbf{E}(|\psi_s|) \right)^2 + \sum_{s,u=1}^{500} w_s w_u \mathbf{Cov}(|\psi_s|\,|\psi_u|). \tag{13}$$

[9] As is customary, we assume that buy orders occur at prices above fundamental values whereas sell orders occur at prices below fundamental values (see, e.g., Roll, 1984).

[10] Here we are, of course, assuming that SPIDERS prices roughly coincide with NAVs (computed at buy or sell prices). Deviations of SPIDERS prices from NAVs would introduce an additional liquidity-related contamination whose contribution is, as said, measurable given our approach (see our comments in the previous section).

Generally speaking, the larger the individual price deviations from fundamental values, the larger the variance of ψ. Importantly, when the price deviations are cross-sectionally correlated, the pair-wise correlations (i.e., the second term on the right-hand side of Eq. (13)) ought to be taken into account. Simply averaging value-weighted firm-specific estimates will underestimate the variance of the aggregate price deviations (and, therefore, the extent of market illiquidity) if these covariances are on average positive (as possibly the case in the presence of aggregate liquidity shocks). This observation, in turn, suggests that first applying the methods to individual stocks (rather than to an index) and subsequently value-weighing the firm-specific estimates might lead to a misleading measure if the ability to buy and sell a large, diversified portfolio near fundamental values is the object of interest, as in our case. In this sense, using an index (and straightforward no-arbitrage reasoning) provides a meaningful solution to the empirical issues that would be posed by the (hardly tractable) computation of moments and cross-moments of individual price deviations for a broad array of stocks using high-frequency data.

5 A look at the factors

5.1 *Market volatility*

Fig. 3 contains 6-hour realized variance estimates ($\sum_{j=1}^{M_t^*} \widetilde{r}_{j,t}^2$) obtained by sampling returns optimally for each day in our sample (Panel a), realized variance estimates constructed using 20-minute intervals (Panel b), and realized variance estimates constructed using 5-minute intervals (Panel c). The 5- and 20-minute frequencies have been widely used in the literature. They are generally employed to reduce the impact of market microstructure noise (ψ, in our notation) on estimates of the fundamental (integrated, over the trading day) return variance. Choosing fixed intervals is, of course, ad-hoc and generally sub-optimal from a finite sample mean-squared-error standpoint (see the Appendix). Not surprisingly, the optimally-sampled realized estimates appear better behaved than the estimates obtained by sampling at fixed intervals. As said, we focus on optimally-sampled (for each day in the sample) realized variances in what follows.

Fig. 5 plots the monthly realized variance estimates constructed using sums of daily realized variance estimates (as in Eq. (5) above) and the monthly variance estimates constructed by summing squared daily returns. The two measures have a correlation of 75%. They both spike during known financial crises such as the Asian crisis (October 1997), the LTCM and Russian debt default (October/November 1998), and so on.

5.2 *Market illiquidity*

Fig. 6 plots the monthly illiquidity estimates FV (constructed using Eq. (4)). The graph suggests a general decline in illiquidity with spikes corresponding, again, to known (il-)liquidity events, such as the Asian crisis, the LTCM collapse and Russian debt default, the 9/11 terrorist attack, and so on. In agreement with what is expected from a proper liquidity measure, the documented decline mirrors well-known downward trends in average bid-ask spreads across stocks. The graph also reflects the increase in liquidity associated with the stock market decimalization (i.e., penny pricing) introduced at the end of 2000, beginning of 2001.

For comparison, in Fig. 7 we report IFV and innovations in the Pástor and Stambaugh liquidity measure over the same period.[11] Pástor and Stambaugh's measure is a price reversal measure. The idea underlying it is that less liquid stocks should have larger price reversals following signed order flow than more liquid stocks. For stock i in month k, liquidity is defined as the least-squares γ estimate from the regression

$$r^e_{i,t+1,k} = \theta_{i,t} + \phi_{i,k} r_{i,t,k} + \gamma_{i,k} sign(r^e_{i,t,k}) v_{i,t,k} + \varepsilon_{i,t+1,k}, \tag{14}$$

where r is a stock return, r^e is an excess stock return, and v is dollar volume. Pástor and Stambaugh expect γ to be negative in general (the price impacts of trades get reversed in the future) and larger in magnitude for less liquid stocks. To construct innovations in aggregate liquidity, they scale the differences in the monthly liquidity measures by relative market size at k and average the differences across stocks with data available in consecutive months, i.e.,

$$\Delta\widehat{\gamma}_k = \left(\frac{m_k}{m_1}\right) \frac{1}{N_k} \sum_{i=1}^{N_k} \left(\gamma_{i,k} - \gamma_{i,k-1}\right). \tag{15}$$

Subsequently, they run the regression

$$\Delta\widehat{\gamma}_k = a + b\Delta\widehat{\gamma}_{k-1} + c\left(\frac{m_{k-1}}{m_k}\right)\widehat{\gamma}_{k-1} + u_k. \tag{16}$$

Finally, innovations in aggregate (il-)liquidity are measured by $PS_k = \frac{u_k}{100}$. The correlation between IFV (our illiquidity proxy) and PS is -0.23 (positive innovations signal possible illiquidity events in our case while negative innovations signal illiquidity events in the case of Pástor and Stambaugh's measure).

The correlation between IFV (PS) and excess market returns is -0.28 (0.23). Similarly, the correlation between IV (innovations in market variance) and excess market returns is -0.41. Subsection 5.3. expands on both observations. Finally, IFV (PS) and IV have a correlation coefficient of 0.57 (-0.31). Thus, market illiquidity and market volatility are higher in times of financial market downturns. In addition, they correlate with each other in potentially important ways. Because they are invariably priced in isolation, both findings justify looking more closely at their individual (and joint) pricing ability.

[11] Data on Pástor and Stambaugh's liquidity measure are downloaded from CRSP.

In our sample, IFV is more highly correlated with SMB (i.e., the difference in returns between small and large firms) and HML (i.e., the difference in returns between high and low book-to-market stocks) than PS (-0.14 and 0.14 versus -0.01 and -0.02). See Table I.

Alternative aggregate liquidity measures have, of course, been proposed. Amihud (2002) recommends using the so-called "illiquidity ratio," ILL. For each stock i and each month k, he computes $\frac{1}{\#Days}\sum_{t=1}^{\#Days}\frac{|r_{i,t,k}|}{v_{i,t,k}}$, where, again, r and v denote stock return and trading volume, respectively. An aggregate measure can then be defined by averaging across stocks for each month k:

$$ILL_k = \frac{1}{N_k}\sum_{i=1}^{N_t}\left(\frac{1}{\#Days}\sum_{t=1}^{\#Days}\frac{|r_{i,t,k}|}{v_{i,t,k}}\right). \tag{17}$$

As earlier, this measure can be re-scaled by $\left(\frac{m_k}{m_1}\right)$. The "illiquidity ratio" looks directly at price impacts. Periods of illiquidity are times when small volumes determine large price moves.

Finally, "share turnover" is sometimes used to quantify aggregate liquidity. For each stock i and each month k, compute $\frac{1}{\#Days}\sum_{t=1}^{\#Days} turn_{i,t,k}$, where $turn$ denotes the ratio between the number of shares transacted and the number of shares outstanding. Aggregate monthly liquidity is then measured by

$$T_k = \frac{1}{N_k}\sum_{i=1}^{N_t}\left(\frac{1}{\#Days}\sum_{t=1}^{\#Days} turn_{i,t,k}\right). \tag{18}$$

Again, this measure can be re-scaled to achieve stationarity. Eckbo and Norli (2002), for example, scale it by a factor $\frac{c_1}{c_t}$ with c_t defined as the 24-month moving average of market turnover (between month $k-24$ and month $k-1$), and c_1 defined as market turnover in the first month in the sample. Fujimoto (2003) contains a thorough discussion of these alternative quantities.

The correlations between IFV and PS, innovations in ILL ($IILL$), and innovations in T (IT) are reported in Table II. Using data between April 93 and December 2002,[12] we find that $IFV, IILL$, and IT correlate with each other in an economically meaningful (and statistically significant) fashion. The correlations between IFV and PS, (un-scaled) $IILL$, and (un-scaled) IT are -0.22, 0.24, and 0.28, respectively. The correlations between IFV, (scaled) $IILL$, and (scaled) IT are 0.21 and 0.26. In our sample, PS is hardly correlated with $IILL$ and IT. The correlation between PS and un-scaled $IILL$ (scaled $IILL$) is -0.06 (-0.02). The correlation between PS and un-scaled IT (scaled IT) is -0.02 (0.002).

5.3 *Illiquidity and volatility premia*

In this subsection we fit simple autoregressive models for V and FV. Subsequently, we regress excess market returns on lagged V and unexpected V (i.e., the residuals from the volatility autoregression).

[12] We thank Akiko Fujimoto for providing the ILL and T data.

We do the same using lagged FV and unexpected FV as regressors. The second-stage regressions are meant to evaluate the presence of market premia associated with expected and unexpected volatility (or liquidity) given a straightforward, but empirically reasonable, dynamic model for the relevant factor proxy.

To this extent, a first-order autoregression of V_t on V_{t-1} gives us

$$V_t = \underset{(4.04)}{0.011} + \underset{(13.18)}{0.74}\, V_{t-1} + \widehat{\varepsilon}_t^V, \tag{19}$$

with an R^2 of 55.7%,while a first-order autoregression of FV_t on FV_{t-1} yields

$$FV_t = \underset{(4.46)}{0.000242} + \underset{(11.77)}{0.71}\, FV_{t-1} + \widehat{\varepsilon}_t^{FV}, \tag{20}$$

with an R^2 of 50.1%. Both volatility and illiquidity are, as expected, highly persistent. A regression of MKT_t (the monthly excess return on the market) on V_{t-1} and the estimated residuals from the previous volatility autoregression gives

$$MKT_t = \underset{(0.43)}{0.35} + \underset{(0.41)}{6.51}\, V_{t-1} - \underset{(-5.51)}{131.05}\widehat{\varepsilon}_t^V + \widehat{u}_t, \tag{21}$$

with an R^2 of 18.2%. Similarly, regressing MKT_t on FV_{t-1} and the estimated residuals from the illiquidity autoregression we obtain

$$MKT_t = \underset{(-0.81)}{-0.80} + \underset{(1.57)}{1729} FV_{t-1} - \underset{(-3.10)}{4691}\, \widehat{\varepsilon}_t^{FV} + \widehat{u}_t, \tag{22}$$

with an R^2 of 7.7%. Finally, the *joint* model yields

$$MKT_t = \underset{(0.11)}{0.133} + \underset{(0.28)}{4.56}\, V_{t-1} + \underset{(0.33)}{363}\, FV_{t-1} - \underset{(-4.2)}{121.13}\widehat{\varepsilon}_t^V - \underset{(-0.57)}{991}\, \widehat{\varepsilon}_t^{FV} + \widehat{u}_t, \tag{23}$$

with an R^2 of 18.4%.

The above results show that (i) market returns are a statistically-significant decreasing function of unexpected volatility and unexpected illiquidity (individually considered), (ii) when volatility and illiquidity are evaluated jointly, unexpected illiquidity is dominated by unexpected volatility (i.e., the latter remains strongly statistically-significant), (iii) both expected liquidity and expected illiquidity hardly affect excess market returns. This last result is of course consistent with the well-established inability to find robust risk-return trade-offs for the market at the monthly frequency.

