Stock Market Liquidity and the Business Cycle

Randi Næs, Johannes A. Skjeltorp and Bernt Arne Ødegaard *

March 2010

Forthcoming, *Journal of Finance*

Abstract

In the recent financial crisis we saw the liquidity in the stock market drying up as a precursor to the crisis in the real economy. We show that such effects are not new, in fact we find a strong relation between stock market liquidity and the business cycle. We also show that the portfolio compositions of investors change with the business cycle and that investor participation is related to market liquidity. This suggests that systematic liquidity variation is related to a "flight to quality" during economic downturns. Overall, our results provide an new explanation for the observed commonality in liquidity.

In the discussion of the current financial crisis, much is made of the apparent causality between a decline in the liquidity of financial assets and the economic crisis. In this paper we show that such effects are not new, changes in the liquidity of the US stock market have been coinciding with changes in the real economy at least since the Second World War. Stock market liquidity is in fact a very good "leading indicator" of the real economy. Using data for the US over the period 1947 to 2008, we document that measures of stock market liquidity contain leading information about the real economy, also after controlling for other asset price predictors.

Figure 1 shows a time series plot of a measure of market liquidity (the Amihud (2002) measure) together with the NBER recession periods (grey bars). This figure serves to illustrate the relationship found between stock market liquidity and the business cycle as liquidity clearly worsens (illiquidity increases) well ahead of the onset of the NBER recessions.

*Randi Næs is at the Ministry of Trade and Industry. Email: ran@nhd.dep.no. Johannes A Skjeltorp is at Norges Bank (the Central Bank of Norway). Email: Johannes-A.Skjeltorp@norges-bank.no. Bernt Arne Ødegaard is at the University of Stavanger and Norges Bank. Email: bernt.a.odegaard@uis.no. We are grateful for comments from an anonymous referee, associate editor, and our editor (Campbell Harvey). We also thank Kristian Miltersen, Luis Viceira and seminar participants at the 4th Annual Central Bank Workshop on the Microstructure of Financial Markets in Hong Kong, Norges Bank, the Norwegian School of Economics and Business Administration (NHH), Statistics Norway (SSB), CREST and the Universities of Oslo, Stavanger and Aarhus (CREATES) for comments. Ødegaard acknowledges funding from "Finansmarkedfondet" (The Finance Market Fund). The views expressed are those of the authors and should not be interpreted as reflecting those of Norges Bank or the Ministry of Trade and Industry.

Figure 1: Liquidity and the business cycle

The figure shows time series plots of the detrended Amihud (2002) illiquidity ratio (*ILR*) for the US over the period 1947-2008. The gray bars indicate the NBER recession periods. The *ILR* is an elasticity (price impact) measures of liquidity and reflects how much prices move as a response to trading volume. The *ILR* is first calculated for each stock for each year. Then the equally weighted cross sectional average for each year is calculated. A more precise definition is found in equation (2) in the paper. Note that the *ILR* reflect illiquidity, so a high value reflect a high price impact of trades(i.e. low liquidity). *ILR* is detrended using a Hodrick-Prescott filter.

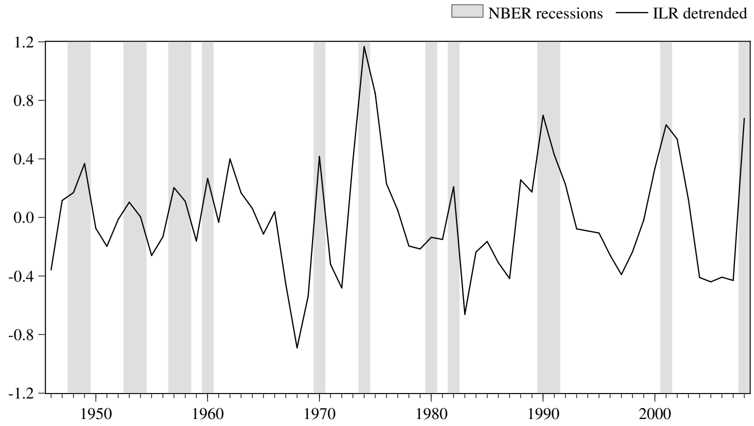

Our results are relevant for several strands of the literature. One important strand is the literature on forecasting economic growth using different asset prices, including interest rates, term spreads, stock returns and exchange rates. The forward-looking nature of asset markets makes the use of these prices as predictors of the real economy intuitive. If a stock price equals the expected discounted value of future earnings, it seems natural that it should contain information about future earnings growth. Theoretically, a link between asset prices and the real economy can be established from a consumption–smoothing argument. If investors are willing to pay more for an asset that pays off when the economy is thought to be in a bad state than an asset that pays off when the economy is thought to be in a good state, then current asset prices should contain information about investors' expectations about the future real economy. In their survey article, Stock and Watson (2003) conclude, however, that there is considerable instability in the predictive power of asset prices.

We shift focus to a different aspect of asset markets, the liquidity of the stock market, i.e. the costs of trading equities. It is a common observation that stock market liquidity tends to dry up during economic downturns. However, we show that the relationship between trading costs and the real economy is much more pervasive than previously thought. A link from trading costs to the real economy is not as intuitive as the link from asset prices, although several possible explanations are suggested in the literature.

One might speculate that the observed effects are the results of aggregate portfolio shifts from individual investors, where changes in desired portfolios are driven by changes in individuals' expectations of the real economy. This is an example of the well known idea of a "flight to quality" or "flight to liquidity," see for instance Longstaff (2004).[1] We

[1]The term "flight to quality" refers to a situation where market participants suddenly shift their

find some empirical evidence consistent with this hypothesis. First, using data for the US, we show that the informativeness of stock market liquidity for the real economy differs across stocks. In particular, the most informative stocks are those of small firms, which are the least liquid. Second, using data for Norway, where we have unusually detailed information about the composition of ownership of the whole stock market, we show that changes in liquidity coincide with changes in portfolio compositions of investors of the hypothesized type. Before economic recessions we observe a "flight to quality", where some investors leave the stock market altogether, and others shift their stock portfolios into larger and more liquid stocks.

Brunnermeier and Pedersen (2009) provide an alternative explanation based on the interaction between securities' market liquidity and financial intermediaries availability of funds. In the model, liquidity providers ability to provide liquidity depends on their capital and margin requirements. During periods of financial stress, a reinforcing mechanism between market liquidity and funding liquidity leads to liquidity spirals. Reduced funding liquidity leads to a flight to quality in the sense that liquidity providers shift their liquidity provision towards stocks with low margins. In our Norwegian data set, we find that mutual funds have a stronger tendency to realize their portfolios in small stocks during downturns than the general financial investor. This result provides some support for the model as mutual funds are most likely to face funding constraints during economic downturns (withdrawals from investors who have to realize their portfolios). The current financial crises has shown that high systemic risk and funding liquidity problems in the financial sector can spread to the real economy.

Another possibility is that stock market liquidity has a causal effect on the real economy, through investment channels. This could for example be that a liquid secondary market makes it easier for investors to invest in productive, but highly illiquid, long-run projects (Levine, 1991; Bencivenga, Smith, and Starr, 1995). Empirical studies provide some support for this hypothesis. In a cross-country regression, Levine and Zervos (1998) find a significant positive correlation between stock market liquidity and current and future rates of economic growth, after controlling for economic and political factors. Moreover, some recent empirical evidence suggests that stock market liquidity is positively related to the costs of raising external capital.[2]

Even though there exist several possible explanations for a link between stock market liquidity and the real economy, it is still puzzling that liquidity measures provide information about the real economy that is not fully captured by stock returns. One explanation of why liquidity seems to be a better predictor than stock price changes is that stock prices contain a more complex mix of information that makes the signals from stock returns more blurred (Harvey, 1988).

Two recent papers that investigate the relationship between equity order flow and macro fundamentals are closely related to our work. Beber, Brandt, and Kavajecz (2010)

---

portfolios towards securities with less risk. In Longstaff (2004) a "flight to liquidity" is defined as a distinct phenomenon where market participants shift their portfolios from less liquid to more liquid bonds with identical credit risk, i.e. from "off the run" to "on the run" Treasuries. We will use the term "flight to quality" throughout the paper, although the portfolio shifts we are assuming are also likely to have elements of a flight to liquidity.

[2]See Lipson and Mortal (2009), which shows a link between capital structure and liquidity. Also, for some direct evidence, see Skjeltorp and Ødegaard (2010), who shows that firms are willing to pay for improved liquidity before seasoned equity issues.

examine the information in order flow movements across equity sectors over the period 1993-2005 and find that an order flow portfolio based on cross-sector movements predicts the state of the economy up to three months ahead. They also find that the cross section of order flow across sectors contains information about future returns in the stock and bond markets. Kaul and Kayacetin (2009) study two measures of aggregate stock market order flow over the period 1988-2004 and find that they both predict future growth rates for industrial production and real GDP. The common theme of these two papers and our research is that the trading process in stock markets contains leading information about the economy. Our results are by far the most robust ones as they are based on a sample period that spans over 60 years and cover 10 recessions. The two order flow papers also find some evidence that order flow contains information about future asset price changes. Kaul and Kayacetin (2009) and Evans and Lyons (2008) argue that the extra information contained in order flow data can be explained by aggregate order flows bringing together dispersed information from heterogeneously informed investors.

A number of other papers are related to our study. Fujimoto (2003) and Söderberg (2008) examine the relationship between liquidity and macro fundamentals. However, they both investigate whether time-varying stock market liquidity has macroeconomic sources. They do not consider the possibility of causality going the other way. Gibson and Mougeot (2004) find some evidence that a time-varying liquidity risk premium in the US stock market is related to a recession index over the 1973-1997 period.

Our paper also contributes to the market microstructure literature on liquidity. Several empirical studies have shown evidence of commonality and time variation in stock market liquidity measures, see Chordia, Roll, and Subrahmanyam (2000), Huberman and Halka (2001) and Hasbrouck and Seppi (2001). It is also well documented that time variation in liquidity affects asset returns, see for example Pastor and Stambaugh (2003) and Acharya and Pedersen (2005). The phenomenon of commonality is, however, so far not fully understood. The Brunnermeier and Pedersen (2009) model discussed above can explain commonality across stocks, although the model is probably most relevant during periods of financial stress.[3] Our finding that time-varying aggregate stock liquidity has a business cycle component is new and quite intriguing. It suggests that pricing of liquidity risk cannot be explained solely by uninformed investors who require a premium for ending up with the stock that the informed investors sell, as suggested in O'Hara (2003). Hence, the traditional arguments why market microstructure matters for asset returns might be too narrow.

By showing that microstructure liquidity measures are relevant for macroeconomic analysis, our paper also enhances our understanding of the mechanism by which asset markets are linked to the macro economy. We show that the predictive power of liquidity holds up to adding existing asset price predictors. Given the documented instability in the predictive power of asset prices, an incremental indicator that might react earlier or in some way differently to shocks in the economy should be useful, also for policy purposes.

The rest of the paper is structured as follows. First, in section I, we look at the data. We define the measures we use, discuss the data sources and give some summary statistics. Next, in section II we document that liquidity is related to the real economy using data

[3]Coughenour and Saad (2004) investigate commonality in liquidity amongst stocks handled by the same NYSE specialist firm and provide some evidence in favor of the Brunnermeier and Pedersen (2009) model.

for the US in the period 1947-2008. In section III we look closer at the causes of this predictability by splitting stocks into size groups and showing that the main source of the predictability is reflected in the liquidity variation of small, relatively illiquid, stocks. In section IV we use Norwegian data to do two things. First, we confirm the US results, that stock market liquidity contains information about the macroeconomy. We go on to show some evidence of the causes of time variation in aggregate liquidity, by linking changes in liquidity to changes in the portfolio composition of all investors at the Oslo Stock Exchange. We construct several measures of changes in the portfolio composition of investors and show that periods when liquidity worsens are the same as periods when there is a "flight to quality" in the stock portfolios of owners. Finally, section V offers some concluding remarks.

I Liquidity measures and data

A Liquidity measures

Given that there are numerous theoretical definitions of liquidity, there are also many different empirical measures used to measure liquidity. Since our focus is on the link between liquidity and the real economy, we are agnostic about this. We use a number of common measures and show that the relevant links are relatively independent of which liquidity measures we employ. Our choices of liquidity measures are driven by our desire for reasonably long time series. Many liquidity measures require intra-day information on trades and orders to be calculated, which is not available for the long time period considered in this paper. We therefore employ measures that can be calculated using data available at a daily frequency. We consider the following four liquidity measures: Relative spread (RS), the Lesmond, Ogden, and Trzcinka (1999) measure (LOT), the Amihud (2002) illiquidity ratio (ILR) and the Roll (1984) implicit spread estimator (*Roll*). The "low-frequency" versions of these liquidity proxies are shown in Goyenko and Ukhov (2009) and Goyenko, Holden, and Trzcinka (2009) to do well in capturing the spread cost and price impact estimated using intra-day data. Note that all the liquidity measures we employ in this study measure illiquidity. Thus, when the measures have a high value, market liquidity is low and it is costly to execute a trade.

Spread costs are observed in dealer markets as well as in limit order markets. The relative spread (RS) is the quoted spread (the difference between the best ask quote and bid quote) as a fraction of the midpoint price (the average of the best ask quote and bid quote) and measures the implicit cost of trading a small number of shares.

Lesmond et al. (1999) suggest a measure of transaction costs (hereafter the LOT measure) that does not depend on information about quotes or the order book. Instead, the LOT measure is calculated from daily returns. It uses the frequency of zero returns to estimate an implicit trading cost. The LOT cost is an estimate of the implicit cost required for a stock's price *not* to move when the market as a whole moves. To get the intuition of this measure, consider a simple market model,

$$R_{it} = a_i + b_i R_{mt} + \varepsilon_{it} \tag{1}$$

where R_{it} is the return on security i at time t, R_{mt} is the market return at time t, a is a constant term, b is a regression coefficient and ε is an error term. In this model, for *any* change in the market return, the return of security i should move according to (1). If it does not, it could be that the price movement that *should* have happened is not large enough to cover the costs of trading. Lesmond et al. (1999) estimate how wide the transaction cost band around the current stock price has to be to explain the occurrence of no price movements (zero returns). The wider this band, the less liquid the security. Lesmond et al. shows that their *LOT* measure is closely related to the bid-ask spread.

