Available online at www.sciencedirect.com
ScienceDirect
Journal of Empirical Finance 15 (2008) 679 – 699
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/jempfin

Liquidity and conditional portfolio choice: A nonparametric investigation ☆

Eric Ghysels [a,1], João Pedro Pereira [b,*,2]

[a] Department of Economics, UNC and Department of Finance, Kenan-Flagler Business School, United States
[b] Finance Department ISCTE Business School - Lisbon, Portugal

Received 1 September 2005; accepted 12 December 2007
Available online 16 January 2008

Abstract

This paper studies the relation between liquidity and optimal portfolio allocations. Given that the portfolio problem of a constant relative risk aversion investor does not have a closed-form solution, we use a nonparametric approach to estimate the optimal allocations. Using a sample of NYSE stocks from 1963–2000, we find that the optimal portfolio weight in small stocks is strongly increasing in liquidity at short daily and weekly horizons. This result is consistent for three different measures of liquidity: price impact, dollar volume, and turnover. However, liquidity does not influence the optimal portfolio choice for large stocks, nor for longer monthly investment horizons.

JEL classification: G11
Keywords: Liquidity; Conditional Portfolio Choice; Nonparametric

1. Introduction

How does an asset's liquidity affect its demand? A substantial amount of research has been devoted to the relation between liquidity and the conditional distribution of returns. However, the question of how liquidity influences portfolio allocations is not easy to address, particularly since for most interesting utility functions — like the widely used power utility — portfolio weights are complex implicit functions of higher-order conditional moments of returns. The goal of this paper is to characterize the dependence of optimal portfolio choices on changes in the liquidity of assets through time.

☆ We are grateful for the helpful comments of Jennifer Conrad, Adam Reed, Harold Zhang, an anonymous referee, and seminar participants at UNC Chapel Hill and at the 2003 Doctoral Student Consortium of the Financial Management Association. We also thank Michael Brandt for providing software code. We thank the Southern Finance Association for awarding this paper the 2003 Best Paper in Investments distinction.
* Corresponding author. ISCTE- Department Financas, Av Forcas Armadas, 1649-026 Lisboa, Portugal. Tel.: +351 217 903 965.
E-mail addresses: eghysels@unc.edu (E. Ghysels), joao.pereira@iscte.pt (J.P. Pereira).
[1] Address: Gardner Hall CB 3305, Chapel Hill, NC 27599-3305. Tel.: +919 966 5325.
[2] Financial support from the Portuguese Science and Technology Foundation is gratefully acknowledged.

0927-5398/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2007.12.002

The dependence of portfolio weights on liquidity is not the only challenging issue, as liquidity itself is a complex and not easy to measure concept. A common definition is to say that an asset is liquid if large quantities can be traded in a short period of time without moving the price too much. Accordingly, several alternative measures of liquidity have been used in the literature, including the price impact of trade, the bid–ask spread, share or dollar volume, and turnover, among others. A security is taken to be more liquid the lower its price impact, the tighter its bid–ask spread, or the higher its volume or turnover.

The standard cross sectional empirical finding is that expected returns are decreasing in liquidity, as measured by the bid–ask spread in Amihud and Mendelson (1986), turnover in Datar, Naik, and Radcliffe (1998), price impact in Brennan and Subrahmanyan (1996), or trading volume in Brennan, Chordia, and Subrahmanyan (1998). The reasoning is that investors anticipate having to pay higher transaction costs when they need to sell the illiquid assets in the future, and thus require a higher expected return to hold them.

However, empirical evidence on the time series relation is quite different. On the one hand,Amihud (2002) finds a long-term (monthly and yearly) negative relation between market liquidity (measured by price impact) and market returns. Using a very long time series, Jones (2002) also finds that high liquidity (high turnover or narrow bid–ask spread) predicts low stock returns 1 year ahead. On the other hand, Gervais, Kaniel, and Mingelgrin (2001) find a short-term (up to 20 days) positive relation between liquidity (measured by volume) and stock returns. Kaniel, Li, and Starks (2005) validate this high volume premium for almost all international developed markets. Hence, it appears that there is a temporal dimension to the effects of changes in liquidity, with expected returns increasing in liquidity at short daily or weekly frequencies, while decreasing at longer monthly or yearly frequencies.

There is little guidance from theory on how *portfolio weights* should respond to liquidity. In a recent paper, Longstaff (2001) associates illiquidity with the notion of thin markets, i.e., with the possibility that sometimes investors may find it impossible to initiate or unwind positions in a given security at any price. Thus, he defines illiquidity as a bound on the amount of shares that can be traded per period. Longstaff shows that in general the investor chooses a lower initial portfolio weight in the presence of liquidity constraints. Pereira and Zhang (2004) define liquidity by the existence of price impact of trading, generalizing Longstaff's model in the sense that the investor is allowed to trade any number of shares as long as he is willing to pay or accept the necessary price. They show that the investor chooses to hold less shares of stocks with high levels of price impact. The intuition is that the buildup and unwinding of a large portfolio is very costly when there is price impact. However, both these papers compare different states of liquidity, which can be interpreted as cross-sectional comparisons between different stocks. None of them addresses directly our question, namely what is the optimal trading strategy that takes advantage of the fact that changes in liquidity through time imply different conditional distributions of future returns.

Our contribution is to analyze the empirical time-series relation between optimal portfolio allocations and different measures of liquidity, at different investment frequencies. We extend the previous empirical papers by studying the optimal portfolio functions directly, instead of focusing only on the first conditional moment of returns. We use two different approaches. First, we adopt the nonparametric method of Brandt (1999) and Aït-Sahalia and Brandt (2001), which allows us to express optimal portfolio choices directly as functions of the state variable (liquidity) without requiring the intermediate step of estimating conditional moments of returns. Second, we specify a parametric estimator based on the nonparametric results. This allows us to confirm the dependence of portfolio weights on liquidity.

Our paper also relates to the growing literature on model and parameter uncertainty as it pertains to portfolio selection. Various approaches have been suggested, ranging from Bayesian methods, to robustness arguments regarding model uncertainty. Since our approach is non-parametric, de facto there is no parameter uncertainty (within a given model). See, for example, Jorion (1986), Kandel and Stambaugh (1996), Xia (2001), Maenhout (2004), Tu and Zhou (2004). Implicitly, our non-parametric approach acknowledges the inability to construct tightly parameterized portfolio allocation rules. The resulting inference is very imprecise, compared to parametric approaches. In this sense our approach demonstrates that, given the amount of data, we have great uncertainty about portfolio rules and their dependence on state variables, in our case liquidity.

We construct two portfolios of large and small New York Stock Exchange (NYSE) stocks for the 1963–2000 period. We test three different measures of liquidity for those portfolios: price impact, dollar volume, and turnover. To address the time dimension of liquidity, we consider the investment decisions of a power-utility investor with three different horizons: 1 day, 1 week, and 1 month. We consider two separate portfolio allocation problems: one where the investor chooses between the portfolio of large stocks and a risk-free asset; another where he chooses between a portfolio of small stocks and a risk-free asset. The focus of the paper is on the time-series relation between liquidity and optimal allocations.

Our main results are as follows. The optimal portfolio weight in small stocks does indeed depend on their liquidity for short daily or weekly investment horizons. More precisely, the portfolio allocations are strongly increasing in liquidity, which is consistent with the short-term price run up documented in Gervais, Kaniel, and Mingelgrin (2001) and Kaniel, Li, and Starks (2005). In contrast, the liquidity of large stocks does not in general contain useful information for portfolio choice. At monthly frequency, liquidity does not determine the portfolio functions for either small or large stocks. These results are consistent across the three liquidity measures (price impact, volume, and turnover). Also, these results are consistent across the nonparametric and parametric estimators. Finally, we find that higher-order moments of returns matter for the optimal allocation with short daily horizons, but not with weekly or monthly horizons. This is consistent with the well-known empirical fact that the departure of stock returns from a normal distribution is stronger at higher frequencies.

The paper proceeds as follows. Section 2 defines the conditional portfolio choice problem and discusses the use of liquidity as conditional information. This section also describes our data. Since the portfolio problem does not have a closed-form solution, we solve it in Section 3 using a nonparametric technique. This section presents the optimal conditional allocations and also discusses the importance of implicitly considering higher-order moments of returns. Section 4 shows that the results are not influenced by non-synchronous trading and that they are consistent with a parametric estimator. Further, it discusses our results in the light of the findings in Amihud (2002). The last section concludes.

2. The conditional portfolio problem

We start by defining the investor's maximization problem. Since this problem does not have a closed-form analytical solution, we delay the discussion of the econometric techniques used to solve it until Section 3. In Section 2.2 we argue why liquidity variables are good candidates for conditioning information. Section 2.3 presents the data used in this paper.

2.1. The standard model

Consider an investor who maximizes the conditional expected utility of next period's wealth, subject to the usual budget constraint:

$$\begin{aligned} &\underset{\alpha_t \in \mathbb{R}}{\text{maximize}} && E[u(W_{t+1})|Z_t] \\ &\text{subject to} && W_{t+1} = W_t\left[R_{f,t+1} + \alpha_t\left(R_{s,t+1} - R_{f,t+1}\right)\right] \end{aligned} \tag{1}$$

where α_t represents the proportion of wealth invested in a risky stock with return $R_{s,t+1}$, $1-\alpha_t$ is the proportion invested in a risk-free asset with return $R_{f,t+1}$, and W_t is the investor's wealth at time t. We consider two separate problems: one where $R_{s,t+1}$ is the return on a portfolio of large stocks; another where it is the return on a portfolio of small stocks (these portfolios are described in Section 2.3). The investor can have three different horizons, i.e., the difference between t and $t+1$ can either be 1 day, 1 week, or 1 month. We assume that the investor is small enough so that his trades do not affect prices and that transaction costs are negligible.[3] The purpose of this paper is to study whether different measures of liquidity, represented here by Z_t, constitute relevant conditioning information.

Assume a utility function with constant relative risk aversion (CRRA):

$$u(W) = \begin{cases} W^{1-\gamma/(1-\gamma)} & \text{if } \gamma > 1 \\ \ln(W) & \text{if } \gamma = 1 \end{cases} \tag{2}$$

where γ represents the coefficient of relative risk aversion.[4] This utility function is attractive in the sense that the portfolio weights are independent of the level of wealth. However, it has the disadvantage of not permitting a closed-form solution to the investor's portfolio choice problem.

[3] Lynch and Balduzzi (2000) show that return predictability continues to have a large effect on portfolio rebalancing even in the presence of transaction costs. Other papers like Constantinides (1986), Heaton and Lucas (1996), and Vayanos (1998) show that transaction costs induce only a negligible utility loss, despite causing the investor to reduce the trading frequency. Nonetheless, it is important to realize that transaction costs may decrease the benefit of following the trading rules described in this paper, especially when trading small stocks daily.

[4] For most of our empirical analysis, we assume a coefficient of relative risk-aversion of $\gamma=10$. Mehra and Prescott (2003) mention that $\gamma=10$ should be considered as an upper bound on the degree of risk-aversion. Even with this number we sometimes find large portfolio weights, but this is just evidence of the well-known equity premium puzzle. For robustness, Section 3.2 also analyzes the results for different degrees of risk aversion.

The first order condition for (1) is:

$$\alpha_t : E[m_t + 1(\alpha_t)|Z_t] \equiv E\left[\frac{du(W_{t+1})}{dW_{t+1}}\left(R_{t+1,s} - R_{t+1,f}\right)|Z_t\right] = 0 \tag{3}$$

Since (3) cannot be solved in closed-form for α_t, one must proceed using the econometric technique described in Section 3. First, we discuss the conditioning information Z_t.