Regressions (20) and (22) are in the spirit of Amihud (2002). He finds that market returns are an increasing function of expected illiquidity (measured by virtue of (17)) and a decreasing function of unexpected illiquidity. These findings may be justified. Being persistent, higher current illiquidity will translate into higher expected illiquidity. If higher expected illiquidity yields higher expected returns, then higher unexpected illiquidity should lead to a drop in prices and, hence, lower realized returns.

Our results are qualitatively similar to those in Amihud (2002). However, the statistical significance of the positive coefficient on FV_{t-1} is, in our case, lower indicating that unexpected illiquidity has a more significant (negative) effect on realized stock market returns than expected illiquidity. Using several aggregate liquidity measures in the literature, including PS and T, Fujimoto (2003) argues in favor of the same conclusion.

We show that, when allowing for unexpected volatility, the statistical significance of unexpected illiquidity decreases drastically. Importantly, this result does not hinge on our assumed illiquidity measure. If we were to employ PS, the slope coefficient from a regression of MKT on PS would be 22 (t-stat = 3.3), the sign being positive in light of the interpretation of PS, as discussed above. If we regressed MKT on PS and unexpected volatility, the coefficient on PS would be 9.8 (t-stat = 1.44), whereas the coefficient on unexpected volatility would continue to be negative and highly statistically significant, -115.5 (t-stat = -4.46).

We now run the same regressions using excess returns on the Fama-French size-sorted decile portfolios. Consider first regressions of the excess portfolio returns on expected and unexpected market volatility (Table III). The coefficients on V_{t-1} are positive but insignificant across portfolios. The coefficients on unexpected volatility, on the other hand, have a negative sign, as in the market case, and are consistently highly statistically significant. More importantly for our purposes, these coefficients decrease (in absolute value) monotonically when going from small cap stocks to large cap stocks. In other words, smaller stocks have more exposure to unexpected volatility risk than larger stocks. This exposure may of course be priced in equilibrium, as we show in the next section. Implementing the same regressions on expected and unexpected illiquidity yields similar results (Table IV). The absolute values of the coefficients on unexpected illiquidity decrease monotonically in the size dimension (again, going from small stocks to large stocks). As earlier in the market case, however, joint consideration of unexpected illiquidity and unexpected volatility (in Table V) leads to parameter estimates on unexpected volatility which remain statistically significant while driving out the statistical significance of the parameter estimates associated with unexpected illiquidity. Having made this point, the effect of the latter on excess portfolio returns remains completely monotonic with smaller stocks being more (negatively) correlated with illiquidity changes than larger stocks. Thus, while the exposures of the size portfolios' returns with respect to illiquidity are such that illiquidity may be priced (with a negative sign), these exposures are poorly estimated in a joint volatility/illiquidity specification.

We conduct the same exercise with the book-to-market decile portfolios. We again find that the loadings of unexpected volatility on excess returns vary roughly monotonically with value, i.e., going from low book-to-market stocks to high book-to-market stocks (see Table VI). However, these loadings decrease (become more negative). In other words, value stocks appear to be more correlated with unexpected market volatility than growth stocks. This pattern may, of course, provide a justification for the well-known value premium: value stocks provide higher average returns since their payoffs are lower in high volatility states. Differently from unexpected volatility, the loadings with respect to unexpected illiquidity increase with value (Table VII). Again, joint consideration of unexpected volatility and unexpected illiquidity tends to drive out the latter in terms of statistical significance (Table VIII). This said, the loadings on volatility changes remain decreasing and significant, while the loadings on illiquidity changes are roughly increasing and insignificant.

Implementing the same regressions using momentum decile portfolios would yield similar findings.[13] Again, when volatility and illiquidity are individually considered, the correlation between momentum portfolios' returns and the two risk proxies is negative and decreasing with momentum. In a joint specification, the illiquidity and volatility exposures follow the same pattern across momentum portfolios but only the volatility exposures are statistically significant, thereby again underlining the superior robustness of the pricing results associated with the volatility proxy. Once more, while illiquidity and volatility may be priced jointly with a negative sign only the volatility factor loadings appear to be fully informative about the extent of distress risk (when, of course, interpreting illiquidity and volatility as proxies for distress risk).

We emphasize that none of these results depends on our proposed illiquidity proxy. Using the PS measure would again yield factor loadings on illiquidity which lose their statistically significance when unexpected volatility is added to the model. It would also, again, lead to factor loadings with respect to unexpected illiquidity which increase with both size and value (thereby calling for a positive risk premium in the value dimension and a negative risk premium in the size dimension, which is of course problematic when pricing the size and book-to-market portfolios jointly, as done in the next section).

In sum, (i) expected illiquidity and expected volatility have a positive, but statistically-insignificant, effect on market and portfolio returns. If there is a premium for expected illiquidity, such a premium is as hard to detect as the more classical premium for expected volatility risk. However, *(ii)* unexpected illiquidity and unexpected volatility are strongly negatively correlated with market and portfolio returns. These correlations are always highly statistically significant for volatility (i.e., in individual and joint models). In the case of illiquidity, they are highly statistically significant only in individual specifications (i.e., when not accounting for volatility). If shocks to volatility are added to the specification, these shocks tend to subsume the information content of illiquidity shocks. Importantly, (iii) with the sole exception of the unexpected illiquidity loadings in the book-to-market case, the exposures (loadings) with respect to illiquidity surprises and volatility surprises

[13] These findings are not reported for brevity. However, they are available from the authors upon request.

align with the returns on the size and book-to-market portfolios thereby suggesting the potential for *negative* illiquidity and volatility risk prices (regardless of whether volatility and illiquidity are evaluated individually or jointly). We now turn to this issue.

6 Cross-sectional pricing

We consider a traditional intertemporal asset-pricing model as in Merton (1973). Denote excess returns on a generic asset i by R_i^e and excess returns on the market by R_m^e. Assume existence of p state variables F_s. Equilibrium expected excess returns are, as always, expressed as linear combinations of the beta of the asset returns with the market return, β_i^m, and the betas of the asset returns with the state variables, β_i^s, namely

$$E(R_i^e) = \lambda_m \beta_i^m + \sum_{s=1}^{p} \lambda_s \beta_i^s. \tag{24}$$

The lambdas have the usual interpretation in terms of prices of risk. Specifically, λ_m is the price of market risk and λ_s is the price of risk associated with the generic factor s.

As typically done in the literature, we employ the Fama-French 25 size- and value-sorted portfolios as test assets (Fama and French, 1992, 1993) and focus on 2- and 3-factor models. In the 2-factor model case, $p = 1$ and the two factors are *MKT* and either volatility measure (*IV* or *IFV*). In the 3-factor model case, $p = 2$ and the three factors are *MKT*, *IV* and *IFV*.

The estimation method consists of two steps. In the first step, we estimate the betas (the factor loadings) for each portfolio by fitting a linear regression model with AR(1)-GARCH(1,1) errors to the time series of each portfolio's excess returns as in Moise (2006). The model is:

$$\begin{array}{l} R_{i,t}^e = \alpha_i + \beta_i^m R_{m,t}^e + \sum_{s=1}^{p} \beta_i^s F_{s,t} + \eta_{i,t} \quad i = 1, ..., 25, \quad t = 1, ..., T \\ \eta_{i,t} = \varepsilon_{i,t} - \gamma_i \eta_{i,t-1}, \\ \varepsilon_{i,t} = \sqrt{h_{i,t}} e_{i,t}, \\ h_{i,t} = \omega_i + \phi_i h_{i,t-1} + \theta_i \varepsilon_{i,t-1}^2, \\ e_{i,t} \sim N(0,1). \end{array} \tag{25}$$

Given the factor loadings, the prices of risk are estimated in the second step by regressing cross-sectionally the portfolios' average excess returns on the factor loadings, as implied by Eq. (24).

6.1 *The pricing of illiquidity and volatility risk*

Figs. 8 and 9 plot the monthly average excess returns of the 25 Fama-French portfolios as well as the factor loadings obtained from 2-factor models (*MKT* and *IV*, *MKT* and *IFV*). Fig. 10 gives the same plots in the 3-factor case (*MKT*, *IV,* and *IFV*). The average excess returns have a familiar pattern: they largely increase in the growth-value dimension (i.e., going from low book-to-market to high book-to-market stocks) and decrease in the size dimension (i.e., going from small cap stocks to large cap stocks). In the 2-factor models and in the 3-factor model, both the IV and the IFV factor loadings decrease with value and increase with size, albeit sometimes not monotonically. The inverse relation between size and factor loadings is generally stronger than the inverse relation between value and factor loadings. This is, of course, fully consistent with our results in Subsection 5.3.

As discussed earlier, because the relation between excess returns and factor loadings is largely negative across size- and value-sorted portfolio, risk may be priced with a negative sign. We test this implication by regressing average excess returns on factor loadings as described in the previous section. Table IX contains the results. In a 2-factor model with MKT and IV, both variables have significant risk prices. The corresponding t-stats are 4 and -2.54. The estimated (yearly) small-minus-big IV risk premia $\left(\widehat{\lambda}_{IV}\left(\widehat{\beta}^{IV}_{Small}-\widehat{\beta}^{IV}_{Big}\right)\right)$ for stocks in the 2nd, 3rd, and 4th book-to-market quintiles are 2.8%, 1%, and 3.1% (Table XI, Panel a). In a 2-factor model with MKT and IFV, the prices of risk are again statistically significant with t-stats equal to 4.05 and -3.45. The estimated (yearly) small-minus-big IFV risk premia $\left(\widehat{\lambda}_{IV}\left(\widehat{\beta}^{IFV}_{Small}-\widehat{\beta}^{IFV}_{Big}\right)\right)$ for stocks in the 2nd, 3rd, and 4th book-to-market quintile are equal to 6.8%, 4.8%, and 8.28% (Table XI, Panel b). When combining MKT, IV, and IFV in a 3-factor model, we again find that all prices of risk are statistically significant with t-stats equal to 4.05, -2.22, and -2.69, respectively. Across the 2nd, 3rd, and 4th book-to-market quintiles, the IV and IFV small-minus-big risk premia are now equal to 1.9%, 0.22%, and 0.94%, and 2.15%, 2%, and 5.52%, respectively (Table XI, Panel c). In sum, for our sample, market illiquidity carries a price of risk which is more statistically significant than the price of risk associated with market volatility. The corresponding risk premia are also larger. Having said this, the factor loadings are, again, more accurately estimated in the market volatility case. Specifically, even though controlling for excess market returns (as done when computing the factor betas in this section) lowers the statistical significant of the loadings associated with both volatility and illiquidity, the volatility betas are generally statistically significant. The illiquidity betas are virtually always insignificant.

Fig. 11 displays the pricing errors. We plot realized average excess returns versus predicted average excess returns for three models, i.e., the CAPM, the Fama-French 3-factor model, and a 3-factor model with MKT, IV, and IFV. As is well-known, the Fama-French 3-factor model strongly dominates the CAPM. Interestingly, we find that a 3-factor model with excess market returns, market volatility, and market illiquidity compares rather favorably with the Fama-French 3-factor model when pricing the cross-section of size- and value-sorted Fama-French portfolios. Not surprisingly, some of the growth portfolios (specifically, the portfolios in the first value quintile and first, second, and third size quintile) are exceptions. This is, of course, easy to explain. The growth portfolios have realized average excess returns which increase, rather than decreasing, across size quintiles (see Fig. 10, for instance). The corresponding IV and IFV loadings increase too. In other words, the become less negative. The combination of negative prices of risk with increasing loadings leads to a divergence between realized average excess returns and excess returns implied by the model: the realized excess returns on these portfolios are excessively low relative to average returns implied by the model.