We also employ as a liquidity measure the Roll (1984) estimate of the implicit spread. This spread estimator, also called the effective bid-ask spread, is measured from the serial covariance of successive price movements. Roll shows that assuming the existence of a constant effective spread s, this can be estimated as $\widehat{s} = \sqrt{-Scov}$ where $Scov$ is the first-order serial covariance of successive returns.[4] We calculate the *Roll* estimator based on daily returns.

Our final liquidity measure, Amihud (2002)'s illiquidity ratio (*ILR*), is a measure of the elasticity dimension of liquidity. Elasticity measures of liquidity try to estimate how much prices move in response to trading volume. Thus, cost measures and elasticity measures are strongly related. Kyle (1985) defines the price impact as the response of price to order flow. Amihud proposes a price impact measure that is closely related to Kyle's measure. The daily Amihud measure is calculated as,

$$ILR_{i,T} = 1/D_T \sum_{t=1}^{T} \frac{|R_{i,t}|}{VOL_{i,t}} \tag{2}$$

where D_T is the number of trading days within a time window T, $|R_{i,t}|$ is the absolute return on day t for security i, and $VOL_{i,t}$ is the trading volume (in units of currency) on day t. It is standard to multiply the estimate by 10^6 for practical purposes. The Amihud measure is called an illiquidity measure since a high estimate indicates low liquidity (high price impact of trades). Thus, *ILR* captures how much the price moves for each volume unit of trades.

B Liquidity data

To calculate the liquidity measures we use data on stock prices, returns, and trading volume. For the US, the data source is CRSP, and the sample we are looking at covers the period 1947 through 2008. To keep the sample as homogeneous as possible through the entire period, we restrict the analysis to the common shares of stocks listed at the New York Stock Exchange (NYSE). For Norway similar data to the CRSP data are obtained from the Oslo Stock Exchange data service.[5] The Norwegian sample covers the period 1980-2008. For both the US and Norwegian sample, we calculate the different liquidity measures each quarter for each security and then take the equally weighted average across securities for each liquidity variable.

[4] This estimator is only defined when $Scov < 0$. Harris (1990) suggests defining the $\widehat{s} = -2\sqrt{Scov}$ if $Scov > 0$, but this would lead to an assumed *negative* implicit spread. A negative transaction cost for equity trading is not meaningful. We therefore only use the Roll estimator for stocks with $Scov < 0$, and leave the others undefined.

[5] We use all equities listed at the OSE with the exception of very illiquid stocks. Our criteria for filtering the data are the same as those used in Næs, Skjeltorp, and Ødegaard (2008), i.e. that we remove years where a stock is priced below NOK 10, and remove stocks with less than 20 trading days in a year.

Table I: Describing liquidity measures

Panels A and B show descriptive statistics for the US liquidity measures. The US sample covers the period from 1947 through 2008. The liquidity measures examined are the relative bid-ask spread (RS), the Lesmond et al. (1999) measure (LOT), the Amihud (2002) illiquidity ratio (ILR) and the Roll (1984) implicit spread estimator ($Roll$). Note that the Relative spread is not universally available, the CRSP database only includes full data on spreads starting in 1980, but there are some observations earlier. The liquidity measures are calculated for each available stock once each quarter. Panel A shows the mean and median of the liquidity measures, the number of securities used, the total number of observations (each security is observed in several quarters), and estimates of average liquidity measures for different subperiods. Panel B shows correlation coefficients between the liquidity measures. The correlations are calculated across all stocks and time, i.e. the liquidity measures are calculated for each available stock once each quarter, and the correlations are pairwise correlations between these liquidity measures. Panels C and D show corresponding statistics for the Norwegian liquidity measures. The Norwegian sample covers the period from 1980 through 2008.

Panel A: Descriptive statistics, US liquidity measures

Liquidity					Means subperiods					
measure	mean	median	no secs	no obs	1947-59	1960-69	1970-79	1980-89	1990-99	2000-08
RS	0.021	0.014	4248	146262	0.021	0.019		0.020	0.027	0.016
LOT	0.035	0.022	5177	340076	0.027	0.031	0.051	0.037	0.040	0.027
ILR	0.657	0.056	5178	340668	1.900	0.818	0.829	0.294	0.366	0.176
$Roll$	0.017	0.013	5141	174326	0.012	0.013	0.015	0.015	0.017	0.018

Panel B: Correlation coefficients, US liquidity measures

	RS	LOT	$Roll$
LOT	0.72		
$Roll$	0.40	0.62	
ILR	0.41	0.38	0.32

Panel C: Descriptive statistics, Norwegian liquidity measures

					Means subperiods		
Liquidity					1980-1989	1990-1999	2000-2008
measure	mean	median	no secs	no obs			
RS	0.042	0.029	788	14942	0.041	0.046	0.040
LOT	0.054	0.039	753	14852	0.055	0.064	0.049
ILR	0.772	0.205	770	15092	1.149	0.875	0.452
$Roll$	0.027	0.021	663	7209	0.027	0.026	0.026

Panel D: Correlation coefficients, Norwegian liquidity measures

	RS	LOT	$Roll$
LOT	0.64		
$Roll$	0.65	0.51	
ILR	0.40	0.34	0.49

In Table I, we give a number of descriptive statistics for these series of liquidity measures. Note that for the US, we do not have complete data for bid-ask spreads and will therefore have to leave these out in our time series analysis for the US.[6] Looking first at the descriptive statistics for the US in panel A of Table I, we see that the average relative spread for the full sample period was 2.1%, while the relative spread of the median firm was 1.4%. Looking at the subperiod statistics, we see that there have been some changes over time across all liquidity measures. Panel B shows the correlations between the liquidity proxies for the US. We see that all the liquidity measures are positively correlated. The lowest correlation is between *ILR* and *Roll*, but the correlation is still as high as 0.32. In addition, the high correlation between *LOT* and *RS* indicates that *LOT* is a good estimator for the actual spread cost.

Panel C of table I gives similar descriptive statistics for the Norwegian sample. The liquidity of the Norwegian market has improved over the sample, but has also varied across subperiods. In Panel D we observe that all the liquidity proxies are strongly positively correlated also for Norway. Overall, the high correlations between these measures suggest they contain some of the same information.

C Macro data

To proxy for the state of the real economy we use real GDP (*GDPR*), unemployment rate (*UE*), real consumption (*CONSR*) and real investment (*INV*).[7] We also use a number of financial variables which are shown in the literature to contain leading information about economic growth. From the equity market we use *Excess market return* ($\mathtt{er_m}$), calculated as the value weighted return on the S&P500 index in excess of the 3-month T-bill rate, and *Market volatility* (*Vola*), measured as the cross-sectional average volatility of the sample stocks, where volatility is calculated as the standard deviation of daily returns over the quarter. We also use the *term spread* (*Term*), calculated as the difference between the yield on a 10-year Treasury bond benchmark and the yield on the 3-month T-bill, and the *credit spread*(*Cred*) measured as the yield difference between the Moody's Baa credit benchmark and the yield on a 30-year government bond benchmark. The Moody's long term corporate bond yield benchmark consists of seasoned corporate bonds with maturities as close as possible to 30 years.[8] We use similar macro series for Norway.[9]

[6]This is due to these not being present in the CRSP data for the whole period. They have been back-filled for the early period, but in the 1950s through the 1970s there is essentially no spread observations in the CRSP data.

[7]The *GDPR* series is the Real Gross Domestic Product, *UE* is the Unemployment rate for fulltime workers, *CONSR* is real Personal Consumption Expenditures, and *INV* is real Private Fixed Investments. All series are seasonally adjusted. *GDPR* and *INV* are from the Federal Reserve Bank of St Louis, *UE* is from the US Bureau of Labor Statistics, and *CONSR* from the US Dept of Commerce.

[8]The source of these variables is Ecowin/Reuters.

[9]*GDPR* is the real Gross Domestic Product for Mainland Norway (excluding oil production). *UE* is the Unemployment Rate (AKU), *CONSR* is the real Households Consumption Expenditure and *INV* is real Gross Investments. All numbers are seasonally adjusted. The data source is Statistics Norway (SSB).

D Time series adjustment of series

The sample period we are looking at covers more than 60 years. Over this long period changes in market structure, competition, technology and activity in financial markets potentially generate non-stationarities in the liquidity series. For this reason, we perform several unit root tests for each series to determine whether the series needs to be transformed to stationary series.

While we want to avoid the risk of obtaining spurious results, we also want to avoid the risk of over-differentiating our variables. We therefore employ two tests. The first test we use is the Augmented Dickey-Fuller (ADF) test with a null that the variable has a unit root. The second test we use is the test proposed by Kwiatkowski, Phillips, Schmidt, and Shin (1992) (KPSS), where the null hypothesis is that the series is stationary. As noted by Kwiatkowski et al., their test is intended to complement unit root tests, such as the ADF test. Among our liquidity proxies, the *Roll* measure is the only variable for which we reject the null of a unit root using the ADF test. We are also unable to reject the null (of stationarity) using the KPSS test. Both the *LOT* and *ILR* series are unit root processes according to the ADF test (both allowing for a drift and deterministic trend), and in both cases the null of stationarity is rejected by the KPSS test.[10]

With respect to the other financial variables we use in the analysis, both the excess market return ($\mathsf{er_m}$), stock market volatility (*Vola*) and the term spread (*Term*) are stationary. However, we cannot reject the null that the credit spread (*Cred*) has a unit root according to the ADF test. In addition, the null of stationarity is rejected by the KPSS test. The result that we cannot reject the null that the credit spread is a unit root has been documented by e.g. Pedrosa and Roll (1998) and Kiesel, Perraudin, and Taylor (2001). Thus, we will transform the *ILR*, *LOT* and *Cred* to preserve stationarity.

Since we are going to perform pseudo out-of-sample tests later in our analysis, we want to be careful when transforming the series and only use information available up to every point in time. For this reason, we report results using a very simple method for making *ILR*, *LOT*, and *Cred* stationary, namely to take log differences.[11] We similarly use a simple differentiation of the macro variables.[12]

Table II shows the contemporaneous correlations between the different variables used in the analysis for the US. All three liquidity measures are negatively correlated with the term structure and positively related to the credit spread. Thus, when market liquidity worsens, the term spread decreases and the credit spread increases. There is a positive correlation between all liquidity measures and market volatility, and a negative correlation between liquidity and the excess return on the market ($\mathsf{er_m}$). Thus, when market liquidity is low, market volatility is high and realized market returns are low. This is

[10] Also, looking at the correlograms for the different series, we see that the autocorrelation function for the Roll measure converges to zero relatively quickly (4 quarters). However, both the *ILR* and *LOT* measures are much more persistent with large and significant autocorrelations up to 24 quarters.

[11] We have also considered two alternative methods for making these three series stationary. One is to *demean* the series relative to a two-year moving average of the series. The other is to use a Hodrick-Prescott filter. In an internet appendix we show that these alternative methods provide similar results.

[12] *dGDPR* is the real GDP growth, calculated as $\mathsf{dGPDR} = \ln(\mathsf{GDPR_t}/\mathsf{GDPR_{t-1}})$. *dUE* is the growth in unemployment rate , calculated as $\mathsf{dUE} = \ln(\mathsf{UE_t}/\mathsf{UE_{t-1}})$, *dCONSR* is the real consumption growth, calculated as $\mathsf{dCONSR} = \ln(\mathsf{CONSR_t}/\mathsf{CONSR_{t-1}})$ and *dINV* is the real growth in investments, calculated as $\mathsf{dINV} = \ln(\mathsf{INV_t}/\mathsf{INV_{t-1}})$.

consistent with the findings in Hameed, Kang, and Viswanathan (2010) that negative market returns decrease stock liquidity. All liquidity variables are negatively correlated with growth in GDP, investments and consumption and positively correlated with the unemployment rate. Note that the macro variables are not known to the market participants before the following quarter, thus, these correlations are a first indication that there is real time information about current underlying economic growth in market liquidity variables. Furthermore, we also see that the term spread has a significant positive correlation with GDP growth and consumption growth, while the credit spread is negatively correlated with GDP growth, investment growth and consumption growth and positively correlated with unemployment. The signs of these correlations are what we would expect. Stock market volatility and returns are not significantly correlated with any of the macro variables, except for consumption growth. Finally, as one would expect, all the macro variables are significantly correlated with each other and have the expected signs.

Table II: Correlations

The table shows the Pearson correlation coefficients between the variables used in the analysis for the US. The associated p-values are reported in parenthesis below each correlation coefficient. *ILR*, *LOT* and *Roll* are the three liquidity measures. The cross sectional liquidity measures are calculated as equally weighted averages across stocks. *Term* is our proxy for the term spread and *Cred* is the credit spread. With respect to additional equity market variables, we examine market volatility (*Vola*) which is calculated as the cross sectional average volatility of all stocks in the CRSP database, and excess market return ($\mathtt{er_m}$) which is the return on the S&P500 index in excess of the risk-free rate (proxied by the 3-month T-bill rate). With respect to macroeconomic variables, *dGDPR* is real GDP growth, *dINV* is growth in investments, *dUE* is growth in the unemployment rate and *dCONSR* is real consumption growth.