2.2. Conditioning information

This section defines and motivates the use of liquidity measures based on existing literature. Note that most of the papers referred below deal with the relation between liquidity and the first or second conditional moments of returns. However, we stress that we are only implicitly interested in conditional moments of returns, since we are able to directly estimate optimal portfolio functions. Nevertheless, these relations can help us rationalize the shape of the portfolio functions found — any risk-averse investor obviously likes higher means and lower variances.

The first measure of liquidity is the price impact of trading. In the microstructure literature, this is usually defined as the price change induced by a given signed (buy/sell) order size. Hence, a more liquid security will have a lower price impact. Brennan and Subrahmanyan (1996) find a positive cross-sectional empirical relation between this measure and stock returns. Pereira and Zhang (2004) study the liquidity premium in a partial-equilibrium portfolio selection model, also showing that there is an increasing and concave relation between price impact and expected returns. Since high frequency data on transactions and quotes is not available for long periods of time, we follow Amihud (2002) and define a daily stock measure of price impact as the ratio of absolute return to dollar volume. For a portfolio with I stocks,

$$\mathrm{pi}_d \equiv \frac{1}{I}\sum_{i=1}^{I} \frac{|R_{id}|}{V_{id}P_{id}}$$

where R_{id} is the return, V_{id} is the share volume, and P_i is the price of stock i on day d. Amihud (2002) computes this measure for the whole market and finds a positive time series relation with stock returns at the monthly and yearly frequencies. Sadka (2004) shows that Amihud's measure is highly correlated with high-frequency measures of price impact.

A second measure of liquidity is dollar volume, defined for a portfolio as

$$\mathrm{dvol}_d \equiv \sum_{i=1}^{I} V_{id}P_{id}.$$

Brennan, Chordia, and Subrahmanyan (1998) find that dollar volume is cross-sectionally negatively related to monthly risk-adjusted returns. In the time series dimension, Gervais, Kaniel, and Mingelgrin (2001) find that stocks which experience unusually high (low) volume over a day or a week tend to appreciate (depreciate) over the subsequent 1, 10, and 20 days. They argue that this finding is consistent with the visibility hypothesis: if higher volume attracts attention to a stock, then the number of potential buyers increases and thus the stock price will also tend to increase. They disaggregate the analysis by market capitalization and find that the effect is stronger for small firms than for large firms. Hence, we expect to see portfolio weights increasing in volume, especially at very short investment horizons.

The third measure of liquidity is turnover, defined for an individual security as V_{id}/N_{id}, where N_{id} is the number of shares outstanding. For a portfolio, we define turnover as the value-weighted average

$$\mathrm{turn}_d \equiv \sum_{i=1}^{I} \omega_{id}\frac{V_{id}}{N_{id}}$$

where $\omega_{id} \equiv N_{id}P_{id} / \sum_{j=1}^{I} N_{jd}P_{jd}$. Datar, Naik, and Radcliffe (1998) propose turnover as a proxy for liquidity and, in an exercise similar to Amihud and Mendelson (1986), find that turnover is cross-sectionally negatively related to monthly returns. Jones (2002) builds a long time series from 1900–2000 of large NYSE stocks and finds that high turnover predicts low stock returns 1 year or more ahead.

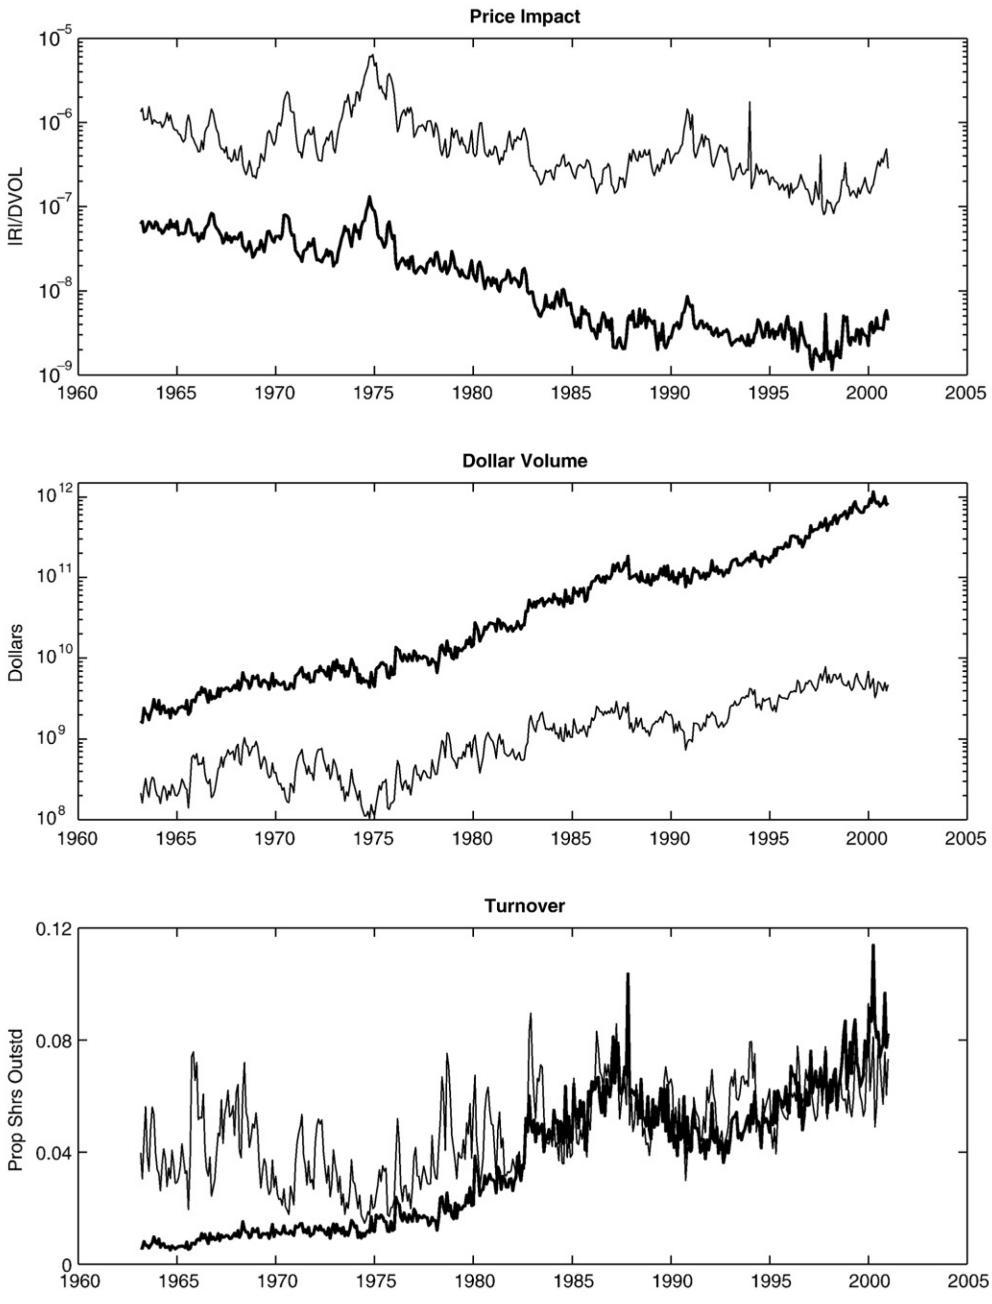

Fig. 1. Time Series of Raw Data. Each plot displays raw (before the transformation in Eq. (4)) monthly data for the indicated liquidity measure. The bold line in each panel represents the portfolio of large stocks (deciles 9 and 10 of the NYSE); the thinner line represents the portfolio of small stocks (deciles 2 and 3 of the NYSE). Note that the vertical axis in the first two plots is in a logarithmic scale.

2.3. Data

We collect daily returns for New York Stock Exchange (NYSE) traded stocks from the Center for Research in Security Prices (CRSP). We sort the stocks on their market capitalization to construct value-weighted returns on two risky assets: a portfolio of small stocks (deciles 2 and 3) and a portfolio of large stocks (deciles 9 and 10). These daily data are then aggregated to weekly and monthly frequencies. The sample is from 1963 to 2000, resulting in sample sizes of 9431 daily observations, 1978 weekly observations, and 455 monthly observations.

We collect samples of one-year Treasury bills from CRSP for the same time period. For most of this study, we assume that Treasury bills are risky free and fix the risk-free rate (R_f) at the Treasury bill's historical average, following Aït-Sahalia and Brandt (2001). By using a constant risk-free rate, we guarantee that any difference in the optimal portfolio functions across frequencies is solely due to the relation between stock returns and liquidity.

Finally, we collect daily volume data from CRSP and compute the measures of liquidity defined above for each of the two stock portfolios. Time-aggregation is done by summing across dates in the cases of turnover and dollar volume, or by averaging across dates in the case of price impact. This follows the conventions in Lo and Wang (2000) and Amihud (2002).

Fig. 1 plots these time series at the monthly frequency. Both price impact and dollar volume show that liquidity has increased through time. Large stocks are at least one order of magnitude more liquid than small stocks and, interestingly, the gap between them has also increased through time (i.e., the liquidity of large stocks has increased faster). Since these measures are functions of stock prices, we expect them to be non-stationary. Indeed, the augmented Dickey–Fuller t-test clearly accepts the null of a unit root in dollar volume for small (DF = − 2.28) and large (DF = 1.71) stocks. For price impact, the statistical evidence of a unit root is less strong: we may accept the unit root for small stocks (DF = −3.89) if we compare it to the 1% critical value (−3.98), but we reject the unit root for large stocks (DF = −4.72).[5] Still, since price impact is a function of dollar volume, which is clearly and sensibly non-stationary, we will also assume that price impact is itself non-stationary. On the other hand, turnover displays very surprising properties: small and large stocks have similar levels of turnover in the second half of the sample; in the first half, small stocks actually have more turnover than large stocks. This confirms the suggestion of Lo and Wang (2000, p. 272) that "smaller-capitalization companies can have high turnover". The Dickey–Fuller test rejects a unit root for the turnover of small stocks (DF = −4.91), but accepts a unit root for large stocks (DF = −2.74). We will assume that turnover is also non-stationary because there are good economic reasons for it: changes in market structure over the sample period are certainly responsible for a secular increase in turnover. In particular, commissions were deregulated during the 1970s, leading to cheaper trades and more volume. Further, the move from eighths to sixteenths in 1977 led to a one-time reduction in spreads and resulted in more trading.[6]

Given that our nonparametric approach requires stationary data, we transform each predictor by taking logs and subtracting a 6-month moving average. For example, we redefine price impact as

$$\mathrm{PI}_t = \ln(\mathrm{pi}_t) - \frac{1}{S}\sum_{s=1}^{S} \ln(\mathrm{pi}_{t-s}) \tag{4}$$

for each of the three frequencies: daily ($t = d$, $S = 125$), weekly ($t = w$, $S = 26$), and monthly ($t = m$, $S = 6$). Stationary measures of dollar volume, DVOL, and turnover, TURN, are obtained through a similar transformation.[7] To facilitate the interpretation of the results, we further rescale all series by subtracting their means and dividing by their standard-deviations. Table 1 presents descriptive statistics of these three standardized predictors. By construction, all measures have mean zero and standard-deviation equal to one. Over 90% of the observations are within 2 standard deviations of zero. The Dickey–Fuller tests strongly reject the null of a unit root, thus indicating that the transformed predictors are stationary.

