

ESSAYS ON STOCK LIQUIDITY AND STOCK RETURN PREDICTABILITY

by

GREGORY W. EATON

(Under the Direction of Bradley Paye and Annette Poulsen)

ABSTRACT

I examine the effects of stock liquidity on asset values and whether aggregate stock liquidity and other forecasting instruments predict stock market returns. In the first chapter, I use tick-size reductions in equity markets as sources of exogenous variation in liquidity to examine the causal effect of transaction costs on firm value. In contrast to the prevailing view, I find that increased liquidity has a marginal or, in some cases, negative impact on firm value. The second chapter evaluates the predictive content of aggregate liquidity for economic activity and stock returns. We decompose illiquidity into a component capturing aggregate volatility and a volatility-adjusted component and find strong evidence that the component of illiquidity uncorrelated with volatility forecasts stock market returns. The third chapter provides new evidence on the stock return forecasting performance of alternative corporate payout yields. We find that the net payout yield forecasts stock returns and generally outperforms the commonly used dividend yield. Additionally, we show that the choice of cash flow used to construct the payout yield is economically significant. An agent relying on the incorrect payout measure as a forecasting instrument is willing to pay an economically significant amount to switch to the optimal policy.

INDEX WORDS: Asset pricing, Stock return predictability, Stock liquidity, Market
microstructure, Net payout yield

ESSAYS ON STOCK LIQUIDITY AND STOCK RETURN PREDICTABILITY

by

GREGORY W. EATON

B.S., University of Missouri, 2007

A Dissertation Submitted to the Graduate Faculty of The University of Georgia
in Partial Fulfillment of the Requirements for the Degree

DOCTOR OF PHILOSOPHY

ATHENS, GEORGIA

2016

©2016

Gregory W. Eaton

All Rights Reserved

ESSAYS ON STOCK LIQUIDITY AND STOCK RETURN PREDICTABILITY

by

GREGORY W. EATON

Approved:

Major Professors: Bradley Paye
 Annette Poulsen

Committee: Stuart Gillan
 Jeffry Netter

Electronic Version Approved:

Suzanne Barbour
Dean of the Graduate School
The University of Georgia
May 2016

Acknowledgments

I would like to express the deepest appreciation to my dissertation chairs Professor Bradley Paye and Professor Annette Poulsen. Their insights and encouragement were instrumental to my successful completion of this dissertation, and without their guidance this dissertation would not have been possible.

I would like to thank my committee members Professor Stuart Gillan and Professor Jeffrey Netter for their valuable comments. My special thanks go to Professor Sara Holland, Professor Harold Mulherin, and Professor Julie Wu for their suggestions.

Table of Contents

Acknowledgements	iv
List of Tables	vii
List of Figures	ix
1 The Effects of Market Liquidity on the Firm: Does Liquidity Impact Firm Value?	1
1.1 Introduction	2
1.2 Empirical Design	7
1.3 Examining the Effect of Liquidity on Firm Value using Tick-Size Reductions in U.S. Equity Markets	12
1.4 Additional Analyses	20
1.5 Concluding Remarks	24
2 Micro(structure) before Macro? The Predictive Power of Aggregate Illiquidity for Stock Returns and Economic Activity	40
2.1 Introduction	42
2.2 Data	47
2.3 Aggregate Illiquidity	51
2.4 The Predictive Power of Aggregate Illiquidity	61
2.5 Concluding Remarks	76

3	Payout Yields and Stock Return Predictability: How Important is the Measure of Cash Flow?	91
3.1	Introduction	93
3.2	Alternative payout yields and cash flow measures	96
3.3	Stock return forecasts and the measure of payout	99
3.4	Economic significance of the payout measure	109
3.5	Concluding Remarks	116
	References	126
	Appendices	139
A	A Timeline of Key Dates for Tick-Size Reductions in U.S. Equity Markets	139
B	Construction of Aggregate Illiquidity and Macroeconomic Variables	142
B.1	Aggregate liquidity and price delay measures	142
B.2	Macroeconomic and financial variables	146
C	Construction of Payout Variables	149
D	Out-of-Sample Forecasting Tests	151

List of Figures

1.1	Effective spreads of NYSE and AMEX stocks surrounding 1997 tick-size reduction	26
1.2	Effective spreads of NASDAQ stocks surrounding 1997 tick-size reduction	27
1.3	Effective spreads of NYSE and AMEX stocks surrounding decimalization	28
1.4	Effective spreads of NASDAQ stocks surrounding decimalization	29
2.1	Alternative effective spread proxies: 1926-2013	78
2.2	Alternative effective spread proxies: 1993-2013	79
2.3	Spread proxies adjusted for aggregate volatility	80
2.4	Out-of-sample Stock Return Forecasting Performance	81
3.1	Time series plots for alternative yields.	118
3.2	Average allocation to stock as a function of investment horizon.	119

List of Tables

1.1	Descriptive statistics of key measures around the 1997 tick-size reduction	30
1.2	NYSE and AMEX stocks sorted by shock to spreads associated with the 1997 tick-size reduction	31
1.3	NASDAQ stocks sorted by shock to spreads associated with the 1997 tick-size reduction	32
1.4	Regression estimates of the impact of a shock to liquidity from the 1997 tick-size reduction on firm value	33
1.5	Descriptive statistics of key measures around the shift to decimalization	35
1.6	Univariate sorts by shock to spreads associated with decimalization	36
1.7	Regression estimates of the impact of a shock to liquidity from decimalization on firm value	37
1.8	Impact of shock to liquidity from decimalization on firm value using alternative measures	38
2.1	Descriptive Statistics: Raw Illiquidity Measures and Selected Macroeconomic and Financial Variables	82
2.2	Extracting a Volatility Component from Illiquidity Measures	83
2.3	Finite-Sample Bias in Predictive Regressions with Long Memory Variables	84
2.4	Forecasting regressions for stock returns: Monthly frequency	85
2.5	Forecasting regressions for macroeconomic activity: Monthly frequency	86
2.6	Long-horizon forecasting regressions	87

2.7	Monthly forecasts including macroeconomic and financial variables	88
2.8	Forecasting regressions: Quarterly frequency	89
2.9	Out-of-sample forecasting performance: monthly frequency	90
3.1	Predictive regressions for stock returns using alternative yields	120
3.2	Multivariate regressions: decomposing the net payout yield	121
3.3	Out-of-sample evidence: stock return forecasts	122
3.4	Joint tests for predictability in stock returns or cash flow growth	123
3.5	VAR calibration for long-horizon portfolio choice analysis	124
3.6	Welfare simulation results by investment horizon and risk aversion level	125

Chapter 1

The Effects of Market Liquidity on the Firm: Does Liquidity Impact Firm Value?