7 Robustness

We evaluate the robustness of our previous findings to (*i*) the use of an alternative volatility measure, (*ii*) the use of an alternative illiquidity measure, and (*iii*) the use of daily frequencies.

7.1 *Low-frequency variance*

We expect monthly volatility measures computed by aggregating intra-daily continuously-compounded squared returns to be considerably less noisy estimates than monthly measures obtained by aggregating squared daily returns. While this is true, in general, it is meaningful to ask whether the reported pricing results hinge on the use of a fairly efficient volatility proxy, as in this paper. We find that this is not the case. Put it differently, all of our findings having to do with individual and joint pricing hold in the presence of a more traditional volatility measure. The volatility factor loadings continue to be accurately estimated (and more accurately estimated than in the illiquidity case). Volatility continues to be priced with a negative sign whether considered in isolation or jointly with our assumed liquidity proxy. Similarly, the illiquidity risk price remains significant, and of comparable magnitude, in the trivariate model.[14]

7.2 *Low-frequency illiquidity*

We previously emphasized that the use of PS as an illiquidity proxy would deliver illiquidity factor loadings for the size- and value-sorted portfolios which fully mimic the behavior of the illiquidity factor loadings obtained by virtue of IFV. The resulting pricing is consistent with this observation. In a bivariate model, PS is priced with a positive sign, which is of course coherent with the interpretation of the measure (Table IX). When controlling for either market volatility or market illiquidity, proxied by IFV, PS becomes insignificant. We do not attribute too much importance to the superior statistical significance of IFV in comparison with PS. While this finding may of course be sample-specific, the relative performance of the two illiquidity measures in relation with market volatility provides further evidence for the importance (and robustness) of market volatility in a model that allows for illiquidity pricing at the monthly level (see, also, Moise, 2006).

[14] For brevity, the corresponding results are not reported. They are available from the authors upon request.

7.3 *Daily frequencies*

As in the monthly case, the portfolios' excess returns largely increase in the value dimension and decrease in the size dimension. Interestingly, contrary to the monthly case, the IV factor loadings increase fairly monotonically with both value and size. The size effect, however, appears stronger, thereby confirming the relation between size and volatility loadings documented by Moise (2006). Table X contains the corresponding prices of risk. The t-stats associated with the market price of risk and the volatility risk are equal to 1.51 and -2.64. Thus, despite the positive relation between value and volatility loadings, and in light of the stronger negative relation between size and volatility loadings, we find a negative price of volatility risk at daily frequency.

Turning to a 2-factor model with MKT and IFV, the IFV factor loadings increase with size and decrease with value, albeit not monotonically. The resulting effect is a strongly significant IFV price of risk in the daily cross-section of Fama-French portfolios. The t-stat for the illiquidity price of risk is equal to -3.54 whereas the market price of risk carries a t-stat equal to 1.08.

When combining MKT, IV and IFV in a 3-factor model, illiquidity largely subsumes the information contained in volatility in terms of resulting pricing. The price of volatility risk becomes insignificant (t-stat of -0.78) while the t-stat associated with the price of illiquidity risk is now -3.82. As in the monthly case, this result completely ignores estimation accuracy associated with the illiquidity and volatility risk exposures. In the joint model, the illiquidity factor loadings nicely align with the portfolios average returns, hence the statistical significance of the resulting prices of illiquidity risk. However, contrary to the univariate model, the illiquidity loadings lose their statistical significance when allowing for exposures to volatility shocks. Thus, even at daily frequencies, the pricing of illiquidity risk is still arguably less robust than the pricing of volatility risk. In the joint model, the significance of the final risk prices is therefore to be taken with caution.

8 Conclusions

Market volatility and market illiquidity have received much attention in recent work on cross-sectional asset pricing. In spite of the substantial correlation between macro volatility and macro illiquidity events, their joint pricing has, however, seldom been considered. We investigate the joint pricing of volatility and illiquidity by using novel proxies extracted from high-frequency SPIDERS transaction data. In particular, aggregate illiquidity is measured by the volatility of the difference between *observed* SPIDERS prices and *unobserved* SPIDERS fundamental values. Market volatility is measured by the volatility of *unobserved* SPIDERS fundamental values. Reliance on high-frequency data makes our market volatility proxy more efficient than proxies filtered from low frequency market returns as in the current asset pricing literature. Similarly, using data sampled at high frequencies and treating equilibrium prices as unobservable has the potential to lead to a more effective quantification of the extent of market frictions (see, e.g., Bandi and Russell, 2005, and the references therein).

We show that innovations in our derived illiquidity proxy correlate with macro illiquidity events in important ways. We also show that, when jointly considered in the context of classical asset pricing paradigms, market volatility and market illiquidity are negatively priced in the cross-section of stock returns. In our sample, the performance of a 3-factor model with market returns, (innovations in) market volatility, and (innovations in) market illiquidity appears to be similar to the performance of the Fama-French 3-factor model when pricing the Fama-French size- and book-to-market-sorted portfolios.

This favorable pricing result is due to return exposures with respect to illiquidity and volatility shocks which align in meaningful ways with observed mean portfolio returns. In particular, portfolios with higher average returns generally correlate more negatively with innovations in volatility and illiquidity. These portfolios pay off less in adverse states of the world, thereby requiring compensation for risk.

Liquidity and volatility are admittedly more fundamental economic variables than portfolio returns, like *SMB* and *HML*, routinely used as systematic factors in the literature. However, they are still likely to be proxies for a more fundamental distress factor. While our results point to the joint pricing ability of correlated risk proxies previously analyzed in isolation, we emphasize that exposures to volatility shocks are more robustly estimated than exposures to illiquidity shocks. Hence, volatility pricing is arguably more robust than liquidity pricing. In particular, as shown, the information contained in illiquidity shocks is largely subsumed by that contained in volatility shocks in joint specifications. Put it differently, once one controls for volatility shocks, the significance of the illiquidity exposures decreases substantially. While this result is robust across measures and does not depend on our assumed illiquidity (or volatility) proxy, it might still hinge on the fact that current illiquidity measures do not capture all relevant facets of illiquidity. It may also be a genuine economic fact speaking to the superior ability of market volatility as a proxy for the type of distress risk investors hedge against.

9 Appendix

This Appendix provides a discussion of the assumptions imposed on the fundamental return process r and price deviations ν justifying our identification approach.

As in much recent work in high-frequency econometrics (see, e.g., the discussions in the review papers of Bandi and Russell, 2008b, and Barndorff-Nielsen and Shephard, 2007), we model the fundamental return process as a stochastic volatility martingale difference sequence driven by Brownian shocks, i.e.,

$$r_{j,t} = \int_{(t-1)h+(j-1)\delta}^{(t-1)h+j\delta} \sigma_s dW_s, \tag{26}$$

where σ_s is a càdlàg spot volatility process bounded away from zero. Spot volatility can therefore display jumps, diurnal effects, high-persistence (possibly of the long memory type), and nonstationarities. The price deviations ψ are assumed to be independent of the fundamental prices and i.i.d. with a bounded fourth moment. In consequence, the return deviations ν follow an MA(1) model with a negative first-order autocovariance equal to $-2\mathbf{E}(\psi^2)$. Since the fundamental returns are uncorrelated, the structure of the return deviations carries over to the observed high-frequency returns. The empirical autocorrelation properties of the high-frequency SPIDERS returns strongly supports this simple specification (see Fig. 1).

Importantly, the deviations ψ are modeled as having a stochastic order of magnitude $O_p(1)$. Price discreteness, as well as the existence of different prices for buyers and sellers, for instance, justify this assumption (Bandi and Russell, 2006). Thus, the friction returns ν do not vanish at high sampling frequencies. Technically, $\nu = O_p(1)$. Differently from the deviation returns, the fundamental return process r has an order of magnitude $O_p(\sqrt{\delta})$ over any sub-interval of length δ, thereby implying that the magnitude of the fundamental price changes decreases with the sampling interval. This assumption, which is standard in continuous-time asset pricing, represents slow accumulation and processing of information leading to negligible fundamental price updates over small time intervals.

Exploiting the different stochastic orders of ν and r, Bandi and Russell (2006, 2008) show that sample moments of high-frequency return data sampled at the highest frequencies at which information arrives identify the deviation moments. Availability of high sampling frequencies is represented here by an asymptotic design which lets the distance between observations δ go to zero in the limit or, equivalently, lets the number of observations M go off to infinity for every trading day. In the second moment case, one obtains

$$\frac{\sum_{j=1}^{M} \widetilde{r}_{j,t}^{2}}{M} \underset{M\to\infty}{\overset{p}{\to}} \mathbf{E}_t(\nu^2) = 2\mathbf{E}_t(\psi^2). \tag{27}$$

The result in Eq. (27) hinges on the fact that the deviation process dominates the fundamental return process at high frequencies. More explicitly, when computing sample moments of the observed return data, the fundamental return component r washes out asymptotically since its stochastic order, $O_p(\sqrt{\delta})$, is smaller than the stochastic order of the frictions ν, $O_p(1)$. Hence, the moments of the observed returns consistently estimate the deviation moments at high frequencies. Importantly, this consistency result would not be affected by the presence of infrequent news arrivals leading to discrete changes in the fundamental price process. In other words, one could easily allow for the presence of a Poisson jump component in the fundamental prices or returns r (see Bandi and Russell, 2005). The estimator, and its consistency properties, would not change. These arguments justify using the estimator in Eq. (27). Because we employ a large number of high-frequency returns per trading day (the average number of intra-daily returns is about 3,000) we expect the consistency result in Eq. (27) to be fairly accurate and the corresponding estimator to be informative.

We now turn to the variance of the fundamental return process, i.e., $V_t = \int_{t-1}^{t} \sigma_s^2 ds$. In the absence of deviations, the sum of the squared intra-daily returns $\sum_{j=1}^{M} \widetilde{r}_{j,t}^{2}$ (realized variance) estimates V_t consistently as $M \to \infty$ (see, e.g., Andersen et al., 2003, and Barndorff-Nielsen and Shephard, 2002). The presence of deviations leads to an important bias-variance trade-off. High sampling frequencies may determine substantial noise accumulation and biased estimates. Low sampling frequencies may lead to (fairly) unbiased but highly volatile estimates. Bandi and Russell (2006, 2008) provide a simple rule-of-thumb to optimize this trade-off and choose the optimal sampling frequency δ (or, equivalently, the optimal number of observations M) for every horizon of interest. For each day in our sample, given the assumed price formation mechanism, the (approximate) optimal number of observations M_t^* may be defined as

$$M_t^* \approx \left(\frac{\widehat{Q}_t}{\widehat{\alpha}_t}\right)^{1/3}, \tag{28}$$

where $\widehat{Q}_t$ is equal to $\frac{\widetilde{M}}{3}\sum_{j=1}^{\widetilde{M}} \widetilde{r}_{j,t}^4$, the quarticity estimator of Barndorff-Nielsen and Shephard (2002) with returns sampled every 15 minutes, and $\widehat{\alpha}_t$ is equal to $\left(\frac{\sum_{j=1}^{M} \widetilde{r}_{j,t}^2}{M}\right)^2$, the friction-in-returns second moment estimator raised to the second power. The optimal number of observations M_t^* can be interpreted as a signal-to-noise ratio. The higher the signal coming from the underlying fundamental price process ($\widehat{Q}_t$ estimates $\int_{t-1}^{t} \sigma_s^4 ds$) relative to the size of the frictions (as represented by $\widehat{\alpha}_t$), the higher the optimal number of high-frequency observations needed for realized variance estimation.

