ECORE
INTERNATIONAL ASSOCIATION FOR RESEARCH AND TEACHING
Economics, Finance, Operations Research, Econometrics and Statistics

++ research ++ teaching ++

ECORE DISCUSSION PAPER

2009/9

Market Liquidity as Dynamic Factors

Marc HALLIN
Charles MATHIAS
Hugues PIROTTE
David VEREDAS

CORE
Université catholique de Louvain

ECARES
Université Libre de Bruxelles

MARKET LIQUIDITY AS DYNAMIC FACTORS

Marc HALLIN[1,2], Charles MATHIAS[3], Hugues PIROTTE[4] and David VEREDAS[5]

January 2009

Abstract

We study market liquidity via daily close relative spreads and daily traded volumes in a sample of 426 S&P500 constituents recorded over the years 2004-2006, a period of "normal" liquidity conditions. We use recent results on the Generalized Dynamic Factor Model (GDFM) with block structure to provide a sound definition of unobservable market liquidity and to assess the complementarity of those two liquidity measures. The advantage of defining market liquidity as dynamic factors is that, contrary to other definitions that can be found in the literature, it tackles time dependence and commonness at the same time, without making any restrictive assumptions on the underlying data generating process. Both relative spread and volume in the dataset under study appear to be driven by the same one-dimensional common shocks, which therefore naturally qualify as the unobservable market liquidity shocks.

Keywords: Commonality, liquidity, equities, factor models, block structure.

JEL classification: C33, C51, G10.

[1]ECARES and Département de Mathématique, Université Libre de Bruxelles; email: mhallin@ulb.ac.be.

[2]Académie Royale de Belgique.

[3]Solvay Brussels School of Economics and Management, Université Libre de Bruxelles; email: chmathia@ulb.ac.be.

[4]Centre E. Bernheim, Solvay Brussels School of Economics and Management, Université Libre de Bruxelles. Luxembourg School of Finance, University of Luxembourg; email: hugues.pirotte@ulb.ac.be.

[5]ECARES, Solvay Brussels School of Economics and Management, Université Libre de Bruxelles; email: dveredas@ulb.ac.be.

Corresponding address: Charles Mathias, Université Libre de Bruxelles, 50 Av F.D. Roosevelt CP139, B1050 Brussels, Belgium. Phone: +3226504133. Fax: +3226504475.

We are grateful to Roberto Pascual for useful comments. Any remaining errors and inaccuracies are ours.

1 Introduction

Of all asset characteristics that are relevant for asset managers, liquidity certainly is among the most elusive. Yet it is ubiquitous in financial practice and theory (for instance, most of the modern applied asset pricing theory lingers on its full presence) and it has a pristine definition: an asset is liquid if it is easily convertible into cash, the reference asset with perfect liquidity. This definition is often rephrased in terms of time, volume, and cost. Indeed, *when people think about liquidity, they may think about trading quickly, about trading large size, or about trading at low cost*.[1] Since Kyle (1985), these three dimensions are well defined. The time dimension refers to *resiliency*—the speed with which pricing errors caused by uninformative order-flow shocks are corrected or neutralized in the market. Cost refers to *tightness*—the accepted price for immediacy in resolving the trade. Last, volume refers to *depth*—the volume that can be traded without price variations.[2]

Though the concept of liquidity is clear, its quantitative evaluation poses a major problem. Liquidity indeed is an unobserved variable, which implies that it has to be evaluated from the measurement of liquidity-related quantities or proxies—call them *liquidity measures*. But this is a delicate task because of the difficulty (i) to capture the three dimensions of liquidity in a single measure and (ii) to reach a consensus on the liquidity measures to be taken into account. This double difficulty seriously challenges the objectivity of any final assessment. The simplest liquidity measures currently considered in the empirical literature cover only one of the three dimensions. Trade durations, for instance, defined as the time intervals between two trades, clearly carry liquidity-related information, but only cover the time dimension, ignoring tightness and depth. Moreover, they require tick-by-tick data: at lower frequencies (such as daily frequency), trade durations cannot be computed, as observations are regularly spaced. Daily close or open bid-ask spreads, defined as the difference between the lowest ask and highest bid prices for an asset at some given point in time, measure liquidity effects as well, but mainly cover tightness. Daily realized volumes also measure the liquidity effects of an asset, but only cover its transacted depth. A number of papers have proposed measures that combine tightness and depth, such as the Hasbrouck and Seppi (2001) *quote slope*, the Domowitz and Wang (2002) *order book integral*, the Amihud (2002) ratio of average volume effect on absolute returns, or the Pastor and Stambaugh (2003) measure for average volume-related return reversal. The former two consider the existing shape of the available order book through time, while the latter rely on transacted prices and volumes. A common drawback of all those liquidity measures is that we do not know up to what extent they capture liquidity dynamics only.[3] To get a deeper understanding of liquidity, we should study interrelations between the various liquidity measures and investigate their relation to liquidity. Literature on this subject is scarce, however: most papers dealing with liquidity perform an analysis based one single liquidity measure at a time.

[1] Harris (2003), p. 394.

[2] See Minguet (2003), O'Hara (1997) or Schwartz (1993), among others, for further details.

[3] For example, when looking at the Spanish stock market, Martinez et al. (2005) found a disconcerting positive correlation between the Amihud measure for illiquidity and the Pastor-Stambaugh measure of liquidity. This indicates that such measures, most likely, are driven by other dynamics than liquidity alone.

The same comments apply in the analysis of *market liquidity*, where the aim is to understand commonness in liquidity across securities. Two fundamental contributions to this domain are Chordia et al. (2000) and Hasbrouck and Seppi (2001). Both find that a common or "market" component is significantly present in various liquidity measures taken over a large cross-section of stocks. Chordia et al. (2000) present evidence of this fact by running regressions of a given liquidity measure, for each individual stock, against its cross-sectional mean. Hasbrouck and Seppi (2001), on the other hand, perform a (classical, hence static) Principal Component Analysis (PCA) on liquidity measures, out of which they consider up to three principal components. Other papers assess the existence of market liquidity risk, that is, they examine whether market liquidity is priced across different stocks. Amihud (2002), Eckbo and Norli (2002), Pastor and Stambaugh (2003) or Acharya and Pedersen (2005) find that market liquidity explains price differences across assets.

The influence of the aforementioned papers should not be underestimated, as they all substantially increase our knowledge and understanding of the role of liquidity in financial markets. However, they rely on the somewhat arbitrary choice of some liquidity measure, without much of an analysis of its net informative contribution in explaining the market liquidity phenomenon. A first attempt to assess market liquidity based on several liquidity measures was made by Korajczyk and Sadka (2008). In a very large sample (more than 4000 stocks followed during 18 years), they use a static factor model method, comparable to that of Hasbrouck and Seppi (2001), to estimate commonness for eight different liquidity measures, which they call "within-measure commonality". Moreover, they also estimate commonness when considering the eight different measures on the same panel, called "across-measure commonality", through the extraction of a common component that is then added to their asset pricing experiment. Just as Hasbrouck and Seppi (2001), the authors extract up to the third principal component for each liquidity measure.

None of these approaches fully exploits the time series nature of the data. In particular, they all overlook the leading/lagging phenomenons that may exist among the various liquidity measures and are particularly relevant here, since liquidity-related data are highly autocorrelated. Taking such time series features into account naturally brings the Generalized Dynamic Factor Model (henceforth GDFM) methods into the picture. These methods were developed in a series of papers by Forni, Hallin, Lippi and Reichlin (2000, 2004, 2005), Forni and Lippi (2001), and Hallin and Liška (2007, 2008). They allow for disentangling commonness (market components) and idiosyncrasy (stock-specific components), not only across panels consisting of some given liquidity measure observed over a large number of stocks, but also across panels juxtaposing several such measures. Contrary to other dynamic factor methods (such as Stock and Watson (2002a and b) or Bai and Ng (2002)), GDFM methods do not impose any restriction (beyond the usual assumptions of second-order stationarity, etc.) on the actual data generating process.

We examine here the complementarity of two simple and widely used liquidity measures: daily close relative bid-ask spread and daily realized dollar volume, for 426 S&P500 listed stocks from January 2004 till December 2006. This is a period characterized by a "normal" state of market liquidity, which is appropriate for the scope of this article. A period with extreme events and/or liquidity crunches, though important, would introduce distortions in the measurement of commonness, which are left for future research.