Abstract

This study provides new evidence on the effects of secondary market liquidity on firm value. The prevailing view is that increased liquidity positively affects firm value. However, a small but growing literature argues that liquidity can have a negligible or, in certain cases, negative impact on the firm. I find empirical support for the latter. Using tick-size reductions in U.S. equity markets as sources of exogenous variation in liquidity, I find that increased liquidity has an economically small or, in some instances, negative impact on firm value. Additionally, I show that the choice of liquidity measure primarily explains the difference between my results and prior findings.

1.1 Introduction

Does liquidity impact firm value? The prevailing view is that increased secondary market liquidity positively impacts firm value. Consistent with this view, Fang, Noe, and Tice (2009) find that liquidity has a positive causal effect on firm value. This result is supported by a large theoretical and empirical literature. Amihud and Mendelson (1986) and Lynch and Tan (2011) theorize that, *ceteris paribus*, liquid assets are priced at a premium relative to illiquid assets due to differences in expectations about returns. Investors require a higher return to hold illiquid assets to offset the higher transaction costs associated with purchasing and selling illiquid investments. Amihud and Mendelson (1986) also conduct an empirical analysis that supports their model’s predictions. Brennan and Subrahmanyam (1996), Hasbrouck (2009), and Corwin and Schultz (2012) are among other studies finding empirical evidence that liquidity level is negatively related to expected return.¹ Butler, Grullon, and Weston (2005b) also examine the effects of liquidity on discount rates, and find that increased liquidity reduces the cost of capital in seasoned equity offerings.

Increased liquidity may also positively impact firm value through corporate governance mechanisms. Kyle and Vila (1991), Kahn and Winton (1998), and Maug (1998) theorize that liquidity facilitates the formation of blocks by reducing the costs of acquiring shares. Blockholders aim to profit from the increase in firm value associated with monitoring activities. Admati and Pfleiderer (2009), Edmans (2009), and Edmans and Manso (2011) present theoretical models where liquidity enables the threat of exit after the block has been formed, as informed blockholders can more easily sell their shares, thereby reducing the stock price. The threat of downward price pressure increases the cost of managerial opportunism in these models, leading the manager to make value-maximizing decisions (to the extent that manager’s wealth is tied to the stock price). Holmström and Tirole (1993) construct a model in which market liquidity facilitates information collection by “speculators,” which leads to more efficient prices and more effective equity-based compensation. In terms of empirical evidence, both Bharath, Jayaraman,

¹However, others find that the liquidity premium has decreased in recent years (Ben-Rephael, Kadan, and Wohl (2015)) or that liquidity is not priced across value-weighted portfolios (Hou, Xue, and Zhang (2014)).

and Nagar (2013) and Edmans, Fang, and Zur (2013) find a positive relation between liquidity and the efficacy of governance, and conclude that the threat of exit is an important driver of this relation. Additionally, Norli, Ostergaard, and Schindele (2015) find a positive relation between liquidity levels and shareholder activism.

In contrast to the aforementioned studies, a small but growing literature argues that increased liquidity has a marginal or, in certain cases, negative impact on firm value. For example, Constantinides (1986) constructs a theoretical model in which investors respond to the presence of transaction costs by rebalancing their portfolios less frequently (since the utility costs of trading less often are small). Therefore, transaction costs have only a second-order effect on asset values. Similarly, Vayanos (1998) predicts that, under realistic parameter values, the effects of transaction costs are small because investors trade only to satisfy life-cycle motives. Additionally, in certain cases, Vayanos (1998) predicts that asset prices increase with transaction costs. This is because trading costs induce investors to hold shares for a longer time period, which can increase total asset demand. Bhidé (1993) and Back, Li, and Ljungqvist (2015) predict that liquidity can be harmful for governance. These studies theorize that greater liquidity reduces the cost of exit, increasing the likelihood that a large investor sells his shares rather than intervene, sometimes referred to as “The Wall Street Walk.”

Empirically, Back, Li, and Ljungqvist (2015) find support for their theoretical prediction. Using unexpected changes in firm information environment or market making operations as exogenous shocks to liquidity, the authors find that liquidity negatively impacts shareholder activism. Fang, Tian, and Tice (2014) find that increased liquidity can adversely impact the firm by negatively affecting future innovation. Barclay, Kandel, and Marx (1998) examine empirically the impact of firms that switch exchanges and find that changes in transaction costs due to exchange reassignment do not impact security prices.

It is challenging to empirically examine the impact of liquidity on firm value due to the endogenous nature of the relationship between market liquidity and the firm. For example, liquidity and firm value are likely to be determined simultaneously in equilibrium, in which case market liquidity could be a function of firm value rather than the other way around. To identify

the relation between the liquidity of a firm’s stock and firm value, I use a plausibly exogenous source of variation in liquidity: tick-size reductions in U.S. equity markets in 1997 and 2001.² Tick-size reductions are appealing instruments because they not only dramatically impacted stock liquidity, but they also generated cross-sectional variation in liquidity improvements.³ Furthermore, tick-size reductions are unlikely to directly affect firm value through channels other than liquidity. A large literature uses tick-size reductions to exploit exogenous variation in liquidity.⁴

Using tick-size reductions in U.S. equity markets as plausibly exogenous shocks to trading costs, I find that increased liquidity (measured by effective spread) generally has a negligible effect on firm equity value; and if the effect is statistically significant, it is significantly negative. For example, following the 1997 tick-size reduction, the average NYSE- or AMEX-listed stock experienced a reduction in spreads of about 11% and a statistically significant, but economically small, one-month drop in firm equity value of approximately 35 basis points, *ceteris paribus*. Analysis of NASDAQ-listed stocks give similar inference. I next examine the effects of a shock to liquidity due to decimalization. The results are qualitatively similar, except that the relation between change in liquidity and change in firm value of equity for the NYSE/AMEX sample is no longer statistically significant. For robustness, I consider alternative time horizons and find that the results are generally unaltered. Additionally, inclusion of other key event dates, such as the date that the tick-size regime change was announced (rather than implemented) does not affect inference.

Liquidity is a multi-dimensional concept that includes the cost, quantity, and time associated with trading. I primarily use effective spread to proxy for liquidity, which measures the cost

²The 2001 tick-size reduction refers to decimalization. Although a vast majority of stocks converted to decimals in 2001, a small group of pilot firms switched to decimal pricing in 2000. Pilot firms are also included in my sample, but for ease of exposition, I often refer to the conversion to decimals as the “2001 tick-size reduction.”

³Studies documenting the effects of tick-size reductions on liquidity include Goldstein and Kavajecz (2000), Jones and Lipson (2001), Bacidore, Battalio, and Jennings (2003), Bessembinder (2003), and Chakravarty, Wood, and Van Ness (2004).