Fig. 2 (Panel a) reports a histogram of the optimal sampling intervals and corresponding descriptive statistics. The average interval is about 29 minutes, the median value is about 14 minutes. Fig. 2 (Panel b) presents a time-series plot of the optimal sampling intervals. The intervals display an obvious downward trend. This trend is due to deviation second moment estimates being relatively higher in the first part of the sample (see Fig. 4). According to the ratio in Eq. (28), in order to achieve deviation reduction, a higher relative deviation component should lead to a smaller optimal number of return observations and, thus, a lower optimal sampling frequency.

Identification of both the deviations second moments and the return variance can be generalized. We could allow for virtually unrestricted dependence in the frictions along similar lines as in Bandi and Russell (2005). As discussed above, this extension is empirically unimportant for our data given its clear MA(1) structure. Consistent (in the presence of price deviations) estimates of the fundamental return variance may be obtained by using kernel estimators such as those proposed by Zhang et al. (2005) and Barndorff-Nielsen et al. (2008). Experimentation with these alternative estimators did not lead to different results.

References

[1] Acharya, V.V., and L.H. Pedersen (2005). Asset pricing with liquidity risk. Journal of Financial Economics 77, 375-410.

[2] Adrian, T., and J. Rosenberg (2006). Stock returns and volatility: pricing the short-run and long-run components of market risk. Journal of Finance, forthcoming.

[3] Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31-56.

[4] Amihud, Y., and H. Mendelson (1980). Dealership markets: market making with inventory. Journal of Financial Economics 8, 31-53.

[5] Amihud, Y. and H. Mendelson (1986). Asset pricing and the bid-ask spread. Journal of Financial Economics 17, 223-249.

[6] Amihud, Y., H. Mendelson, and L.H. Pedersen (2005). Liquidity and asset prices. Foundations and trends in Finance 1, 269-264. NOW publishers.

[7] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. Econometrica 71, 579-625.

[8] Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. Journal of Finance 61, 259-299.

[9] Bandi, F.M., and J.R. Russell (2005). Full-information transaction costs. *Working paper.*

[10] Bandi, F.M., and J.R. Russell (2006). Separating microstructure noise from volatility. Journal of Financial Economics 79, 655-692.

[11] Bandi, F.M., and J.R. Russell (2008). Microstructure noise, realized variance, and optimal sampling. Review of Economic Studies 75, 339-369.

[12] Bandi, F.M., and J.R. Russell (2008b). Volatility. In J. R. Birge and Vadim Linetsky (eds.), *Handbook of Financial Engineering.* Amsterdam: North-Holland.

[13] Barndorff-Nielsen, O.E., and N. Shephard (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. Journal of the Royal Statistical Society, Series B, 64, 253-280.

[14] Barndorff-Nielsen, O.E., and N. Shephard (2007). Variation, jumps, market frictions and high-frequency data in financial econometrics. In R. Blundell, P. Torsten and W. K. Newey (Eds.)

Advances in Economics and Econometrics. Theory and Applications. Ninth World Congress. Econometric Society Monographs, Cambridge University Press.

[15] Barndorff-Nielsen, O.E., P. Hansen, A. Lunde and N. Shephard (2008). Designing realized kernels to measure ex-post variation of equity prices in the presence of noise. Econometrica, forthcoming.

[16] Boehemer, B., and E. Boehemer (2002). Trading your neighbour's ETFs: competition or fragmentation? *Working paper.*

[17] Brennan, M.J., and A. Subrahmanyam (1996). Market microstructure and asset pricing: on the compensation for illiquidity in stock returns. Journal of Financial Economics 41, 441-464.

[18] Cherkes, M., J. Sagi, and R. Stanton (2006). A liquidity-based theory of closed-end funds. *Working paper.*

[19] Cochrane, J.H. (2005). NBER asset pricing program review: liquidity, trading, and asset prices. *Working paper.*

[20] Datar, V.T., N.Y. Naik, and R. Radcliffe (1998). Liquidity and stock returns: an alternative test. Journal of Financial Markets 1, 203-219.

[21] Eckbo, E.B., and O. Norli (2002). Pervasive liquidity risk. *Working paper.*

[22] Eleswarapu, V.R. (1997). Cost of transacting and expected returns in the NASDAQ market. Journal of Finance 52, 2113-2127.

[23] Elton, E.J., M.J. Gruber, G. Comer, and K. Li (2002). Spiders: where are the bugs? Journal of Business 75, 453-472.

[24] Engle, R.E., and D. Sarkar (2002). Pricing exchange traded funds. *Working paper.*

[25] Fama, E., and J. MacBeth (1973). Risk, return, and equilibrium: empirical tests. Journal of Political Economy 81, 607-636.

[26] Fama, E., and K. French (1992). The cross-section of expected stock returns. Journal of Finance 47, 427-465.

[27] Fama, E., and K. French (1993). Common risk factors in returns on stocks and bonds. Journal of Financial Economics 33, 3-56.

[28] Fleming, J., C. Kirby, and B. Ostdiek (2001). The economic value of volatility timing. Journal of Finance, 56, 329-352.

[29] Fleming, J., C. Kirby, and B. Ostdiek (2003). The economic value of volatility timing using realized volatility. Journal of Financial Economics, 67, 473-509.

[30] Fujimoto, A. (2003). Liquidity and expected market returns: an alternative test. Working paper.

[31] Gastineau, G.L. (2001). Exchange-traded funds: an introduction. Journal of Portfolio Management 27, 88-96.

[32] Hasbrouck, J. (2002). Intraday price formation in US equity index markets. *Working paper.*

[33] Jones, C. (2002). A century of stock market liquidity and trading costs. *Working paper.*

[34] Merton, R.C. (1973). An intertemporal capital asset pricing model. Econometrica 41, 867-887.

[35] Moise (2006). Stochastic volatility risk and the size anomaly. *Working paper.*

[36] O'Hara (1994). *Market microstructure theory*. Blackwell.

[37] Pastor, L., and R. Stambaugh (2003). Liquidity risk and expected stock returns. Journal of Political Economy 111, 642–685.

[38] Poterba, J.M., and J.B. Shoven (2002). Exchange traded funds: a new investment option for taxable investors. American Economic Review 92, 422-427.

[39] Sadka, R. (2003). Liquidity risk and asset pricing. *Working paper.*

[40] Spence, J. (2002). Clearing up an ETF liquidity myth. *Editorial.*

[41] Stoll, H.R. (2000). Friction. Journal of Finance 55, 1479-1514.

[42] Zhang, L., P. Mykland, and Y. Aït-Sahalia (2005). A tale of two time scales: determining integrated volatility with noisy high-frequency data. Journal of the American Statistical Association 100, 1394-1411.

```
 Lag   Covariance   Correlation   -1 9 8 7 6 5 4 3 2 1 0 1 2 3 4 5 6 7 8 9 1      Std Error

   0   3.26278E-7       1.00000   |                    |********************|             0
   1   -7.1485E-8       -.21909   |                ****| .                  |      0.044811
   2   -1.7989E-9       -.00551   |                  . | .                  |      0.046913
   3   -4.2528E-9       -.01303   |                  . | .                  |      0.046914
   4   -3.2565E-8       -.09981   |                  **| .                  |      0.046921
   5   3.18552E-9       0.00976   |                  . | .                  |      0.047346
   6   1.33231E-8       0.04083   |                  . |*.                  |      0.047350
   7   -2.3546E-8       -.07216   |                  .*| .                  |      0.047420
   8   3.62875E-9       0.01112   |                  . | .                  |      0.047640
   9   -2.1113E-8       -.06471   |                  .*| .                  |      0.047646
  10   3.64909E-9       0.01118   |                  . | .                  |      0.047822
  11   5.31233E-9       0.01628   |                  . | .                  |      0.047827
  12   1.66286E-8       0.05096   |                  . |*.                  |      0.047838
  13   -2.4186E-9       -.00741   |                  . | .                  |      0.047947
  14     -9.81E-9       -.03007   |                  .*| .                  |      0.047949
  15    1.6344E-8       0.05009   |                  . |*.                  |      0.047987
  16   -3.012E-10       -.00092   |                  . | .                  |      0.048092
  17   -4.2426E-9       -.01300   |                  . | .                  |      0.048092
  18   -3.0794E-9       -.00944   |                  . | .                  |      0.048099
  19   2.92423E-9       0.00896   |                  . | .                  |      0.048103
  20   -1.3661E-8       -.04187   |                  .*| .                  |      0.048106
  21   1.65097E-8       0.05060   |                  . |*.                  |      0.048179
  22   1.56147E-8       0.04786   |                  . |*.                  |      0.048286
  23   8.77186E-9       0.02688   |                  . |*.                  |      0.048381
  24   1.17495E-8       0.03601   |                  . |*.                  |      0.048411

"." marks two standard errors
```

Figure 1. Autocorrelation function of the observed SPDR returns. High-frequency transaction prices on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. The data are collected from the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific, and Philadelphia.

Panel (a)

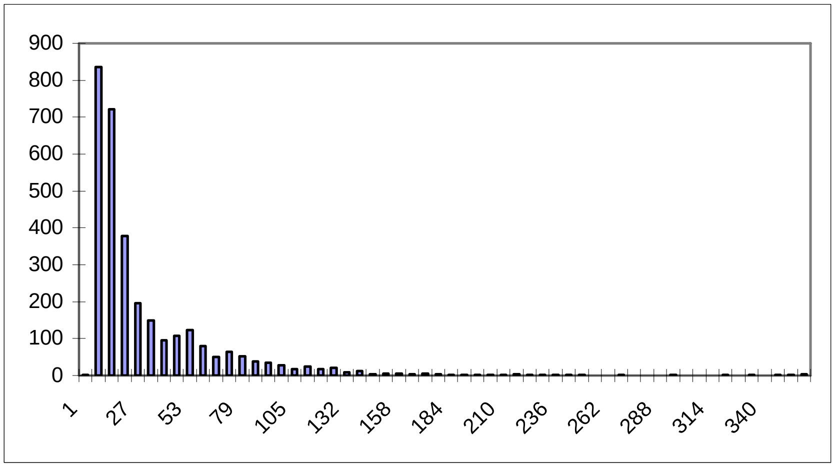

Mean	28.59
Median	13.99
Std. Dev.	36.87
Kurtosis	17.97
Skewness	3.37
Minimum	1.00
Maximum	360.00

Panel (b)