3. Nonparametric estimation of conditional portfolio choices

We start with a short discussion of the nonparametric technique. Then, Section 3.2 provides the mains results on optimal conditional portfolios. Section 3.3 studies the importance of higher-order moments for portfolio selection.

3.1. Nonparametric technique

The portfolio choice model in Section 2.1 does not have a closed-form solution. We solve this problem by using the nonparametric estimation technique of Brandt (1999) and Aït-Sahalia and Brandt (2001). It consists of replacing

[5] All the augmented Dickey-Fuller tests in this paper follow case 4 in table 17.3 of Hamilton (1994). That is, we report the OLS t test of $\rho = 1$ from the regression $yt = \rho y_{t-1} + c_0 + c_1 t + b_1 \Delta y_{t-1} + \ldots + b_p\ \Delta y_{t-p} + \varepsilon_t$. Following Mills (1999, p.74), we set $p = T$ 0.25. The 1% critical values are computed from MacKinnon response surfaces (see Mills, 1999, p.81).

[6] We thank the referee for pointing this out.

[7] Different studies use different transformations of these variables. For turnover for example, Campbell, Grossman, and Wang (1993) take logs of daily data and subtract a 1-year moving average (m.a.), Llorente, Michaely, Saar, and Wang (2002) subtract a 200-day m.a., Lo and Wang (2000) consider several different transformations, including the ratio to a 4-week m.a. Using daily volume, Andersen (1996) considers the ratio to a centered 2-year m.a.

Table 1
Descriptive statistics

	Small stocks portfolio				Large stocks portfolio			
	TURN	DVOL	PI	R_{t+1}	TURN	DVOL	PI	R_{t+1}
Panel A: Daily frequency								
Mean	0.00	0.00	0.00	1.00	0.00	0.00	0.00	1.00
Stdev	1.00	1.00	1.00	0.01	1.00	1.00	1.00	0.01
Skewness	0.26	0.26	0.03	−0.57	0.18	0.15	0.77	−0.93
Kurtosis	3.27	3.27	3.06	23.84	3.69	3.63	4.81	28.72
Percentiles								
5%	−1.58	−1.57	−1.64	0.99	−1.55	−1.57	−1.47	0.99
50%	−0.05	−0.03	−0.01	1.00	−0.03	−0.02	−0.08	1.00
95%	1.73	1.74	1.64	1.01	1.70	1.68	1.71	1.01
Cross correlations								
TURN	1.00				1.00			
DVOL	0.97	1.00			0.97	1.00		
PI	−0.60	−0.68	1.00		−0.30	−0.40	1.00	
R_{t+1}	0.13	0.12	−0.11	1.00	0.02	0.02	0.01	1.00
Auto correlations								
1-day lag	0.75	0.81	0.70	0.32	0.59	0.64	0.40	0.11
2-day lag	0.64	0.72	0.66	0.08	0.41	0.48	0.37	−0.03
5-day lag	0.57	0.66	0.65	0.08	0.34	0.41	0.36	−0.00
22-day lag	0.30	0.39	0.47	−0.00	0.10	0.16	0.22	−0.01
Unit root test								
Dickey–Fuller *t*-stat	−13.22	−11.62	−9.44	−25.25	−17.55	−15.79	−13.28	−30.72
1% critical value	−3.96	−3.96	−3.96	−3.96	−3.96	−3.96	−3.96	−3.96
Panel B: Monthly frequency								
Mean	0.00	0.00	0.00	1.03	0.00	0.00	0.00	1.01
Stdev	1.00	1.00	1.00	0.06	1.00	1.00	1.00	0.04
Skewness	0.30	0.30	−0.06	0.66	0.42	0.28	0.06	−0.10
Kurtosis	3.37	3.21	2.92	11.30	3.30	3.12	3.14	4.92
Percentiles								
5%	−1.56	−1.65	−1.71	0.93	−1.55	−1.60	−1.74	0.95
50%	−0.05	−0.03	−0.01	1.03	−0.05	−0.05	−0.04	1.02
95%	1.78	1.72	1.66	1.12	1.81	1.63	1.68	1.08
Cross correlations								
TURN	1.00				1.00			
DVOL	0.96	1.00			0.93	1.00		
PI	−0.63	−0.74	1.00		−0.30	−0.52	1.00	
R_{t+1}	0.08	0.05	−0.07	1.00	−0.04	−0.04	0.03	1.00
Auto correlations								
1-month lag	0.55	0.63	0.70	0.15	0.22	0.36	0.61	−0.02
3-month lag	0.01	0.08	0.15	−0.04	0.07	0.09	0.11	−0.00
6-month lag	−0.26	−0.23	−0.28	−0.02	−0.26	−0.26	−0.25	−0.05
12-month lag	0.04	0.02	0.19	0.16	0.23	0.21	0.14	0.06
Unit root test								
Dickey–Fuller *t*-stat	−10.37	−9.39	−9.37	−8.58	−11.52	−10.99	−9.95	−8.73
1% critical value	−3.98	−3.98	−3.98	−3.98	−3.98	−3.98	−3.98	−3.98

Descriptive statistics for the portfolio measures of liquidity and returns. The liquidity measures TURN, DVOL, and PI are constructed as described in Section 2.3, hence having a mean of 0 and a standard deviation of 1 by construction. The sample is from 1963 to 2000.

the conditional expectations with sample analogues. To estimate the optimal portfolio weight in a given reference state $Z=\overline{z}$, we find the number $\alpha(\overline{z})$ that solves the sample equivalent to the optimality conditions in (3):

$$\alpha(\overline{z}) : \frac{1}{\mathrm{Th}}\sum_{t=1}^{T} m_{t+1}(\alpha(\overline{z}))k(Z_t;\overline{z},h)\frac{1}{\sum_{t=1}^{T}k(Z_t;\overline{z},h)} = 0. \tag{5}$$

The kernel function $k(Z_t; \bar{z}, h)$ measures how far each sample observation is from the reference point $\bar{z}$, giving more weight to observations close to $\bar{z}$. We use a normal kernel: $k\,(Z_t;\, \bar{z},\, h)=(2\pi)^{-1/2}\, \exp(-d_t^2\,/\,2)$, with $d_t=(Z_t-\bar{z})\,/\,h$, where h is the bandwidth. We apply a standard exactly identified GMM to (5), i.e., $\alpha(\bar{z})$ is the number that sets the square of the left-hand side of (5) to zero.

The choice of the bandwidth h is crucial. A larger h implies averaging across more data points, thus reducing the variance but increasing the bias; a smaller h makes the estimator differentiate more between observations in different states and use fewer points, thus reducing the bias but increasing the variance. The conventional solution to optimize this tradeoff between variance and bias is to choose a bandwidth of the form $h=\lambda\sigma\,(Z)\,T^{-1\,/\,(\mathrm{K}+4)}$, where $\sigma\,(Z)$ is the standard deviation of the predictor Z, T is the sample size, K is the dimension of Z (one in our case), and λ is a constant.

There is no good guidance in the literature for bandwidth selection with non i.i.d. observations, as is the case with our data. For instance, Chaudhuri and Marron (1999) advocate a simple trial and error approach. In our application, a small λ usually produces very noisy portfolio weights that vary a lot with even small changes in liquidity, thus not making much economic sense; a big enough λ will eventually result in a flat portfolio weight, equal to the unconditional estimator. The results in Section 3.2 below are for values of λ of 9, 6, and 3 for respectively the daily, weekly, and monthly frequencies.[8] These values eliminate local noisy fluctuations, while still keeping the basic shape of the portfolio function. The parametric analysis in Section 4.2 ensures that the results do not depend on the particular values of λ used here.

The main advantage of Brandt's (1999) approach is that it permits estimating portfolio weights directly from the predictor variables. Whereas traditional studies in conditional portfolio choice usually start by estimating a model relating returns to forecasting variables and then find the portfolio weights given the conditional distribution of returns, we perform a single nonparametric estimation relating predictors directly to portfolio weights. The advantage is that we avoid the introduction of additional noise and potential misspecifications from the intermediate step of estimating the return distribution. Being nonparametric, this approach also has the advantage of giving a consistent estimator of the portfolio weights, but has the disadvantage of having higher variance than a correctly specified parametric estimator would have. However, lack of closed-form solutions prevent us from following the latter approach.

3.2. Conditional portfolio weights

We now present the main results of this paper: portfolio weights as functions of liquidity. Each one of the different predictors — price impact, dollar volume, and turnover — is analyzed in turn. For each of these measures, we present the results in a figure and a corresponding table.

Fig. 2 shows portfolio choices conditional on *price impact*. The bold line in each plot shows the proportion allocated to risky assets, α, as a function of the state variable Z (price impact in this case). The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (left column) or a portfolio of large stocks (right column). In the first row the portfolio is rebalanced daily; in the second row, weekly; and in the third, monthly.

The thin line in each plot represents the optimal unconditional portfolio weight. This is computed by applying a standard GMM procedure to the unconditional first order conditions:

$$\alpha : \frac{1}{T}\sum_{t=1}^{T} m_{t+1}(\alpha) = 0. \tag{6}$$

When the risky asset is a portfolio of large stocks, the allocation to stocks is close to 0.5 regardless of the investment horizon. More precisely, the point estimates (and standard errors) are: 0.51 (0.13) for daily, 0.50 (0.12) for weekly, and 0.53 (0.13) for monthly. For the case of small stocks, the allocation to the risky asset decreases from 1.29 (0.26 standard error) at the daily frequency, to 0.84 (0.17) at the weekly frequency, to 0.56 (0.10) at the monthly frequency.

[8] As an example to further interpret these numbers, with weekly data (T=1952, λ=6), Z=PI for small stocks, when computing the optimal weight at the mean, $\alpha(0)$, the weight given to an observation located exactly at the mean is $k(0;0,1.32)/\sum_{t=1}^{1952} k(0;0,1.32) = 6.42\times10^{-4}$, while an observation one standard deviation away from 0, $Z_t=\pm1$, gets a smaller weight of 4.81×10^{-4}, and for $Z_t=\pm2$ the weight is even smaller, 2.04×10^{-4}.