⁴Fang, Noe, and Tice (2009) use decimalization as an exogenous shock to liquidity to examine the relation between liquidity and Tobin’s Q, a proxy for firm value. Bharath, Jayaraman, and Nagar (2013) use tick-size reductions to measure how increased liquidity impacts the relation between blockholdings and firm value. Edmans, Fang, and Zur (2013) and Gerken (2014) use decimalization to identify the effects of liquidity on blockholders’ decisions or actions. Fang, Tian, and Tice (2014) use tick-size reductions to identify the impact of liquidity on firm innovation.

of a round-trip trade. Although this proxy is commonly used in studies that employ tick-size reductions as instruments, it does not provide a complete measure of liquidity. I attempt to temper this concern by including quoted depth, which is designed to measure the quantity dimension of liquidity, in the analysis. While quoted depth generally falls (liquidity worsens) following tick-size reductions, there is considerable variation in change in depth. Indeed, average quoted depth for NASDAQ-listed stocks following the 1997 tick-size reduction slightly increases (consistent with evidence in Van Ness, Van Ness, and Pruitt (2000)). Importantly, on average, shock to quoted depth does not offset shock to effective spread's effect on firm value. Still, effective spread and quoted depth do not fully measure liquidity, and it is possible that other (latent or difficult to measure) aspects of liquidity provide differential effects on firm value.

In a related study, Fang, Noe, and Tice (2009) find a positive relation between stock liquidity and firm value. They also use exogenous shocks to liquidity to examine the causal impact of trading costs on firm value. However within this general framework, I make several different empirical design choices. First, I measure liquidity using effective bid-ask spread, rather than relative effective spread (sometimes referred to as percent effective spread), which is effective spread scaled by the bid-ask midpoint (or true price). I do not use relative effective spread because scaling the liquidity measure by contemporaneous price induces a relationship between liquidity and concurrent firm market value by construction: holding effective spread constant, an increase in price causes relative effective spread to decrease and firm value to increase. Second, I measure change in firm value using stock returns instead of Tobin's Q.⁵ Third, I compare changes in liquidity and firm value in the months, rather than the years, surrounding shocks to liquidity because using a longer window increases the likelihood that liquidity or firm value change for reasons other than an exogenous shock to trading costs. Fourth, I examine multiple shocks to liquidity instead of focusing solely on decimalization.

⁵There is a tradeoff between using stock returns versus Tobin's Q to measure change in firm value. Stock returns measure the change in market value of equity, while Tobin's Q is designed to measure total (debt plus equity) firm value. However, many consider Tobin's Q to be a noisy proxy for firm value as it may measure other aspects of the firm, such as growth opportunities (see e.g., Roberts and Whited (2013) for a discussion of measurement error and valuation ratios). Therefore, I use stock returns to measure change in firm (equity) value following a shock to liquidity.

I show that the choice of bid-ask spread measure is particularly important in explaining differences between my results and previous findings. I find that a decrease in effective spread (or increase in liquidity) has a negligible, if not negative, impact on firm value, while a decrease in relative effective spread, where spread is scaled by concurrent price, has a significantly positive impact on firm value. However, if I slightly alter the definition of relative effective spread by scaling spread with the stock price from the preceding period, I find that liquidity has a marginal or negative effect on firm value. This result suggests that it is not the choice of effective spread versus relative effective per se that leads to different inference, rather the date used to obtain the stock price deflator primarily determines the differences in results.

This study's primary contribution is to provide new evidence on the impact of liquidity on firm value. Using tick-size reductions in U.S. equity markets as plausibly exogenous shocks to trading costs, I find that increased liquidity has a small impact on firm equity value on average, and when the relation is significant, it is significantly negative. Additionally, I show that my results primarily differ from prior findings due to the choice of liquidity measure. I provide evidence that scaling trading costs by current stock price induces a mechanical relation between shock to liquidity and change in firm value. Overall, my results are consistent with theory suggesting that liquidity has a second-order effect on firm value (see e.g., Constantinides (1986) and Vayanos (1998)), or with arguments that liquidity can, in some instances, negatively impact the firm through such channels as reduced blockholder activism (Back, Li, and Ljungqvist (2015)) or reduced innovation output (Fang, Tian, and Tice (2014)). Since the topic of stock liquidity and firm value is relevant for firm management tasked with maximizing firm value and policymakers and regulators who shape liquidity through financial market laws and regulations, this subject is of interest to a broad audience. The JOBS Act of 2012 illustrates the importance of this issue as it required the U.S. Securities and Exchange Commission (S.E.C.) to examine the effects of decimalization on firms and to implement a change in minimum price variation for certain groups of stocks if deemed appropriate.⁶

⁶Although the S.E.C. initially recommended to not increase tick sizes, in May of 2015, it finalized a pilot program in which small-cap stocks will trade in increments of five cents for two years beginning in 2016. See <http://www.sec.gov/divisions/marketreg/tick-size-pilot-plan-final.pdf>

1.2 Empirical Design

1.2.1 Methodology

It is challenging to determine empirically the effect of market liquidity on firm value due to the endogenous nature of liquidity. To illustrate, consider the following linear model:

$$V_i = \alpha + \beta L_i + \epsilon_i, \quad (1.1)$$

where V_i denotes firm value, L_i represents liquidity, ϵ_i represents an error term, and α and β represent the parameters to be estimated.

To consistently estimate β , the liquidity measure needs to be exogenous, or orthogonal to the error term. If liquidity is endogenous, the estimate for β is unidentified and we cannot determine the causal impact of liquidity on firm value. As pointed out by Fang, Noe, and Tice (2009), liquidity in equation (1.1) is likely to be endogenous. For example, it is plausible that liquidity and firm value are both determined by an unobservable omitted variable, such as manager ability. Additionally, liquidity and firm value may be determined simultaneously in equilibrium, in which case market liquidity could be a function of firm value rather than the other way around.

To overcome these identification issues, I employ an empirical design that isolates exogenous variation in liquidity. Specifically, I use tick-size reductions in U.S. equity markets as plausibly exogenous shocks to liquidity. The major U.S. stock exchanges reduced minimum price variation from eighths to sixteenths in 1997 and further reduced tick size to decimals in 2000 and 2001. Tick-size reductions dramatically impacted stock liquidity and are unlikely to directly affect firm value through channels other than liquidity. As noted in the introduction, a number of other studies use changes in minimum price variation to exploit exogenous variation in liquidity; these studies include Fang, Noe, and Tice (2009), Bharath, Jayaraman, and Nagar (2013), Edmans, Fang, and Zur (2013), Fang, Tian, and Tice (2014), and Gerken (2014).

Using tick-size reductions as shocks to liquidity, I evaluate the impact of liquidity on firm value by estimating the following model:

$$\Delta V_i = \alpha + \beta \Delta L_i + \delta' C_i + \epsilon_i, \quad (1.2)$$

where ΔL_i represents the change in liquidity due to the tick-size reduction, C_i represents a vector of firm control variables, and ΔV_i denotes the change in firm value. I describe how I measure these variables in the next section.