We consider the method proposed by Hallin and Liška (2008) for the analysis of large panels with block structure, where the blocks represent the two subpanels of volume and relative spread, respectively. The method permits to identify, estimate and compare the factors driving each subpanel and the factors driving the joint panel. In particular, it allows us to assess up to what extent commonality in volume coincides with commonality in relative spread. As volume and relative spread cover different aspects of liquidity (depth and tightness, respectively), they are a priori unlikely to carry exactly the same information: it could be that some features of liquidity are explained by realized volume but not by relative spread, and conversely. In GDFM terms, this means that some common spread shocks might a priori be idiosyncratic to volume and vice versa. Moreover, the analysis using GDFM takes into account the dynamic interactions between the measures: some liquidity features may be leading in volume while lagging in spread, and vice versa.

Our findings mainly go into three directions. First, it appears that the common relative spread and common volume spaces coincide, and have dynamic dimension one. This means that, although relative spread and volume cover different aspects of liquidity, their market or common components have the same origin and thus carry the same information. Moreover, that common space being one-dimensional, it is driven by a unique shock, which therefore strongly qualifies as the unobservable market liquidity shock. This suggests some homogeneity when markets are confronted with liquidity and that therefore there should be no distinct market liquidity effects originating, for instance, from different sectors or different types of investors. Second, on average, market related shocks account for 12% of the total variation of a stock's relative spread and for 18% of the total variation of its volume. This may seem a rather low proportion, but is not surprising if compared to the variance decompositions obtained in Hasbrouck and Seppi (2001) and Chordia et al. (2000), even though we should be careful with such comparisons, given that databases and the nature of measures differ. Third, we observe a significant difference between the autocorrelograms of idiosyncratic spread and volume. On average, idiosyncratic relative spread components are only weakly autocorrelated, but they are persistant. By contrast, volume components exhibit higher autocorrelations, with much faster decay. This difference can be explained by the fact that S&P500 constituents are highly traded stocks with thick limit order books. Relative spreads for such stocks do not move quickly, so that the impact of a market shock stays for long. Volumes, on the contrary, are more flexible, and their adaptation to changing market circumstances is much faster.

The outline of the article is as follows. Section 2 frames the previous contributions on liquidity with respect to the present one, and explains the opportunities offered by GDFM for liquidity econometrics. Section 3 explains the building blocks of GDFM. Section 4 gives information about the dataset used and comments on the liquidity measures considered. The main results are presented in Section 5, and Section 6 concludes.

2 Commonness in liquidity

The analysis of liquidity is deeply rooted in market microstructure theory; its origins can be traced back to Kyle (1985) and Amihud and Mendelson (1986), among others. Most of the models developed under the market microstructure perspective focus on the liquidity of individual securities, and give little attention to the common determinants of liquidity across the market. On the other hand, modern finance emerged through the study of portfolio theory and the benefits of risk diversification by exploiting return volatility risk and its market component, paying little attention to liquidity risk, though.

The present article builds further on recent contributions with respect to the identification of market liquidity. The main idea is that, while agreeing on the existence of standard market illiquidity (i.e. that in "normal" market conditions), the liquidity of every asset can be seen as the sum of a common (hence non diversifiable) and an idiosyncratic (hence diversifiable) component, in the same sense as for return volatility.

The contributions and methodology of three leading articles in the identification of common versus idiosyncratic liquidity are discussed in the next paragraphs.

Chordia et al. (2000) look at all NYSE transactions in 1992 and analyze commonality in daily percentage changes of five order book-related liquidity measures (quoted and effective spread, both in absolute and in proportional terms, and quoted depth)[4] by considering regression equations of the form

$$\Delta \mathrm{LIQ}_t^i = \beta_i \Delta \mathrm{LIQ}_t^M + \xi_t^i, \tag{1}$$

where LIQ_t^i denotes the value taken by one of the five liquidity measures at time t for stock i, LIQ_t^M is the average liquidity over all stocks except i, and $\Delta = 1 - L$ where L denotes the lag operator. The first term of the right hand side accounts for market related variations. The second term, $\xi_t^i = \alpha_i + \varepsilon_t^i$, includes the intercept of the regression, α_i, and a white noise term ε_t^i . The authors find that the β_i's are significantly different from zero, which indicates the presence of common underlying determinants of liquidity.[5] However, they obtain low R^2 for all measures (about 4%).

Hasbrouck and Seppi (2001) perform a similar analysis on four 15-minute interval order flow measures (share volume, dollar volume and square root of dollar volume, both signed and in absolute values)[6] of 30 Dow Jones stocks traded during 1994. Instead of the differences $\Delta \mathrm{LIQ}_t^i$, they consider the levels LIQ_t^i of each liquidity measure. Moreover, they use PCA instead of linear regressions, by defining LIQ_t^M as the principal component of the order flows of the 30 stocks:

$$\mathrm{LIQ}_t^i = \beta_i \mathrm{LIQ}_t^M + \xi_t^i. \tag{2}$$

[4] The *quoted spread* is the difference between the best ask and bid quotes. The *effective spread* is twice the observed deviation of the price at which the transaction took place and the *midquote*, which is the average between highest bid and lowest ask. *Proportional* (or *relative*) *spreads* are quoted spreads divided by the midquote. *Quoted depth* is the depth at the best quotes.

[5] The authors perform two distinct analyses, based on unweighted and value-weighted averages, respectively.

[6] The sign of a liquidity measure is negative or positive, according as the trade takes place at a price which is below or above the midquote.

They find that, for signed dollar volume, 8% of the total variance is explained by the first common factor, and empirically suggest that the second and third common components could be negligible.

Korajczyk and Sadka (2008) use the asymptotic principal component method of Connor and Korajczyk (1986) and the EM algorithm to identify market liquidity. First they apply both methods to eight different liquidity measures independently: the Amihud (2002) montlhly average effect of volume on absolute value return, turnover (i.e. the ratio of montly volume and shares outstanding), quoted percentage spread, effective percentage half-spread and four parametric estimates of price impact components. As Hasbrouck and Seppi (2001), the authors consider up to the third principal component. They obtain very different R^2 for the different measures: 7% for turnover versus 24% for effective spread, when looking at the first principal component. The major contribution of Korajczyk and Sadka (2008) is that, for their pricing experiment, they also compute common liquidity by including their eight different liquidity measures into a single joint panel, as they argue that this would eliminate some liquidity measurement bias.

Although these articles significantly contribute to the study and understanding of commonalities in liquidity, neither of them take into account the time series nature of the various liquidity measures. Assuming the ξ_t^i's in (1) and (2) to be serially uncorrelated clearly is unrealistic; it requires, for instance, LIQ_t^M in (2) to account for *all* dynamic aspects of *all* LIQ_t^i's. Lagged influence of unobserved common liquidity factors is also precluded: LIQ_t^M is a purely static principal component which only depends on contemporaneous LIQ_t^i's, with an implicit and questionable assumption that the liquidity characteristics of all stocks are perfectly synchronized.

The Generalized Dynamic Factor Model (GDFM) estimates commonality in a spirit which is somewhat similar to Hasbrouck and Seppi (2001) and Korajczyk and Sadka (2008) in the sense that it also seeks for a variance maximizing linear combination of observations. The main difference, however, is that, by allowing for lagged loadings (instead of contemporaneous) and autocorrelated idiosyncratics (rather than white noise residuals), it does not force any model on the data, while fully exploiting their time series nature. In the present context, this means that GDFM tackles persistence and co-movement at the same time, as it estimates the effect of common shocks. This also implies that all common components are orthogonal, at all leads and lags, to all the idiosyncratic ones, while allowing for mild cross-sectional correlation among the idiosyncratic components of distinct individual stocks (typically, the idiosyncratic component of a given stock may yield autocorrelation with a *finite* number of other, closely related, idiosyncratic components). The latter is an attractive property in the study of liquidity, as it provides us with a clear and formal distinction between commonness and idiosyncrasy. Another advantage of the GDFM theory is that it identifies the dimension of the common space, as opposed to Hasbrouck and Seppi (2001) and Korajczyk and Sadka (2008), who look at the first three principal components without having a rigorous criterion of whether they all significantly contribute to commonality. Finally, the GDFM method allows for a global analysis of (arbitrarily many) different liquidity measures. The next section describes in detail the building blocks of the GDFM method.

3 Dynamic factors and commonness in liquidity

First consider a panel of n stocks, for which some liquidity measure has been recorded over a time period of length T. Denote by LIQ_t^i, $i = 1, \ldots, n$, $t = 1, \ldots, T$ the observation made at time t for stock i. These observations are treated as finite realizations of a double-indexed zero mean second order stationary stochastic process $\{\text{LIQ}_t^i : i \in \mathbb{N},\ t \in \mathbb{Z}\}$. Both n and T, in the sequel, are assumed to be large, and asymptotic statements are made as n and T tend to infinity.