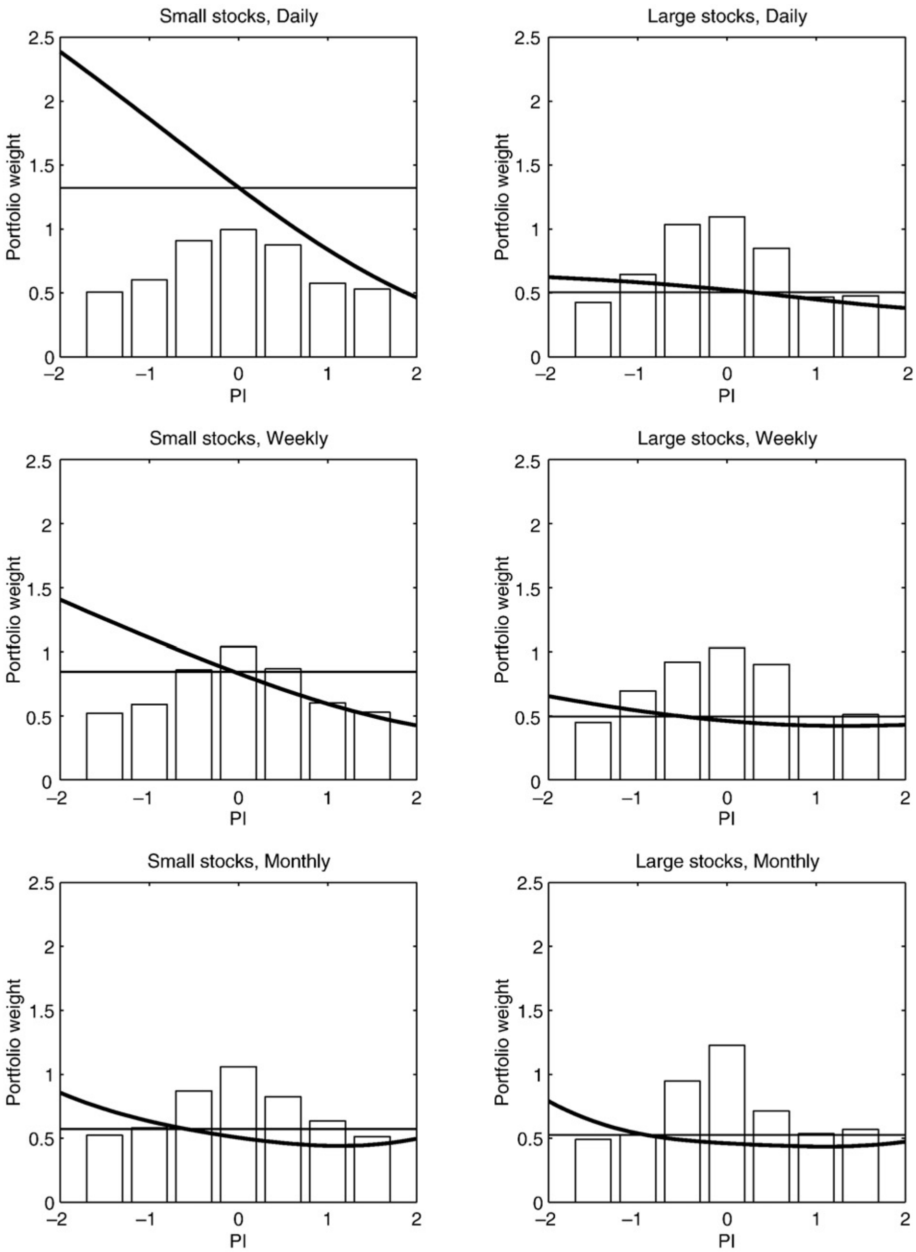

Fig. 2. Optimal Portfolio Weights as a Function of Price Impact. The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (left column) or a portfolio of large stocks (right column). In the first row, both the investment horizon and the rebalancing frequency are 1 day; in the second row, 1 week; and in the third, 1 month. The bold line in each panel represents the optimal fraction of wealth allocated to the respective stock portfolio as a function of Price Impact. Each point in this function is the solution to Eq. (5). The thin horizontal line represents the optimal unconditional allocation, which is given by Eq. (6). The bars in the background represent the histogram (scaled to add up to 5) of Price Impact. The sample is from 1963 to 2000.

Table 2 is the companion table to Fig. 2. Each panel shows $\alpha(Z)$ for the corresponding plot in Fig. 2. Standard errors are obtained with the stationary bootstrap of Politis and Romano (1994).[9] This technique accounts for the

[9] The algorithm consists of: (1) let the time indices I_1 , I_2 , ... be a sequence of i.i.d. random variables with a discrete uniform distribution on $\{1, ..., T\}$; (2) let the block lengths L_1 , L_2 , ... be a sequence of i.i.d. random variables with a geometric distribution; (3) generate pseudo-time series Z^* and R^* by stacking the overlapping blocks, $Z^* = [Z_{I_1}, \ldots, Z_{I_1+L_1-1}, Z_{I_2}, \ldots, Z_{I_2+L_2-1}, \ldots]$ and similarly for R^*, stopping when the length of these new series reaches T; (4) estimate α; (5) repeat B times; (6) compute the standard deviation of the B numbers α. We need to specify the mean of the geometric distribution in step 2, which determines the average block length. Following Horowitz (2000), we set it to $T^{1/5}$. We perform $B = 1000$ iterations.

Table 2
Optimal portfolio weights as a function of price impact

	Small stocks					Large stocks				
	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$
Panel A: Daily frequency										
Weight	2.39	1.86	1.33	0.84	0.46	0.63	0.58	0.53	0.45	0.38
	(0.32)	(0.28)	(0.25)	(0.23)	(0.23)	(0.16)	(0.14)	(0.14)	(0.15)	(0.17)
Slope	−0.51	−0.54	−0.52	−0.44	−0.31	−0.04	−0.05	−0.07	−0.08	−0.05
	{−5.36}	{−6.76}	{−7.61}	{−6.56}	{−4.38}	{−0.63}	{−0.98}	{−1.38}	{−1.45}	{−0.83}
Panel B: Weekly frequency										
Weight	1.40	1.11	0.83	0.59	0.42	0.65	0.54	0.46	0.42	0.43
	(0.23)	(0.20)	(0.18)	(0.17)	(0.17)	(0.15)	(0.13)	(0.12)	(0.13)	(0.14)
Slope	−0.31	−0.29	−0.26	−0.21	−0.13	−0.13	−0.10	−0.06	−0.01	0.02
	{−4.14}	{−5.21}	{−5.05}	{−3.67}	{−2.10}	{−1.98}	{−1.89}	{−1.25}	{−0.28}	{0.34}
Panel C: Monthly frequency										
Weight	0.85	0.63	0.50	0.44	0.50	0.79	0.54	0.46	0.43	0.47
	(0.29)	(0.16)	(0.13)	(0.13)	(0.15)	(0.24)	(0.17)	(0.15)	(0.16)	(0.21)
Slope	−0.29	−0.16	−0.10	−0.02	0.14	−0.38	−0.14	−0.04	−0.01	0.09
	{−0.74}	{−1.63}	{−1.08}	{−0.21}	{1.45}	{−2.08}	{−1.21}	{−0.39}	{−0.06}	{0.63}

The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (columns under “Small Stocks”) or a portfolio of large stocks (columns under “Large Stocks”). In panel A, both the investment horizon and the rebalancing frequency are 1 day; in panel B, 1 week; and in C, 1 month. Each panel presents the optimal allocation to the respective stock portfolio conditional on the value of Price Impact indicated in the column heading. The optimal weight is the solution to Eq. (5). Standard-errors from the bootstrap detailed in Section 3.2 are in parenthesis. The slope at each value of the predictor variable is obtained from Eq. (7). Bootstrap *t*-statistics for a zero slope are in curly braces. The sample is from 1963 to 2000.

autocorrelation in the data and has the advantage, over a simple block bootstrap, of generating a stationary resampled pseudo-time series.[10] The standard errors are presented only to gauge the precision of the nonparametric method used; it is obviously meaningless to test whether a portfolio weight is statistically away from zero. The important question is whether the weights respond to changes in liquidity. Hence Table 2 also presents the first derivative of $\alpha(Z)$, approximated by the finite difference

$$\frac{\partial \alpha(Z)}{\partial Z}\Big|_{Z=\bar{z}} \cong \frac{\alpha(\bar{z}+0.1)-\alpha(\bar{z}-0.1)}{0.2}. \tag{7}$$

Stationary bootstrap *t*-statistics for a zero slope are presented in curly braces below the point estimates.

Several results emerge. First, at the shorter *daily* horizon, an increase in liquidity (i.e., a decrease in price impact) is accompanied by a strong increase in the optimal allocation to small stocks: from $\alpha=0.46$ for $\mathrm{PI}=2$, the optimal weight rises to $\alpha=2.39$ for $\mathrm{PI}=-2$, with all slopes statistically significant at the 1% level or better. For large stocks, the optimal weight does not seem to depend on price impact: the conditional weight is very close to the unconditional and the slopes are not significantly different from zero. Second, at the weekly frequency, the optimal allocation to small stocks still responds strongly to changes in price impact: the optimal weight rises from $\alpha=0.42$ for $\mathrm{PI}=2$, to $\alpha=1.40$ for $\mathrm{PI}=-2$, with all slopes statistically significant. Third, at the monthly frequency, the optimal allocation to either small or large stocks is mostly irresponsive to price impact.[11]

[10] For another application of the stationary bootstrap see Sullivan, Timmermann, and White (1999). Even though there is some recent research on the asymptotic properties of the bootstrap applied to GMM estimators (for example, Horowitz (2000) and Inoue and Shintani (2001) present cases where the bootstrap provides asymptotic refinements over first-order approximations), little is known about the bootstrap properties when GMM estimators run on top of a kernel function (which is our case).

[11] The results are robust to reasonable changes in the bandwidth. For example, with weekly data for small stocks, if we increase the bandwidth (h) by 25%, the optimal weights range from $\alpha(\mathrm{PI}=-2)=1.24$ to $\alpha(\mathrm{PI}=+2)=0.52$, whereas if we decrease h by 25%, the optimal weights range from $\alpha(\mathrm{PI}=-2)=1.68$ to $\alpha(\mathrm{PI}=+2)=0.32$. In all cases, the slopes are negative at the 1% significance level or better (except for an insignificant slope at PI=2 with the smaller bandwidth). Furthermore, the portfolio functions in this section are consistent with the parametric estimators in Section 4.2, which do not depend on any bandwidth.

The second measure of liquidity is *dollar volume*. The results are in Fig. 3 and Table 3. Overall, the results for dollar volume are consistent with the previous findings for price impact. First, at the *daily* frequency the optimal weight in small stocks increases strongly with liquidity, i.e. with dollar volume. The optimal weight rises from $\alpha = 0.43$ for DVOL$=-2$ to $\alpha = 1.80$ for DVOL $= 2$. This effect is stronger in the negative range of dollar volume: the