Since tick-size reductions are applied equally to the entire sample of firms, one concern is that these events do not create exogenous variation in liquidity. However as I discuss below when describing sample characteristics, reductions in minimum-price variation generate significant cross-sectional variation in liquidity improvements. Some firms experience large increases in liquidity while others' liquidity hardly changes, which allows comparison of the effects of liquidity on firm value across firms. Still, it is possible that shock to liquidity is not randomly assigned across firms. It is therefore paramount that the vector of control variables is included in Equation (1.2) to control for differences in observed firm characteristics. Nevertheless, it is possible that an unobserved variable affects liquidity shock assignment. However, this unobservable would also have to be correlated with change in firm value to affect identification. To reduce concerns about proper identification, I measure change in firm value over relatively short windows (as short as one day in robustness tests) to decrease the likelihood that change in firm value is systematically affected by variables other than an exogenous shock to liquidity. Additionally, analysis of multiple tick-size reductions provides robustness because even if an omitted variable that is correlated with both shock to liquidity and change in firm value exists for one of the tick-size regime changes, it is less likely that the omitted variable persists for both events.

1.2.2 Variable construction

It can be difficult to precisely measure firm value and liquidity. Tobin’s Q is commonly used to proxy for firm value. However, valuation ratios are subject to measurement error as they may proxy for investment opportunities or something else other than firm value (see e.g., Roberts and Whited (2013) for a discussion of measurement error and valuation ratios). To address these concerns, I measure change in firm value as the change in stock price. Specifically, I use stock returns to measure the price reaction associated with a liquidity shock. I primarily measure stock returns using compound raw returns, but I obtain qualitatively similar results if I use cumulative raw returns instead.

Liquidity is also difficult to measure as it is a multi-dimensional concept that incorporates the cost, quantity, and time associated with trading. It is common to proxy for stock liquidity using bid-ask spread (see e.g., Fang, Noe, and Tice (2009)), and I primarily measure liquidity with effective spread, which measures the cost of a round-trip trade.⁷ Effective spread is defined as

$$ESPRD_{i,k} = 2|P_k - M_k|. \quad (1.3)$$

Effective spread for the i th firm’s k th trade is twice the absolute value of the difference between the trade price and the prevailing bid-ask midpoint. To construct effective spread, I follow the procedures described in Chordia, Roll, and Subrahmanyam (2001), Fang, Noe, and Tice (2009), and Fang, Tian, and Tice (2014). Prior to 1999, I match each trade with the prevailing quote five seconds before the trade. Beginning in 1999, I match each trade with the prevailing quote that immediately precedes the trade. Additionally, I drop trades occurring out of sequence, trades recorded outside of normal trading hours, and trades with special settlement conditions. Also, I require that the ask quote is no more than 1.5 times greater than the bid quote and that quote conditions are normal, which excludes quotes associated with trading halts.⁸ I construct each firm’s daily spread as the value-weighted average spread across trades, where each trade

⁷As described below, I also consider a measure of depth, which is designed to measure the quantity associated with trading.

⁸Specifically, I retain TAQ quotes with the following quote conditions: 1, 2, 3, 6, 10, and 12.

is weighted by its trade size in dollars.⁹ I compute each firm’s monthly effective spread as the equally weighted average of daily spreads.¹⁰

To compute the shock to liquidity due to the reduction in minimum-price variation, I construct each firm’s percentage change in effective spread ($\Delta ESPRD$) from the month before the tick-size reduction (month $t - 1$) to the month in which the tick-size reduction occurred (month t), where month $t - 1$ contains the four calendar weeks immediately preceding the tick-size reduction and month t begins on the date of the tick-size-reduction implementation and ends four calendar weeks later. I compute the percentage change in effective spread, rather than simply the change in spread, to mitigate concerns about differences in scale across firms. To mitigate the effects of outliers, I winsorize percentage change in effective spread at 100%.

In addition to effective spread, I also construct a depth measure, which is designed to proxy for the quantity associated with trading. It is important to consider depth because there is evidence that depth fell following tick-size reductions (see e.g., Goldstein and Kavajecz (2000) and Bessembinder (2003)), although others find that depth increased in some cases (Van Ness, Van Ness, and Pruitt (2000)). I measure quoted depth as the total number of shares quoted at the best bid and the ask, divided by two. To construct each firm’s daily quoted depth, I compute weighted-average depth, where each quote is weighted by the amount of time that the quote was in force. I compute each firm’s monthly quoted depth as the equally weighted average of daily depth values. Finally, to examine how quoted depth changed surrounding tick-size reductions, I construct each firm’s percentage change in depth ($\Delta DEPTH$) from the month before the tick-size reduction to the month in which the tick-size reduction occurred. As with effective spread, I winsorize percentage change in depth at 100%.

In addition to market liquidity and price reaction, I construct several control variables. I construct measures that prior research suggests are related to firm equity value. Firm size ($SIZE$) is the natural log of market value of equity, and market-to-book (MB) is the natural log of market value of equity divided by book value of equity (Compustat item CEQ). See e.g.,

⁹For robustness, I also construct average daily spreads where each trade is equally weighted, and find qualitatively similar results.

¹⁰I require at least 15 days of available observations per month.

Fama and French (1992, 1993) for evidence on the relation between *SIZE* and *MB* and asset prices. I measure profitability as return-on-assets (*ROA*), which is defined as operating income before depreciation (Compustat item EBITDA) over book assets (Compustat item AT). I also construct investment (*INVST*), defined as the percentage change in annual book assets. See e.g., Hou, Xue, and Zhang (2014) and Fama and French (2015) for discussion and evidence on the effects of profitability and investment on asset prices. I construct a momentum proxy (*MOM*) designed to capture the effects of past stock performance on firm value. This measure is the natural log of one plus compound stock returns, where returns are constructed over a three month period prior to the tick-size reduction (see e.g., Jegadeesh and Titman (1993) for evidence on the effect of past returns on firm equity value). I also construct effective spread level (*ESPRD*), which is an equally weighted average of daily effective spread, averaged over a three month window prior to the tick-size reduction.

Since construction of *ROA* and *INVST* relies on financial statements data that are only available at an annual or quarterly frequency, I construct these measures as of the fiscal year-end preceding the tick-size reduction. *SIZE*, *MOM*, and *ESPRD* are computed as of the day prior to the announcement date for the 1997 tick-size reduction and as of the period (i.e., one month) prior to implementation of tick-size reduction for decimalization. Construction of *MB* relies on both equity market data and financial statements data. Market value is computed on the same day that *SIZE* is computed and book value is computed as of the preceding fiscal year end.