Denote by $\mathbf{\Sigma}_n(\theta)$ the $n \times n$ spectral density matrix of the the n-dimensional vector process $\{\mathbf{LIQ}_{\mathbf{n},t} := (\text{LIQ}_t^1, \ldots, \text{LIQ}_t^n)'; t \in \mathbb{Z}\}$, and assume that, for all $n \in \mathbb{N}$, $k \in \{1, \ldots, n\}$ and some $c_k > 0$, $\sup_\theta(\mathbf{\Sigma}_n(\theta))_{kk} \leq c_k$. For any $\theta \in [-\pi, \pi]$, let $\lambda_{n,k}(\theta)$ be $\mathbf{\Sigma}_n(\theta)$'s k-th eigenvalue (in decreasing order of magnitude).

Denote by q the number of diverging such eigenvalues, that is, define q as $q := \min\{k \in \mathbb{N} : \sup_n \|\lambda_{n,k}(\theta)\| < \infty\ \theta - \text{a.e.}\} - 1$, and assume that $q < \infty$. Theorem 2 in Forni and Lippi (2001) then establishes the existence of a unique decomposition of LIQ_t^i into

$$\text{LIQ}_t^i = \chi_t^i + \xi_t^i = \mathbf{B}_i'(L)\mathbf{u}_t + \xi_t^i \quad \text{for all } i \in \mathbb{N},\ t \in \mathbb{Z}, \tag{3}$$

where χ_t^i and ξ_t^i are mutually orthogonal at all leads and lags, $\mathbf{u}_t := (u_{1t}, \ldots, u_{qt})'$ is q-dimensional orthonormal white noise, and $\mathbf{B}_i(L) := (B_{i1}(L), \ldots, B_{iq}(L))'$ is a vector of square-summable filters (the decomposition into $\chi_t^i + \xi_t^i$ is unique; the filters $\mathbf{B}_i(L)$ and the $\mathbf{u}_t$'s are not). Equation (3), with unspecified q, thus is not a *statistical model*, but a *canonical representation* of the panel under study—contrary to (1) or (2). That representation is called the *dynamic factor representation* of LIQ_t^i; the χ_t^i's are the *common*, and the ξ_t^i's the *idiosyncratic* components, respectively, of LIQ_t^i. The process $\{\chi_{i_0,t}\}$ is cross-correlated with infinitely many liquidity measureprocesses $\{\text{LIQ}_t^i\}$, $i \neq i_0$, as $n \to \infty$ and therefore can be identified as the component of $\{\text{LIQ}_{i_0,t}\}$ which is driven by the market, while $\xi_{i_0,t}$ is specific to stock i_0 (market-uncorrelated), and presents cross-correlations with a finite number of related cross-sectional processes only. The Hilbert space spanned by the χ_t^i's is called the *common* space. It has dynamic dimension q and its elements are *market liquidity variables*. The corresponding innovation process $\{\mathbf{v}_t : t \in \mathbb{Z}\}$ (namely, any orthonormal white noise such that the Hilbert space generated until time s coincides with the Hilbert space generated up to s by all χ_t^i's, $i \in \mathbb{N}$) naturally are interpreted as the *market liquidity shocks*.

Forni, Hallin, Lippi and Reichlin (2000) show how the common and idiosyncratic components χ_t^i and ξ_t^i can be consistently reconstructed from the observed LIQ_t^i's, along with estimators of their respective variances. The variance decomposition

$$\text{Var}[\text{LIQ}_t^i] = \text{Var}[\chi_t^i] + \text{Var}[\xi_t^i] \tag{4}$$

for given i (because of stationarity, these variances do not depend on t) of course indicates how common or idiosyncratic the liquidity of a particular stock i is.

The Hallin and Liška (2008) method for the analysis of panel data with block structure very much relies on the Hallin and Liška (2007) procedure for the identification of the number of dynamic factors. That identification procedure consists in tuning the penalty term of an information-theoretic criterion by a positive factor c. A grid (n_k, T_k), $k = 1, \ldots, K$ of increasing n and/or T values ($n_K = n$; $T_K = T$) is considered, and, for each value of c and k, a number $q_k(c)$ of factors is selected as the value of $q \in \mathbb{N}$ minimizing the information criterion with tuning constant c, computed from a panel consisting of the series $1, \ldots, n_k$ observed over $t = 1, \ldots, T_k$. A particular value c^* then is chosen as the second smallest value of c for which the K selected $q_k(c)$'s are stable across the (n_k, T_k) grid. In practice, this is achieved by examining a double plot. In the first one, the empirical variance $V(c)$ of the K-tuple $(q_1(c), \ldots, q_K(c))$ is plotted against c; the second plot provides the corresponding final selection $q_K(c)$ as a function of c. The number of factors which is ultimately selected then is $\hat{q} := q_K(c^*)$, where c^* belongs to the second interval of c values for which the empirical variance $V(c)$ takes value zero. Hallin and Liška provide several versions of their criterion, which all yield essentially the same results. The version adopted here is their IC_2 log criterion, together with their penalty $p_1(n, T)$—see Section 3.2 of Hallin and Liška (2007) for details.

No method so far allows for reconstructing the market liquidity *fundamental* shocks (its innovation process $\mathbf{v}_t$); however, see Forni and Lippi (2009) for significant progress in that direction. When results are available on that point, the GDFM definitely will be the ideal tool for identifying, estimating, and forecasting market liquidity, as well as its impact on individual stocks' liquidity.

A major advantage of the GDFM in the analysis of market liquidity is that several liquidity measures can be handled, either jointly or separately, via the Hallin and Liška (2008) methodology for dynamic factor models in the presence of block structure. The blocks here are the (sub)panels associated with a given liquidity measure—here, relative spread and volume. The method, as we shall see, provides interesting insights into the interrelations between those two measures, and answers such questions as "do relative spread and volume convey the same information about market liquidity?" "should we choose one of them, or rather combine them?" "is there an optimal way to do so?" etc.

For $K = 2$ blocks (in order to fix the ideas, call them relative spread and volume, respectively, and denote by $\mathbf{SPR}_n$ and $\mathbf{VOL}_n$ the corresponding subpanels), the method actually decomposes the Hilbert space spanned by all variables in the joint panel $\mathbf{LIQ}_n$ (consisting of all SPR_t^i's and VOL_t^i's, for $i = 1, \ldots, n$ and $t = 1, \ldots, T$) into a product of four mutually orthogonal Hilbert subspaces, spanned (for the sake of simplicity, we write "spread-common" instead of "relative-spread-common" etc.) by the spread- and volume-common, spread-common and volume-idiosyncratic, spread-idiosyncratic and volume-common, and spread- and volume-idiosyncratic components, respectively. Projecting SPR_t^i and VOL_t^i onto those four subspaces yields the following refinements of (3):

$$\text{SPR}_t^i = \phi^i_{\text{SPR};t} + \psi^i_{\text{SPR};t} + \nu^i_{\text{SPR};t} + \xi^i_{\text{SPR};t}, \qquad (5)$$
$$\text{VOL}_t^i = \phi^i_{\text{VOL};t} + \psi^i_{\text{VOL};t} + \nu^i_{\text{VOL};t} + \xi^i_{\text{VOL};t}. \qquad (6)$$

The $\phi^i{}_{;t}$'s and $\xi^i{}_{;t}$'s are called *strongly common* and *strongly idiosyncratic*, the $\psi^i{}_{;t}$'s and $\nu^i{}_{;t}$'s *weakly common* and *weakly idiosyncratic* components, respectively. This decomposition is shown in Figure 1; see Hallin and Liška (2008) for details.

Figure 1: Schematic representation of the Hilbert space decomposition for two blocks

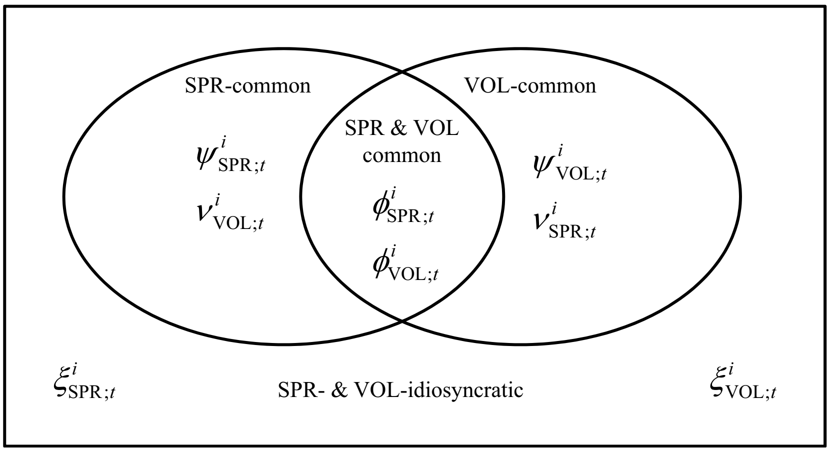

Now, if (as will appear in Section 5) all spread-common and volume-idiosyncratic and all volume-common and spread-idiosyncratic components are zero, that is, if (5) and (6) boil down to

$$\mathrm{SPR}^i_t = \phi^i_{\mathrm{SPR};t} + \xi^i_{\mathrm{SPR};t} \qquad \text{and} \qquad \mathrm{VOL}^i_t = \phi^i_{\mathrm{VOL};t} + \xi^i_{\mathrm{VOL};t}, \tag{7}$$

spread and volume are driven by the same common shocks, which unambiguously can be interpreted as *the* market liquidity shocks.