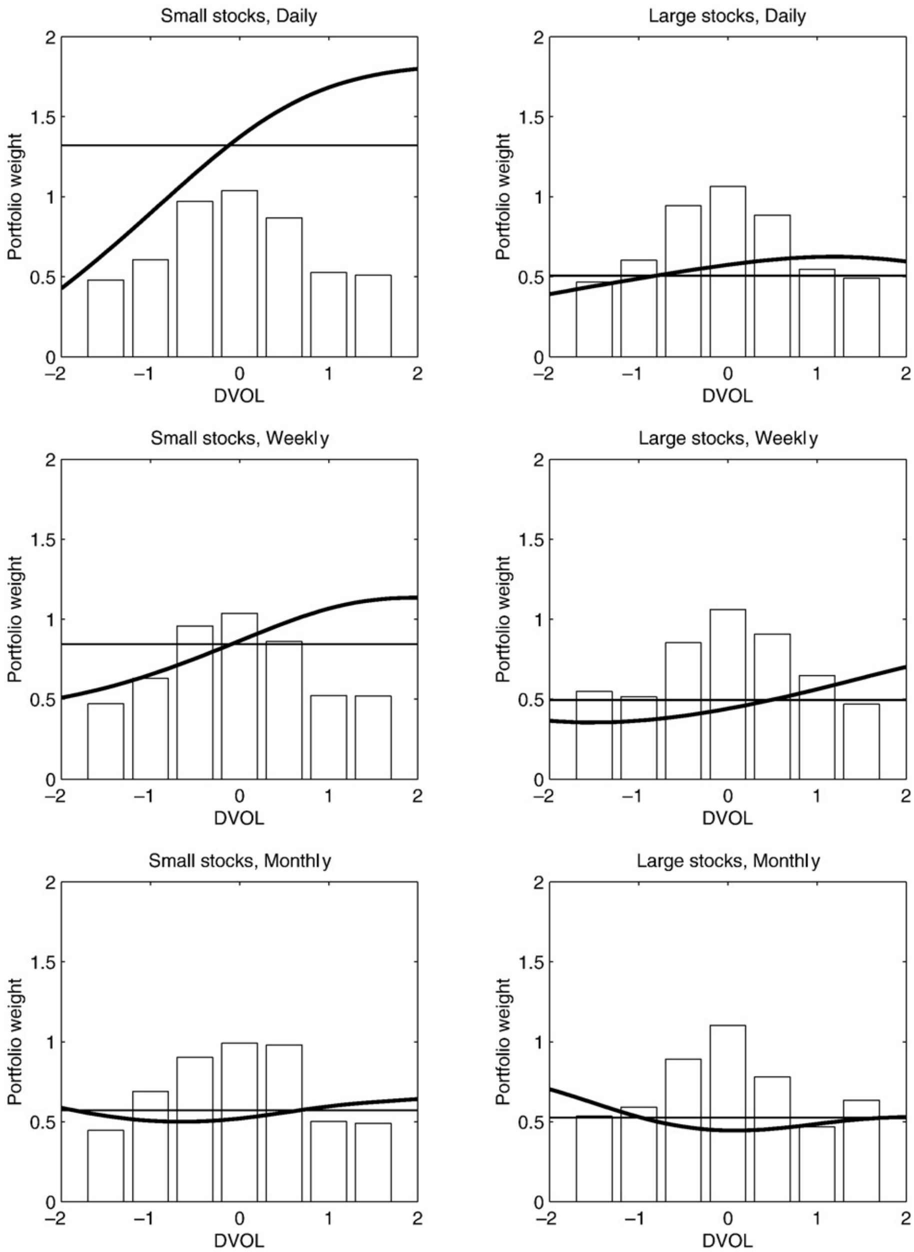

Fig. 3. Optimal Portfolio Weights as a Function of Dollar Volume. The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (left column) or a portfolio of large stocks (right column). In the first row, both the investment horizon and the rebalancing frequency are 1 day; in the second row, 1 week; and in the third, 1 month. The bold line in each panel represents the optimal fraction of wealth allocated to the respective stock portfolio as a function of Dollar Volume. Each point in this function is the solution to Eq. (5). The thin horizontal line represents the optimal unconditional allocation, which is given by Eq. (6). The bars in the background represent the histogram (scaled to add up to 5) of Dollar Volume. The sample is from 1963 to 2000.

Table 3
Optimal portfolio weights as a function of dollar volume

	Small stocks					Large stocks				
	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$
Panel A: Daily frequency										
Weight	0.43	0.90	1.37	1.68	1.80	0.39	0.49	0.57	0.62	0.60
	(0.22)	(0.21)	(0.24)	(0.34)	(0.53)	(0.15)	(0.13)	(0.13)	(0.14)	(0.19)
Slope	0.43	0.50	0.42	0.20	0.05	0.10	0.09	0.07	0.02	−0.07
	{5.88}	{6.24}	{3.12}	{1.00}	{0.25}	{1.74}	{1.92}	{1.47}	{0.23}	{−0.78}
Panel B: Weekly frequency										
Weight	0.51	0.65	0.86	1.07	1.13	0.36	0.37	0.44	0.56	0.70
	(0.17)	(0.16)	(0.17)	(0.19)	(0.29)	(0.14)	(0.13)	(0.12)	(0.13)	(0.15)
Slope	0.10	0.18	0.23	0.15	−0.01	−0.05	0.04	0.10	0.14	0.13
	{2.11}	{3.89}	{3.99}	{1.54}	{−0.05}	{−0.73}	{0.92}	{2.40}	{2.93}	{1.84}
Panel C: Monthly frequency										
Weight	0.59	0.51	0.52	0.60	0.64	0.70	0.53	0.45	0.49	0.53
	(0.17)	(0.14)	(0.13)	(0.15)	(0.26)	(0.20)	(0.15)	(0.14)	(0.16)	(0.24)
Slope	−0.12	−0.04	0.06	0.07	0.05	−0.15	−0.15	−0.01	0.07	0.00
	{−0.97}	{−0.47}	{0.75}	{0.50}	{0.21}	{−1.04}	{−1.49}	{−0.12}	{0.58}	{0.01}

The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (columns under "Small Stocks") or a portfolio of large stocks (columns under "Large Stocks"). In panel A, both the investment horizon and the rebalancing frequency are 1 day; in panel B, 1 week; and in C, 1 month. Each panel presents the optimal allocation to the respective stock portfolio conditional on the value of Dollar Volume indicated in the column heading. The optimal weight is the solution to Eq. (5). Standard-errors from the bootstrap detailed in Section 3.2 are in parenthesis. The slope at each value of the predictor variable is obtained from Eq. (7). Bootstrap *t*-statistics for a zero slope are in curly braces. The sample is from 1963 to 2000.

slopes are significantly positive from DVOL=−2 up to DVOL = 0. For large stocks, there is weak evidence of a small increase in the portfolio weight over the negative range of DVOL. Second, at the weekly frequency, the optimal allocation to small stocks still increases with DVOL, particularly over the negative range. For large stocks, again we find some evidence of a small increase, particularly in the center of the distribution. Third, at the monthly frequency optimal portfolio weights for both classes of stocks do not seem to respond to changes in volume.

Fig. 4 and Table 4 present the results for the third liquidity measure, turnover. Again, the results are consistent with the previous liquidity measures. First, at the daily horizon, turnover is a strong determinant of the allocation to small stocks, with the optimal weight ranging from 0.46 at TURN=−2 to 1.74 at TURN = 2. The increase of α with liquidity is particularly strong over the negative range of turnover, i.e., the slopes are all positive and statistically significant from TURN=−2 to TURN = 0. Second, at the weekly frequency, we again see a positive relation in small stocks (significant slopes from TURN=−2 to TURN = 1). Interestingly, the allocation to large stocks also seems to increase somewhat at the weekly frequency (significant slopes on the positive range of turnover). Third, at the monthly frequency, the optimal portfolio weight is no longer sensitive to changes in turnover.

To analyze whether these results are robust to different degrees of risk aversion (γ), we re-estimate the optimal portfolio weights for $\gamma = 3$, $\gamma = 5$, and $\gamma = 20$. Table 5 presents the results for price impact.[12] We observe that different degrees of risk-aversion change mainly the level of the portfolio function (lower γ, higher α), having little effect on the shape of this function. Hence, our conclusions regarding the relation between optimal weights and liquidity variables are not determined by the level of risk aversion.

[12] Tables for the other liquidity variables are available from the authors upon request.

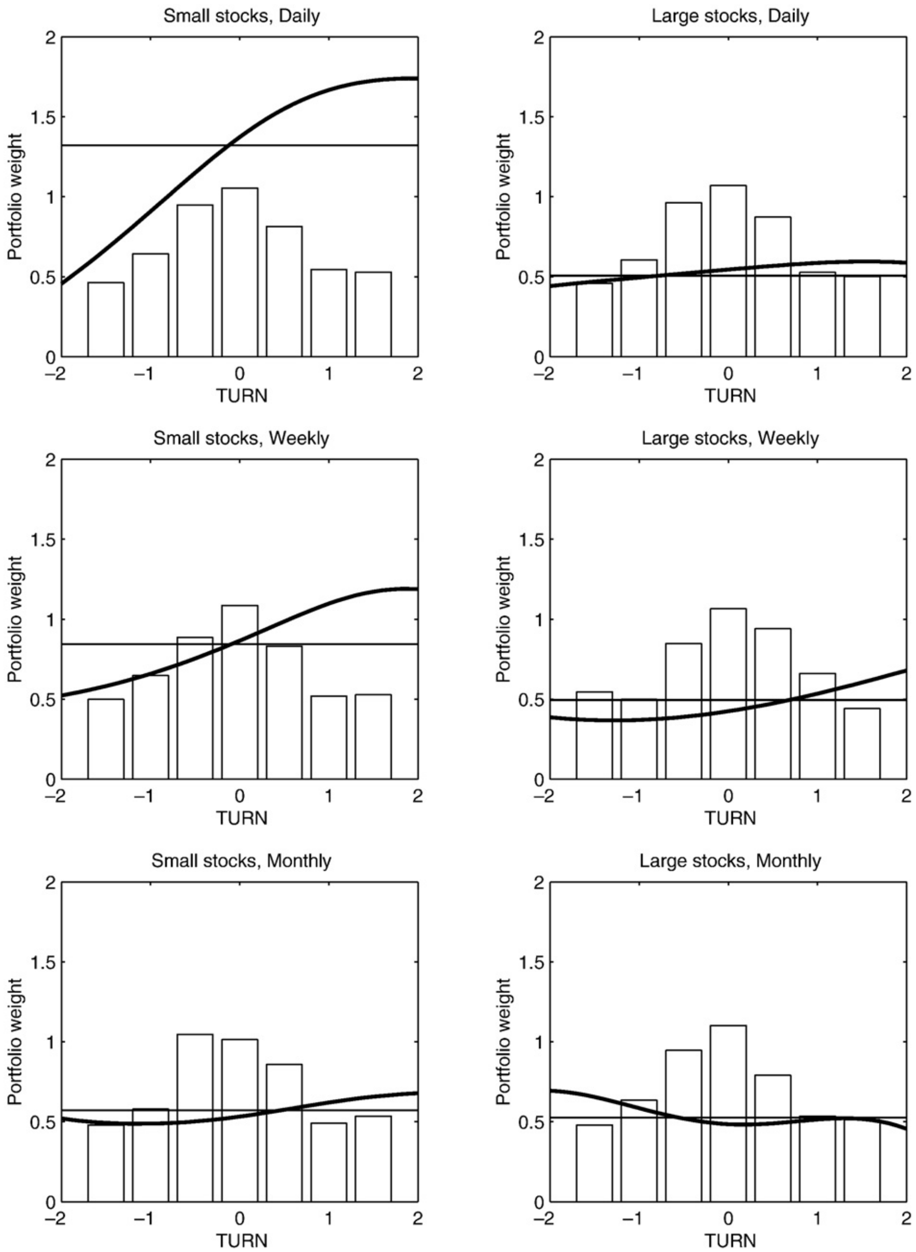

Fig. 4. Optimal Portfolio Weights as a Function of Turnover. The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (left column) or a portfolio of large stocks (right column). In the first row, both the investment horizon and the rebalancing frequency are 1 day; in the second row, 1 week; and in the third, 1 month. The bold line in each panel represents the optimal fraction of wealth allocated to the respective stock portfolio as a function of Turnover. Each point in this function is the solution to Eq. (5). The thin horizontal line represents the optimal unconditional allocation, which is given by Eq. (6). The bars in the background represent the histogram (scaled to add up to 5) of Turnover. The sample is from 1963 to 2000.

To summarize the results in this section, the three liquidity measures consistently show that the optimal allocation to small stocks increases in liquidity at the shorter daily and weekly frequencies. For large stocks, the optimal weight is much less sensitive to liquidity. At the longer monthly frequency, liquidity does not influence the optimal weight. Finally, we stress that the main findings in this section regard the shape of the optimal portfolio functions; the actual level obviously depends on the degree of risk aversion.