	Market variables							**Macro variables**		
	dILR	*dLOT*	*Roll*	*Term*	*dCred*	*Vola*	$\mathtt{er_m}$	*dGDPR*	*dINV*	*dCONSR*
Term	-0.17	-0.14	-0.04							
	(0.00)	(0.04)	(0.55)							
dCred	0.32	0.34	0.42	-0.21						
	(0.00)	(0.00)	(0.00)	(0.00)						
Vola	0.30	0.57	0.47	-0.15	0.42					
	(0.00)	(0.00)	(0.00)	(0.02)	(0.00)					
$\mathtt{er_m}$	-0.53	-0.19	-0.35	0.33	-0.17	-0.33				
	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)				
dGDPR	-0.16	-0.10	-0.31	0.16	-0.27	0.01	0.09			
	(0.02)	(0.15)	(0.00)	(0.02)	(0.00)	(0.87)	(0.19)			
dINV	-0.16	-0.17	-0.40	0.18	-0.26	-0.07	0.09	0.73		
	(0.02)	(0.01)	(0.00)	(0.00)	(0.00)	(0.27)	(0.21)	(0.00)		
dCONSR	-0.27	-0.15	-0.38	0.21	-0.34	-0.08	0.16	0.68	0.57	
	(0.00)	(0.02)	(0.00)	(0.00)	(0.00)	(0.24)	(0.01)	(0.00)	(0.00)	
dUE	0.16	0.15	0.33	-0.10	0.28	0.08	-0.04	-0.65	-0.62	-0.56
	(0.01)	(0.03)	(0.00)	(0.14)	(0.00)	(0.21)	(0.58)	(0.00)	(0.00)	(0.00)

E Norwegian ownership data

An important reason for including Norwegian data in the paper is the availability of data on stock market ownership for all investors at the Oslo Stock Exchange, which we use to investigate aggregate patterns in stock ownership.

Our data on stock ownership is from the centralized records on stock ownership in Norway. All ownership of stocks at the Oslo Stock Exchange is registered in a single, government-controlled entity, the Norwegian Central Securities Registry (VPS). From this source we have access to monthly observations of the equity holdings of the complete stock market. At each date we observe the number of stocks held by every owner. Each owner has a unique identifier which allows us to follow each owner's holdings over time. For each owner the data also includes a sector code that allows us to distinguish between such types as mutual fund owners, financial owners (which include mutual funds), industrial (nonfinancial corporate) owners, private (individual) owners, state owners and foreign owners. This data set allows us to construct the actual monthly portfolios of all investors at the stock exchange. We can also calculate, for each stock, measures of ownership concentration and fractions held by different owner types.[13] Table III shows some descriptive statistics for the stock ownership data at the Oslo Stock Exchange.

Table III: Descriptive statistics for the Norwegian ownership data

The table shows some summary statistics for the Norwegian ownership data. For each stock we calculate the fraction of the stock held by its largest owner (Largest owner) and three largest owners (Three largest). We also calculate two Herfindahl indices; the sum of squared ownership fractions of all the firms' owners (Herfindahl index), and the sum of squared ownership fractions of all but the three largest owners (Herfindahl index all but 3 largest). We also show the total number of owners, only counting owners owning more than 100 shares. (Total no owners > 100 shares), and the fraction of the firm held by the five different mutually exclusive owner types: State, foreign, nonfinancial (industrial), individual (private) and financial owners. Finally, in the last line we show the fraction owned by the subgroup of financial owners which are mutual funds (Note that these mutual funds are contained in the total holdings of financials in the line above.) Data from 1989–2007. (Annual 1989–1992, monthly 1993-2007.)

	1989–2007			1989–1994			1995–1999			2000–2007		
	average		med	average		med	average		med	average		med
	vw	ew		vw	ew		vw	ew		vw	ew	
Largest owner	37.2	27.5	21.1	28.4	26.2	20.8	29.4	27.0	21.0	44.8	28.2	21.3
Three largest	50.9	44.1	41.9	45.1	43.4	38.5	44.8	43.4	41.8	56.6	44.7	43.4
Herfindahl Index	0.22	0.15	0.08	0.15	0.14	0.08	0.15	0.15	0.08	0.29	0.16	0.09
Herfindahl (all but three largest)	0.03	0.04	0.03	0.03	0.04	0.03	0.03	0.04	0.03	0.02	0.04	0.02
Total no owners(>100 shares)	13965	2330	851	7861	1853	654	7511	1847	814	19902	2781	967
Fraction State Owners	27.0	6.2	0.5	21.2	6.5	1.0	19.6	6.3	0.4	33.4	6.0	0.4
Fraction Foreign Owners	31.6	22.6	12.6	29.3	20.5	13.3	33.4	22.5	13.7	31.2	23.4	11.2
Fraction Nonfinancial Owners	19.1	35.1	28.9	25.6	41.0	40.8	20.9	33.6	28.8	16.0	34.2	28.0
Fraction Individual Owners	7.5	19.7	13.3	10.9	18.3	12.4	8.8	20.0	13.0	5.7	19.9	13.7
Fraction Financial Owners	16.8	18.7	16.6	18.5	20.6	18.1	20.5	21.0	19.4	13.9	16.8	14.2
Fraction Mutual Fund Owners	5.5	6.8	4.9	4.5	5.8	5.2	6.6	7.2	6.1	5.0	6.8	4.4

II Predicting US economic growth with market illiquidity

A In-sample evidence

We start by assessing the in-sample predictive ability of market illiquidity. The models we examine are predictive regressions on the form:

$$y_{t+1} = \alpha + \beta LIQ_t + \gamma' \mathbf{X}_t + u_{t+1}, \quad (3)$$

where y_{t+1} is the realized growth in the macro variable over quarter $t+1$, LIQ_t is the market illiquidity measured for quarter t, and $\mathbf{X}_t$ contains the additional control variables

[13]More details about this data can be found in e.g. Bøhren and Ødegaard (2001), Bøhren and Ødegaard (2006) and Ødegaard (2009).

(*Term*, *dCred*, *Vola*, er_m and the lag of the dependent variable) observed at t, and γ' is the vector of coefficient estimates for the control variables. We use three different proxies for equity market illiquidity; *ILR*, *LOT* and *Roll*. Our main dependent variable (y_{t+1}) is real GDP growth. However, we also examine three additional macro variables related to economic growth; growth in the unemployment rate (*dUE*), real consumption growth (*dCONSR*) and real growth in private investments (*dINV*).

Table IV summarizes the results from the various regression specifications. The first specification only includes the liquidity variable and one lag of the dependent variable.[14] We see that the coefficient on market illiquidity ($\hat{\beta}$) is highly significant for most models regardless of which illiquidity proxy we use. An increase in market illiquidity predicts lower real GDP growth (*dGDPR*), an increase in unemployment (*dUE*) and a slowdown in consumption (*dCONSR*) and investment (*dINV*).

To give some more information about the significance of the liquidity variable we report the $\bar{R}^2$ for models estimated with and without liquidity in the columns on the right of the table. So for example, adding liquidity to the regression forecasting *dGDPR* improves the $\bar{R}^2$ from 3% to 13%.

It is at this point useful to interpret the coefficients to get at the magnitude of the estimated effects. Starting with the regression predicting changes in GDP as a function of changes in *ILR*, we ask how much does growth change? Let us look at a one standard deviation change in *dILR*. The standard deviation of *dILR* is 0.26. Multiplying this with the estimated coefficient for *dILR* of -0.013, we would predict a change in *dGDPR* of -0.003, i.e. a fall in quarterly GDP growth of 0.3%, for a one standard deviation increase in *ILR*. During this period, average GDP growth was 0.8% per quarter. The predicted change in GDP is thus about a third of average quarterly growth. Doing a similar exercise for the *LOT* variable, the model predicts a change in GDP growth of -0.2% (-0.00192) for a one standard deviation increase in *LOT*. Similarly, a one standard deviation increase in *Roll* predicts a change in GDP growth of -0.8% (-0.00796).

In sum the results indicate that market illiquidity contains economically significant information about future economic growth. When market liquidity worsens, this is followed by a significant slowdown in economic growth.

Several other financial variables have been found to contain information about future macroeconomic conditions. We therefore also consider regression specifications where we control for these variables. Table II shows that our liquidity proxies are correlated with the term spread, the credit spread as well as the market return and volatility. This is what we would expect, since one hypothesis is that variations in market liquidity captures changes in expectations about future growth which should also be reflected in other financial variables. The main purpose of adding other financial control variables to the models is to determine whether liquidity provide an additional (or less noisy) signal about future macro fundamentals. We start by including two non-equity control variables (in addition to the lag of the dependent variable). The control variables we include are the term spread (*Term*) and credit spread (*Cred*). Harvey (1988) shows that (*Term*) is a strong predictor of consumption growth and a superior predictor of growth in *GNP* relative to stock returns (Harvey, 1989). With respect to *Cred*, Gilchrist, Yankov, and Zakrajsek (2009) show that credit spreads contain substantial predictive power for

[14] We have also estimated the models with different lag specifications with up to four lags of the dependent variable and the liquidity variables. This does not materially affect the results.

Table IV: In-sample prediction of macro variables

The table shows the results from predictive regressions where we regress next-quarters growth in different macro variables on three proxies for market illiquidity for the period 1947-2008. Market illiquidity (LIQ) is proxied by one of three illiquidity measures: the Amihud Illiquidity ratio (*ILR*), the *LOT* measure and the Roll measure (*Roll*). We use the log difference in *ILR* and *LOT* to preserve stationarity, while the Roll measure is not differenced. The crossectional liquidity measures are calculated as equally weighted averages across stocks. The model estimated is $y_{t+1} = \alpha + \beta^{LIQ} LIQ_t + \gamma' \mathbf{X}_t + u_{t+1}$ where y_{t+1} is one of real GDP growth (*dGDPR*), growth in the unemployment rate (*dUE*), real consumption growth (*dCONSR*) or growth in private investments (*dINV*). We also include one lag of the dependent variable (y_t) and *Term*, *dCred*, *Vola* and er_m as control variables. The Newey-West corrected t-statistics (with 4 lags) are reported in parentheses below the coefficient estimates, and $\bar{R}^2$ is the adjusted R^2. The column on the far right, labeled "ex.liq. $\bar{R}^2$," gives the adjusted R^2 for a model *without* the liquidity variable.

Panel A: *ILR* liquidity measure

Dependent variable (y_{t+1})	$\hat{\alpha}$	$\hat{\beta}^{LIQ}$	$\hat{\gamma}^{y}$	$\hat{\gamma}^{Term}$	$\hat{\gamma}^{dCred}$	$\hat{\gamma}^{Vola}$	$\hat{\gamma}^{er_m}$	$\bar{R}^2$	ex.liq. $\bar{R}^2$
dGDPR	0.006	-0.013	0.224					0.13	0.03
	(7.58)	(-5.38)	(3.68)						
dUE	0.003	0.074	0.300					0.13	0.07
	(0.61)	(3.68)	(5.14)						
dCONSR	0.006	-0.006	0.305					0.11	0.08
	(7.08)	(-3.33)	(4.46)						
dINV	0.006	-0.034	0.265					0.15	0.06
	(2.95)	(-6.19)	(3.97)						
dGDPR	0.006	-0.011	0.207	0.001	-0.012			0.17	0.10
	(5.14)	(-4.59)	(3.48)	(0.95)	(-2.91)				
dUE	0.014	0.055	0.298	-0.009	0.089			0.18	0.15
	(1.88)	(3.10)	(5.09)	(-2.61)	(3.01)				
dCONSR	0.004	-0.005	0.303	0.001	-0.003			0.13	0.12
	(3.81)	(-2.79)	(4.41)	(2.23)	(-0.94)				
dINV	0.002	-0.027	0.239	0.004	-0.035			0.23	0.17
	(0.57)	(-5.27)	(3.79)	(2.41)	(-3.93)				
dGDPR	0.006	-0.008	0.196	0.000	-0.012	0.000	0.015	0.17	0.15
	(5.82)	(-3.87)	(3.38)	(0.72)	(-2.99)	(0.07)	(1.95)		
dUE	0.005	0.021	0.302	-0.007	0.097	-0.033	-0.228	0.22	0.22
	(0.75)	(1.17)	(6.05)	(-2.44)	(3.16)	(-0.93)	(-4.54)		
dCONSR	0.005	-0.001	0.301	0.001	-0.003	0.002	0.026	0.17	0.18
	(4.65)	(-0.35)	(4.36)	(2.19)	(-1.21)	(0.39)	(3.38)		
dINV	0.003	-0.020	0.236	0.003	-0.037	0.007	0.045	0.24	0.22
	(1.21)	(-3.81)	(3.70)	(2.37)	(-3.87)	(0.50)	(2.02)		

Panel B: *LOT* liquidity measure

Dependent variable (y_{t+1})	$\hat{\alpha}$	$\hat{\beta}^{LIQ}$	$\hat{\gamma}^{y}$	$\hat{\gamma}^{Term}$	$\hat{\gamma}^{dCred}$	$\hat{\gamma}^{Vola}$	$\hat{\gamma}^{er_m}$	$\bar{R}^2$	ex.liq. $\bar{R}^2$
dGDPR	0.007	-0.017	0.168					0.06	0.03
	(7.52)	(-2.78)	(2.59)						
dUE	0.003	0.129	0.261					0.10	0.07
	(0.47)	(3.14)	(4.42)						
dCONSR	0.006	-0.009	0.282					0.09	0.08
	(7.04)	(-1.74)	(3.86)						
dINV	0.007	-0.039	0.218					0.07	0.06
	(3.04)	(-2.56)	(3.21)						
dGDPR	0.006	-0.012	0.160	0.001	-0.014			0.11	0.10
	(5.20)	(-2.11)	(2.52)	(1.06)	(-3.48)				
dUE	0.014	0.088	0.269	-0.009	0.098			0.16	0.15
	(1.76)	(2.53)	(4.58)	(-2.73)	(3.26)				
dCONSR	0.004	-0.006	0.285	0.001	-0.004			0.12	0.12
	(3.94)	(-1.30)	(3.95)	(2.32)	(-1.21)				
dINV	0.002	-0.021	0.200	0.004	-0.043			0.18	0.17
	(0.71)	(-1.61)	(3.17)	(2.61)	(-4.60)				
dGDPR	0.007	-0.012	0.155	0.000	-0.014	0.006	0.028	0.16	0.15
	(6.29)	(-2.13)	(2.64)	(0.60)	(-3.48)	(1.03)	(3.63)		
dUE	0.004	0.110	0.285	-0.007	0.098	-0.085	-0.261	0.23	0.22
	(0.61)	(2.73)	(5.83)	(-2.32)	(3.17)	(-2.02)	(-5.44)		
dCONSR	0.005	-0.006	0.290	0.001	-0.003	0.005	0.027	0.18	0.18
	(4.94)	(-1.18)	(4.26)	(2.16)	(-1.22)	(0.90)	(4.41)		
dINV	0.005	-0.024	0.207	0.003	-0.041	0.017	0.075	0.22	0.22
	(1.67)	(-1.80)	(3.21)	(2.33)	(-4.38)	(1.14)	(3.85)		