4 Data

We consider $n = 426$ S&P500 constituents that listed from Monday January 5th, 2004, till Friday December 29th, 2006, and that were still listed in November 2008. This is a period characterized by "standard" market illiquidity, i.e. without extreme illiquidity conditions, which is appropriate for the scope of this article. From *Reuters 3000 Xtra*, we extracted, for each of these stocks, the daily close best ask, daily close best bid and the daily realized dollar volume, from which we constructed two liquidity measures. The first one is the relative spread, defined as

$$\mathrm{SPR}^i_t := \frac{\mathrm{ask}^i_t - \mathrm{bid}^i_t}{\mathrm{mq}^i_t},$$

where ask^i_t (respectively bid^i_t) is the daily close best ask (respectively bid), $\mathrm{mq}^i_t := (\mathrm{ask}^i_t + \mathrm{bid}^i_t)/2$ is the midquote of stock i at day t. We denote the spread subpanel by $\mathbf{SPR}_n := \{(\mathrm{SPR}^1_t, \dots, \mathrm{SPR}^n_t)'; t = 1, \dots, T\}$. The second measure is the realized dollar volume, denoted by VOL^i_t. The corresponding subpanel (or volume subpanel) is denoted by $\mathbf{VOL}_n$, the total panel by $\mathbf{LIQ}_n := (\mathbf{SPR}'_n, \mathbf{VOL}'_n)'$.

There are several reasons for choosing these two measures. First, both are simple and widely used in practice. Second, each of them covers a different dimension of liquidity: VOL^i_t is a proxy for depth and SPR^i_t a proxy for tightness. Third, they cover different aspects of the trading process: SPR_t^i is a pre-trading measure, conveying information about the state of the limit-order book and the immediacy cost (measuring liquidity *ex ante*), whereas VOL_t^i is a post-trading measure, conveying information about the actual trade (measuring liquidity *ex post*).

Prior to estimation, we applied some algorithms to clean the data. First, days on which trading for more than 80% of the stocks was suspended were eliminated from the analysis. Second, days for which at least one stock showed negative spread also were eliminated. Third, missing spread or volume values were interpolated. In total, this leaves us with $T = 747$ observation dates. Fourth, SPR_t^i is multiplied by 10^3 and VOL_t^i by 10^{-6}. Relative spread for S&P500 constituents is very small, which may entail numerical problems. Likewise, traded volumes are very large, which may entail numerical problems as well.

We checked for the presence of weekly seasonality. We therefore regressed the demeaned relative spread and volume, denoted generically by $\text{LIQ}_t^{i,*}$, on a set of dummy variables that account for the different days of the week

$$\text{LIQ}_t^{i,*} = \sum_{k=1}^{5} \beta_k^i D_{k,t} + \varepsilon_t^i \, , \tag{8}$$

where $D_{1,t}$ stands for the Monday indicator function, $D_{2,t}$ for the Tuesday one, etc. We thus obtained 426 least squares estimates of β_k^i for a given day k. The plots in Figure 2 show the estimated parameters for all the stocks and the 5% confidence bands for relative spread.[7] In very few cases there are parameters significantly different from zero. Table 1 shows the percentage of stocks with parameters that are significantly different from zero for the different days of the week and for three probability levels. Overall results indicate that very few stocks show significant weekly seasonal patterns, either for relative spread or for volume.

Table 1: Percentage of significant seasonal parameters

	level α	Mon	Tue	Wed	Thu	Fri
Relative Spread	$\alpha = 10\%$	7.51	5.63	9.86	9.62	11.5
	$\alpha = 5\%$	5.63	3.76	5.63	6.34	6.81
	$\alpha = 1\%$	1.41	0.70	2.82	2.82	3.29
Volume	$\alpha = 10\%$	7.51	8.69	5.16	8.69	2.35
	$\alpha = 5\%$	5.87	6.57	4.23	7.04	1.88
	$\alpha = 1\%$	3.05	3.05	3.29	3.99	1.64

Percentage of parameters in (8) that are significantly different from zero for the two liquidity measures, at three different probability levels α, across all stocks.

Table 2 shows descriptive statistics of the means ($\bar{x}$) and variances (s_x^2) of the liquidity measures per stock over the considered period. Their cross-sectional means (denoted $m_{\bar{x}}$ and $m_{s_x^2}$ respectively) and interquartile ranges ($\text{IQR}_{\bar{x}}$ and $\text{IQR}_{s_x^2}$) are computed for the complete

[7] Similar results for volume are available under request.