I collect intra-day trade and quote data from TAQ, stock price data from CRSP, and financial statements data from COMPUSTAT. The sample consists of firms at the intersection of these three data sets. I require that the sample firms be U.S. common stocks listed on the NYSE, AMEX, or NASDAQ. I choose to separate NYSE/AMEX listings from NASDAQ listings due to differences in institutional details and attributes of the listings across the exchanges.¹¹ The final sample of stocks for the 1997 tick-size reduction consists of 1,723 NYSE- and AMEX-

¹¹It is common to perform separate analyses on a NYSE and AMEX sample and a NASDAQ sample (see e.g., Asparouhova, Bessembinder, and Kalcheva (2010)). Results are unchanged if I further separate NYSE listings from AMEX listings.

listed stocks and 1,468 NASDAQ-listed stocks. The final sample for the 2001 tick-size reduction includes 1,565 NYSE and AMEX stocks and 2,352 NASDAQ stocks. I winsorize all measures at the 1st and 99th percentile each year. I elaborate further on sample characteristics below when I discuss the analysis of the 1997 and 2001 tick-size reductions, respectively.

1.3 Examining the Effect of Liquidity on Firm Value using Tick-Size Reductions in U.S. Equity Markets

1.3.1 Tick-size reductions in U.S. equity markets: The setting

In the mid- to late-1990s, U.S. stock exchanges began facing pressure to reduce minimum price variation. At the time, U.S. stock exchanges quoted prices in increments of eighths, but policymakers and some in the investment community were calling for reduced tick size to sixteenths or to decimals. Appendix A presents a timeline of key dates for tick-size reductions in U.S. equity markets. On March 13, 1997, Congressman Michael Oxley introduced a bill to the U.S. House of Representatives directing the SEC to require decimal pricing for all equity securities. On the same day, the AMEX's board of directors approved a change in minimum price variation to sixteenths. Later that month, NASDAQ's board made a similar announcement. AMEX and NASDAQ implemented a tick-size reduction from eighths to sixteenths on May 7, 1997 and June 2, 1997, respectively. On June 5, the NYSE announced plans to convert to decimal pricing as soon as the essential systems were in place. As an interim step, the NYSE approved a switch from an eighths regime to a sixteenths regime, a change that it implemented on June 24.

Later in 1997, the exchanges decided to delay conversion to decimal pricing until 2000 so that the industry could first focus on year 2000 (Y2K) issues. By August of 2000, the exchanges began phasing in decimal pricing for groups of pilot firms. On January 29, 2001, the NYSE and AMEX fully converted to decimal pricing and the NASDAQ followed suit on April 9, 2001.

Since tick-size reductions in U.S. equity markets dramatically impacted stock liquidity and generated cross-sectional variation in liquidity improvements, I use these events to examine how

a shock to liquidity affects firm value. Although the shock to liquidity occurs on the day that the exchange implements a tick-size reduction, it is possible that the market reaction occurred on the day that the change in minimum price variation was announced rather than implemented. I therefore examine price reactions over various windows following the implementation and announcement dates. As the timeline in Appendix A suggests, it is relatively straightforward to determine announcement and implementation dates for the 1997 tick-size reduction. In each case the exchange’s board of directors announced plans for a tick-size reduction and the exchange implemented the change a short time later. However, determining announcement dates for decimalization is more problematic. For example, the NYSE announced plans to adopt decimal pricing on the same day it announced plans to switch to a sixteenths regime, which was nearly three and a half years before the NYSE actually implemented decimal pricing. Since the appropriate announcement date is unclear or confounded by other events, in the analysis surrounding decimalization, I primarily focus on effects following the implementation date.

1.3.2 Exploiting the 1997 tick-size reduction

Table 1.1 reports descriptive statistics for NYSE/AMEX (Panel A) and NASDAQ (Panel B) stocks around the 1997 tick-size reduction. Consistent with prior findings, effective spreads decreased on average following the tick-size regime change, as the mean change was -11.123% for NYSE and AMEX stocks and -5.530% for NASDAQ stocks. Additionally, the percentage change in effective spread varies considerably across firms. $\Delta ESPRD$ ’s standard deviation for NYSE and AMEX stocks is 21.503%, and its values at the 25th and 75 percentiles are -24.118% and -0.593%, respectively. $\Delta ESPRD$ for NASDAQ stocks exhibits similar variation.¹² The existence of cross-sectional variation in change in effective spread is important, as it facilitates analysis of the effects on firm value from a large shock to trading costs compared to effects on firm value from a small shock to trading costs, *ceteris paribus*.

¹²Table 1.1 reveals that spreads following the tick size regime change from eighths to sixteenths increased for some firms. For robustness, I repeat the analysis after removing firms that experienced an increase in spreads following tick-size reductions and I obtain qualitatively similar results.

Table 1.1 also reports summary statistics for change in quoted depth ($\Delta DEPTH$) surrounding the 1997 tick-size reduction. As with change in effective spread, $\Delta DEPTH$ exhibits considerable cross-sectional variation. Additionally, quoted depth fell on average for NYSE and AMEX stocks, but slightly increased on average for NASDAQ stocks (consistent with findings in Van Ness, Van Ness, and Pruitt (2000)). The summary statistics suggest that although lower depth may have at least partially offset liquidity gains from reduced spreads for some sample firms, it is plausible that depth did not erode liquidity gains from reduced spreads for a large number of sample firms. Finally, Table 1.1 reveals that stocks in the NYSE/AMEX sample are larger, have fewer growth opportunities (as proxied by MB), are more profitable, and have lower effective spread levels than NASDAQ sample. These differences in firm characteristics stress the importance of analyzing the samples separately.

In an initial examination of the impact of a shock to effective spread on firm value, I sort firms into terciles based on the percentage change in effective spread ($\Delta ESPRD$) from the month before tick-size reduction to the month in which the firm's tick-size reduction occurred. Since spread is inversely related to liquidity, a large spread reduction is equivalent to a large liquidity increase. Stocks in tercile 1 are those with the largest spread reduction, while stocks in tercile 3 are those with the smallest spread reduction (or spread increase in some cases). Both Figure 1.1 (NYSE and AMEX sample) and Figure 1.2 (NASDAQ sample) illustrate the impact that the regime change from eighths to sixteenths had on spreads. Trade day 0 in the figures represents the day that the stock's exchange implemented the tick-size reduction. In both figures, stocks in tercile 1 experienced a dramatic, immediate drop in effective spreads, while stocks in tercile 3 experienced an increase in spreads. Stocks in the middle tercile experienced a more gradual reduction in spreads.