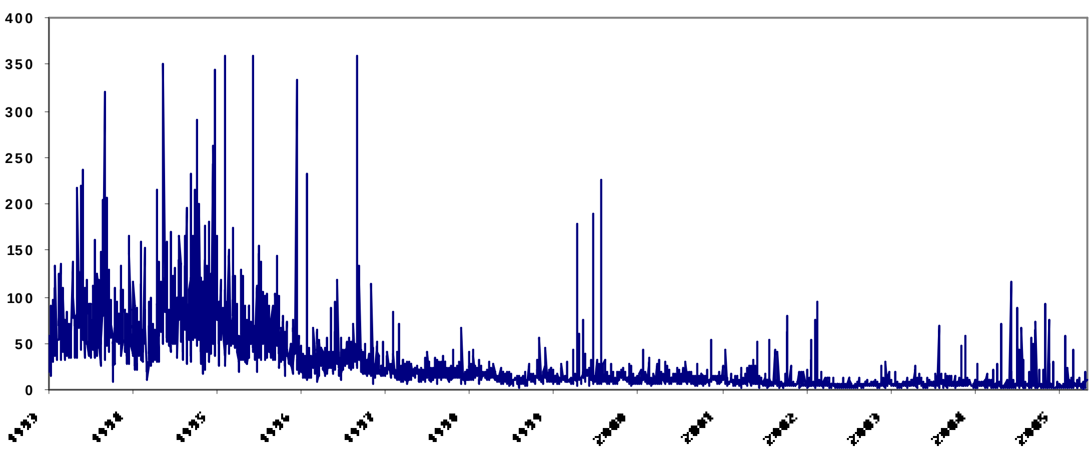

Figure 2. Optimal sampling frequencies. Panel (a) plots a histogram of daily (MSE-based) optimal sampling frequencies for the realized variance estimator $\sum_{j=1}^{M^*} \tilde{r}_j^2$ constructed using SPDR return data. The optimal frequencies are estimated as $\delta=1/M^*$, where $M^* = \left(\hat{Q}/\hat{\alpha}\right)^{1/3}$, with the numerator representing the 15-minutes quarticity estimator $\hat{Q} = \left(\sum_{j=1}^{\tilde{M}} \tilde{r}_j^4\right)\tilde{M}/3$ and the denominator being defined as $\hat{\alpha} = \left(\sum_{j=1}^{\tilde{M}} \tilde{r}_j^2 / M\right)^2$. $\tilde{M}$ and M denote the number of 15-minute SPDR returns and the number of observed SPDR returns over the trading day. On average, M=3,000. Panel (b) plots the (daily) time series plot of the optimal intervals. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected for the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia.

Panel (a) Optimally-sampled realized variances

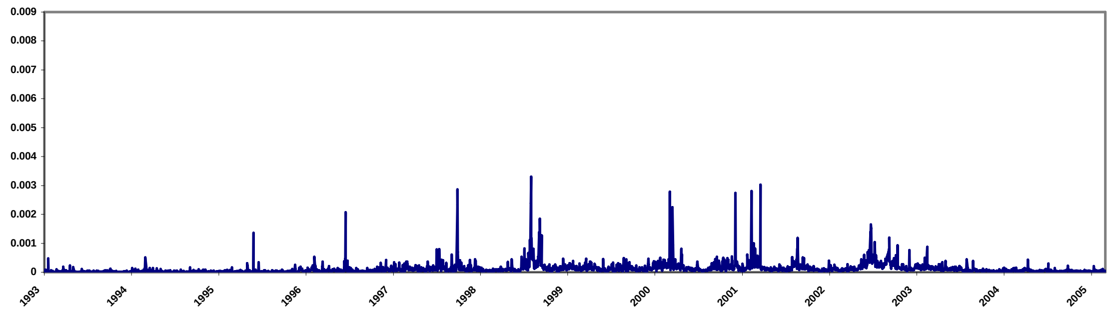

Panel (b) 20-minute realized variances

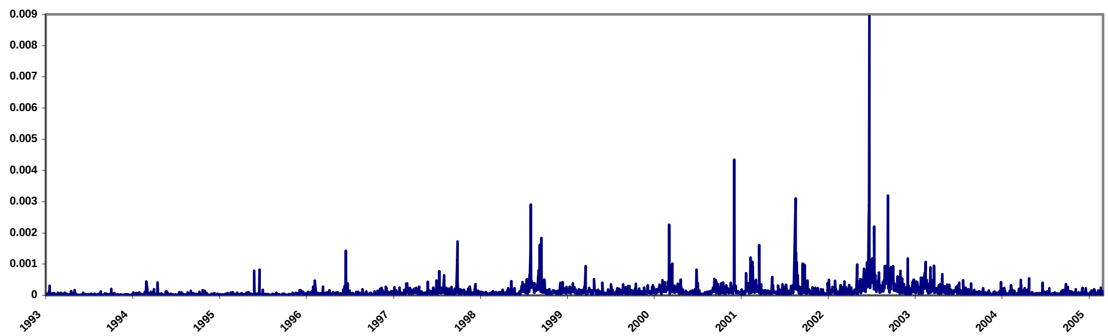

Panel (c) 5-minute realized variance

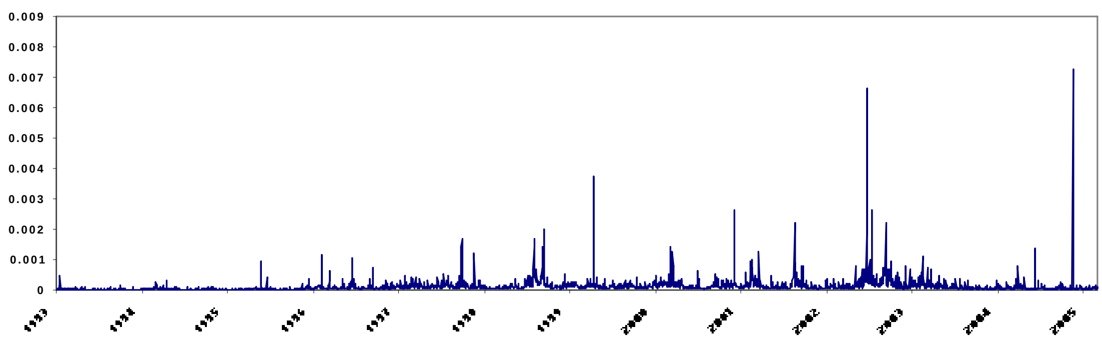

Figure 3. Daily realized variance estimates. Panel (a) plots the time series of daily realized variance estimates constructed using optimally-sampled SPDR returns. Panel (b) plots the time series of daily realized variance estimates constructed using 20-minute SPDR returns. Panel (c) plots the time series of daily realized variance estimates constructed using 5-minute SPDR returns. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected for the consolidated market, which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia.

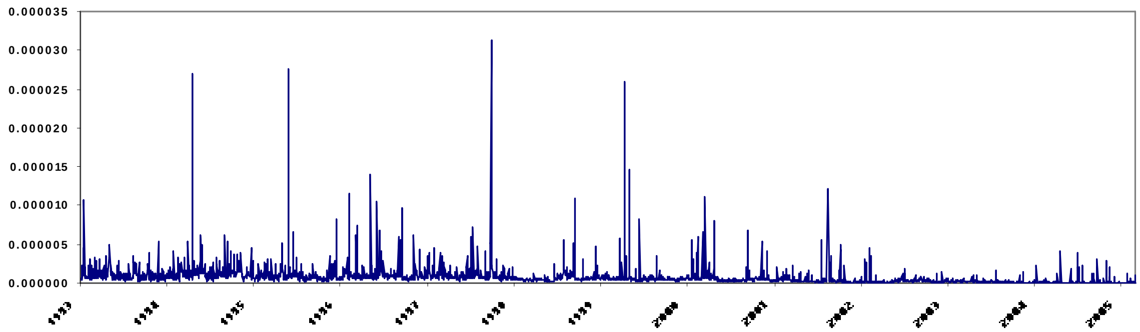

Figure 4. Price deviation second moment estimates. The figure plots the daily time series of price deviation second moment estimates constructed using SPDR returns. For each day in our sample, the estimates are obtained by computing un-centered second moments of the SPDR high-frequency tick-by-tick returns. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected for the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia.

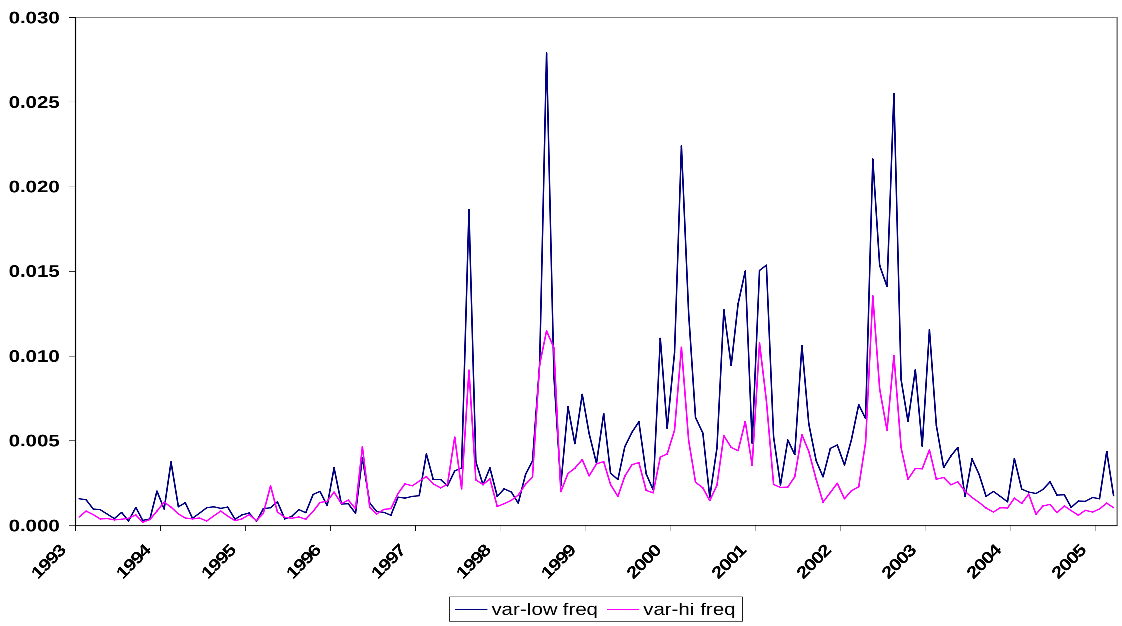

Figure 5. Monthly market variances. We plot monthly variances computed using sums of daily optimally-sampled realized variance estimates over the month (var-hi freq) and monthly variances computed by summing up squared daily S&P500 returns (var-low freq). The daily optimally-sampled realized variances sum up squared optimally-sampled intra-daily SPDR returns. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected for the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia.