Figure 2: Weekly seasonal parameters

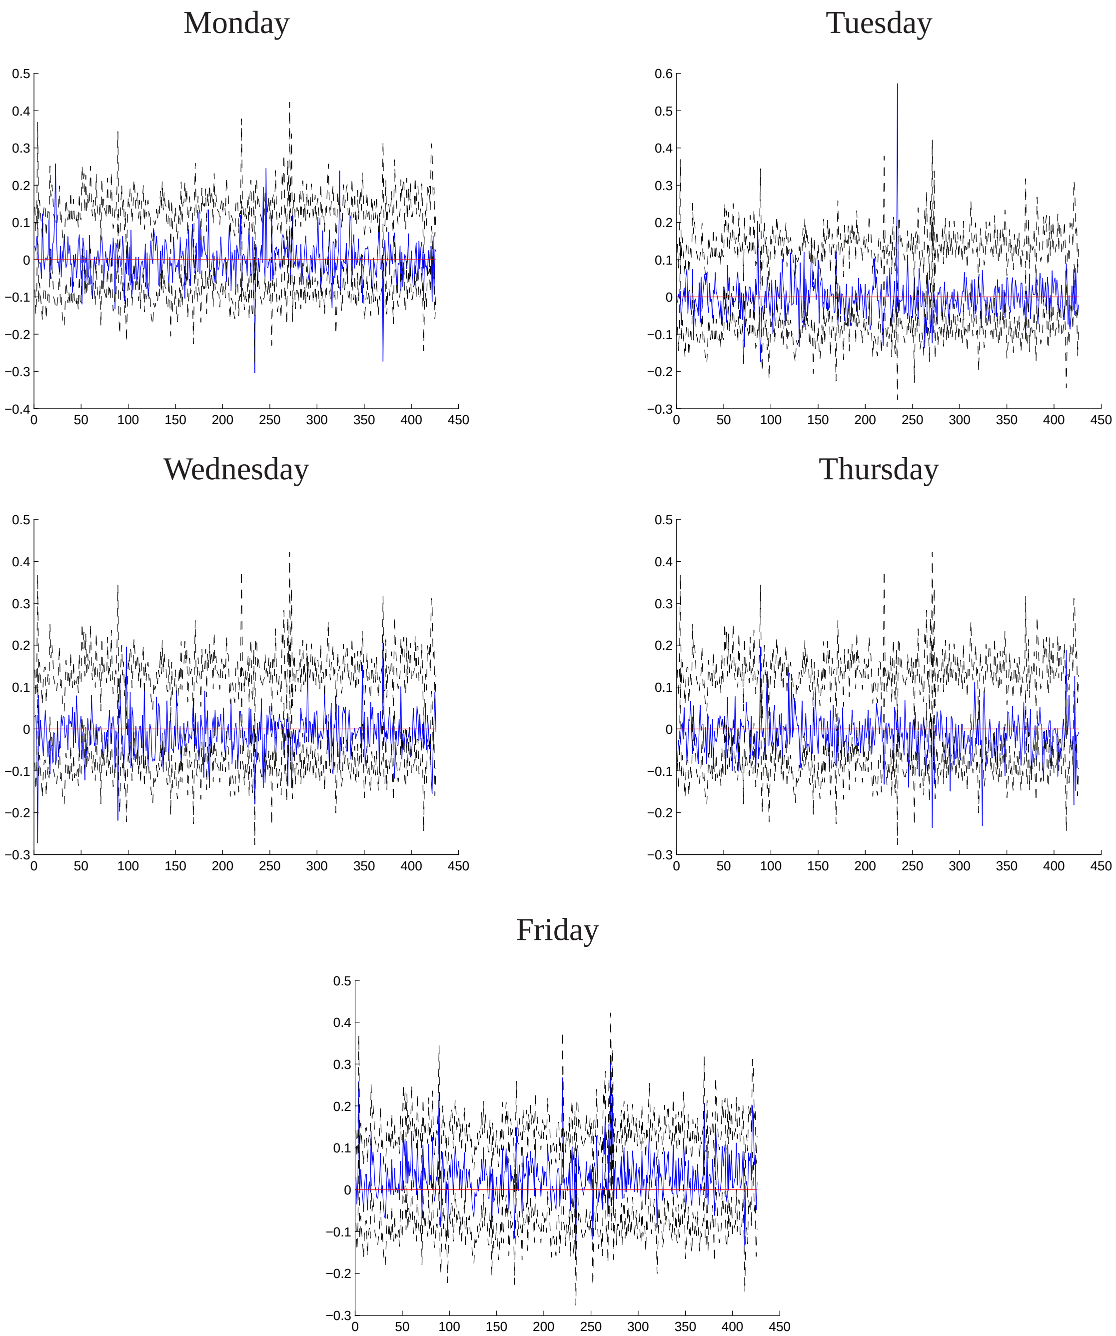

Estimated weekly seasonal parameters (see (8)) for relative spreads (solid lines), along with the corresponding 5% confidence bands (dashed lines); on the horizontal axis, the 426 stocks.

collection of all stocks (first row) and for subsets of them. The subsets are defined according to the 25, 50 and 75% quantiles of relative spreads and volumes. So, for instance, row $[Q^{0.25}, Q^{0.5})$ shows the mean and interquartile range of the means and variances of the stocks whose average relative spread and volume (over the period under study) are between the 25 and 50% quantiles. Two conclusions can be drown from this table. First, stocks with large relative spread and volume (columns $m_{\bar{x}}$) also have large variances (columns $m_{s_x^2}$). Very liquid stocks have tight spreads that do not move much, as the limit order book is very thick. However, since these are the stocks that are the most scrutinized, they are the first to react if an event arrives to the market (such as a bad macroeconomic news) and hence will experiment large changes in trading volume. On the other hand, the less liquid stocks take, in principle, longer to react to events, smoothing the traded volume, but, since the limit order book is thinner, an unexpected increase in the trading activity entails an unexpected variation in spread. Second, the interquantile range tells us that the stocks that are on the tails of the distributions of relative spread and volume show the largest discrepancies at both the average and variance. This suggests that, on one hand, there are large differences in the traded volume between the most liquid stocks. In a sense, there are stocks that are "more blue chips than others". On the other hand, the heterogeneity in the relative spread of the less liquid stocks suggests that differences in the thickness and depth of the limit order books are larger than among the more liquid stocks. All in one, there are significant differences in the liquidity measures between the stocks. As far as this article is concerned, the question is whether, regardless of these differences, the stocks share a common component and up to what extent it drives the relative spread and volume across all stocks. We provide an answer to this question in Section 5.

Table 2: Descriptive statistics

	Relative Spread				Volume			
	$m_{\bar{x}}$	$\text{IQR}_{\bar{x}}$	$m_{s_x^2}$	$\text{IQR}_{s_x^2}$	$m_{\bar{x}}$	$\text{IQR}_{\bar{x}}$	$m_{s_x^2}$	$\text{IQR}_{s_x^2}$
All stocks	0.96	0.35	0.58	0.29	2.59	2.07	6.70	2.23
$[0, Q^{0.25})$	0.62	0.08	0.26	0.15	0.56	0.31	0.16	0.15
$[Q^{0.25}, Q^{0.5})$	0.78	0.06	0.33	0.12	1.18	0.34	0.69	0.57
$[Q^{0.5}, Q^{0.75})$	0.95	0.08	0.42	0.13	2.16	0.57	1.82	1.10
$[Q^{0.75}, +\infty)$	1.46	0.45	1.26	0.74	6.14	3.60	18.51	15.54

Descriptive statistics for the relative spread and volume for all stocks, and for subsets thereof classified by quantile ranges (subsequent rows); Q^{α} stands for the quantile of order α and $[Q^{\alpha_1}, Q^{\alpha_2})$ for all the stocks with liquidity measure (either relative spread or volume) lying between the α_1- and α_2- quantiles. The column $m_{\bar{x}}$ (respectively $m_{s_x^2}$) shows the sample mean of individual stock means (respectively variances), and the column $\text{IQR}_{\bar{x}}$ (respectively $\text{IQR}_{s_x^2}$) the interquantile range of individual stock means (respectively variances).

The left plots in Figure 3 show the evolution of the averaged relative spread (top plot) and volume (bottom plot) over the 426 stocks. Visual inspection does not suggest any violation of the assumption of second-order stationarity. Yet, it clearly reveals some heteroskedasticity. Note that all these plots show averaged values. The assumption for applying GDFM is that all the relative spread and volume series in $\mathbf{SPR}_n$ and $\mathbf{VOL}_n$ are second order stationary. Because of obvious space reasons we do not show the 852 series, but visual inspection also suggests that they all fulfill the GDFM assumption of second order stationarity.[8] Right plots of Figure 3 show the autocorrelation functions of the averaged relative spread (top plot) and volume (bottom plot). They confirm the well-known stylized fact that liquidity time series are strongly autocorrelated. The averaged relative spread seems to be more persistent than averaged volume, as the autocorrelations decay more slowly. This may be due to the fact that the S&P500 constituents we consider are in general heavily traded with small spreads near the tick, which implies persistent autocorrelations. In the spirit of Chordia et al. (2000), these averaged relative spread and volume series can be seen as estimates of market liquidity. Their strong and persistent autocorrelations give more credit to our claim that liquidity needs models that can handle time dependencies.

Figure 3: Averaged data and their autocorrelograms

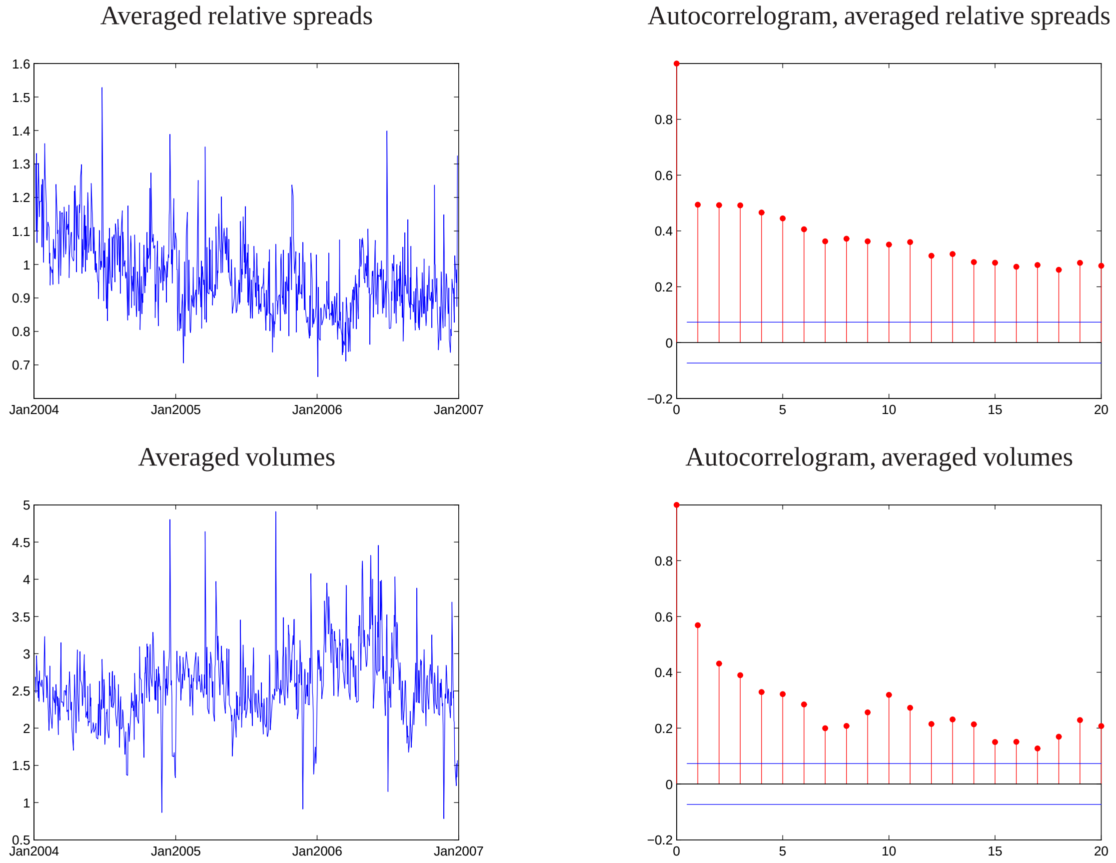

Averaged (over all stocks) observed liquidity measures (left), along with their autocorrelograms (right).