Table 4
Optimal portfolio weights as a function of turnover

	Small stocks					Large stocks				
	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$
Panel A: Daily frequency										
Weight	0.46	0.91	1.37	1.67	1.74	0.44	0.50	0.54	0.59	0.59
	(0.23)	(0.22)	(0.23)	(0.32)	(0.50)	(0.16)	(0.14)	(0.14)	(0.14)	(0.17)
Slope	0.40	0.48	0.41	0.17	−0.00	0.06	0.05	0.05	0.03	−0.03
	{5.36}	{6.17}	{3.31}	{0.87}	{−0.02}	{1.00}	{1.07}	{1.05}	{0.48}	{−0.35}
Panel B: Weekly frequency										
Weight	0.52	0.66	0.87	1.10	1.19	0.39	0.37	0.42	0.53	0.68
	(0.17)	(0.16)	(0.16)	(0.18)	(0.26)	(0.14)	(0.13)	(0.13)	(0.13)	(0.15)
Slope	0.10	0.17	0.23	0.20	−0.04	−0.06	0.02	0.09	0.13	0.15
	{2.00}	{3.67}	{4.55}	{2.44}	{−0.19}	{−0.81}	{0.42}	{1.95}	{3.06}	{2.46}
Panel C: Monthly frequency										
Weight	0.52	0.49	0.53	0.62	0.68	0.69	0.58	0.48	0.51	0.46
	(0.17)	(0.15)	(0.14)	(0.16)	(0.26)	(0.19)	(0.14)	(0.14)	(0.16)	(0.23)
Slope	−0.08	0.01	0.08	0.09	0.03	−0.02	−0.15	−0.03	0.05	−0.21
	{−0.72}	{0.14}	{0.86}	{0.70}	{0.14}	{−0.16}	{−1.40}	{−0.36}	{0.43}	{−1.10}

The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (columns under "Small Stocks") or a portfolio of large stocks (columns under "Large Stocks"). In panel A, both the investment horizon and the rebalancing frequency are 1 day; in panel B, 1 week; and in C, 1 month. Each panel presents the optimal allocation to the respective stock portfolio conditional on the value of Turnover indicated in the column heading. The optimal weight is the solution to Eq. (5). Standard-errors from the bootstrap detailed in Section 3.2 are in parenthesis. The slope at each value of the predictor variable is obtained from Eq. (7). Bootstrap *t*-statistics for a zero slope are in curly braces. The sample is from 1963 to 2000.

3.3. The importance of higher-order moments

Our analysis so far implicitly assumes that the investor cares about higher order moments of the return distribution. To assess the relevance of this assumption, we now compute the optimal portfolio holdings for an

Table 5
Optimal portfolio weights as a function of price impact for different degrees of risk aversion

	Small stocks					Large stocks				
	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$	$Z=-2$	$Z=-1$	$Z=0$	$Z=1$	$Z=2$
Panel A: Daily frequency										
$\gamma=3$	5.99	5.22	4.08	2.72	1.53	2.07	1.93	1.73	1.48	1.26
$\gamma=5$	4.31	3.51	2.58	1.66	0.93	1.25	1.16	1.05	0.89	0.76
$\gamma=10$	2.39	1.86	1.33	0.84	0.46	0.63	0.58	0.53	0.45	0.38
$\gamma=20$	1.24	0.95	0.67	0.42	0.23	0.31	0.29	0.26	0.23	0.19
Panel B: Weekly frequency										
$\gamma=3$	3.98	3.33	2.62	1.92	1.39	2.16	1.79	1.52	1.40	1.42
$\gamma=5$	2.65	2.13	1.63	1.17	0.84	1.30	1.08	0.92	0.84	0.86
$\gamma=10$	1.40	1.11	0.83	0.59	0.42	0.65	0.54	0.46	0.42	0.43
$\gamma=20$	0.72	0.56	0.42	0.30	0.21	0.33	0.27	0.23	0.21	0.21
Panel C: Monthly frequency										
$\gamma=3$	2.50	1.93	1.58	1.40	1.58	2.55	1.76	1.50	1.42	1.54
$\gamma=5$	1.63	1.22	0.99	0.87	0.98	1.56	1.07	0.91	0.86	0.94
$\gamma=10$	0.85	0.63	0.50	0.44	0.50	0.79	0.54	0.46	0.43	0.47
$\gamma=20$	0.44	0.32	0.25	0.22	0.25	0.40	0.27	0.23	0.22	0.24

The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (columns under "Small Stocks") or a portfolio of large stocks (columns under "Large Stocks"). In panel A, both the investment horizon and the rebalancing frequency are 1 day; in panel B, 1 week; and in C, 1 month. Each panel presents the optimal allocation to the respective stock portfolio conditional on the value of price impact indicated in the column heading. The optimal weight is the solution to Eq. (5). The coefficient γ defines the degree of risk aversion of the CRRA utility function in Eq. (2). The sample is from 1963 to 2000.

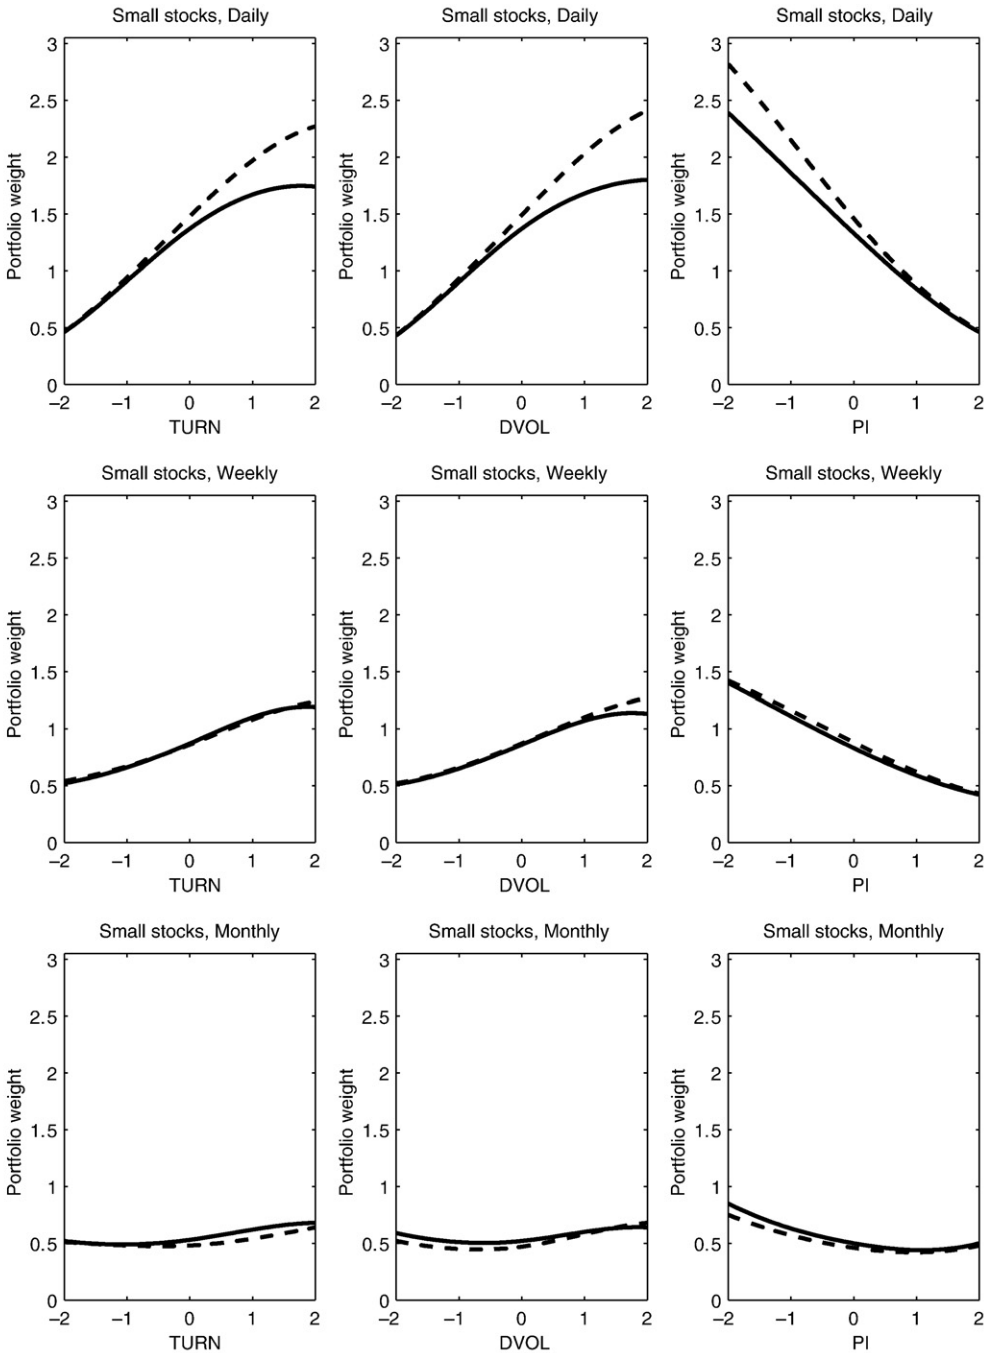

Fig. 5. Portfolio weights for lower-order moments. Optimal portfolio allocation to small stocks as a function of three different liquidity measures. In the first row, the investment horizon is 1 day; in the second row, 1 week; third row, 1 month. The solid lines represent the optimal nonparametric functions for a CRRA utility function, as computed in Section 3.2. The dashed lines represent the optimal functions for an investor that only cares about the first two moments of returns, as in Eq. (8). The sample is from 1963 to 2000.

investor with simplified preferences. In particular, we assume that the investor only cares about the first two moments of returns.

Consider a Taylor expansion of the utility function around $R_f : u(W) = u(R_f) + u'(R_f)(W - R_f) + \frac{1}{2}u''(R_f)(W - R_f)^2 \ldots$ Ignoring higher-order terms, imposing the budget constraint, and normalizing $W_t = 1$, the investor's goal now becomes to maximize$_{\alpha_t}\left\{u(R_f)+u'(R_f)E\left[R_{t+1,s}-R_f|Z_t\right]\alpha_t+\frac{1}{2}u''(R_f)E\left[\left(R_{t+1,s}-R_f\right)^2|Z_t\right]\alpha_t^2\right\}$, which has the following first-order condition:

$$\alpha_t=\frac{R_f}{\gamma}\frac{E\left[R_{t+1,s}-R_f|Z_t\right]}{E\left[\left(R_{t+1,s}-R_f\right)^2|Z_t\right]}. \tag{8}$$

The two conditional expectations in (8) are estimated nonparametrically. Fig. 5 compares this new portfolio function with the one obtained above in Section 3.2. We present the optimal portfolio allocations to small stocks, conditional on the three liquidity measures defined above. At the daily frequency, we observe that the two portfolio functions become substantially different for states of high liquidity, with (8) being larger than the CRRA weight. For example, for TURN = 1 the optimal allocation increases from 1.67 to 1.97 when the investor cares only about the first two moments. This is because high values of turnover forecast "adverse values" for higher-order moments. For example, conditional on TURN > 1, the stock portfolio return has skewness = −1.96 and kurtosis = 59.12, which compares with unconditional values of −0.57 and 23.84, respectively (see Table 1). Since the mean and variance are not greatly changed, the portfolio weight given by (8) is substantially larger than the CRRA weight. At the weekly and monthly frequencies however the results show that both portfolio functions are quite similar. Hence, for weekly and monthly horizons, higher moments of returns do not seem to matter much.