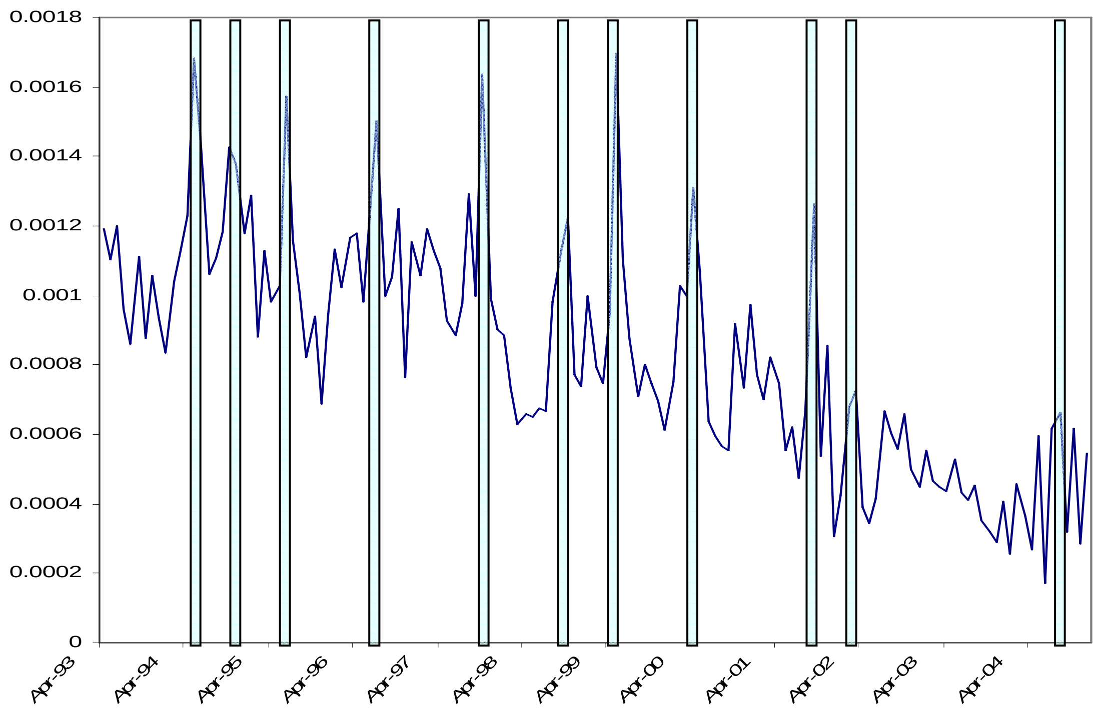

Figure 6. Monthly illiquidity. This figure plots the time series of monthly illiquidity measures obtained from high--frequency SPDR returns. For each day in our sample, we compute un-centered second moments of SPDR high-frequency tick-by-tick returns. The monthly illiquidity estimates are calculated by taking the square root of the averages of the daily estimates over the month. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected from the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia. The months highlighted below are:

May 1994: rate hike (The Credit Union Accountant)
October/November 1994: significant liquidity shortage in the market associated with exacerbated volatile market conditions and the Peso crisis (The Financial Post);
June 1995: Fed announced carrying out weekend system repos since the market needed extra liquidity (AFX News);
July 1996: weaker-than-expected employment report (Investors Chronicle);
October 1997: Asian crisis;
October/November 1998: LTCM crisis and the Russian debt default;
May 1999: uncertainty over the direction of monetary policy in Argentina and anticipation of a rate hike by US Fed (Emerging Markets Debt Report);
April 2000: considerable increase in oil price (Financial Director);
September 2001: 9/11 attack;
March 2002: worries about the Iraq war (Financial Times);
August 2004: worries about oil prices and global liquidity (Business Line).

Panel (a) Innovations in illiquidity

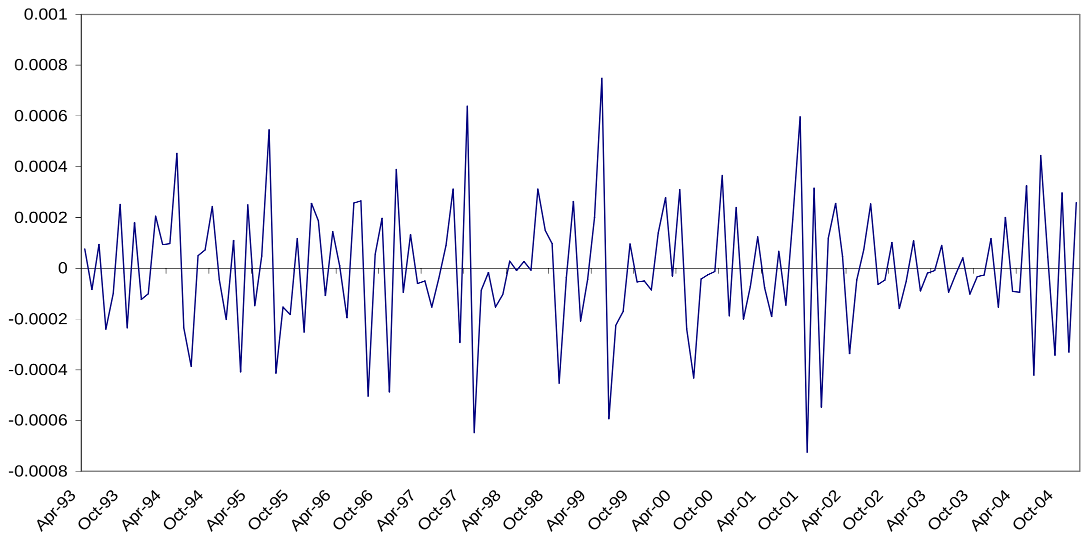

Panel (b) Innovations in the liquidity factor of Pastor and Stambaugh (2003)

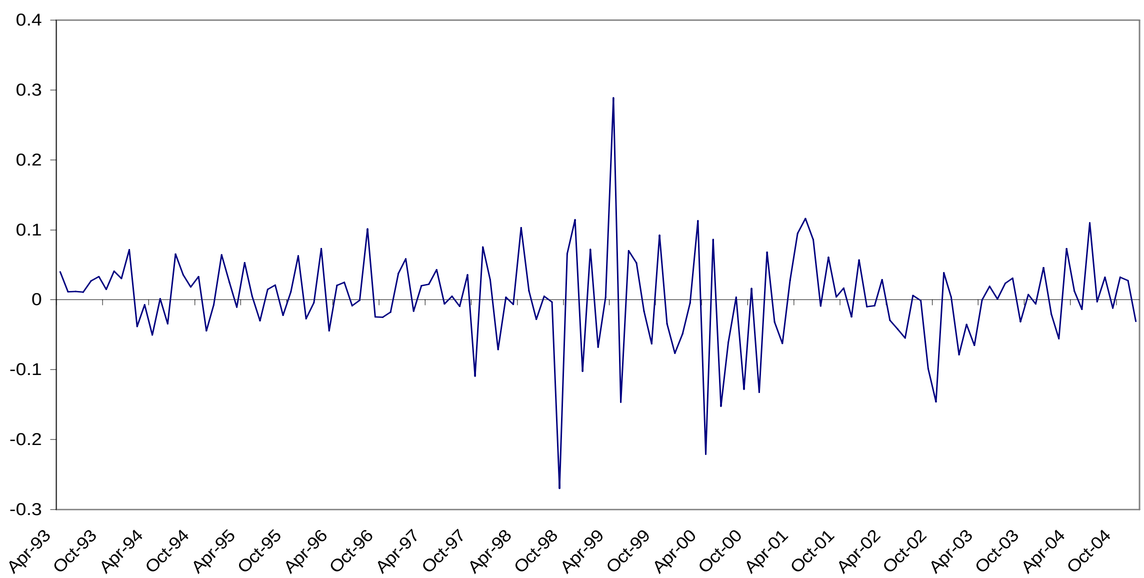

Figure 7. Two illiquidity proxies. Panel (a) plots innovations in the monthly time series of illiquidity estimates constructed using SPDR returns. For each day in our sample, we compute un-centered second moments of the SPDR high-frequency tick-by-tick returns. The monthly illiquidity estimates are calculated by taking the square root of the averages of the daily estimates over the month. High-frequency transaction price data on the Standard & Poor's depository receipts (SPDRs) are downloaded from the Trade and Quote (TAQ) database in CRSP for the period February 1993 – March 2005. Trade prices on SPDRs are collected from the consolidated market which comprises the following exchanges: AMEX, CBOE, NASDAQ, NYSE, Boston, Cincinnati, Midwest, Pacific and Philadelphia. Panel (b) plots innovations in Pastor and Stambaugh's (2003) liquidity measure over the same period. Pastor and Stambaugh's measure is downloaded from CRSP.

Panel (a) Average monthly excess returns   Panel (b) Volatility loadings $-\hat{\beta}_i^{IV}$

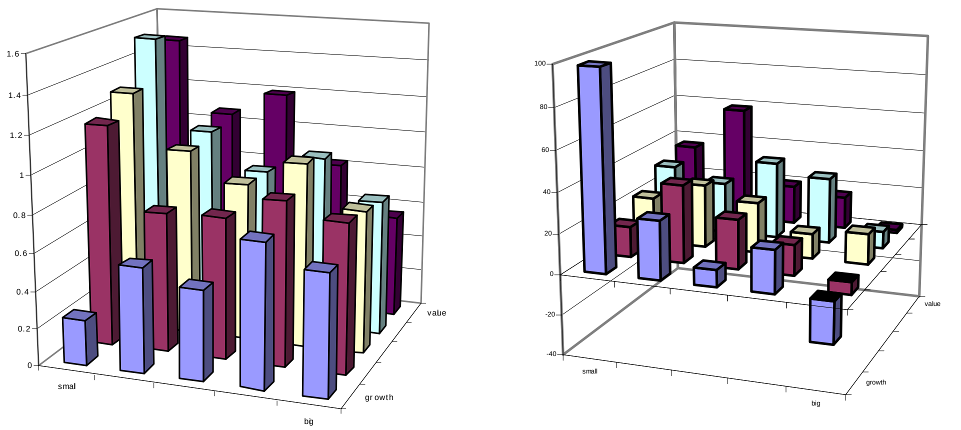

Figure 8. Monthly average returns and optimally-sampled volatility loadings. Panel (a) plots monthly average excess returns for the 25 Fama-French size- and value-sorted portfolio. The return data are collected for the period February 1993 – March 2005. Panel (b) plots the volatility factor loadings (with a minus sign) associated with innovations in optimally-sampled realized variance (IV) from the regression:

$$
\begin{gathered}
R_{i,t}^{e} = \alpha_i + \beta_i^m R_{m,t}^{e} + \beta_i^{IV} IV_t + \eta_{i,t} \\
\eta_{i,t} = \varepsilon_{i,t} - \gamma_i \eta_{i,t-1} \\
\varepsilon_{i,t} = \sqrt{h_{i,t}}\, e_{i,t} \\
h_{i,t} = \omega_i + \theta_{1,i}\varepsilon_{i,t-1}^2 + \phi_{1,i} h_{i,t-1} \\
e_{i,t} \approx N(0,1)
\end{gathered}
$$

$i = 1,...25$, $t = 1,...T$, where $R_{i,t}^{e}$ denotes excess returns on portfolio i and $R_{m,t}^{e}$ denotes excess returns on the market portfolio.

Panel (a) Average monthly excess returns

Panel (b) Illiquidity loadings $-\hat{\beta}_i^{IFV}$

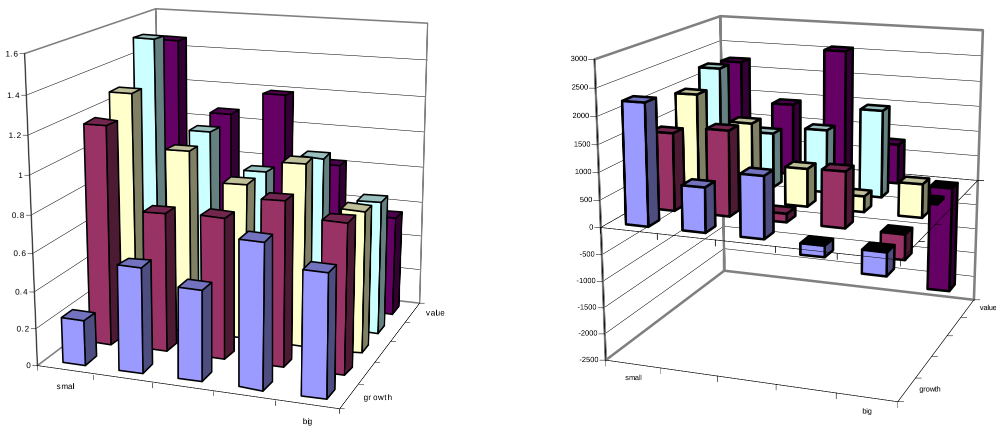

Figure 9. Monthly average returns and illiquidity loadings. Panel (a) plots monthly average excess returns for the 25 Fama-French size- and value-sorted portfolio. The return data are collected for the period February 1993 – March 2005. Panel (b) plots the illiquidity factor loadings (with a minus sign) associated with innovations in illiquidity (IFV) from the regression:

$$R_{i,t}^e = \alpha_i + \beta_i^m R_{m,t}^e + \beta_i^{IV} IFV_t + \eta_{i,t}$$
$$\eta_{i,t} = \varepsilon_{i,t} - \gamma_i \eta_{i,t-1}$$
$$\varepsilon_{i,t} = \sqrt{h_{i,t}}\, e_{i,t}$$
$$h_{i,t} = \omega_i + \theta_{1,i}\varepsilon_{i,t-1}^2 + \phi_{1,i} h_{i,t-1}$$
$$e_{i,t} \approx N(0,1)$$

$i = 1,...25$, $t = 1,...T$, where $R_{i,t}^e$ denotes excess returns on portfolio i and $R_{m,t}^e$ denotes excess returns on the market portfolio.