The left plot in Figure 4 presents a scatter plot of the contemporaneous average spread and volume time series. This plot shows a cloud with undefined principal directions, indicating

[8]Results are available under request.

that relative spread and volume are not contemporaneously correlated. Correlations however are found when looking at lagged relationships. The right plot shows the cross-correlations (y-axis) for different orders (x-axis). Negative orders stand for the relation between lagged volume and lead relative spread. So, for instance, the correlation of order -5 explains the relationship between volume five days ago and today's relative spread. We observe that, although there is some cross-correlation at the negative orders, the bulk of it is on the positive orders: relative spread leads volume. The conclusions drawn upon these cross-correlations go in the same direction as for the autocorrelograms of Figure 3: liquidity panels show important time series features that cannot be explained nor exploited via static factor models. In the next section, we show the results of a GDFM analysis.

Figure 4: Scatter plot and cross-correlogram

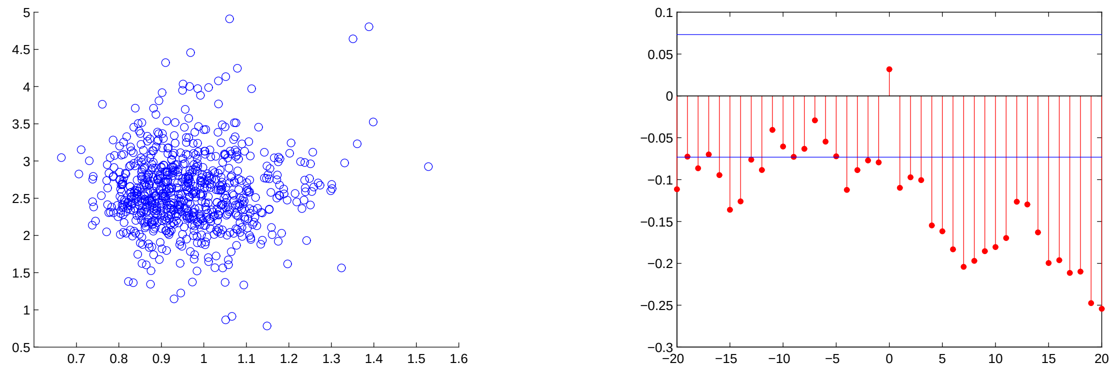

Left plot shows the scatter plot of relative spread against volume. Right plot shows the cross-correlogram of averaged relative spreads against averaged volumes. Negative lags correspond to cross-correlations between leading volumes and lagging relative spreads.

5 Main results

As explained in Section 3, we use the Hallin and Liška (2007) information criterion to identify the numbers of factors, that is, the dynamic dimensions of the common spaces of the three (sub)panels $\mathbf{SPR}_n$, $\mathbf{VOL}_n$ and $\mathbf{LIQ}_n$. Identification is based on a visual inspection of the three double plots of Figure 5. For each (sub)panel, the figure shows a measure (the variance $V(c)$, dashed line) of the instability of the selection associated with various values of the tuning factor c, along with the final selection associated with the same value of c (solid line). The procedure then consists in spotting the second interval (starting from the left) of c values over which the dashed line touches the horizontal axis (hence, $V(c) = 0$); the number of factors to be selected then is obtained by reading, on the solid line curve, the corresponding value shown by the solid line.

Each of the three plots leads to a selection of one single factor. This and Lemma 1 in Hallin and Liška (2008) implies the existence of a unique strongly common factor driving the common components of both subpanels, thus yielding the particular case, described in (7), of empty weakly common and weakly idiosyncratic spaces: see Figure 6.[9] Two important conclusions can be drawn from this result. First, $\mathbf{SPR}_n$ and $\mathbf{VOL}_n$ share the same common space, meaning that the shocks driving commonness in relative spread and commonness in volume are the same. This result supports the conjecture that, in a liquid market such as S&P500, a single liquidity measure (either relative spread or volume) suffices to understand market liquidity dynamics. Second, this common space presents a dynamic dimension one, suggesting some homogeneity when markets are confronted with liquidity. It may be an indication that no market liquidity effects are originating from, for instance, different sectors or different types of investors but only from the market itself.[10] The innovation of the one-dimensional factor driving both the common spread and the common volume components therefore strongly qualifies as the unobservable market liquidity shock.

We find that, on average, market liquidity accounts for 12% of total variations of relative spread and for 18% of total variations of volume, as shown in Figure 6. These proportions are much larger than those of Chordia et al. (2000) (about 2-4%), larger also than those of Hasbrouck and Seppi (2001) (about 8-14%), and they are comparable to those of Korajczyk and Sadka (2008). These differences can be due to an array of reasons: different databases, different time frequencies, different liquidity measures, and different methods to extract commonality in liquidity. Yet, our results support previous studies and offer an alternative based on a representation result, rather than on model assumptions, and on a rigorous methodology that leaves little room to subjectivism. Idiosyncratic components on the other hand account for 88% and 82% of total variations of relative spread and volume. As mentioned earlier, this large proportion does not mean that relative spread and volume are very noisy measures. Since GDFM allows for mild correlation among the idiosyncratic components of individual stocks, some groups of stocks may share liquidity drivers that are uncorrelated with market liquidity. Further, idiosyncratic terms in the GDFM may be autocorrelated, i.e. an idiosyncratic component at time t may contain information about its future values.

All these percentages and variance decompositions are proportions averaged over the panel. Individual decompositions, however, also may reveal interesting features. The plots in Figure 7 show the proportions explained by the common components for each individual stock (left plot is for relative spread, right plot for volume). Stocks are ordered from smallest to largest relative spread and volume, respectively. Vertical lines divide the stocks according to the same quantile ranges as in Table 2. In that table we found significant differences for those stocks with the largest relative spread and volume. Similar features are observed in Figure 7. For stocks with average relative spread and volume smaller than the 75% quantile, the proportions of total variance explained by market dynamics do not exhibit any clear pattern, even though showing significant differences among them. However, beyond the 75% quantile, the market impact on total variances seems to be increasing with the liquidity measure value (this is very clear for volume, somewhat less so for relative spread since there is also a large

[9]In a way, this result is in line with the empirical intuition provided by Hasbrouck and Seppi (2001), who argue that only the first principal component, in their PCA approach, should be taken into account.

[10]Such effects may exist; but then they only have an impact on idiosyncratic components, hence on a limited number of stocks.

Figure 5: Implementation of the Hallin-Liška identification method

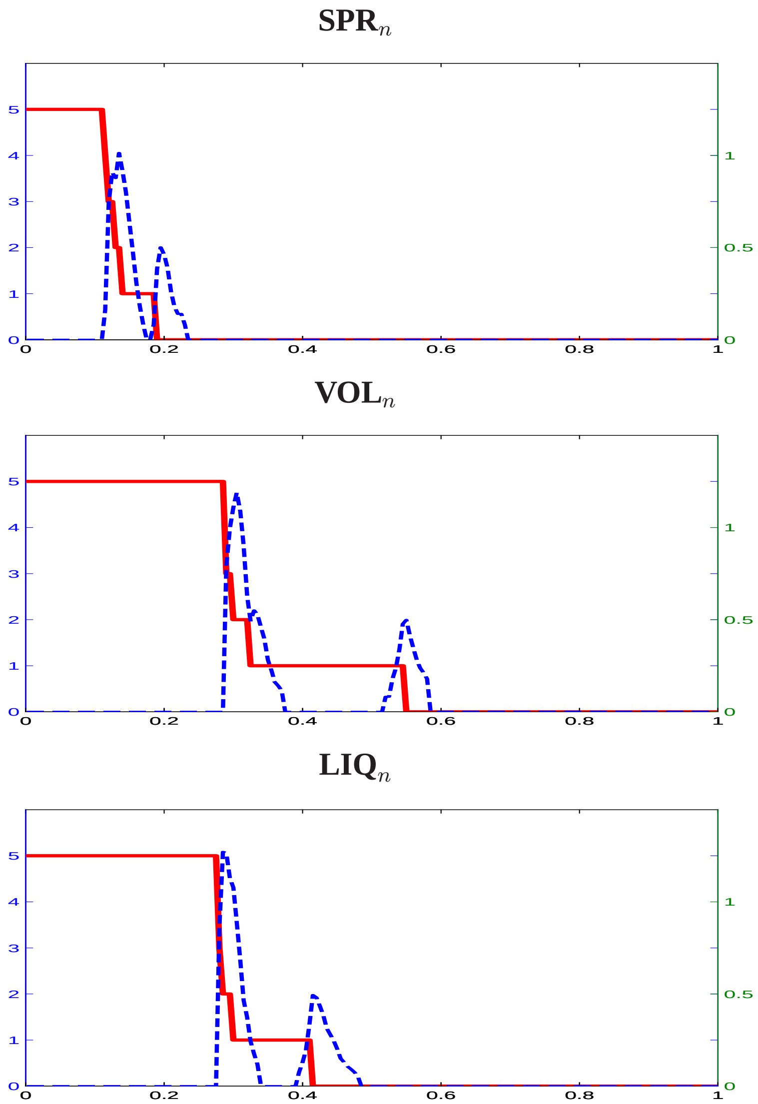

Application of the Hallin and Liška (2007) information criterion in the identification of the dynamic dimensions of the common spaces of the various (sub)panels (top: relative spreads; middle: volumes; bottom: joint panel), as described in Section 5. Dashed lines are a measure of the instability of the selection associated with various values of the tuning factor c. Solid lines are the final selected number of factors as a function of the tuning factor c.