Alternatively, we could start by assuming mean-variance preferences (i.e., a quadratic utility function), which would also lead to a simple expression relating α_t to the first two moments of returns. The optimal portfolio functions in this case have quite similar shapes to the ones from (8), but usually not the same level (not shown, available upon request). The Taylor series approach followed above preserves the interpretation of the coefficient of relative risk aversion, γ, therefore making the results directly comparable to the fully nonparametric portfolio weights.

4. Robustness and discussion

4.1. Non-synchronous trading

The return series for the small-stock portfolio may be influenced by non-synchronous trading. For example, Boudoukh, Richardson, and Whitelaw (1994) find that non-synchronous trading induces serial autocorrelation in weekly returns of small-stock portfolios. The basic intuition is that most stocks respond in the same direction to economic news, but since not all small stocks trade each day, the news will take a few days to be incorporated into all stock prices, leading to spurious return autocorrelation at the aggregated portfolio level.[13]

Therefore, one possible concern with our findings is that non-synchronous trading might somehow explain the optimal portfolio allocation to small stocks at the daily and weekly horizons (see Section 3.2). We address this concern in three ways.

First, note that we are using Liquidity to predict future returns, not returns today to predict returns tomorrow. This might still be an issue if liquidity was just a proxy for returns. However, the contemporaneous cross-correlation between daily liquidity (Z_d) and daily small stocks returns (R_d) is low: −0.16 for PI, 0.20 for DVOL, and 0.19 for TURN. Hence, it does not seem that liquidity today is proxying for returns today.

Second, we follow Chordia and Swaminathan (2000) and exclude the return (R_d) and liquidity (Z_d) of any stock that did not trade at day d or $d-1$.[14] This ensures that we exclude returns that may be due to information released yesterday.

[13] The literature is not unanimous on the importance of non-synchronous trading. Lo and MacKinlay (1990) and other references cited in Boudoukh, Richardson, and Whitelaw (1994) argue that only a small part of the autocorrelations can be explained by non-synchronous trading.

[14] This resulted in only a very small loss in valid observations: from 3,213,649 to 3,095,766 daily individual stock returns.

The optimal allocation to this new portfolio of small stocks (deciles 2 and 3 of NYSE) as a function of PI is the following (daily horizon):

	PI = −2	PI = −1	PI = 0	PI = 1	PI = 2
Weight	2.36	1.86	1.35	0.88	0.51
(standard error)	(0.31)	(0.28)	(0.25)	(0.23)	(0.23)
Slope	−0.48	−0.51	−0.50	−0.43	−0.31
{t-stat}	{−5.17}	{−6.58}	{−7.44}	{−6.56}	{−4.52}

This new portfolio function is very similar to the one in Section 3.2 and still strongly decreasing in PI, thus suggesting that non-synchronous trading is not the reason for our results. The portfolio functions for DVOL and TURN are also consistent with our previous results.

Third, as an alternative test, we follow Gervais, Kaniel, and Mingelgrin (2001) and redefine the small stock portfolio to include deciles 2 through 5. Adding deciles 4 and 5, which represent larger and likely more frequently traded stocks, should reduce the relative importance of any non-synchronous trading. The optimal daily allocation is the following:

	PI = −2	PI = −1	PI = 0	PI = 1	PI = 2
Weight	2.31	1.80	1.28	0.81	0.43
(standard error)	(0.34)	(0.29)	(0.26)	(0.24)	(0.24)
Slope	−0.49	−0.52	−0.51	−0.43	−0.32
{*t*-stat}	{−4.93}	{−5.80}	{−6.61}	{−6.02}	{−4.33}

Again, this portfolio function is consistent with the results above, further stressing that non-synchronous trading is not the reason for our results. The results for DVOL and TURN are similar.

4.2. Parametric estimation of conditional portfolio functions

The nonparametric approach used so far has the big advantage of producing a consistent estimator of the portfolio functions (see Section 3). Therefore, we can safely draw conclusions regarding the overall shape of these functions. However, the procedure is subject to high sampling error, as can be seen in the relatively high standard deviations of the conditional allocations. Furthermore, little is known about the properties of the bootstrap when applied to “GMM-with-Kernel” estimators.

In view of these considerations, we now proceed to estimate parametric portfolio functions. This will provide more precise estimators to confirm the results above. Given the smooth shapes of the portfolio functions observed in Section 3.2, we assume that α(Z) can be approximated by a third degree polynomial in Z:

$$\alpha_{par}(Z) \equiv c_0 + c_1 Z + c_2 Z^2 + c_3 Z^3 \tag{9}$$

This specification is parsimonious and is able to capture all the interesting features re-ported in Section 3.2. Namely, (9) can produce increasing, decreasing, concave, or convex portfolio functions. Though we could conjecture more general parametric functions, for example capable of having multiple inversion points, there would be little economic motivation for doing so. In fact, given that our three predictors (price impact, dollar volume, and turnover) are continuous variables, it seems reasonable to assume that optimal policies should not have abrupt changes (discontinuities), kinks, or too much variation. The constants $c_0 ... c_3$ are estimated through GMM on the moment conditions

$$\frac{1}{T}\sum_{t=1}^{T} m_{t+1}\left(\alpha_{par}(Z_t)\right) \otimes g(Z_t) = 0 \tag{10}$$

with $g(Z_t) = [1, Z_t, Z_t^2, Z_t^3]$.

Table 6
Parametric estimators of conditional portfolio weights

	Small stocks				Large stocks			
	c_0	c_1	c_2	c_3	c_0	c_1	c_2	c_3
Panel A: Daily frequency								
Z=PI	1.66	−1.86	0.12	0.15	0.57	−0.31	−0.04	0.05
	(6.10)	(−6.24)	(1.11)	(2.33)	(3.48)	(−1.93)	(−0.41)	(1.61)
Z=DVOL	1.35	0.86	−0.21	−0.04	0.66	0.26	−0.10	−0.02
	(5.19)	(3.21)	(−2.42)	(−1.00)	(4.56)	(1.33)	(−2.02)	(−0.63)
Z=TURN	1.39	1.12	−0.22	−0.11	0.61	0.22	−0.06	−0.01
	(5.08)	(3.35)	(−2.73)	(−2.02)	(4.00)	(1.33)	(−1.09)	(−0.63)
Panel B: Weekly frequency								
Z=PI	0.87	−0.95	0.17	0.08	0.38	−0.06	0.16	−0.05
	(4.56)	(−4.83)	(2.13)	(1.59)	(2.53)	(−0.40)	(1.91)	(−1.44)
Z=DVOL	1.00	0.68	−0.03	−0.11	0.41	0.53	0.12	−0.09
	(5.59)	(3.40)	(−0.46)	(−2.16)	(3.07)	(2.77)	(1.80)	(−2.23)
Z=TURN	0.95	0.88	0.07	−0.13	0.38	0.40	0.11	−0.06
	(5.26)	(4.07)	(1.02)	(−3.25)	(2.71)	(2.23)	(1.71)	(−1.64)
Panel C: Monthly frequency								
Z=PI	0.36	−0.22	0.32	−0.03	0.33	−0.05	0.24	−0.07
	(3.08)	(−1.15)	(3.44)	(−0.46)	(2.19)	(−0.25)	(2.26)	(−1.26)
Z=DVOL	0.48	0.01	0.09	−0.01	0.42	−0.13	0.10	0.03
	(3.58)	(0.08)	(1.03)	(−0.11)	(2.62)	(−0.65)	(1.04)	(0.43)
Z=TURN	0.50	0.10	0.07	−0.02	0.54	−0.05	−0.01	−0.02
	(4.02)	(0.61)	(0.99)	(−0.57)	(3.78)	(−0.28)	(−0.09)	(−0.34)

The investor allocates his wealth between a risk-free asset and either a portfolio of small stocks (columns under "Small Stocks") or a portfolio of large stocks (columns under "Large Stocks"). In panel A, both the investment horizon and the rebalancing frequency are 1 day; in panel B, 1 week; and in C, 1 month. Each panel gives estimates of the parameters c_i for the optimal portfolio function defined in Eq. (9). These estimates are obtained through GMM on the moment conditions (10), using the Newey–West estimator of the spectral density matrix. t-statistics are under the point estimates within parenthesis. The sample is from 1963 to 2000.

Table 6 shows the estimation results. We focus the analysis on the slope of the portfolio function in the center of the distribution, $\alpha_{\text{par}}{}' (Z)|_{z=0}=c_1$. First, at the *daily* frequency we observe that the optimal portfolio allocation to small stocks increases with liquidity. In particular, for price impact the optimal portfolio allocation has a strong downward slope (high and strongly significant $c_1=-1.86$). For DVOL and TURN, the c_1 coefficient is positive and also strongly significant. Second, at the weekly frequency the results are similar: the optimal portfolio functions increase with liquidity. Third, at the monthly the optimal weights no longer respond to changes in liquidity.

Hence, these parametric results support the nonparametric findings in Section 3.2. We conclude that the optimal allocation to small stocks increases with liquidity at the shorter daily and weekly frequencies, but does not respond to liquidity at the longer monthly frequency. The allocation to large stocks is mostly irresponsive to liquidity.

4.3. Forecasting the first two moments of returns

Given that Amihud (2002) finds that liquidity forecasts expected returns 1 month ahead (particularly for small stocks), it is perhaps surprising that the optimal portfolio allocations for a 1-month horizon do not depend on liquidity. To better understand this result, we proceed in two steps.

First, we estimate the optimal portfolio allocation to small stocks at the monthly frequency using the precise measure of price impact used in Amihud (2002). Amihud's measure differs from ours in that he uses only $\ln(pi_m)$; he does not subtract a moving average as we do in Eq. (4).[15] The results are the following:

	PI = −2	PI = −1	PI = 0	PI = 1	PI = 2
Weight	0.56	0.53	0.56	0.62	0.66
(standard error)	(0.25)	(0.17)	(0.13)	(0.14)	(0.24)
Slope	−0.04	−0.00	0.05	0.06	0.01
{*t*-stat}	{−0.08}	{−0.05}	{0.57}	{0.59}	{0.01}

Hence, the portfolio function is also not responsive to Amihud's measure of liquidity (all slopes are statistically insignificant). This suggests that if price impact is forecasting an higher mean return, it must also be forecasting a higher variance (or other higher-order moments that the investor dislikes).

Second, to analyze this hypothesis, we estimate forecasting regressions of the first two moments of monthly returns of the small stocks portfolio. Again, in addition to our measure of price impact, we also study the precise measure of price impact used in Amihud (2002). The simplest forecasting regression is

$$y_{t+1} = \beta_0 + \beta_1 Z_t + \varepsilon_{t+1}$$

The dependent variable can either be $y_{t+1} = R^{e}_{t+1}$ or $y_{t+1} = (R^{e}_{t+1})^2$, where R^{e}_{t+1} is the excess return on the small-stock portfolio, $R^{e}_{t+1} \equiv R_{t+1,\text{small}} - R_{t+1,f}$, and the return on a risk-free 1-month Treasury bill, $R_{t+1,f}$, is known at time t. Z_t can either be Amihud's or our measure of price impact. In either case, since both Z_{t+1} and y_{t+1} are functions of prices at time $t+1$, we have that $\mathrm{E}[\varepsilon_{t+1}|Z_{t+1}, Z_t] \neq 0$. This leads to the ordinary least squares (OLS) estimator being biased in finite samples, as shown in Stambaugh (1999).