Tables 1.2 and 1.3 present the sorting results. Panel B in the tables presents mean and median price reaction at key event dates, sorted by the percentage change in effective spread. I use compound stock returns to measure the change in firm value. Since the liquidity shock occurs on the date that the tick-size reduction is implemented, I primarily focus on the price reaction over various windows following the regime change implementation date. However, since it's

possible that market participants price in the effects of the reduced tick size at announcement, I also examine price reaction following announcement and over a window that combines the implementation and announcement dates.

The results provide univariate evidence that liquidity (measured using effective spread) is marginally, if not negatively, related to firm value. The results are similar for both NYSE/AMEX and NASDAQ stocks. Firms with the largest reduction in spreads (tercile 1) experienced the smallest one-month price reaction following the tick-size reduction implementation date, while firms with the smallest reduction or a slight increase in spreads (tercile 3) experienced the largest one-month price reaction following the implementation date. This pattern of results does not reverse at three- and six-month horizons. I next consider price reactions over a three-day and one-week window following the tick-size reduction announcement. Average returns are relatively flat following the announcement date, providing evidence that announcement returns do not offset implementation date returns. The last column of Panel B in the tables makes this point explicit by combining returns over the three-day announcement window with returns over the one month implementation window. Since announcement returns are relatively flat across terciles, the returns from the implementation window drive the patterns of returns for the combined window. Price reaction over the combined window is generally negatively related to shock to liquidity.

Panel A in Table 1.2 and Table 1.3 presents means and medians of key firm measures, sorted by the percentage change in spread. Quoted depth falls (liquidity worsens) for NYSE and AMEX stocks following the tick-size reduction, however the change in depth is relatively flat across terciles. This suggests that the change in depth is not strongly related to change in effective spread and that I can control for change in depth while still exploiting variation in effective spread. Furthermore, quoted depth slightly improves for NASDAQ stocks, and this improvement is strongest for firms with the largest spread reduction (tercile 1). Thus, for

NASDAQ stocks, change in depth does not offset the liquidity gain from a spread reduction and, if anything, improves liquidity further.¹³

Panel A also report means and medians for other key firm characteristics. While some firm characteristics, such as firm size, appear to be related to $\Delta ESPRD$, most of the measures are relatively flat across groups. For example, log market-to-book and investment do not appear to be strongly related to change in effective spread. These results suggest that key observable firm characteristics are generally not highly correlated with $\Delta ESPRD$. Nevertheless, to reduce concerns about omitted variable bias, I next examine the impact of a shock to liquidity on firm value, while controlling for firm characteristics that could be related to value. To do so, I estimate the following regression model:

$$RET_i = \alpha + \beta \Delta ESPRD_i + \delta' C_i + \epsilon_i, \quad (1.4)$$

where RET_i represents stock i 's return following either the implementation date, the announcement date, or both, $\Delta ESPRD_i$ represents stock i 's percentage change in spreads due to the tick-size reduction, and C_i represents a vector of control variables described in Section 1.2.2. I present regression estimates of equation (1.4) in Table 1.4; Panel A includes regression estimates for NYSE and AMEX stocks and Panel B includes those for NASDAQ stocks. I examine the price reaction to the shock to liquidity over a one-, three-, or six-month window following tick-size implementation, over a three-day or one-week window following tick-size announcement, and over a period that combines the three-day price reaction following announcement with the one-month price reaction following implementation. For each return window, I estimate a univariate regression model, where δ is set to zero, and a multivariate regression model that includes all of the controls. I present estimated slope coefficients and the associated t -statistics in parentheses.

¹³In unreported results, I compute Pearson correlation coefficients. Consistent with the patterns observed in the sorting results, the correlation between $\Delta ESPRD$ and $\Delta DEPTH$ is low for the NYSE/AMEX sample (correlation of -0.0418) and is negative for the NASDAQ sample (correlation of -0.180).

The regression estimates confirm the evidence from the univariate sorting results. Although the relation between change in liquidity (as proxied by change in effective spread) and change in firm value is negative, the estimated effect for the average firm is economically small. For example, in column two of Panel A, the estimated slope coefficient on $\Delta ESPRD$ is 0.032. This suggests that the average firm in the NYSE/AMEX sample experienced a reduction in spread of 11.123% (from Table 1.1), but a one-month drop in firm equity value of only 35 basis points, *ceteris paribus*. Three-month and six-month returns are also generally insignificant. Additionally, the price reaction following the tick-size reduction announcement is insignificantly related to $\Delta ESPRD$, suggesting that the market did not significantly react to the tick-size regime change announcement date. The effects for the average firm in the NASDAQ sample, while larger, are also economically small.

I also include change in quoted depth ($\Delta DEPTH$) in the regression specifications. Consistent with the effects of shock to effective spread, change in quoted depth is generally associated with an insignificant price reaction following tick-size implementation. The estimated coefficients on $\Delta DEPTH$ are close to 0 for the NYSE/AMEX sample and are larger in magnitude but negative for the NASDAQ sample. The negative estimated coefficient on change in depth suggests that increased liquidity (increased depth) is associated with a negative change in firm value.

In summary, shock to liquidity due to the 1997 tick-size reduction has a small effect on firm value on average, and when the effect of increased liquidity is significant, it is significantly negative. These results are consistent with the small but growing theoretical and empirical literature that suggests that the effects of liquidity on the firm are marginal or, in some cases, negative.

1.3.3 Exploiting the 2001 tick-size reduction

Table 1.5 provides descriptive statistics of key measures around the shift to decimalization. The shift to decimalization had an even larger impact on spreads than did the 1997 tick-size reduction. Average spread reduction is over 18% for NYSE/AMEX stocks and over 21%

for NASDAQ stocks. As with the 1997 tick-size reduction, the effect of decimalization on spreads varies considerably in the cross-section. Quoted depth also fell dramatically following decimalization, and in contrast to the 1997 tick-size reduction, depth fell on average for both the NASDAQ and NYSE/AMEX samples. This suggests that the effects of quoted depth on overall liquidity and on firm value should be considered in the analysis that follows.

As with the analysis of the 1997 tick-size reduction, I sort firms into terciles based on the percentage change in effective spread from the month before decimalization to the month in which decimalization occurred. Figure 1.3 and Figure 1.4 illustrate that decimalization had a significant impact on effective spreads and that the impact varied considerably in the cross-section. Spreads dramatically drop following the decimalization implementation date (trade day 0) for stocks in tercile 1 and tercile 2, while spreads for stocks in tercile 3 are relatively flat.