Figure 6: Spaces, factors and variance decompositions

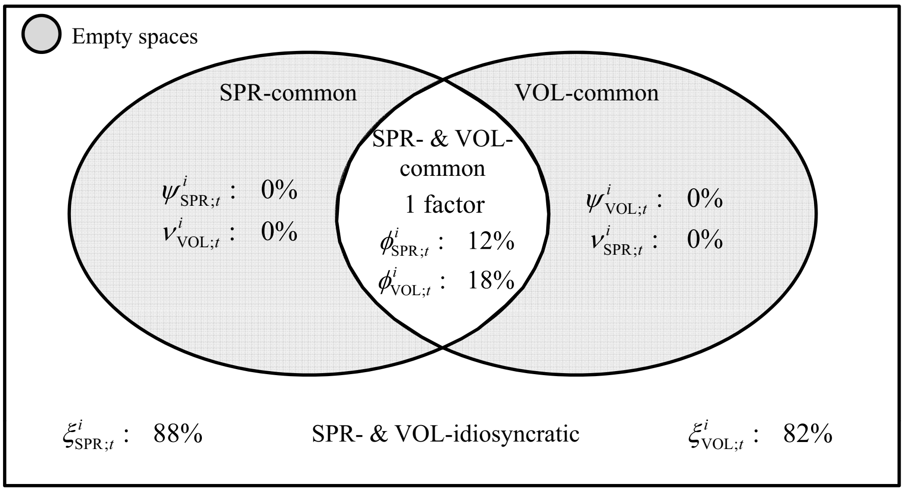

increase in variability). This fact may be explained with similar arguments as in the discussion of Table 2. Relative spreads lying beyond the 75% quantile are more sensitive to market conditions, as their limit order books are thinner. A shock in market liquidity should imply a reaction in relative spread that is more important than it is for more liquid stocks. On the other hand, stocks with the largest volumes, the blue chips, are the driving forces of the market, which means that a shock in market liquidity entails a reaction in volume which is larger than for the less liquid stocks.

Figure 7: Variances of common components of invidual stocks

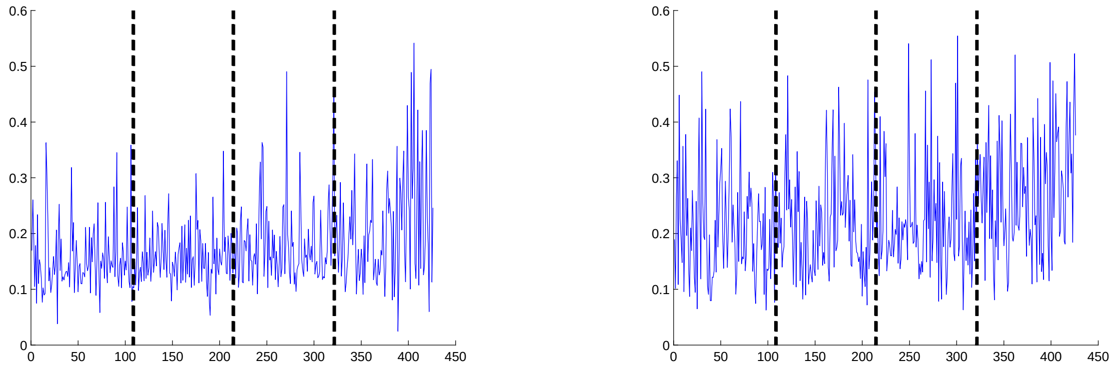

Proportions of variance explained by the common components for all individual stocks and for relative spread (left plot) and volume (right plot). Stocks are ordered from smallest relative spread (respectively volume) to largest.

Figure 8 shows a plot of the a cross-sectional average of relative spread (top left) and volume (top right) autocorrelations, along with the corresponding plots for their common (middle plots) and idiosyncratic components (bottom plots). These plots reveal quite interesting differences between those components.

First, note that the autcorrelations of averaged relative spreads, as shown in Figure 3, strikingly differ from these averaged autocorrelations of observed series and common components. Averaging all spreads and volumes indeed cancels out all leading/lagging dynamics that may exist between stocks, and therefore provides a biased picture of reality. These important differences, especially for relative spread, demonstrate the danger of using averaged relative spreads or volumes as a proxy for market liquidity, and, more particularly, the danger of defining market liquidity shocks as the innovation of cross-sectionally averaged spread of volume series.

Second, by the fundamental property of GDFM that common and idiosyncratic components are mutually orthogonal at all leads and lags, the average autocorrelations of observed series are a linear combination of the common and idiosyncratic average autocorrelations, with coefficients given by the variance ratios. As a consequence of idiosyncratic predominance in variance decompositions, the autocorrelation profiles of observed series look closer to those of the corresponding idiosyncratic components than to those of the common ones.

Third, the autocorrelations for the common components look very different for relative spread and volume. This indicates that the ways market liquidity shocks are transferred to relative spread and to volume also are quite different. While the impact of a market liquidity shock on volume is instantaneously very significant, it also vanishes relatively fast. By contrast, the same market liquidity shock has a rather weak impact on relative spread, but that impact decays very slowly with time. This difference in the way liquidity shocks are loaded by the two liquidity measures under study can be explained by the fact that, as mentioned in Section 4, the limit order book, in the stocks we consider, is relatively thick and hence, on average, the spread does not move quickly, so that the impact of a market shock stays for long. But this is not the case for volume, a measure that adapts much faster to market variations. A more precise discussion of this, would require a comparative analysis of the loading filters—which is impossible at this stage.

Finally, observe that idiosyncratic volume components are much more autocorrelated than the idiosyncratic relative spread ones. This indicates that observed persistence in relative spread almost entirely originates in market dynamics, whereas serial autocorrelations for observed volume clearly both have market-wide and idiosyncratic origins. A possible explanation for this is that, while the relative spread essentially is a bounded variable, and less dependent on stock specificities, trading volumes are clearly connected to the size of the firms, so that huge cross-sectional magnitude discrepancies may exist.

Summing up, although relative spread and volume provide equivalent characterizations of market liquidity as dynamic factors, one should not conclude that they constitute equivalent liquidity measures, since the ways they react to market liquidity shocks are drastically different.