We use the method of Amihud and Hurvich (2004) to correct this bias.[16] The method consists of starting by estimating an AR(1) for the regressor:

$$Z_{t+1} = \alpha + \rho Z_t + v_{t+1}$$

The OLS estimator $\hat{\rho}$ is then used to produce a second-order bias corrected estimator, $\hat{\rho}^{c} = \hat{\rho} + (1+3\hat{\rho})/T + 3(1+3\hat{\rho})/T^2$. This $\hat{\rho}^{c}$ is used to generate corrected residuals of the AR(1) process, v_t. Finally, these corrected residuals, v_t, are added to the forecasting regression,

$$y_{t+1} = \beta_0 + \beta_1 Z_t + \beta_2 v_{t+1} + \varepsilon_{t+1}$$

The OLS estimator of β_1 is the reduced-bias estimator of Amihud and Hurvich (2004). The reduced-bias estimator of the standard-error of $\hat{\beta}_1$ is given by their Eq. (10). Further, for comparison with Amihud (2002), we also consider a specification with a January dummy to account for the well-known January effect:

$$y_{t+1} = \beta_0 + \beta_1 Z_t + \beta_2 v_{t+1} + \beta_3 \mathrm{JanDum}_{t+1} + \varepsilon_{t+1}.$$

Table 7 shows the results. Regarding our measure of price impact, we find that it is at best a weak predictor of the mean (not significant when we consider the January dummy) and that it is not a significant predictor of the second moment. Hence, the portfolio function at the monthly frequency does not depend on PI. On the contrary, we find that Amihud's measure has a strongly significant positive coefficient in both the mean and the variance equation. Hence, the reason why the portfolio allocation does not change with Amihud's price impact is very different: a higher mean that the investor likes is being compensated by an also higher variance that the investor dislikes. The two effects cancel each other out.

[15] It is not clear whether Amihud's measure is stationary. The Augmented Dickey-Fuller test for a unit root produces a t-statistic of −3.39, which slightly fails to reject at the 5% critical level (−3.42). Even though we are not able to strongly reject a unit root, we still proceed with our nonparametric estimation.

[16] This method is straightforward to implement and can be easily extended to multivariate regressions.

Table 7
Forecasting the first two moments of returns

	Z=PI			Z=Amihud		
	OLS	Reduc-bias	Reduc-bias	OLS	Reduc-bias	Reduc-bias
Panel A: Dependent variable $y_{t+1}=R^e_{t+1,\ small}$						
c	2.036	2.036	1.643	−3.312	−3.312	−3.987
	(7.33)	(8.19)	(6.52)	(−1.60)	(−1.74)	(−2.16)
Z_t	−0.458	−0.484	−0.273	0.870	0.800	0.844
	(−1.65)	(−1.74)	(−1.00)	(2.62)	(2.40)	(2.62)
v_{t+1}		−3.685	−3.549		−8.212	−7.916
		(−10.6)	(−10.5)		(−8.82)	(−8.79)
JanDum_{t+1}			4.751			5.187
			(5.34)			(5.81)
R^2	0.006	0.207	0.255	0.015	0.160	0.218
D–W	1.74	2.21	2.15	1.70	2.12	2.10
Panel B: Dependent variable $y_{t+1}=(R^e_{t+1,\ small})^2$						
c	38.78	38.78	30.05	−145.4	−145.4	−159.4
	(6.95)	(6.95)	(5.31)	(−3.57)	(−3.57)	(−4.04)
Z_t	−0.619	−0.696	3.996	29.87	29.84	30.74
	(−0.11)	(−0.12)	(0.73)	(4.56)	(4.55)	(4.84)
v_{t+1}		−11.01	−7.983		−3.622	2.506
		(−1.41)	(−1.05)		(−0.18)	(0.13)
JanDum_{t+1}			105.6			107.4
			(5.28)			(5.62)
R^2	0.000	0.005	0.063	0.440	0.044	0.107
D–W	1.95	1.98	1.95	2.05	2.05	2.04

This table presents several cases of the general forecasting regression $y_{t+1}=\beta_0+\beta_1 Z_t+\beta_2 v_{t+1}+\beta_3 \text{JanDum}_{t+1}+\varepsilon_{t+1}$. The dependent variable can either be $y_{t+1}=R_{t+1,\text{small}}{}^e$, the excess return on the small-stocks portfolio (panel A), or its square (panel B). The main variable of interest is Z_t and we consider two different cases: $Z=PI$ is the price impact measure as constructed in this paper; Z=Amihud is Amihud's (2002) measure of price impact. The column under "OLS" is a simple OLS estimate. The columns under "reduc-bias" show the reduced-bias estimates using the method of Amihud and Hurvich (2004). T-statistics are within parenthesis. The sample is from 1963 to 2000 with monthly observations.

For comparison purposes, we also estimated forecasting regressions for dollar volume (DVOL) and turnover (TURN).[17] Similarly to the results for PI, we find that neither DVOL nor TURN are significant predictors of monthly excess returns on small stocks. This is consistent with optimal portfolio allocations to small stocks at the monthly horizon not depending on liquidity, as described above.

5. Conclusion

This paper documents the empirical relation between optimal portfolio weights and liquidity. In a fist step, we find the optimal portfolio function of a CRRA investor using a nonparametric technique. In a second step, we use the nonparametric findings to define a parametric estimator.

We find that the relevance of liquidity for conditional portfolio choice depends on both the asset and the investment horizon. The optimal portfolio allocation to small stocks is strongly increasing in liquidity at the daily or weekly frequency. These results cannot be explained by non-synchronous trading. However, liquidity does not matter for large stocks nor for longer monthly horizons. Furthermore, the optimal allocations for the CRRA investor depend on higher-order moments of returns only at the daily horizon.

Liquidity is non-stationary over our sample period. Similarly to other studies, we find that one must be careful in choosing a particular transformation to make liquidity stationary. We find that our measure of price impact and the one used in Amihud (2002) have different predictive power of the first two moments of returns. Interestingly however, they result in optimal portfolio allocations with the same properties.

In short, this paper suggests that it may be fruitful to consider the information contained in liquidity for short-horizon trading in small stocks.

[17] Not shown, but available upon request.

References

Aït-Sahalia, Y., Brandt, M.W., 2001. Variable selection for portfolio choice. Journal of Finance 56 (4), 1297–1351.

Amihud, Y., 2002. Illiquidity and stock returns: cross-section and time-series effects. Journal of Financial Markets 5, 31–56.

Amihud, Y., Hurvich, C.M., 2004. Predictive regressions: a reduced-bias estimation method. Journal of Financial and Quantitative Analysis 39 (4), 813–841.

Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid–ask spread. Journal of Financial Economics 17, 223–249.

Andersen, T.G., 1996. Return volatility and trading volume: an information flow interpretation of stochastic volatility. Journal of Finance 51 (1), 169–204.

Boudoukh, J., Richardson, M.P., Whitelaw, R.F., 1994. A tale of three schools: insights on the autocorrelations of short-horizon stock returns. Review of Financial Studies 7 (3), 539–573.

Brandt, M.W., 1999. Estimating portfolio and consumption choice: a conditional Euler equations approach. Journal of Finance 54 (5), 1609–1646.

Brennan, M.J., Chordia, T., Subrahmanyan, A., 1998. Alternative factor specifications, security characteristics, and the cross-section of expected stock returns. Journal of Financial Economics 49, 345–373.

Brennan, M.J., Subrahmanyan, A., 1996. Market microstructure and asset pricing: on the compensation for illiquidity in stock returns. Journal of Financial Economics 41, 441–464.

Campbell, J.Y., Grossman, S.J., Wang, J., 1993. Trading volume and serial correlation in stock returns. Quarterly Journal of Economics 905–939.

Chaudhuri, P., Marron, J., 1999. SiZer for exploration of structures in curves. Journal of the American Statistical Association 94 (447), 807–823.

Chordia, T., Swaminathan, B., 2000. Trading volume and cross-autocorrelations in stock returns. Journal of Finance 55 (2), 913–935.

Constantinides, G.M., 1986. Capital market equilibrium with transaction costs. Journal of Political Economy 94 (4), 842–862.

Datar, V.T., Naik, N.Y., Radcliffe, R., 1998. Liquidity and stock returns: an alternative test. Journal of Financial Markets 1, 203–219.

Gervais, S., Kaniel, R., Mingelgrin, D.H., 2001. The high-volume return premium. Journal of Finance 56 (3), 877–919.

Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press.

Heaton, J., Lucas, D.J., 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. Journal of Political Economy 104 (3), 443–487.

Horowitz, J.L., 2000. The Bootstrap. University of Iowa. Working Paper.

Inoue, A., Shintani, M., 2001. Bootstrapping GMM Estimators for Time Series. North Carolina State University. Working Paper.

Jones, C.M., 2002. A Century of Stock Market Liquidity and Trading Costs. Columbia University. Working paper.

Jorion, P., 1986. Bayes–Stein estimation for portfolio analysis. Journal of Financial and Quantitative Analysis 21 (3), 279–292.

Kandel, S., Stambaugh, R.F., 1996. On the predictability of stock returns: an asset-allocation perspective. Journal of Finance 51 (2), 385–424.

Kaniel, R., Li, D., Starks, L., 2005. Investor Visibility Events: Cross-Country Evidence. University of Texas. Working Paper.

Llorente, G., Michaely, R., Saar, G., Wang, J., 2002. Dynamic volume-return relation of individual stocks. Review of Financial Studies 15 (4), 1005–1047.

Lo, A., MacKinlay, C., 1990. An econometric analysis of nonsynchronous trading. Journal of Econometrics 45, 181–211.

Lo, A.W., Wang, J., 2000. Trading volume: definitions, data analysis, and implications of portfolio theory. Review of Financial Studies 13 (2), 257–300.

Longstaff, F.A., 2001. Optimal portfolio choice and the valuation of illiquid securities. Review of Financial Studies 14 (2), 407–431.

Lynch, A.W., Balduzzi, P., 2000. Predictability and transaction costs: the impact on rebalancing rules and behavior. Journal of Finance 55 (5), 2285–2309.

Maenhout, P.J., 2004. Robust portfolio rules and asset pricing. Review of Financial Studies 17 (4), 951–983.

Mehra, R., Prescott, E.C. 2003, "The equity premium in retrospect," NBER Working Paper No. 9525.

Mills, T.C., 1999. The Econometric Modelling of Financial Time Series. Cambridge University Press.

Pereira, J.P., Zhang, H.H., 2004. The Liquidity Premium in a Dynamic Model with Price Impact. ISCTE Business School and UT Dallas. Working Paper.

Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. Journal of the American Statistical Association 89 (428), 1303–1313.

Sadka, R., 2004. Informational versus Non-Informational Aspects of Liquidity Risk: The Pricing of Momentum Stocks. University of Washington Business School. Working paper.

Stambaugh, R.F., 1999. Predictive regressions. Journal of Financial Economics 54, 375–421.

Sullivan, R., Timmermann, A., White, H., 1999. Data-snooping, technical trading rule performance, and the bootstrap. Journal of Finance 54 (5), 1647–1691.

Tu, J., Zhou, G., 2004. Data-generating process uncertainty: what difference does it make in portfolio decisions? Journal of Financial Economics 72, 385–421.

Vayanos, D., 1998. Transaction costs and asset prices: a dynamic equilibrium model. Review of Financial Studies 11 (1), 1–58.

Xia, Y., 2001. Learning about predictability: the effects of parameter uncertainty on dynamic asset allocation. Journal of Finance 56 (1), 205–246.