Table IV: (Continued)

Panel C: *Roll* liquidity measure

Dependent variable (y_{t+1})	$\hat{\alpha}$	$\hat{\beta}^{LIQ}$	$\hat{\gamma}^{y}$	$\hat{\gamma}^{Term}$	$\hat{\gamma}^{dCred}$	$\hat{\gamma}^{Vola}$	$\hat{\gamma}^{er_m}$	$\bar{R}^2$	ex.liq. $\bar{R}^2$
dGDPR	0.019	-0.811	0.136					0.10	0.03
	(5.94)	(-4.11)	(2.16)						
dUE	-0.074	5.206	0.236					0.12	0.07
	(-3.07)	(3.29)	(4.23)						
dCONSR	0.013	-0.436	0.269					0.11	0.08
	(4.23)	(-2.28)	(3.47)						
dINV	0.039	-2.192	0.188					0.13	0.06
	(4.26)	(-3.61)	(3.08)						
dGDPR	0.016	-0.716	0.133	0.001	-0.012			0.157	0.104
	(5.29)	(-3.79)	(2.15)	(1.91)	(-2.84)				
dUE	-0.051	4.639	0.248	-0.011	0.083			0.189	0.153
	(-2.23)	(3.15)	(4.64)	(-3.62)	(2.60)				
dCONSR	0.011	-0.465	0.268	0.001	-0.002			0.150	0.121
	(3.98)	(-2.54)	(3.47)	(3.00)	(-0.68)				
dINV	0.030	-2.007	0.177	0.005	-0.034			0.248	0.187
	(3.85)	(-3.80)	(3.25)	(3.56)	(-3.89)				
dGDPR	0.016	-0.614	0.135	0.001	-0.013	0.006	0.021	0.18	0.15
	(4.78)	(-3.03)	(2.30)	(1.39)	(-3.07)	(1.12)	(2.74)		
dUE	-0.044	3.559	0.270	-0.009	0.091	-0.065	-0.219	0.23	0.22
	(-1.80)	(2.25)	(5.90)	(-3.05)	(2.84)	(-1.71)	(-4.74)		
dCONSR	0.010	-0.318	0.282	0.001	-0.002	0.004	0.023	0.19	0.18
	(3.63)	(-1.76)	(4.03)	(2.71)	(-0.93)	(0.92)	(3.66)		
dINV	0.030	-1.895	0.179	0.005	-0.037	0.028	0.055	0.28	0.23
	(3.81)	(-3.43)	(3.17)	(3.20)	(-4.11)	(2.34)	(2.84)		

economic activity.

These regression specifications are also listed in table IV. Looking first at the estimation results for GDP growth, we see that while *dCred* enters significantly in all three models, the coefficients on liquidity retains their level, sign and significance. Interestingly, the coefficient on the term spread ($\hat{\gamma}^{Term}$) is not significant in the models that include *dILR* or *dLOT*. In unreported specifications we find that excluding the liquidity variables in these models restores the significance of *Term*. The results for the other macro variables yield the same results. The coefficients on liquidity are robust to the inclusion of the term spread and credit spread in the models. However, the results suggest that both the term spread and credit spread are important predictor variables, and a model that contains the two bond market variables in addition to liquidity has higher adjusted R-squared compared to the model just containing liquidity and the lag of the dependent variables.

As a final exercise, we include the equity market variables excess market return (er_m) and volatility (*Vola*) into the models in addition to the term spread and credit spread. In the models for GDP growth, we find that while market volatility is insignificant, market return enters significantly with a positive coefficient. However, this does not affect the significance of any of the liquidity coefficients. Thus, market liquidity retains its predictive power for real GDP growth. In the models for the unemployment rate, the results are more mixed. In the model with *dILR*, we see that adding market return renders the *dILR* coefficient insignificant. However, in the models with *Roll* and *dLOT*, the coefficients are unaffected. In the models for real consumption growth, we see that market liquidity (regardless of liquidity measure) is rendered insignificant when the excess return on the market is included in the model. Finally, in the models for investment growth, the liquidity coefficients are unaffected by the inclusion of market return.

Overall, the results show that while other financial variables are clearly useful for predicting future economic growth, we find that there is additional information in market illiquidity, even after controlling for well known alternative variables. Market liquidity seems to be a particularly strong and robust predictor of real GDP growth, unemployment and investment growth. For future real consumption growth, however, there does not seem to be additional information in liquidity that is not already reflected in the term spread and market return.

A.1 Causality

We are primarily interested in predicting macroeconomic conditions with liquidity, but there is also the possibility of causality going in the opposite direction, i.e. that changes in economic conditions affect market illiquidity. We know from earlier studies that monetary policy shocks have an effect on stock and bond market illiquidity (see e.g. Söderberg (2008) and Goyenko and Ukhov (2009)), while there is no effect of shocks to real economic variables on stock market illiquidity. On the other hand, neither of these studies considers the reverse causality from market liquidity to real economic variables. We therefore look directly at this issue by performing Granger causality tests. We return to the specification with only liquidity and real variables and perform Granger causality tests between the different illiquidity proxies and real GDP growth.[15] Table V reports the results from these tests. The tests are done in a Vector Auto Regression (VAR) framework. We perform the tests for the whole sample and for different sub-samples. We both split the sample period in the middle and into five 20 year sub-periods (overlapping by 10 years). The first row of Table V shows the number of quarterly observations in each sample period, and the second row shows the number of NBER recessions that occurred within each sample period. In part (a) of the table we run Granger causality tests between *dILR* and *dGDPR*. Looking first at the column labeled "Whole sample," we see that the null hypothesis that GDP growth *does not* Granger cause *dILR* ($dGDPR \nrightarrow dILR$) cannot be rejected, while the hypothesis that *dILR does not* Granger cause GDP growth ($dILR \nrightarrow dGDPR$) is rejected at the 1% level. For the different sub-periods we see that the relation is remarkably stable. Thus, part (a) of the table shows a strong and stable one way Granger causality from market illiquidity, proxied by *dILR*, to *dGDPR*, while there is no evidence of a reverse causality from *dGDPR* to *dILR*. In parts (b) and (c) of the table, we perform the same tests for the *dLOT* and the *Roll* measures. For the full sample period, we find support for a Granger causality from *dLOT* and *Roll* to GDP growth, while there is no evidence of a reverse causality. Also for the sub-periods, we find support for a one-way Granger causality from the *Roll* measure to *dGDPR*, except for the first 20-year period where we are only able to reject the null that the Roll measure does not Granger cause real GDP growth at a 10% significance level. Based on the sub-sample results for the *dLOT* measure we cannot reject the null that *dLOT does not* Granger

[15] Results from a much more comprehensive VAR specification are reported and discussed in an internet appendix. There we also examine the dynamic linkages between the other financial variables and liquidity as well as testing for Granger causality between all the variables used in the analysis. Furthermore, we analyze the robustness of the response function of *dGDPR* to a shock in *dILR* for different orderings of the endogenous variables.

cause $dGDPR$ in the second half of the sample. One potential reason why the LOT measure has become less informative over the sample period is the increase in trading activity. Recall that the LOT measure uses zero return days to identify the implicit transaction cost for a stock. Thus, if the number of zero return days has decreased at the same time as the trading activity has increased, the LOT measure may have become a more noisy estimator of actual transaction costs in the last part of the sample.

Table V: Granger causality tests

The table shows Granger causality tests between quarterly real GDP growth ($dGDPR$) and the (a) Amihud Illiquidity ratio (ILR), (b) the LOT measure and (c) the Roll measure. The crossectional liquidity measures are calculated as equally weighted averages across stocks. The test is performed for the whole sample period and different subperiods. For each measure we first test the null hypothesis that real GDP growth *does not* Granger cause market illiquidity and then whether market illiquidity *does not* Granger cause real GDP growth. We report the χ^2 and p-value (in parenthesis) for each test. We choose the optimal lag length for each test based on the Schwartz criterion. For each illiquidity variable the test is performed on the whole sample period (1947q1-2008q4), the first (1947q1-1977q4) and second half (1978q1-2008q4) of the sample, and for rolling 20-year subperiods overlapping by 10 years. The first two rows report the number of quarterly observations covered by each sample period and the number of NBER recession periods within each sample. ** and * denotes a rejection of the null hypothesis at the 1% and 5% level, respectively.

	Whole sample	First half	Second half	20-year subperiods				
	1947-2008	1947-1977	1977-2008	1950-1970	1960-1980	1970-1990	1980-2000	1990-2008
N (observations)	243	119	124	84	84	84	84	76
NBER recessions	11	6	5	5	4	4	2	3
(a) *ILR*								
H_0: $dGDPR \nrightarrow dILR$								
χ^2	4.08	1.66	3.13	3.66	3.56	3.35	2.83	2.66
p-value	0.13	0.44	0.21	0.16	0.17	0.19	0.24	0.26
H_0: $dILR \nrightarrow dGDPR$								
χ^2	31.97**	19.01**	14.50**	15.81**	8.89**	11.7**	11.64**	11.85**
p-value	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.00
(b) *LOT*								
H_0: $dGDPR \nrightarrow dLOT$								
χ^2	2.21	1.77	1.13	2.20	1.48	1.21	0.06	1.05
p-value	0.14	0.18	0.29	0.14	0.22	0.27	0.80	0.31
H_0: $dLOT \nrightarrow dGDPR$								
χ^2	9.55**	13.37**	1.45	8.24**	7.7**	6.81**	1.22	0.99
p-value	0.00	0.00	0.23	0.00	0.01	0.01	0.27	0.32
(c) *Roll*								
H_0: $dGDPR \nrightarrow Roll$								
χ^2	0.086	0.305	0.745	0.270	0.012	2.300	1.332	0.014
p-value	0.77	0.58	0.39	0.60	0.91	0.13	0.25	0.91
H_0: $Roll \nrightarrow dGDPR$								
χ^2	15.96**	5.56*	10.80**	2.95	10.74**	9.31**	4.43*	10.18**
p-value	0.00	0.02	0.00	0.09	0.00	0.00	0.04	0.00

A.2 Market liquidity and NBER recessions

The in-sample results on the predictive content of liquidity for macro variables can be visualized by a form of "event study." We use the onset of a recession as the "event date," and show the evolution of the various series of interest around this date in a plot. In panel A of figure 2 we plot changes in liquidity relative to the onset of a recession, as defined by the NBER. For each NBER recession, we first calculate the quarterly GDP growth starting 5 quarters before ($\mathsf{t} = -5\mathsf{Q}$) the first NBER recession quarter (NBER1) and ending 5 quarters after the end of each NBER recession ($\mathsf{t} = 5\mathsf{Q}$). Next, we average the GDP growth for each quarter across all recessions, and then accumulate the average GDP growth over the event window. Then we do the same for the *ILR* measure. Thus, the figure shows the average pattern in *ILR* before, during and after US recessions averaged across all the 10 NBER recessions (shaded area) in our sample from 1947-2008.[16] This style of analysis also lets us give some informative comparisons of the informational content of the different predictive variables. Panel B of figure 2 shows similar plots, where we also add the financial control variables term spread, credit spread, excess market return and volatility. Looking first at the term spread (dotted line) in picture (a), we see that there is a systematic decline in the term spread in all the quarters prior to the first NBER recession quarter (NBER1). This is consistent with the notion that the yield curve has a tendency to flatten and invert before recessions. We also see that the term spread increases again already during the first quarters of the recession, predicting the end of the recession and increased growth. Thus, before the recession, the signal from both the term spread and market liquidity (solid line) seems to capture similar information about GDP growth. For the credit spread in picture (b), both market liquidity and the credit spread seems to share a very similar path, although the liquidity series changes earlier than the credit spread. As we will see later in the out-of-sample analysis, the credit spread and market liquidity have very similar out-of-sample performance when predicting GDP growth. In picture (c) we see that the accumulated excess market return is relatively stable until the quarter just before the NBER recession starts. Thus, it seems to be responding later than the other variables. Finally, in picture (d), we see that volatility increases in the quarter just before the NBER recessions starts. However, consistent with the regression results, the information in market volatility seems small compared to the other variables.

B Out-of-sample evidence for the US

In the previous section, we found that market illiquidity had predictive power for economic growth for the whole sample period, for subperiods, and when controlling for other financial variables that are found in the literature to be informative about future economic growth. However, in-sample predictability does not necessarily mean that the predictor is a useful predictor out of sample. In this section, we therefore examine whether market illiquidity is able to forecast quarterly real GDP growth out-of-sample.

[16]Note that some NBER recessions only last for 3 quarters (e.g. 1980Q1-1980Q3), while there are some recessions that last up to 6 quarters (e.g. 1973Q4-1975Q1 and 1981Q3-1982Q4). However, the most important point of the figure is that all NBER recessions are aligned to start at the same point (NBER1) in event time.