Figure 8: Averaged autocorrelograms: Observed, Common and Idiosyncratic

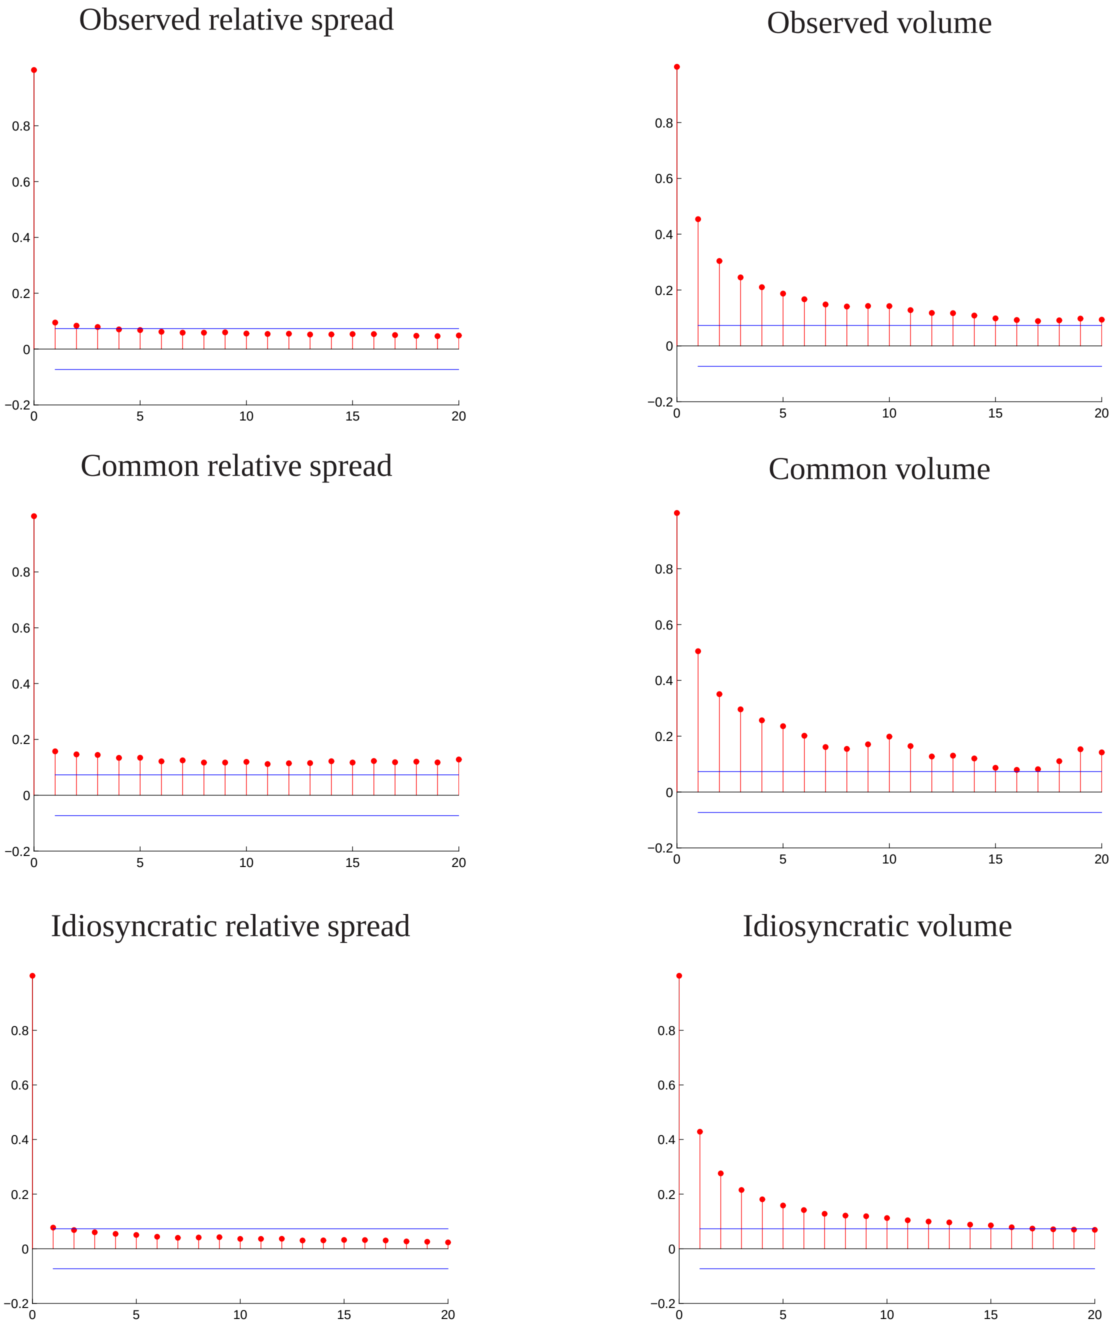

Averaged (over all stocks) autocorrelograms of observed liquidity measures (top), their common (middle) and idiosyncratic (bottom) components; to be contrasted with Figure 3, where autocorrelograms of averaged (over all stocks) liquidity measures, common and idiosyncratic components, are shown.

6 Conclusions

The GDFM presents a number of advantages in the identification, analysis and forecasting of market liquidity dynamics. First, unlike its competitors, it is based on a general representation result which is free of restrictive model assumptions. Second, it tackles co-movement and time dependencies, two stylized facts of liquidity time series. Third, it provides a clear distinction between commonness and idiosyncrasy. Fourth, it allows to estimate the dimension of the common space. Finally, it allows identifying commonality over different liquidity measures on a global analysis. An application of GDFM to panels of relative spreads and volumes suggests that these two liquidity-related quantities actually convey the same information about market liquidity. The one-dimensional common shocks driving these two panels therefore strongly qualify as the unobservable market liquidity shocks. Such results of course are calling for more extensive and detailed empirical investigations, involving larger databases and further liquidity measures.

Extensions of the present paper go hand in hand with further developments in the theory of GDFM. For instance, a most attractive research direction would consist in a comparative study, based on their dynamic factor loading filters, of stocks and liquidity measures. These filters indeed characterize the way these stocks and liquidity measures react to market liquidity shocks. Such a study however requires being able to recover the fundamental liquidity shocks; Forni and Lippi (2009) provide a significant progress into that direction. A clean identification of the impact of liquidity shocks on various liquidity measures also would lead to a better assessment of the links between liquidity and asset pricing, and a better analysis of the macroeconomic drivers of liquidity. Uncertainty in liquidity could be dealt with by applying the Dynamic Factor GARCH model introduced by Alessi, Barigozzi and Capasso (2007), while the Eichler, Motta and von Sachs (2008) extension of GDFM to non stationary time series opens the door, for instance, to a better understanding of financial crises.

References

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

[2] Alessi, L., Barigozzi, M., Capasso, M., 2007. Dynamic Factor GARCH: multivariate volatility forecast for a large number of series. LEM Working Paper.

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

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

[5] Bai, J. , Ng, S., 2002. Determinig the number of factors in approximate factor models. *Econometrica* **79**, 191-221.

[6] Chordia, T., Roll, R., Subrahmanyam, A., 2000. Commonality in liquidity. *Journal of Financial Economics* **56**, 3-28.

[7] Domowitz, I., Wang, X., 2002, Liquidity, liquidity commonality and its impact on portfolio theory. Penn State University Working Paper.

[8] Eckbo, B.E., Norli, O., 2002. Pervasive Liquidity Risk. Working Paper.

[9] Eichler, M., Motta, G., von Sachs, R., 2008. Fitting dynamic factor models to non-stationary time series. UCL Working Paper.

[10] Forni, M., Lippi, M., 2001. The generalized dynamic factor model: representation theory. *Econometric Theory* **17:6**, 1113-1141.

[11] Forni, M., Lippi, M., 2009. The unrestricted dynamic factor model: representation results. Working Paper.

[12] Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2000. The generalized dynamic factor model: identification and estimation. *The Review of Economics and Statistics* **82**, 540-554.

[13] Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2004. The generalized dynamic factor model: consistency and rates. *Journal of Econometrics* **119**, 231-255.

[14] Forni, M., Hallin, M., Lippi, M., Reichlin, L., 2005. The generalized dynamic factor model: one-sided estimation and forecasting. *Journal of the American Statistical Association* **100**, 830-840.

[15] Hallin, M., Liška, R., 2007. The generalized dynamic factor model: determining the number of factors. *Journal of the American Statistical Association* **102**, 103-117.

[16] Hallin, M., Liška, R., 2008. Dynamic factors in the presence of block structure. EUI ECO Working Paper 2008/22.

[17] Harris, L., 2003. *Trading and Exchanges. Market Microstrucrure for Practitioners*. Oxford University Press, Oxford.

[18] Hasbrouck, J., Seppi, D., 2001. Common factors in prices, order flows, and liquidity. *Journal of Financial Economics* **59**, 383-411.

[19] Korajczyk, R.A., Sadka, R., 2008. Pricing the commonality across alternative measures of liquidity. *Journal of Financial Economics* **87**, 45-72

[20] Kyle, A., 1985. Continuous auctions and insider tradings. *Econometrica* **53**, 1315-1335.

[21] Martinez, M.A., Nieto, B., Rubio, G., Tapia, M., 2005. Asset pricing and systematic liquidity risk: an empirical investigation of the Spanish stock market. *International Review of Economics and Finance* **14**, 81-103.

[22] Minguet, A., 2003. *La Microstructure des Marchés d'Actions*. Economica, Paris.

[23] O'Hara, M., 1998. *Market Microstructure Theory*. J. Wiley, New York.

[24] Pastor, L., Stambaugh, R.F., 2003. Liquidity risk and expected stock returns. *Journal of Political Economy* **111**, 642-685.

[25] Schwartz, R., 1993, *Reshaping Equity Markets: a Guide for the 1990's*. Business One Irwin, Homewood, Illinois.

[26] Stock, J., Watson, M., 2002a. Macroeconomic forecasting using diffusion indexes. *Joural of Business and Economic Statistics* **20**, 147-162.

[27] Stock, J., Watson, M., 2002b. Forecasting using principal components from a large number of predictors. *Journal of the American Statistical Association* **47**, 1167-1179