Panel (a) Average monthly excess returns

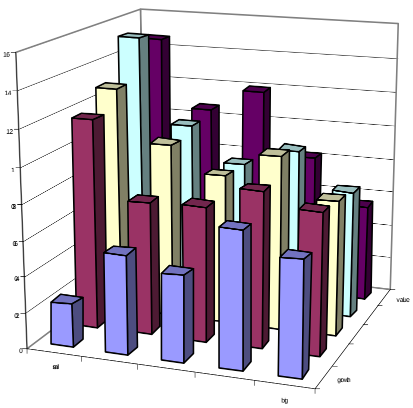

Panel b) Illiquidity loadings $-\hat{\beta}_i^{IFV}$

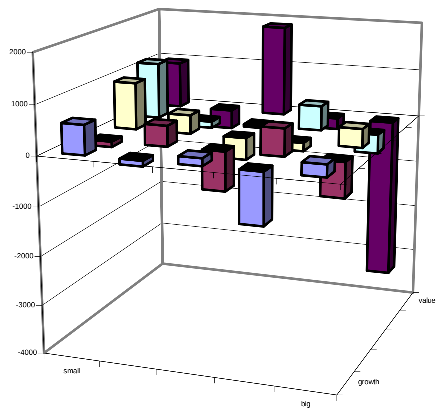

Panel (c) Volatility loadings $-\hat{\beta}_i^{IV}$

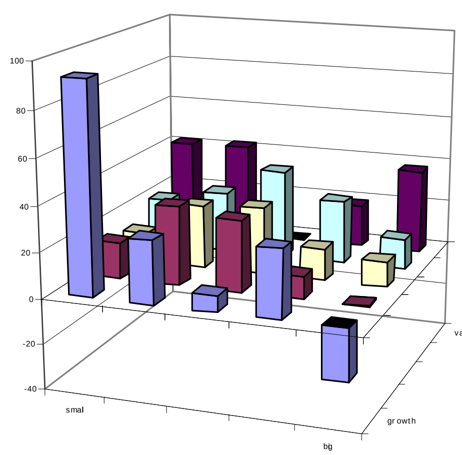

Figure 10. Monthly average returns and factor loadings. Panel (a) plots monthly average excess returns for the 25 Fama-French size- and value-sorted portfolio. The return data are collected for the period February 1993 – March 2005. Panels (b) and (c) plot illiquidity factor loadings (with a minus sign) associated with innovations in illiquidity (IFV) and market volatility factor loadings (with a minus sign) associated with innovations in market volatility (IV) from the regression

$$R_{i,t}^e = \alpha_i + \beta_i^m R_{m,t}^e + \beta_i^{IV} IV_t + \beta_i^{IFV} IFV_t + \eta_{i,t}$$
$$\eta_{i,t} = \varepsilon_{i,t} - \gamma_i \eta_{i,t-1}$$
$$\varepsilon_{i,t} = \sqrt{h_{i,t}}\, e_{i,t}$$
$$h_{i,t} = \omega_i + \theta_{1,i}\varepsilon_{i,t-1}^2 + \phi_{1,i} h_{i,t-1}$$
$$e_{i,t} \approx N(0,1)$$

$i = 1,...25$, $t = 1,...T$, where $R_{i,t}^e$ denotes excess returns on portfolio i and $R_{m,t}^e$ denotes excess returns on the market portfolio.

Panel (a) CAPM

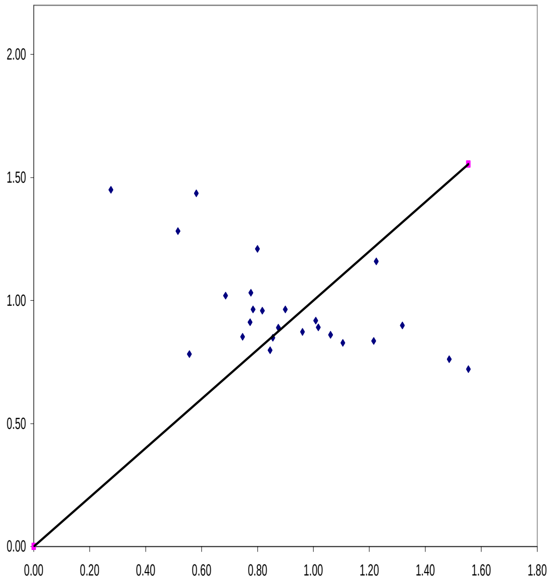

Panel (b) FF-3 Factor model (R_m^e, SMB, HML)

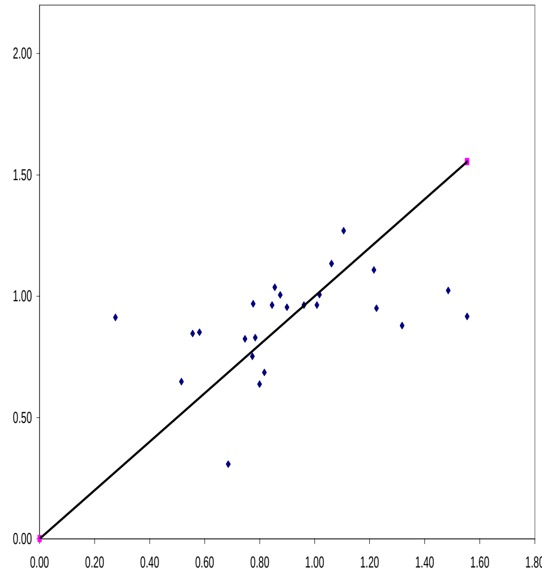

Panel (c) 3-Factor model (R_m^e, IV, IFV)

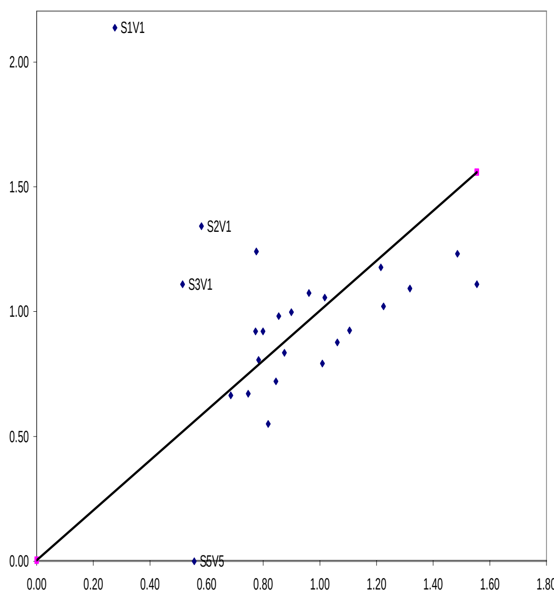

Figure 11. Predicted versus realized monthly average excess returns – 25 Fama-French portfolios. Predicted mean excess returns are plotted on the vertical axis. Mean realized excess returns are plotted on the horizontal axis. We use 3 asset pricing models: the CAPM, a 3-factor model consisting of market excess return, innovations in market volatility (IV) and innovations in illiquidity (IFV), and the Fama-French 3 factor model, FF-3. In Panel (c), S_iV_j stands for the portfolio in the i^{th} size quintile and in the j^{th} value quintile. We use value-weighted portfolios sorted on size and book-to-market equity covering the period February 1993 – March 2005.

Table I
Cross-Correlation Matrix of the Risk Factors

IV represents innovations in market volatility, IFV denotes innovations in illiquidity (as estimated in this paper), $XMKT$ refers to excess market returns, while SMB and HML are the Fama-French size and book-to-market factors. PS denotes innovations in Pastor and Stambaugh's illiquidity measure. Data cover the period February 1993 – March 2005.

	$XMKT$	IV	IFV	PS	SMB	HML
$XMKT$	1.00					
IV	-0.41	1.00				
IFV	-0.28	0.57	1.00			
PS	0.23	-0.31	-0.23	1.00		
SMB	0.19	-0.19	-0.14	-0.01	1.00	
HML	-0.54	0.13	0.14	-0.02	-0.51	1.00

Table II
Cross-Correlation Matrix of the Illiquidity Proxies

Panel (a). IFV denotes innovations in illiquidity (as estimated in this paper). $IILL$, PS, and IT denote innovations in Amihud's (2002) illiquidity measure, innovations in Pastor and Stambaugh's (2003) illiquidity measure, and innovations in aggregate turnover. $IILL$ and IT are un-scaled. Data cover the period February 1993 – December 2002.

	IFV	PS	$IILL$	IT
IFV	1.00			
PS	-0.22	1.00		
$IILL$	0.24	-0.06	1.00	
IT	0.28	-0.02	-0.12	1.00

Panel (b). IFV denotes innovations in illiquidity (as estimated in this paper). $IILL$, PS, and IT denote innovations in Amihud's (2002) illiquidity measure, innovations in Pastor and Stambaugh's (2003) measure, and innovations in aggregate turnover. $IILL$ and IT are scaled. Data cover the period February 1993 – December 2002.