Figure 2: Market illiquidity around NBER recessions

The figure in panel A shows the accumulated quarterly growth in *ILR* (solid line) and accumulated quarterly GDP growth (bars) averaged in event time across different NBER recession periods. All recession periods are aligned to start at NBER1, the first NBER recession quarter. The figure shows the results when looking at all 10 NBER recessions during the full sample period 1947-2008. In Panel B we show similar figures, adding evolutions of the cumulative average changes in (a) term spread, (b) credit spread, (c) excess market return and (d) volatility.

Panel A: Liquidity evolution approaching recessions.

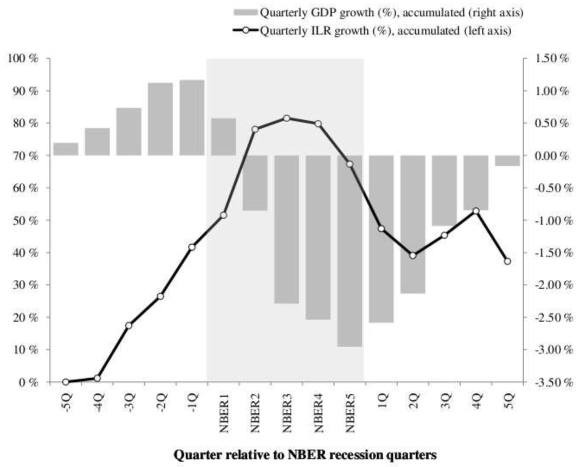

Panel B: Comparing to other financial variables.

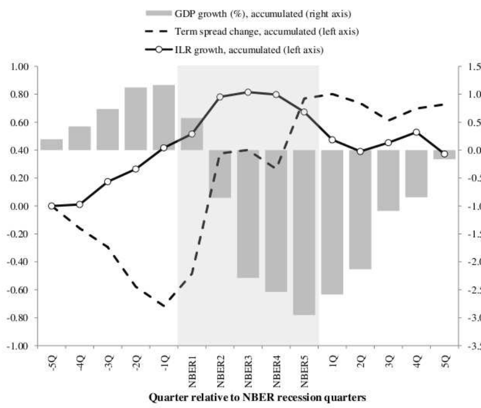

(a) Term spread (*Term*)

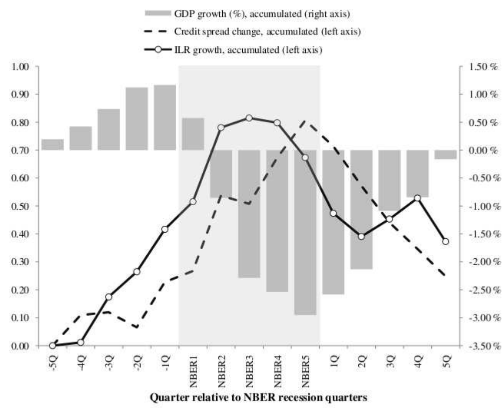

(b) Credit spread (*Cred*)

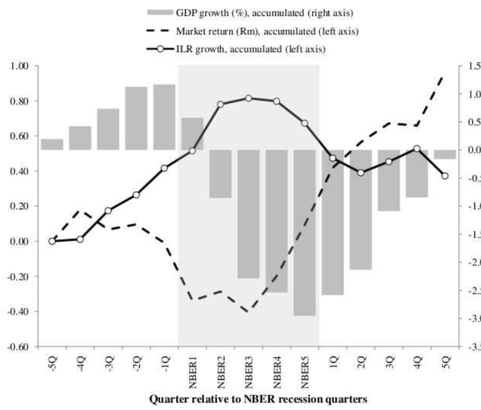

(c) Market return (er_m)

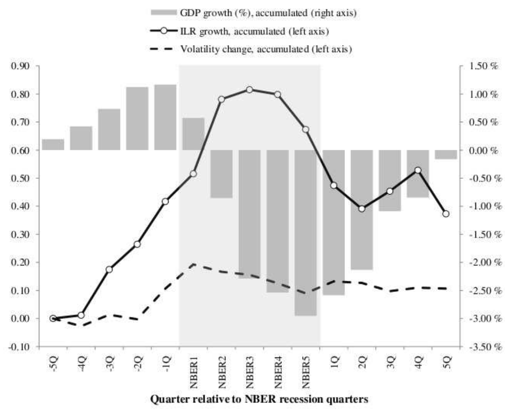

(d) Volatility (*Vola*)

B.1 Methodology and timing of information

When setting up our out-of-sample procedure, we need to be careful about the timing of the data so it reflects what would have been available to the forecaster when a forecast is made. While the illiquidity variables and the other financial variables are observable in real-time without revisions, real GDP growth is not. First, there is a publication lag of one quarter for GDP.[17] Secondly, there is an issue of later revisions in most macro variables. While the publication lag is easily accounted for, the revisions are more tricky. Basically, the question is whether we want to forecast the first or final vintage of GDP growth. This depends on the question we are asking. If we were using macro variables to predict financial variables (e.g. returns), we would want to use the first vintage (real time version) of the macro variable since the later vintages (revised figures) would not be known to the forecaster (investor) when making his forecast. However, since the question we are asking is whether financial variables contain information about expected economic growth, we want to forecast the last vintage. The argument for this is that since the revisions are mainly due to measurement errors in the first/early vintage series, market participants' expectations about underlying economic growth should be unrelated to ("see through") measurement errors in the first vintages. Thus, we want to forecast the most precisely measured version of the macro variable, i.e. the last vintage series.

In our out-of-sample analysis we consider a rolling estimation scheme with a fixed width of 20 quarters (5 years). For all models, our first out-of-sample forecast is made at the end of the first quarter of 1952 for GDP growth for the second quarter of 1952. At this point we estimate each model using data from the first quarter of 1947 through the fourth quarter of 1951 (which is then the most recent GDP observation available to us). We then produce a forecast of real GDP growth for the second quarter of 1952 based on the estimated model coefficients and the most recent observation of the predictor variable. In the case the predictor variables is market liquidity or any of the other financial variables, these are observed for the same quarter as we construct our forecast for next quarter, i.e. first quarter of 1952. Next, we move the window forward by one quarter, re-estimate the models, and produce a new forecast for the next quarter, and so on. The last forecast is made at the fourth quarter of 2008 for GDP growth for the first quarter of 2009.

We compare the performance of a model with market liquidity as the predictor against models with other financial variables both individually, as well as looking at the contribution of adding liquidity to a benchmark model that contains all the financial market variables used in the previous analysis. We also compare the illiquidity model against a benchmark model where we forecast GDP growth using an autoregressive model. In that case, the most recent observation of GDP available to us at the end of the first quarter of 1952, when we produce the first forecast of GDP growth for the second quarter of 1952, is GDP for the fourth quarter of 1951. Thus, we estimate the autoregressive model for GDP growth with data including the fourth quarter of 1951 and construct a forecast for the second quarter of 1952 based on the estimated coefficients and the most recent GDP observation available, which is the final figure for GDP growth for the fourth quarter of 1951.

[17]The Bureau of Economic Analysis releases the *final* GDP figure for quarter $t-1$ in the last month of the following quarter (t). However, they also release an "advance" estimate in the first month of the following quarter as well as a "preliminary" release in the second month of the following quarter. Thus, at the end of t, a forecaster has the "final" number available for $t-1$ GDP growth.

B.2 Out-of-sample performance of different liquidity measures

We begin by evaluating univariate forecast models for real GDP growth using the three different liquidity proxies. The models are evaluated by comparing the mean squared forecast error (MSE) from the series of one-quarter ahead forecasts. Since we compare models for the same dependent variable, but with different predictor variables, the models are non-nested. We use two statistics to compare the out-of-sample performance of the different liquidity measures; the mean-squared forecasting error (MSE) ratio and the modified Diebold-Mariano (MDM) encompassing test proposed by Harvey, Leybourne, and Newbold (1998), which has greater power than the original Diebold and Mariano (1995) test, especially in small samples. In addition, Harvey et al. (1998) advocate comparison of the MDM statistic with critical values from the Student's t distribution, instead of the standard normal distribution.

The Diebold and Mariano (1995) statistic (hereafter DM) is calculated in the following way: Suppose we have a candidate predictor i and a competing predictor k. We want to test the null hypothesis of equal predictive accuracy that $E[\bar{d}] = 0 \; \forall \; t$, where $\bar{d} = P^{-1} \cdot \sum_t (\varepsilon^2_{k,t+1} - \varepsilon^2_{i,t+1})$, P is the number of rolling out-of-sample forecasts, and $\varepsilon^2_{k,t+1}$ and $\varepsilon^2_{i,t+1}$ are the squared forecast errors from the two models. The DM statistic is calculated as:

$$DM = \frac{\bar{d}}{(\sigma^2_{\bar{d}}/P)^{1/2}}, \tag{4}$$

and the modified DM statistic is calculated as:

$$MDM = \left[\frac{P + 1 - 2h + P^{-1}h(h-1)}{P}\right]^{1/2} DM, \tag{5}$$

where DM is the original statistic, P is the number of out-of-sample forecasts and h is the forecast horizon (overlap). The MDM statistic is compared with critical values from the Student's t distribution with $(P-1)$ degrees of freedom.

Panel A in Table VI shows the results when we compare different forecasting models for quarterly GDP growth using different proxies for market liquidity. The liquidity variables labeled in the first row (under Model 1) constitute the respective candidate variable (i), and the liquidity variables labeled in the first column (under Model 2) are the competing variables (k). For example, the first pair of numbers compares the MSE from a model (Model 1) that uses *dILR* as predictor variable against a model (Model 2) that uses *dLOT* as the predictor variable. The first number shows the relative MSE between the two models, which is 0.89. This means that the model with *dILR* as a predictor variable has a lower MSE than the model that uses *dLOT*. The second number shows the modified Diebold/Mariano statistic (MDM) which provides a statistic to test for whether the MSE of model 1 is significantly different from that of Model 2. The last row in the table shows the MSE for each model specification labeled under Model 1. Looking first at the last row, we see that the model with *dILR* has the lowest MSE across the models. Also, when comparing the forecast performance of the different models against each other we see that the model with *dILR* in all cases has a significantly lower MSE compared to models with *dLOT* and *Roll* as predictor variables. The model with *dLOT* as the predictor variable has a lower MSE than the *Roll* model. The MDM statistic cannot however reject the null that the MSE of the *dLOT* model is not significantly different from the MSE of the *Roll* model.

Overall, the results in panel A of table VI show that *dILR* has the lowest forecast error for GDP growth among the three liquidity proxies we examine. This is consistent with the in-sample results where *dILR* was the strongest and most robust predictor of GDP growth. In the rest of the out-of-sample analysis we therefore use the *dILR* as our liquidity predictor variable.

B.3 Out-of-sample performance of illiquidity versus other variables

We next want to evaluate the out-of-sample predictive ability of *dILR* against different baseline models. We assess the out-of-sample performance of *dILR* against two types of baseline models. The first set of baseline models are models where GDP growth is regressed on *one* of the financial control variables (*Term*, *dCred*, *Vola*, er_m) that we used in the in-sample analysis. Each of these models is then a restricted (nested) version of a larger model where GDP growth is regressed on the control variable *in addition* to *dILR*. We also look at the performance of a more comprehensive restricted model for *dGDPR* containing all the financial control variables, which we compare to an unrestricted model where we add *dILR*. The second type baseline model that we compare *dILR* to is an autoregressive model for GDP growth. In that case, the autoregressive GDP model is the restricted version of a model where we include both lagged GDP growth and *dILR* as predictor variables for next quarter GDP growth. We also compare the models with the other financial variables to the restricted autoregressive model for GDP growth.

We evaluate forecast performance using two test statistics. The first test is an encompassing test (ENC-NEW) proposed by Clark and McCracken (2001). The ENC-NEW test asks whether the restricted model (the model that does not include *dILR*), encompasses the unrestricted model that includes *dILR*. If the restricted model *does not* encompass the unrestricted model, that means that the additional predictor (*dILR*) in the larger, unrestricted, model improves forecast accuracy relative to the baseline. Clark and McCracken (2001) shows that the ENC-NEW test has greater power than tests for equality of MSE. The ENC-NEW test statistic is given as

$$\text{ENC-NEW} = (P - h + 1) \cdot \frac{P^{-1} \sum_t \left[\varepsilon_{r,t+1}^2 - \varepsilon_{r,t+1} \cdot \varepsilon_{u,t+1}\right]}{MSE_u}, \tag{6}$$

where P is the number of out-of-sample forecasts, $\varepsilon_{r,t+1}$ denotes the rolling out-of-sample errors from the restricted (baseline) model that excludes *dILR*, $\varepsilon_{u,t+1}$ is the rolling out-of-sample forecast errors from the unrestricted model that includes *dILR*, and MSE_u denotes the mean squared error of the unrestricted model that includes *dILR*.