Table 1.6 reports univariate sorting results. Panel A reports results for NYSE/AMEX stocks and Panel B presents results for NASDAQ stocks. The last three columns report the change in firm value of equity over various windows following decimalization.¹⁴ I focus on stock returns following the decimalization implementation date rather than the announcement date because, as can be seen in the timeline in Appendix A, the decimalization announcement date is either unclear or confounded by other events.¹⁵

There is not a strong univariate relation between change in effective spread following decimalization and change in one-month firm equity value. One-month change in firm value of equity is relatively flat across the terciles. This relation, however, strengthens at longer horizons, particularly for the six-month window in the NASDAQ sample. NASDAQ stocks in tercile 1 experience a change in firm value of equity of -6.407% over the six months following decimalization, while those in tercile 3 experience six-month returns of 7.380%. However, longer run returns should be interpreted with caution as it is more likely that returns at longer horizons are influenced by factors other than the exogenous shock to liquidity. Indeed, decimalization occurred during what is often referred to as the “technology bubble,” which was a time of high

¹⁴Note that one- and three-month returns in the NASDAQ sample are abnormally large. This is due to the high volatility experienced by NASDAQ stocks during the sample period.

¹⁵For robustness, I consider price reactions at the various possible announcement dates and the results that follow are generally unchanged.

volatility for NASDAQ issues. It is therefore possible that stocks in the NASDAQ sample were systematically impacted by events other than decimalization.

Table 1.6 also reports means and medians of other key firm characteristics. Quoted depth falls following decimalization. However, change in depth ($\Delta DEPTH$) is relatively flat across the $\Delta ESPRD$ terciles. In unreported results, the Pearson correlation between $\Delta ESPRD$ and $\Delta DEPTH$ is 0.040 for the NYSE/AMEX sample and -0.035 for the NASDAQ sample. The low correlation suggests that it is possible to hold $\Delta DEPTH$ constant while exploiting variation in $\Delta ESPRD$, and vice versa. Other firm characteristics such as *SIZE* and *MOM* vary across terciles.¹⁶ In an attempt to control for differences in key firm characteristics, I next conduct a regression analysis to estimate the impact of a shock to liquidity due to decimalization on firm value, holding key observables constant.

Table 1.7 presents regression estimates. I estimate the same regression model described in Section 1.3.2. I examine the price reaction to the shock to liquidity over a one-, three-, or six-month window following decimalization. For each return window, I estimate a univariate regression model and a multivariate regression that includes controls. The results are consistent with earlier evidence: the relation between change in liquidity and change in firm value of equity is either insignificant or significantly negative. For the NYSE/AMEX sample, the estimated slope coefficient of 0.26 on $\Delta ESPRD$ in column two suggests that the average firm experienced a spread reduction of about 18% and a one-month drop in firm value of equity of only 47 basis points. With a *t*-statistic of 1.32, this effect is insignificantly different from zero. The estimated effect for the NASDAQ sample is larger and statistically significant. The estimated slope coefficient of 0.045 on $\Delta ESPRD$ in column eight suggests that the average NASDAQ-listed stock experienced a drop in spreads of more than 21% and a one-month drop in firm value of equity of 0.98%.¹⁷ Results are qualitatively similar at longer horizons.

¹⁶The past stocks returns measure (*MOM*) is particularly variable for NASDAQ-listed stocks. In untabulated results, I find qualitatively similar results if *MOM* is excluded from the regression analysis that follows.

¹⁷Proper identification of the estimated effect of liquidity on firm value assumes that an unobservable does not exist that is correlated with both shock to liquidity and change in firm value. Given the volatility associated with NASDAQ stocks during the decimalization sample period, it is possible that this assumption is violated. I therefore view results from the NASDAQ-decimalization analysis as a complement to the other results in the study.

I next consider the effect of change in depth ($\Delta DEPTH$) on change in firm value. An increase in depth (increase in liquidity) tends to have an insignificant effect on firm value, and when the effect is significant it is significantly negative. This effect is consistent with previous results.

The analysis in this section, which uses decimalization as a shock to liquidity, provides similar results to those obtained from the analysis using the 1997 tick-size reduction. An increase in liquidity generally has an insignificant impact on firm value, and, if anything, the impact is negative. The results continue to support the literature that suggests that liquidity has a negligible, or in certain cases, negative impact on the firm.

1.4 Additional Analyses

1.4.1 Comparison of results to prior findings

In this section, I examine why I find that liquidity has a marginal or, in some instances, negative effect on firm value, while a related work by Fang, Noe, and Tice (2009) shows that liquidity positively impacts firm value. Fang, Noe, and Tice (2009) also use a tick-size reduction in U.S. equity markets to exploit plausibly exogenous variation in liquidity. However, I make several different empirical design choices. First, I use effective spread rather than relative effective spread, which is effective spread scaled by the bid-ask midpoint (or true price), to proxy for liquidity. Second, I use stock returns instead of change in Tobin's Q to measure the change in firm value following tick-size reduction. Third, I compare changes in liquidity and firm value in the months, rather than the years, surrounding shocks to liquidity. Fourth, I use the 1997 and 2001 tick-size reductions instead of focusing solely on decimalization. In this section, I examine how these different research design choices impact the results.

As before, I examine the effect of liquidity on firm value by estimating the following model:

$$\Delta V_i = \alpha + \beta \Delta L_i + \delta' C_i + \epsilon_i, \quad (1.5)$$

where ΔL_i represents firm i 's change in liquidity due to the tick-size reduction, ΔV_i denotes firm i 's change in firm value, and C_i represents a vector of firm control variables.

I consider two proxies for change in firm value: stock returns following tick-size reduction and change in Tobin's Q . Stock returns (RET), as described in Section 1.2.2, are compound returns over the period following tick-size reduction. I follow Fang, Noe, and Tice (2009) and construct Tobin's Q as book assets (Compustat item AT) plus market equity (market price (PRCC) multiplied by shares outstanding (CSHO)) minus book equity (CEQ) minus deferred taxes (TXDB), all over book assets. Following Fang, Noe, and Tice (2009), I measure change in Tobin's Q (ΔQ) from the fiscal year prior to tick-size reduction to the fiscal year after tick-size reduction. Since Q is constructed using annual financial statement data, I cannot construct it at a monthly frequency. I can, however, construct stock returns for both the month and fiscal year following tick-size reduction.

I also consider multiple liquidity proxies. I primarily focus on variants of the bid-ask spread. As described in Section 1.2.2, I define effective spread ($ESPRD$) as twice the absolute value of the difference between the trade price and the prevailing bid-ask midpoint. In addition, I construct two versions of relative effective spread. I construct the first, $RESPRD_1$, following Fang, Noe, and Tice (2009). It is defined as

$$RESPRD_1 = \frac{2|P_k - M_k|}{M_k}. \quad (1.6)$$

Relative effective spread for each firm's k th trade is twice the absolute value of the difference between the trade price and the prevailing bid-ask midpoint, scaled by the prevailing midpoint. I apply the same trade-quote matching procedure and filters used to construct $ESPRD$. Daily $RESPRD_1$ is the value-weighted average spread across trades, where each trade is weighted by dollar trade size. Monthly or yearly $RESPRD_1$ is the equal-weighted average of daily spreads over the respective period.