	IFV	PS	$IILL$	IT
IFV	1.00			
PS	-0.22	1.00		
$IILL$	0.21	-0.02	1.00	
IT	0.26	0.00	-0.19	1.00

Table III

Time Series Regressions of Excess Returns on Expected and Unexpected Volatility – Size-Sorted Portfolios

We run regressions of excess portfolio returns on lagged market volatility, V, and the residuals from an $AR(1)$ regression of V on lagged V, namely $R_{i,t}^{e} = \alpha^{i} + \beta^{i} V_{t-1} + \lambda^{i} \varepsilon_{t} + u_{t}$, where $\varepsilon_{t} = V_{t} - \hat{\delta}_{0} - \hat{\delta}_{1} V_{t-1}$. V_{t-1} and ε_{t} represent expected and unexpected volatility, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on size. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
β	19.24	15.52	11.80	18.60	5.53
	(0.28)	(0.72)	(0.66)	(1.04)	(0.34)
λ	-203.60	-177.20	--159.60	-148.10	-113.10
	(-5.70)	(-5.70)	(-5.90)	(-5. 50)	(-4.59)

Table IV

Time Series Regression of Excess Returns on Expected and Unexpected Illiquidity – Size-Sorted Portfolios

We run regressions of excess portfolio returns on lagged illiquidity, FV, and the residuals from an $AR(1)$ regression of FV on lagged FV, namely, $R_{i,t}^{e} = \alpha^{i} + \beta^{i} FV_{t-1} + \lambda^{i} \varepsilon_{t} + u_{t}$, where $\varepsilon_{t} = FV_{t} - \hat{\delta}_{0} - \hat{\delta}_{1} FV_{t-1}$. FV_{t-1} and ε_{t} represent expected and unexpected illiquidity, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on size. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
β	1231	1418	1302	1390	2311
	(0.74)	(0.99)	(1.03)	(1.11)	(2.07)
λ	-8704	-6710	-5902	-5640	-3709
	(-3.70)	(-3.40)	(-3.30)	(-3.70)	(-2.35)

Table V

Time Series Regression of Excess Returns on Expected and Unexpected Volatility and Illiquidity – Size-Sorted Portfolios

We run regressions of excess portfolio returns on lagged illiquidity, FV, lagged market volatility, V, the residuals from an $AR(1)$ regression of FV on lagged FV and the residuals from an $AR(1)$ regression of V on lagged V, namely $R^e_{i,t} = \alpha^i + \beta^i_V V_{t-1} + \lambda^i_V \varepsilon_t + \beta^i_{FV} FV_{t-1} + \lambda^i_{FV} u_t + v_t$, where $\varepsilon_t = V_t - \hat{\delta}_0 - \hat{\delta}_1 V_{t-1}$, and $u_t = FV_t - \hat{\delta}_0 - \hat{\delta}_1 FV_{t-1}$. V_{t-1} and ε_t represent expected and unexpected volatility, respectively. FV_{t-1} and u_t represent expected and unexpected illiquidity, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on size. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
Average portfolio returns	1.115	0.762	0.622	0.771	0.608
$\boldsymbol{\beta_V}$	12.86	12.31	9.30	16.23	4.68
	(0.53)	(0.58)	(0.5)	(0.87)	(0.28)
$\boldsymbol{\lambda_V}$	-185.90	-168.62	-153.05	-140.42	-101.52
	(-4.30)	(-4.50)	(-4.70)	(-4.30)	(-3.40)
$\boldsymbol{\beta_{FV}}$	-858	-475	-419	-178	1167
	-(0.53)	-(0.34)	-(0.34)	-(0.14)	(1.03)
$\boldsymbol{\lambda_{FV}}$	-2919	-1461	-1163	-1117	-592
	(-1.11)	(-0.65)	(-0.59)	(-0.56)	(-0.33)

Table VI

Time Series Regression of Excess Returns on Expected and Unexpected Volatility – BE/ME-Sorted Portfolios

We run regressions of excess portfolio returns on lagged market volatility, V, and the residuals from an $AR(1)$ regression of V on lagged V, namely $R_{i,t}^e = \alpha^i + \beta^i V_{t-1} + \lambda^i \varepsilon_t + u_t$, where $\varepsilon_t = V_t - \hat{\delta}_0 - \hat{\delta}_1 V_{t-1}$. V_{t-1} and ε_t represent expected and unexpected volatility, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on book-to-market. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
β	9.84 (0.59)	8.18 (0.52)	8.25 (0.54)	14.00 (1.0)	1.75 (0.10)
λ	-115.50 (-4.60)	-130.30 (-5.50)	-126.60 (-5.50)	·131.10 (-6. 20)	-135.40 (-5.30)

Table VII

Time Series Regression of Excess Returns on Expected and Unexpected Illiquidity – BE/ME-Sorted Portfolios

We run regressions of excess portfolio returns on lagged illiquidity, FV, and the residuals from an $AR(1)$ regression of FV on lagged FV, namely, $R_{i,t}^e = \alpha^i + \beta^i FV_{t-1} + \lambda^i \varepsilon_t + u_t$, where $\varepsilon_t = FV_t - \hat{\delta}_0 - \hat{\delta}_1 FV_{t-1}$. FV_{t-1} and ε_t represent expected and unexpected illiquidity, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on book-to-market. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
β	1796 (1.60)	1116 (1.01)	1994 (1.89)	977 (0.98)	213 (0.18)
λ	-4424 (-2.70)	-4098 (-2.60)	-4670 (-3.10)	-4801 (-3.40)	-3592 (-2.13)

Table VIII

Time Series Regression of Excess Returns on Expected and Unexpected Volatility and Illiquidity – BE/ME-Sorted Portfolios

We run a regression of excess portfolio returns on lagged illiquidity, FV, lagged market volatility, V, the residuals from an $AR(1)$ regression of FV on lagged FV and the residuals from an $AR(1)$ regression of V on lagged V, namely $R^e_{i,t} = \alpha^i + \beta^i_V V_{t-1} + \lambda^i_V \varepsilon_t + \beta^i_{FV} FV_{t-1} + \lambda^i_{FV} u_t + v_t$, where $\varepsilon_t = V_t - \hat{\delta}_0 - \hat{\delta}_1 V_{t-1}$, and $u_t = FV_t - \hat{\delta}_0 - \hat{\delta}_1 FV_{t-1}$. V_{t-1} and ε_t represent expected and unexpected volatility, respectively. FV_{t-1} and u_t represent the expected and unexpected illiquidity, respectively. We use monthly value-weighted return data collected for the period February 1993 – March 2005 for the 10 Fama-French portfolios sorted on book-to-market. We present results for the 2nd, 4th, 6th, 8th, and 10th decile portfolios. T-statistics are reported in parentheses.

	2nd	4th	6th	8th	10th
Average portfolio returns	0.688	0.862	0.798	0.791	0.962
β_V	7.45	8.08	6.01	13.26	3.4
	(0.43)	(0.50)	(0.38)	(0.85)	(0.20)
λ_V	-101.93	-132.62	-113.01	-127.88	-152.50
	(-3.30)	(-4.60)	(-4.10)	(-5.04)	(-5.03)
β_{FV}	651	-375	722	-454	-1511
	(0.57)	-(0.35)	(0.69)	-(0.47)	-(1.31)
λ_{FV}	-1245	9	-1186	-761	1025
	(-0.68)	(0.10)	(-0.70)	(-0.50)	(0.56)

Table IX

Prices of Risk - monthly

We report the estimated factors' risk prices associated with the asset pricing model $E(R_{i,t+1}^{e}) = \lambda_m \beta_i^m + \sum_{s=1}^{p} \lambda_s \beta_i^s$. The left-hand side variable is the vector of mean excess returns on the Fama-French 25 portfolios sorted on size and book-to-market. The betas represent factors' risk loadings estimated from the corresponding time-series model. The factors of interest are: excess return on the market portfolio, IV (innovations in market volatility), IFV (innovations in illiquidity), SMB (the Fama-French size factor), HML (the Fama-French value factor), and PS (innovations in the illiquidity factor of Pastor and Stambaugh, 2003). The time period is February 1993 – March 2005. Estimation is conducted using GMM. T-statistics are reported in parentheses.

$\hat{\lambda}_m$	$\hat{\lambda}_{IV}$	$\hat{\lambda}_{IFV}$	$\hat{\lambda}_{PS}$	$\hat{\lambda}_{SMB}$	$\hat{\lambda}_{HML}$	J_T	p	df
0.7397	-0.0111					9.9368	0.9917	23
(4.00)	(-2.54)							
0.6999		-0.0003				7.3829	0.9992	23
(4.05)		(-3.45)						
0.8724			0.0519			6.7968	0.9996	23
(4.97)			(2.23)					
0.7691	-0.0100	-0.0003				9.6164	0.9895	22
(4.05)	(-2.22)	(-2.69)						
0.6698	-0.0118		0.0375			8.6923	0.9948	22
(3.63)	(-2.81)		(1.55)					
0.6839		-0.0003	0.0282			8.7111	0.9947	22
(3.91)		(-2.90)	(1.12)					
0.6963	-0.0110	-0.0002	0.0322			9.3396	0.9863	21
(3.62)	(-2.57)	(-2.39)	(1.23)					
0.5913				0.2754	0.3740	4.6856	0.9999	22
(2.69)				(0.80)	(1.71)			

Table X

Prices of Risk – daily

We report the estimated factors' risk prices associated with the asset pricing model $E(R_{i,t+1}^{e}) = \lambda_m \beta_i^m + \sum_{s=1}^{p} \lambda_s \beta_i^s$. The left-hand side variable is the vector of mean excess returns on the Fama-French 25 portfolios sorted on size and book-to-market. The betas represent factors' risk loadings estimated from the corresponding time-series model. The factors of interest are: excess return on the market portfolio, IV (innovations in market volatility), IFV (innovations in illiquidity), SMB (the Fama-French size factor), and HML (the Fama-French value factor). The time period is February 1993 – March 2005. Estimation is conducted using GMM. T-statistics are reported in parentheses.

$\hat{\lambda}_m$	$\hat{\lambda}_{IV}$	$\hat{\lambda}_{IFV}$	$\hat{\lambda}_{SMB}$	$\hat{\lambda}_{HML}$	J_T	p	df
0.0198	-0.0033				31.977	0.1006	23
(1.51)	(-2.64)						
0.0141		-0.0006			22.8503	0.4695	23
(1.08)		(-3.54)					
0.0201	-0.0011	-0.0005			23.8349	0.9863	22
(1.54)	(-0.78)	(-3.82)					
0.0239			0.0099	0.0240	6.7489	0.9992	22
(1.80)			(0.55)	(2.88)			

Table XI
Risk Premia

The differences in risk premia between small-cap and large-cap stocks are estimated as the product between the factors' risk prices and the differences in the factors' loadings. IV represents innovations in market volatility. IFV represents innovations in illiquidity. Results are reported in percentages, per month.

Panel (a) The model is $E(R^e_{i,t+1}) = \lambda_m \beta_i^m + \lambda_{IV} \beta_i^{IV}$.

	Book-to-Market Equity (BE/ME) Quintiles				
	Low	**2**	**3**	**4**	**High**
Average Excess Returns	0.5339	0.8708	0.9759	1.0131	1.0215
$\hat{\lambda}_{IV}(\hat{\beta}^{IV}_{Small} - \hat{\beta}^{IV}_{Big})$	1.3230	0.2359	0.0791	0.2616	0.3863

Panel (b) The model Is $E(R^e_{i,t+1}) = \lambda_m \beta_i^m + \lambda_{IFV} \beta_i^{IFV}$.

	Book-to-Market Equity (BE/ME) Quintiles				
	Low	**2**	**3**	**4**	**High**
Average Excess Returns	0.5339	0.8708	0.9759	1.0131	1.0215
$\hat{\lambda}_{IFV}(\hat{\beta}^{IFV}_{Small} - \hat{\beta}^{IFV}_{Big})$	0.7979	0.5788	0.4019	0.6941	1.2971

Panel (c) The model Is $E(R^e_{i,t+1}) = \lambda_m \beta_i^m + \lambda_{IV} \beta_i^{IV} + \lambda_{IFV} \beta_i^{IFV}$.

	Book-to-Market Equity (BE/ME) Quintiles				
	Low	**2**	**3**	**4**	**High**
Average Excess Returns	0.5339	0.8708	0.9759	1.0131	1.0215
$\hat{\lambda}_{IV}(\hat{\beta}^{IV}_{Small} - \hat{\beta}^{IV}_{Big})$	1.1584	0.1580	0.0187	0.0797	0.0461
$\hat{\lambda}_{IFV}(\hat{\beta}^{IFV}_{Small} - \hat{\beta}^{IFV}_{Big})$	0.1071	0.1793	0.1730	0.4612	1.2812