The second test statistic we examine is an F-type test for equal MSE between two nested models proposed by McCracken (2007), termed MSE-F. This test is given by

$$\text{MSE-F} = (P - h + 1) \cdot \frac{MSE_r - MSE_u}{MSE_u}, \tag{7}$$

where MSE_r is the mean squared forecast error from the restricted model that excludes *dILR*, and MSE_u is the mean squared forecast error of the unrestricted model that includes *dILR*. Both the ENC-NEW and MSE-F statistics are nonstandard and we use the bootstrapped critical values provided by Clark and McCracken (2001).[18]

[18]The bootstrapped critical values are available at
`http://www.kansascityfed.org\Econres\addfiles\criticalvalues_tec.xls`

Table VI: Results of out-of-sample tests

Panel A reports the results of one-quarter ahead, non-nested, forecast comparisons of models with different liquidity proxies. The variable being forecast is quarterly GDP growth (*dGDPR*). Each pair of numbers compares two alternative univariate forecast models (which includes a constant term). The table compares the out-of-sample MSE of a model that uses one of the liquidity variables labeled under Model 1 as a predictor, with a model that uses one of the variables labeled in the first column under Model 2. For each model pair, the table shows the relative MSE between model 1 and model 2, and the modified Diebold/Mariano test statistic (labeled MDM). The null hypothesis for the MDM test is that the MSE of Model 2 and Model 1 are equal against the alternative that the MSE for model 1 is less than that of model 2. An MDM statistic with ** or * denotes a rejection of the null hypothesis of equal forecast accuracy at the 1% and 5% level, respectively. Panel B reports the results from nested model comparisons for predicting quarterly real GDP growth out-of-sample one quarter and two quarters ahead. The first column shows which variables are included in the unrestricted model, and the second column shows which variables that are included in the restricted (baseline) model. Columns 3 to 5 show the relative MSE, the MSE-F test for equality of MSE and the ENC-NEW test for the one quarter ahead forecast. Columns 6 to 8 show the test statistics for the two-quarter-ahead forecasts. The last row in panel B report the out-of-sample results when we compare a restricted model containing all the financial variables with an unrestricted model where we add *dILR* to the model. ** and * denotes a rejection of the null hypothesis (at the 1% and 5% level, respectively) of equal forecast precision for the MSE-F test, while it denotes a rejection of the null that the restricted model encompasses the unrestricted model for the ENC-NEW test. Panel C shows the model comparison results when the baseline model is an autoregressive model (of order 1) for GDP growth. In that case the unrestricted model adds *dILR* and each of the other financial variables to the restricted model.

Panel A: Choosing liquidity variable: Predicting GDP growth with different liquidity proxies

		Model 1		
Model 2	Statistic	*dILR*	*dLOT*	*Roll*
dLOT	MSE_1/MSE_2	0.89	-	
	MDM	1.74*	-	
Roll	MSE_1/MSE_2	0.82	0.91	-
	MDM	1.89*	0.47	-
	MSE (x10^3)	0.088	0.099	0.108

Panel B: Forecasting real GDP growth: Illiquidity (*dILR*) versus other financial variables

		1 quarter-ahead forecasts			2 quarters-ahead forecasts		
Unrestricted model	Restricted model	$\frac{MSE_u}{MSE_r}$	MSE-F	ENC-NEW	$\frac{MSE_u}{MSE_r}$	MSE-F	ENC-NEW
dILR, *Term*	*Term*	0.917	20.95**	41.96**	0.927	18.09**	31.49**
dILR, er_m	er_m	0.976	5.69**	14.39**	1.003	-0.59	12.33**
dILR, *dCred*	*dCred*	1.000	0.02	18.73**	0.964	8.53**	22.86**
dILR, *Vola*	*Vola*	0.889	28.76**	50.91**	0.895	26.88**	35.98**
dILR, *Term*, er_m, *dCred*, *Vola*	*Term*, er_m, *dCred*, *Vola*	1.016	-3.58	7.27**	1.030	-6.79	10.35**

Panel C: Forecasting real GDP growth: Financial variables versus an autoregressive model

		1 quarter-ahead forecasts			2 quarters-ahead forecasts		
Unrestricted model	Restricted model	$\frac{MSE_u}{MSE_r}$	MSE-F	ENC-NEW	$\frac{MSE_u}{MSE_r}$	MSE-F	ENC-NEW
dILR, *dGDPR*	*dGDPR*	0.849	41.16**	60.17**	0.803	56.36**	40.60**
Term, *dGDPR*	*dGDPR*	0.988	2.91	34.75**	0.866	35.44**	28.99**
er_m, *dGDPR*	*dGDPR*	0.905	24.20**	45.54**	0.850	40.66**	30.91**
dCred, *dGDPR*	*dGDPR*	0.838	44.63**	51.37**	0.850	40.54**	28.77**
Vola, *dGDPR*	*dGDPR*	1.109	-22.77	9.92*	1.049	-10.81	1.26

Panel B of table VI provides the results for nested model comparisons of one-quarter ahead and two-quarter-ahead out-of-sample forecasts of GDP growth for the full sample period 1947-2008. The first column shows which variables are included in the unrestricted model, and the second column shows which variable constitutes the restricted (baseline) model. In column three to five we report the relative mean squared error between the unrestricted ($\mathrm{MSE_u}$) and restricted model ($\mathrm{MSE_r}$), the MSE-F test statistic and the ENC-NEW statistic for the one-quarter-ahead forecasts. In the last three columns we report the same test statistics for the two-quarters-ahead forecasts.

Looking first at the one-quarter-ahead forecasts in panel B of table VI, we see that the relative MSE is less than one for all model comparisons except in the case when the baseline model is the credit spread (*dCred*). The MSE-F test for equal MSE between the unrestricted and restricted model reject the null of equal MSE in favor of the $\mathrm{MSE_u}$ being lower than $\mathrm{MSE_r}$ for all models except in the case when credit spread constitutes the baseline model. Based on the ENC-NEW test, we reject the null that the unrestricted models are encompassed by the restricted model at the 1% significance level for all cases. These results provide strong support that *dILR* improves forecast accuracy relative to all of the baseline models. For the two-quarters-ahead forecasts, we get similar results, although based on the MSE-F test we cannot reject the null that the MSE of a model with *dILR* and $\mathsf{er_m}$ has lower MSE than a model with only $\mathsf{er_m}$. The ENC-NEW test, however, supports the claim that *dILR* contains additional information to $\mathsf{er_m}$.

In the last row in panel B we examine the effect on forecast performance of adding *dILR* to a more comprehensive restricted model that contains all the financial variables examined earlier. For both the one-quarter and two-quarter ahead forecasts we cannot reject the null that the $\mathrm{MSE_u}$ is equal to $\mathrm{MSE_r}$, suggesting that there is no value adding *dILR* tho the restricted model. This is not surprising since we saw that adding *dILR* to a model with only *dCred*, did not change the MSE. Thus, we would expect a similar result for a larger model containing *dCred* as one of the predictor variables in the restricted model. However, the ENC-NEW test still rejects, at the 1% level, the null that the restricted model encompasses the unrestricted model, suggesting that adding *dILR* to the restricted model improves forecast performance both at the one-quarter and two-quarter horizon.

In Panel C of table VI we change the baseline model to an autoregressive model for GDP growth and test whether adding *dILR* (or any of the other financial variables) improves forecast accuracy of GDP growth relative to an autoregressive model for GDP growth. Looking first at the one-quarter-ahead forecasts, we find that *dILR*, $\mathsf{er_m}$ and *dCred* significantly improve the MSE relative to the baseline model. Adding the term spread or volatility to the model does not significantly reduce the MSE. The more powerful ENC-NEW test rejects the null that the baseline model encompass the unrestricted model at the 1% level for all variables except for market volatility, where the null is rejected at the 5% level.

For the two-quarters-ahead forecasts, all variables except market volatility improve the forecast accuracy of the autoregressive baseline model. Note also that the unrestricted model that includes *dILR* shows the greatest improvement in MSE over the baseline model when giving two-quarters-ahead forecasts. One final observation from Panel C is worth noting. The model that adds the term spread does not improve the MSE relative to the restricted autoregressive model in the one-quarter-ahead forecast comparison. However, when we look at the two-quarter ahead forecast comparison, the performance of the unrestricted model that adds *Term* to restricted model is greatly improved. *Term* thus has better performance for longer term forecasts.

III Firm size and the information content of liquidity

Small firms are relatively more sensitive to economic downturns than large firms. Therefore firm size might be of particular interest for the purpose of this paper. If the business cycle component in liquidity is caused by investors moving out of assets that have a tendency to perform particularly poorly in recessions, we would expect that the liquidity of small firms reflects this effect most strongly. Thus, we would expect the liquidity variation of small firms to be higher than the liquidity variation of large firms, and also the liquidity of small firms to be more informative about future macro fundamentals. To examine this more closely we run in-sample predictive regressions with liquidity variables constructed for different firm size quartiles. Firms are assigned into size quartiles at the beginning of the year based on their market capitalization the last trading day of the previous year. We construct two version of each liquidity variables, one calculated for the 25% smallest firms (LIQ^{small}) and one for the 25% largest firms (LIQ^{large}).

Table VII reports the results from regression models where we predict GDP growth using liquidity proxies calculated separately for small and large firms, and where we include the different control variables used earlier.[19] We find that the liquidity of small firms has a significant coefficient ($\hat{\beta}_S^{\mathsf{LIQ}}$) for all three liquidity proxies. The liquidity of large firms has an insignificant coefficient ($\hat{\beta}_L^{\mathsf{LIQ}}$) for all liquidity proxies in all models. We make a similar conclusion from the comparison of the $\bar{R}^2$ of the different specifications, reported on the right in the table. A regression specification with *only* the liquidity of the large firms has no R^2 improvement relative to models without liquidity; all the improvement in R^2 comes from the liquidity of small firms. This result is also confirmed in panel B in the table, which shows the results from Granger causality tests between the liquidity proxies for small and large firms and GDP growth. In the second and third column we report the χ^2 statistic and associated p-value from the test of the null that GDP growth does not Granger cause the respective liquidity variable. We cannot reject the null for any of the models. In the two last columns, we test the null that the liquidity variable does not Granger cause GDP growth. For all liquidity measures sampled for the small firms, we reject the null at the 5% level or better.

Overall the results in table VII suggest that the illiquidity of smaller firms is most informative about future economic conditions. We view this result as consistent with our conjecture that variation in market liquidity is caused by portfolio shifts, from illiquid more risky assets into safer more liquid assets, due to changing expectations about economic fundamentals or binding funding constraints.

Finally, if investors have a tendency to move out of small firms and this causes activity to drop and liquidity to worsen, we would expect this to show up in the trading activity

[19] In an internet appendix, we report the results for all three liquidity variable and the other macro variables.

Table VII: Predicting macro with market liquidity - size portfolios

The table in panel A shows the multivariate OLS estimates from regressing next-quarter GDP growth on current market illiquidity of small and large firms and four control variables. We examine three different proxies for market illiquidity, sampled for small and large firms. The crossectional liquidity measures are calculated as equally weighted averages across stocks. The estimated model is $y_{t+1} = \alpha + \beta_S^{LIQ} LIQ_t^{small} + \beta_L^{LIQ} LIQ_t^{large} + \gamma \mathbf{X}_t + u_{t+1}$, where y_{t+1} is real GDP growth, LIQ^{small} is the respective illiquidity proxy sampled for the 25% smallest firms and LIQ^{large} is the illiquidity of the 25% largest firms, $\mathbf{X_t}$ contains the additional control variables (*Term*, *dCred*, *Vola* and er_m) and γ' is the vector of the coefficient estimates for the control variables. The Newey-West corrected t-statistics (with 4 lags) is reported in parentheses below the coefficient estimates, and $\bar{R}^2$ is the adjusted R^2. The three last columns report the adjusted $\bar{R}^2$ for the models estimated without any liquidity measures (ex.LIQ $\bar{R}^2$), only including the liquidity sampled for the 25% largest firms (ex.$LIQ^S \bar{R}^2$), and the model only including the liquidity sampled for the 25% smallest firms (ex.$LIQ^L \bar{R}^2$). Panel B shows the results of Granger causality tests between real GDP growth and the illiquidity of small and large firms for the three different illiquidity proxies. The first column denotes the liquidity variable, columns two and three show the χ^2 and associated p-value for Granger causality tests where the null hypothesis is that GDP growth *does not* Granger cause the liquidity variables. Similarly, columns four and five show the results when the null hypothesis is that the liquidity variable *does not* Granger cause GDP growth. ** and * denotes a rejection of the null hypothesis at the 1% and 5% level, respectively.

Panel A: Predicting GDP with various liquidity measures

Liquidity variable	Const.	β_S^{LIQ}	β_L^{LIQ}	$\hat{\gamma}_1{}^{Term}$	$\hat{\gamma}_2{}^{dCred}$	$\hat{\gamma}_3{}^{Vola}$	$\hat{\gamma}_4{}^{er_m}$	$\bar{R}^2$	ex.LIQ $\bar{R}^2$	ex.LIQ^S $\bar{R}^2$	ex.LIQ^L $\bar{R}^2$
dILR	0.008	-0.008	0.003	0.000	-0.014	0.001	0.021	0.14	0.12	0.12	0.14
	(7.64)	(-3.74)	(1.09)	(0.54)	(-3.16)	(0.21)	(2.31)				
dLOT	0.009	-0.014	0.000	0.000	-0.015	0.009	0.029	0.14	0.12	0.12	0.15
	(7.52)	(-2.12)	(-0.06)	(0.42)	(-3.61)	(1.58)	(3.55)				
Roll	0.017	-0.306	-0.251	0.001	-0.013	0.007	0.022	0.15	0.12	0.14	0.15
	(5.14)	(-2.38)	(-0.91)	(1.39)	(-3.12)	(1.29)	(2.74)				

Panel B: Granger Causality tests

Liquidity variable (LIQ)	*dGDPR* $\nrightarrow$ LIQ χ^2	p-value	LIQ $\nrightarrow$ *dGDPR* χ^2	p-value
dILR ^{S}	4.34	(0.23)	10.33*	(0.02)
dILR ^{L}	6.86	(0.08)	1.32	(0.72)
dLOT ^{S}	3.19	(0.07)	9.83**	(0.00)
dLOT ^{L}	0.20	(0.65)	0.03	(0.87)
Roll ^{S}	0.67	(0.72)	6.44*	(0.04)
Roll ^{L}	0.19	(0.91)	5.60	(0.06)

of these firms. We have actually investigated whether trading volume predict economic growth, and found it to be less informative than other liquidity measures about real variables,[20] but looking at volume may still help in our understanding of the mechanisms. In figure 3 we therefore examine whether the change in turnover before and during NBER recessions is different for small and large firms. Turnover is measured as the shares traded divided by the number of outstanding shares. We sort firms into size quartiles at the end of each year and calculate the equally weighted average turnover for the first quartile and fourth quartile. As before, the bars show the cumulative average quarterly growth in real GDP and the solid line the cumulative average change in *ILR*. The dashed line shows the cumulative average change in turnover for small firms, and the dotted line shows the same series for large firms. The results in the figure indicate a striking systematic difference in the development of trading activity in small and large firms before recessions. While the turnover for large firms is essentially unchanged before the first recession quarter, the turnover for small firms falls steadily already from four quarters before the first NBER recession quarter (NBER1). Furthermore, the turnover for both small and large firms starts increasing already in the middle of the NBER recessions. Since this pattern is strongest for small firms, it indicates that investors increase their demand for equities in general, and for smaller firms in particular, when they start expecting future economic conditions to improve.

IV Systematic liquidity variations and portfolio shifts - evidence from Norway

In the introduction, we conjectured that the systematic liquidity variations found are linked to portfolio shifts and changes in market participation during economic downturns, i.e. that investors seek to move away from equity investments in general and from small illiquid stocks in particular. Using special data on stock ownership from the Oslo Stock Exchange (OSE), we can examine this conjecture. In addition, the Norwegian data set provides a valuable robustness check of our results from the US market.

A The Norwegian evidence of predictability

We first check that we get similar results on predictability as in the US case. For brevity we do not report the Norwegian results on predictability, only summarize the results[21] and start by assessing the in-sample predictive ability of market liquidity for the macro variables real GDP growth ($dGDPR$), growth in the unemployment rate (dUE), real consumption growth ($dCONSR$) and growth in investments ($dINV$). We use the Amihud illiquidity ratio (ILR) and relative spread (RS) as our liquidity proxies.[22]

We look at two model specifications. In the first specification, we use only market liquidity and the lagged dependent variable as predictors for next quarter growth in

[20] In an internet appendix, we report the results from a comprehensive VAR specification which includes turnover as an alternative explanatory variable. We find that turnover has no predictive ability for $dGDPR$.

[21] The results for the Norwegian sample are reported in an internet appendix to the paper.

[22] Both the ILR and RS pass the stationarity tests in the Norwegian sample, so we do not transform any of the liquidity series.

Figure 3: Market illiquidity and trading activity (turnover) around NBER recessions
The figure shows the accumulated average growth in *ILR* (solid line) and accumulated average GDP growth (bars) averaged in event time before, during and after NBER recession periods. In addition, the dashed line shows the accumulated average change in turnover for the 25% smallest firms and the dotted line shows the accumulated average change in turnover for the 25% largest firms. Turnover is measured as the shares traded divided by the number of outstanding shares. All the NBER recession periods are aligned to start at NBER1.

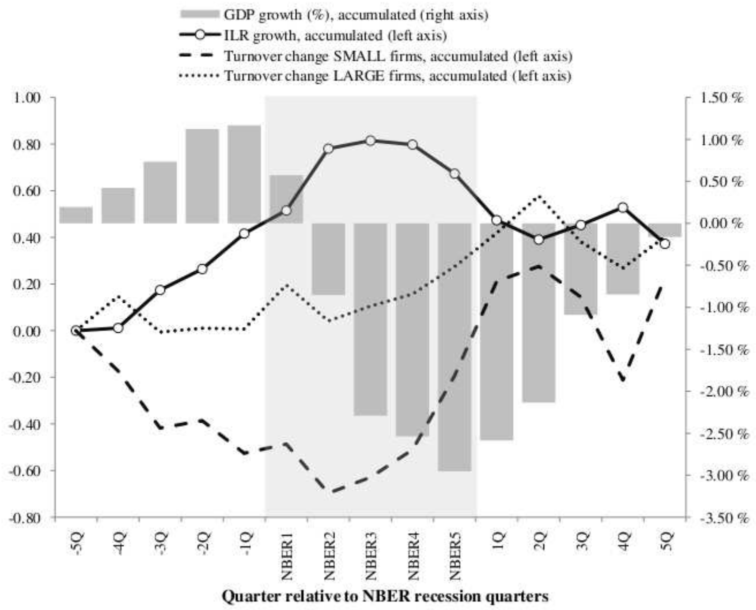

the respective macro variable. We find that regardless of choice of liquidity proxy, the coefficient on market liquidity is highly significant across all models and has the expected signs. A worsening of market liquidity (increase in *RS* or *ILR*) predicts a decrease in next quarter GDP growth, consumption growth, investment growth and an increase in the unemployment rate.

In the second model specification, we control for other variables. In the US analysis, we used four financial control variables; the term spread, credit spread, market returns and market volatility. In Norway, no credit spread series are available for the length of our sample period. This is mainly due to a historically very thin credit market in Norway. Thus, we are only able to control for the other three variables. The results from regressions based on this specification show that the coefficient on market liquidity is highly significant for all models except when the dependent variable is real consumption growth. This is basically the same result we found for the US; after controlling for the term spread and stock market returns, the coefficient on *ILR* was rendered insignificant in the equation for *dCONSR*. However, none of the other financial variables have significant coefficients. It should also be noted that if we exclude the relative spread, the term spread enters significantly into the models for *dGDPR* and *dUE*, although the adjusted R-squared of the models is more than halved. Thus, although *Term* is highly correlated with our liquidity proxies, there seem to be a significant amount of additional information in market liquidity. We also perform Granger causality tests for the Norwegian sample, between *dGDPR* and *RS* and *ILR*. In that analysis, we are unable to reject the null that GDP growth does not Granger cause *RS*, while we reject the reverse hypothesis at the 1% level. This result is similar when we use the *ILR* as our liquidity proxy.

We also perform an out-of-sample analysis for Norway. In nested model comparisons between *RS* or *ILR* and the other financial control variables (*Term*, $\mathtt{er}_\mathtt{m}$, *Vola*), the MSE-F test suggests that the MSE of an unrestricted model (including *RS* as a predictor) has a significantly lower MSE across all models. When we use *ILR* as the liquidity proxy, we are only able to reject the null of equal forecast accuracy in the model where $\mathtt{er}_\mathtt{m}$ is the competing predictor variable. Both for the *RS* and *ILR*, the results are weaker with respect to the ENC-NEW test, and much weaker compared to the results for the US. We are only able to reject the null at the 5% level, that *RS* is encompassed by a model with $\mathtt{er}_\mathtt{m}$ or *Vola*. For *ILR*, we only reject the null of encompassing when the restricted model contains $\mathtt{er}_\mathtt{m}$.

Similar to US out-of-sample analysis, we also compare the out-of-sample forecast performance of liquidity to an autoregressive model for GDP growth. Adding either *RS* or *ILR* to the autoregressive GDP model significantly improves the MSE. In addition, the null that the restricted GDP model encompasses the unrestricted model that adds either *RS* or *ILR* is rejected at the 1% level.

Finally, we examine whether the informativeness of the liquidity about future GDP growth differs between small and large firms also in Norway. We sort firms on the OSE into four groups based on their market capitalization at the end of the previous year, and calculate the average liquidity for each size group. We use the liquidity series for the smallest and largest group as explanatory variables. The results are very similar to what we found for the US in table VII. Also, in the Granger causality tests, we reject the null hypothesis that both RS^{S} and ILR^{S} sampled for the small firms *does not* Granger cause *dGDPR*, while we are unable to reject the null when using the liquidity measured for the largest firms.

In summary, while the in-sample results and Granger causality tests for Norway are very similar to the US results, the out-of-sample results are a bit weaker for Norway. Note however, that the Norwegian sample is much shorter, and covers only about three business cycles. Overall, the results for Norway indicate that the result that stock market liquidity is related to future economic growth is robust to change of market, market structure and trading system.

B Portfolio shifts and liquidity

A possible channel through which the documented relationship between stock market liquidity and business cycles may work is changes in portfolio compositions. In this section, we therefore investigate whether investors do in fact tilt their portfolios towards more liquid assets in economic downturns. Our Norwegian data set includes monthly ownership of all investors in all Norwegian companies listed on the Oslo Stock Exchange (OSE) over the period 1992-2007. The challenge lies in constructing aggregate measures of changes in portfolio composition. We do this in two different ways. First we focus on market participation and look at the full portfolio of each investor. Then we look at concentration and movements between owner types for individual stocks, without controlling for the portfolios *across* stocks.

B.1 Market participation on an investor-by-investor basis

Our ownership data lets us construct the actual portfolios of all investors at the monthly frequency and also the changes in portfolio composition over time. We want a variable that can be informative about both the degree to which investors move in and out of the stock market and the degree to which the structure of their stock portfolios change. The measure should mainly be influenced by actual changes in stock ownership. This rules out measures based on wealth changes, since such measures have the undesirable characteristic that wealth can change due to stock price changes, even if investors do not make any active portfolio changes. We therefore use the *number of shares* owned by an investor as the basic piece of data. We can not sum the number of shares across stocks, since this is again sensitive to price differences across shares. Instead, we simply ask: When does an owner realize the portfolio? Obviously when he sells *all* his stocks. Our measure of aggregate changes uses these cases to identify aggregate movements in and out of the market or a group of stocks, such as a size portfolio.

Our time series is constructed by comparing the set of participants at two following dates. The set of investors which were present at the first date, but not on the second date, is the set of investors *leaving* the market entirely. Similarly, we count the number of investors present at the second date, but not at the first. This is the number of investors *entering* the market. The net change in investors is the number of investors entering the market less the number of investors leaving the market. This number is used as a measure of the change in portfolio composition. The net change in investors is calculated for all owners as well as for each of the owner types (personal, foreign, financial, nonfinancial(corporate) and state owners).[23] Panel A of Table VIII shows some

[23]In implementing the calculation, we attempt to reduce noise by removing trivial holdings of less than

descriptive statistics for the net change in portfolio compositions at the annual level. On average about 15 thousand investors enter the market from one year to the next, which is about a quarter of the investors present at the beginning of the year. The net change is positive, which indicates that on average the number of investors on the exchange has been increasing over the sample period. Panel A also shows the average number of investors leaving and entering the market within each owner type. Note that in the calculations for different owner types, we only consider owners *of the given type*, i.e. the fraction of investors is conditioned on the type. For example, the average of 51 financial owners entering corresponds to 14% of financial investors. As is clear from the table, the most common investor type is personal investors.[24]

As we saw for both the US and Norway, the time series of small firms' liquidity have more predictive content than the time series of large firms' liquidity. To look into such issues we therefore construct measures of changes in participation for different size quartiles, i.e. we sort the stocks at the OSE based on size and each year construct four size-based stock portfolios. We then calculate the same participation measure, the net number of new owners, but now *only* for the stocks in each size portfolio. So, if an investor only had holdings in small stocks, but moved them to large stocks, we would count this as leaving the small stock portfolio and entering the large stock portfolio.

Panel B of Table VIII shows the correlations between liquidity, measured by the relative bid-ask spread (RS), and portfolio changes for various owner types. If liquidity worsens (RS increases) when the number of participants in the market falls, we should expect a negative correlation between RS and changes in the number of investors. This relationship should be strongest for the least liquid stocks. That is exactly what we find. For the portfolio of the smallest stocks on the OSE, there is a significantly negative correlation between relative spreads and changes in participation. The correlation becomes smaller in magnitude when we move to portfolios of larger firms, the correlation being smallest in magnitude for the portfolio of the largest firms.

B.2 Movements between owner types for individual stocks

A problem with the measure of participation above may be that it *only* considers cases of complete withdrawal from the market. We therefore complement the analysis by looking at a couple of alternative, related measures, namely owner concentration and owner types. These measures are much simpler to calculate than the previous, as they can be found on a stock-by-stock basis, without looking at the full portfolio of individual investors. We first look at the concentration of ownership, measured by such standard measures as the fraction of the company owned by the largest owner, and a couple of Herfindahl measures of concentration. We also look at the total number of owners. Ownership concentration is related to participation by the simple book-keeping argument that since all stocks must be held by somebody, if participation declines, the number of

a hundred shares, since this is the minimum lot size at the Oslo Stock Exchange.

[24]There is an institutional reason for the decrease in foreign investors. It is a reflection of the increased ownership through nominee accounts, where foreign owners register through a nominee account. The Norwegian Central Securities Registry do not have details on nominee ownership, they only have data on the total held in nominee accounts. The number of foreign investors we are using is the number of directly registered foreign owners, which has decreased, although the fraction of OSE held by foreigners has increased throughout the period.

Table VIII: Changes in portfolio composition and liquidity

The table in panel A describes changes in ownership participation measured at an annual frequency. Each year in the sample we calculate the number of investors leaving the market totally, entering the market, and the net change. We also normalize the numbers by calculating them as a fraction of owners at the beginning of the period. Panel B presents correlations between stock market liquidity measured by the average relative bid ask spread in a period and the changes in stock market participation in the period. Change in stock market participation is the change in the number of investors in the stock market, or the given portfolio, of the specified types. Numbers in parenthesis are p-values. Panel C shows correlation between changes in measures of ownership concentration and changes in liquidity. We calculate four concentration measures: The size of the largest owner and the total number of owners, as well as two Herfindahl indices. The Herfindahl index is the sum of squared ownership fractions. In the first version we include all owners. In the second version we exclude the three largest owners. The numbers in the tables are the correlations between the concentration measures and the liquidity, measured by the relative spread, for individual stocks. Numbers in parenthesis are p-values. Panel D shows correlations between liquidity and changes in aggregate fraction of the firm owned by the five different owner types. The numbers in the tables are the correlations between the ownership fraction by the given type, and liquidity, measured by the relative spread, for individual stocks. Numbers in parenthesis are p-values. In the tables we show data for five mutually exclusive owner types: Individual (private), nonfinancials (corporate), state, foreign and financial owners. In some tables we also show data for mutual funds, which is a subgroup of financials, and included in the financials. For annual data we use each year from 1990 to 2006, giving 16 observations. For the calculations with quarterly data we use data between 1993:1 to 2006:12, giving 56 quarterly observations.

Panel A: Describing annual changes in portfolio composition

Investor	Number of investors			Fraction of investors		
type	entering	leaving	net	entering	leaving	net
All	15220	11934	3286	24.1	18.5	5.6
Personal owners	13445	10087	3358	24.3	17.5	6.8
Foreign owners	862	1119	-256	33.7	35.3	-1.6
Financial owners	51	44	6	14.8	12.4	2.4
Nonfinancial owners	1013	838	175	24.4	19.6	4.8
State owners	14	11	3	20.8	15.1	5.7

Panel B: Correlation liquidity and change in stock market participation (quarterly)

	All firms		Firm size quartiles: Q1 (smallest firms)		Q2		Q3		Q4 (largest firms)	
All owners	-0.07	(0.32)	-0.35	(0.00)	-0.10	(0.22)	-0.20	(0.07)	-0.11	(0.22)
Personal owners	-0.02	(0.45)	-0.33	(0.01)	-0.09	(0.25)	-0.18	(0.09)	-0.08	(0.28)
Foreign owners	-0.18	(0.09)	-0.30	(0.01)	-0.16	(0.12)	-0.25	(0.03)	-0.23	(0.04)
Financial owners	-0.06	(0.33)	-0.11	(0.21)	0.01	(0.46)	-0.09	(0.25)	-0.08	(0.27)
Nonfinancial owners	-0.16	(0.12)	-0.35	(0.00)	-0.11	(0.21)	-0.21	(0.06)	-0.20	(0.06)
State owners	-0.06	(0.34)	-0.20	(0.07)	0.19	(0.08)	-0.10	(0.23)	-0.06	(0.34)

Panel C: Correlation change in liquidity and change in ownership concentration (quarterly)

Concentration measure	All firms	Firm Size Quartile: Q1 (smallest firms)	Q2	Q3	Q4 (largest firms)
largest owner	0.07 (0.30)	0.13 (0.15)	0.13 (0.16)	0.09 (0.25)	-0.06 (0.31)
Herfindahl	0.09 (0.24)	0.20 (0.06)	0.10 (0.22)	0.18 (0.08)	-0.12 (0.18)
No owners	0.37 (0.00)	-0.09 (0.23)	-0.22 (0.04)	-0.27 (0.02)	0.37 (0.00)
Herfindahl (ex 3 largest)	0.18 (0.08)	0.29 (0.01)	0.23 (0.04)	-0.07 (0.29)	-0.05 (0.36)

Panel D: Correlation change in liquidity and movement across owner types (quarterly)

Owner type	All firms	Firm Size Quartile: Q1 (smallest firms)	Q2	Q3	Q4 (largest firms)
Financial fraction	-0.12 (0.18)	-0.14 (0.14)	-0.10 (0.21)	-0.07 (0.29)	0.24 (0.03)
Mutual fund fraction	-0.06 (0.32)	-0.13 (0.16)	-0.00 (0.49)	0.04 (0.37)	-0.18 (0.08)
Individual fraction	0.05 (0.35)	-0.03 (0.42)	-0.14 (0.13)	0.01 (0.46)	0.06 (0.32)
Nonfinancial fraction	-0.06 (0.34)	0.10 (0.22)	0.05 (0.36)	-0.14 (0.13)	-0.17 (0.09)
Foreign fraction	-0.08 (0.26)	-0.16 (0.11)	-0.04 (0.38)	-0.07 (0.29)	0.21 (0.05)
State fraction	-0.09 (0.23)	-0.30 (0.01)	-0.18 (0.08)	-0.07 (0.29)	0.22 (0.05)

owners declines and ownership concentration increases.

In panel C of table VIII we show the results of looking at correlations between changes in liquidity and ownership concentration. The interesting numbers are the differences between the portfolio of small firms (quartile 1) and large firms. We see that when for example the spread increases, the concentration increases for the portfolio of small stocks (positive correlation), but decreases for the portfolio of large stocks. Similarly, when the spread increases the number of owners decreases for the portfolio of small stocks, but increases for the large stocks.

The changes in ownership participation may also be related to the prevalent type of owner. We therefore calculate the aggregate fractions of companies owned by five different owner types, and relate changes in these fractions to changes in liquidity. The results are given in panel D of the table. There are a number of interesting patterns in the table. First, we see that when liquidity worsens, this coincides with a movement *into* large stocks by individual investors, which is consistent with a portfolio rebalancing by that type of investor, and a "flight to quality" type of behavior. Secondly, it is also interesting to note what type of investors "take up the slack" in small firms. One obvious candidate is financial investors, but that turns out not to be the case. When the spread increases (liquidity worsening) financials tend to also decrease their stake in small stocks.

A potential explanation of this result concerns the issue of funding problems discussed earlier. Included in the group of financial investors are mutual funds. Since these have a tendency to experience outflows of funds in economic downturns, as investors realize some of their portfolios to fund consumption, mutual funds are faced with a funding problem and have to realize a part of their portfolios. If this is the case we would expect outflow from small stocks to be more prevalent among mutual funds than other financial investors. We are able to investigate this conjecture, as the database on ownership identifies which of the financial owners are mutual funds. We therefore redo the calculation *only* for those financial owners that are mutual funds. The results are shown as the bottom line in panel D. The results show that mutual funds have a *stronger* tendency to realize their holdings of small stocks. This is consistent with an explanation based on funding problems.

We observe that the group that seems to "take up the slack" in small firms (although this number is not significant) is foreign investors. This group, which for example includes large international funds, seems to buy small stocks when liquidity is worsening.

To sum up, using various measures of changes in portfolio compositions, we find evidence consistent with our hypothesis that liquidity changes are related to portfolio shifts.

V Conclusion

The prime contribution of this paper is to provide two empirical observations. First, we show that stock market liquidity contains useful information for estimating the current and future state of the economy. These results are shown to be remarkably robust to our choice of liquidity proxy and sample period. The relationship is also very similar for two different markets, the US and Norway. Second, we find evidence that time variation in equity market liquidity is related to changes in participation in the stock market, especially for the smallest firms. Participation in small firms decreases when the economy (and market liquidity) worsens. This is consistent with a "flight-to-quality" effect and with the finding that the liquidity of the smallest firms contains most information about future economic conditions. In addition to suggesting a new financial market-based predictor, our results provide a new explanation for the observed commonality in liquidity.

There are a number of interesting ways to follow up our results. First, our results showing that (Granger) causality goes from the stock market to the real economy has interesting implications for prediction, particularly in a policy context. The ability to improve forecasts and "nowcasts" (Giannone, Reichlin, and Small, 2008) of such central macroeconomic variables as unemployment, GDP, consumption and the like will be particularly interesting for central banks and other economic planners. For such purposes it would be interesting to do more extensive comparisons of the predictive power of different liquidity proxies, or combinations of proxies. Second, while we have found evidence of the link from observed liquidity to the economy using data for the US and Norway, it would be interesting to also look at a larger cross-section of stock markets. Finally, our finding that stock market participation is related to time variation in liquidity should be important input to asset-pricing theorists attempting to understand why liquidity seems to be priced in the cross-section of stock returns.

References

Viral A Acharya and Lasse H Pedersen. Asset pricing with liquidity risk. *Journal of Financial Economics*, 77:375–410, 2005.

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

Alessandro Beber, Michael W Brandt, and Kenneth A Kavajecz. What does equity sector orderflow tell us about the economy? Unpublished Working paper, University of Amsterdam, February 2010.

Valerie R. Bencivenga, Bruce D. Smith, and Ross M. Starr. Transactions costs, technological choice, and endogenous growth. *Journal of Economic Theory*, 67:153–177, 1995.

Øyvind Bøhren and Bernt Arne Ødegaard. Patterns of corporate ownership: Insights from a unique data set. *Nordic Journal of Political Economy*, pages 57–88, 2001.

Øyvind Bøhren and Bernt Arne Ødegaard. Governance and performance revisited. In Paul Ali and Greg Gregouriu, editors, *International Corporate Governance after Sarbanes-Oxley*, pages 27–64. Wiley, February 2006.

Markus K Brunnermeier and Lasse H Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22:2201–2238, 2009.

Tarun Chordia, Richard Roll, and Avanidhar Subrahmanyam. Commonality in liquidity. *Journal of Financial Economics*, 56:3–28, 2000.

Todd E Clark and Michael W McCracken. Tests of equal forecast accuracy and encompassing for nested models. *Journal of Econometrics*, 105: 85–110, 2001.

Jay F Coughenour and Mohsen M Saad. Common market makers and commonality in liquidity. *Journal of Financial Economics*, 73:37–69, 2004.

Francis X Diebold and Roberto S Mariano. Comparing predictive accuracy. *Journal of Business and Economic Statistics*, 13:253–263, 1995.

Martin DD Evans and Richard K Lyons. The term structure as a predictor of real economic activity. *Journal of Financial Economics*, 88:26–50, 2008.

Akiko Fujimoto. Macroeconomic sources of systematic liquidity. Unpublished Working Paper, Yale University, October 2003.

Domenico Giannone, Lucrezia Reichlin, and David Small. Nowcasting: The real-time informational content of macroeconomic data. *Journal of Monetary Economics*, 55:665–676, 2008.

Rajna Gibson and Nicolas Mougeot. The pricing of systematic liquidity risk: Empirical evidence from the US stock market. *Journal of Banking and Finance*, 28:157–178, 2004.

Simon Gilchrist, Vladimir Yankov, and Egon Zakrajsek. Credit market shocks and economic fluctuations: Evidence from corporate bond and stock markets. *Journal of Monetary Economics*, 56:471–493, May 2009.

Ruslan Y Goyenko and Andrey D Ukhov. Stock and bond market liquidity: A long-run empirical analysis. *Journal of Financial and Quantitative Analysis*, 44:189–212, 2009.

Ruslan Y Goyenko, Craig W Holden, and Charles A Trzcinka. Do liquidity measures measure liquidity? *Journal of Financial Economics*, 92(2):153 – 181, 2009.

Allaudeen Hameed, Wenjin Kang, and S. Viswanathan. Stock market declines and liquidity. *Journal of Finance*, 65:257–293, 2010.

Lawrence Harris. Statistical properties of the Roll serial covariance bid/ask spread estimator. *Journal of Finance*, 45:579–590, 1990.

Campbell R Harvey. The real term structure and consumption growth. *Journal of Financial Economics*, 22:305–333, 1988.

Campbell R Harvey. Forecasts of economic growth from the bond and stock markets. *Financial Analysts Journal*, pages 38–45, September-October 1989.

David I Harvey, Stephen J Leybourne, and Paul Newbold. Tests for forecast encompassing. *Journal of Business and Economic Statistics*, 16:254–259, 1998.

Joel Hasbrouck and Duane Seppi. Common factors in prices, order flows, and liquidity. *Journal of Financial Economics*, 59:383–411, 2001.

Gur Huberman and Dominika Halka. Systematic liquidity. *Journal of Financial Research*, XXIV (2):161–178, 2001.

Aditya Kaul and Volkan Kayacetin. Forecasting economic fundamentals and stock returns with equity market order flows: Macro information in a micro measure? Unpublished Working Paper, University of Alberta, February 2009.

Rudiger Kiesel, William Perraudin, and Alex Taylor. The structure of credit risk: spread volatility and ratings transitions. Bank of England Working paper 131, 2001.

Denis Kwiatkowski, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? *Journal of Econometrics*, 54:159–178, 1992.

Albert Kyle. Continous auctions and insider trading. *Econometrica*, 53:1315–35, 1985.

David A Lesmond, Joseph P Ogden, and Charles A Trzcinka. A new estimate of transaction costs. *Review of Financial Studies*, 12:1113–1141, 1999.

Ross Levine. Stock markets, growth, and tax policy. *Journal of Finance*, 46:1445–1465, 1991.

Ross Levine and Sara Zervos. Stock markets, banks, and economic growth. *American Economic review*, 88:537–558, 1998.

Marc L. Lipson and Sandra Mortal. Liquidity and capital structure. *Journal of Financial Markets*, 12, 2009.

Francis A Longstaff. The flight-to-quality premium in U.S. treasury bond prices. *Journal of Business*, 77(3):511–525, 2004.

Michael W McCracken. Asymptotics for out-of-sample tests for Granger causality. *Journal of Econometrics*, 140:719–752, 2007.

Randi Næs, Johannes Skjeltorp, and Bernt Arne Ødegaard. Liquidity at the Oslo Stock Exchange. Working Paper Series, Norges Bank, ANO 2008/9, May 2008.

Bernt Arne Ødegaard. Who moves equity prices? Monthly evidence. Unpublished Working Paper, University of Stavanger, February 2009.

Maureen O'Hara. Presidential adress: Liquidity and price discovery. *Journal of Finance*, LVIII (4), August 2003. 1335-1354.

Lubos Pastor and Robert F Stambaugh. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3):642–685, June 2003.

Monica Pedrosa and Richard Roll. Systematic risk in corporate bond credit spreads. *Journal of Fixed Income*, pages 7–26, 1998.

Richard Roll. A simple implicit measure of the effective bid–ask spread in an efficient market. *Journal of Finance*, 39(4):1127–1139, September 1984.

Johannes Skjeltorp and Bernt Arne Ødegaard. Why do firms pay for liquidity provision in limit order markets? Unpublished Working Paper, University of Stavanger, January 2010.

Jonas Söderberg. Do macroeconomic variables forecast changes in liquidity? An out-of-sample study on the order-driven stock markets in Scandinavia. Working Paper 2009:10, University of Växjö, December 2008.

James H Stock and Mark W Watson. Forecasting output and inflation: The role of asset prices. *Journal of Economic Literature*, 41:788–829, September 2003.