Alternatively, I construct a second relative effective spread measure, $RESPRD_2$, defined as monthly or annual $ESPRD$, scaled by stock price from the period before tick-size reduction.¹⁸ This convention deflates spread by price, but removes the potential of creating a relationship between change in relative effective spread and change in concurrent firm value by construction. I compute the change in each of the three spread measures from the fiscal year (month) prior to tick-size reduction to the fiscal year (month) following tick-size reduction.

I employ the same control variables described in Section 1.2.2.¹⁹ For the analysis in which change in liquidity and change in firm value are constructed at an annual horizon, the control variables are constructed as of the year-end preceding the period in which the liquidity and firm value proxies are computed. For brevity, I report results using only decimalization as an exogenous shock to liquidity. However, results are qualitatively similar for the 1997 tick-size reduction. I present results associated with decimalization to offer a direct comparison to the findings in Fang, Noe, and Tice (2009).

Table 1.8 presents regression estimates of equation (1.5) using the alternative variable definitions. Results are separated into a NYSE/AMEX sample and a NASDAQ sample. Since results are qualitatively similar across these samples, I primarily focus my discussion to stocks listed on the NYSE or AMEX. Panel A presents regression estimates using annual change in liquidity and firm value. The first column presents estimates from a regression of stocks returns over the fiscal year following decimalization on the change in $ESPRD$ from the fiscal year prior to the fiscal year following decimalization. Consistent with results in Section 1.3, the coefficient on $\Delta ESPRD$ is positive. It is also significant, suggesting that measuring liquidity and firm value over longer horizons magnifies this relation.

Column two presents regression estimates of a specification that is identical to that in column one, except $\Delta RESPRD_1$ replaces $\Delta ESPRD$. Recall that $\Delta RESPRD_1$ is constructed

¹⁸The date used to obtain the stock price deflator depends on whether $RESPRD_2$ is constructed at a monthly or annual frequency. If constructed at a monthly frequency, the stock price deflator is obtained one month (four calendar weeks) before the tick-size reduction. If $RESPRD_2$ is constructed at an annual frequency, the stock price deflator is obtained from the fiscal year-end preceding the period in which change in $ESPRD$ is constructed.

¹⁹The lone exception is I do not include momentum (MOM) in the analysis where change in liquidity and firm value are measured at an annual frequency. This is because MOM is designed to measure the effects of stock performance from the preceding period, but in some cases MOM is nearly three years removed from dates used in the construction of change in liquidity and firm value proxies.

such that effective spread is deflated by concurrent midpoint, or true price. The results suggest that substituting $\Delta RESPRD_1$ for $\Delta ESPRD$ is far from innocuous. The estimated coefficient on $\Delta RESPRD_1$ is negative and significant, suggesting that, in contrast to earlier findings, increased spread (reduced liquidity) has a negative impact on firm value. However, this result is sensitive to the definition of relative effective spread. Column three presents estimates of a model that replaces $\Delta RESPRD_1$ with $\Delta RESPRD_2$. Recall that $\Delta RESPRD_2$ deflates change in effective spread using prior-period stock price rather than concurrent price. The estimated coefficient on $\Delta RESPRD_2$ flips back to positive and significant, consistent with the estimate on $\Delta ESPRD$. This result suggests that it is not the use of relative effective spread per se that leads to different inference, rather the date used to measure the stock price deflator is critical. Deflating effective spread by concurrent price can lead to a mechanical relation between change in liquidity and change in firm equity value.

Columns four through six report regression estimates of models that are identical to the first three columns except change in Tobin's Q replaces stock returns as the dependent variable. The results are qualitatively similar. The estimated coefficients are positive and significant on $\Delta ESPRD$ and $\Delta RESPRD_2$ and are negative and significant on $\Delta RESPRD_1$. This result suggests that the choice of RET versus ΔQ does not dramatically impact the results.

Panel B of Table 1.8 presents regression estimates using monthly, rather than yearly, stock returns and change in liquidity. The same pattern of results observed in Panel A, also holds in Panel B. The estimated coefficients on both $\Delta ESPRD$ and $\Delta RESPRD_2$ are consistent with one another, as both suggest that shock to liquidity has an insignificant effect on firm value. In contrast, the estimated coefficients on $\Delta RESPRD_1$ suggest a different result, a negative and statistically significant effect on firm value.²⁰ The effects of a shock to liquidity on firm value tend to be tempered when using monthly data instead of yearly data. The estimated effect of liquidity on firm value measured with annual data (Panel A) is significant in every case, while the estimated effect using monthly data tends to be insignificant (although, the estimated coefficient on $\Delta RESPRD_1$ is always significant and negative).

²⁰Note that the estimates in columns one and four are taken from results in Table 1.7. I present these estimates again for ease of reference.

In summary, the results in Table 1.8 suggest that the choice of liquidity proxy is particularly important in explaining differences between my results and previous findings. Inference is dramatically impacted when $\Delta ESPRD$ rather than $\Delta RESPRD_1$ is used to proxy for a shock to liquidity. The firm value proxies considered do not dramatically alter results. Inference is unchanged if RET instead of ΔQ is used to measure change in firm value. Measuring change in liquidity and firm value at yearly rather than monthly horizons can also impact results. Estimated effects are stronger at longer horizons, although estimated effects tend to have the same sign regardless of horizon. Since estimation over longer windows increases the likelihood that liquidity or firm value change for reasons other than an exogenous shock to spreads, I interpret the results using annually constructed measures with caution.

1.4.2 Robustness tests

In addition to the results reported above, I conduct a number of robustness tests. For brevity, I briefly discuss the results rather than present additional tables. With the exception of the analysis conducted in Section 1.4.1, I rely on the one-month change in effective spread to proxy for a shock to liquidity. For robustness, I consider one-day, three-day, and one-week windows. For example, I regress one-day stock returns at tick-size implementation on the one-day change in effective spread, plus controls. I conduct analogous tests at three-day and one-week horizons. It is useful to consider shorter windows to reduce the likelihood that an unobservable drives both shock to liquidity and change in firm value. The robustness test results are consistent with results presented in the study. Regardless of horizon, change in effective spread has an insignificant impact on firm value, and the estimated effect of shock to liquidity is generally negative.

1.5 Concluding Remarks

This study provides new evidence on the effects of stock liquidity on firm value. Using tick-size reductions in U.S. equity markets to exploit plausibly exogenous variation in stock liquidity,

I find that increased liquidity has a negligible or, in some instances, negative effect on firm equity value. For both the 1997 and 2001 tick-size reductions, the relation between change in liquidity and one-month change in firm value is negative. However, the average estimated effect is generally economically small. These effects are unaltered at alternative horizons and are unchanged by the inclusion of other key event dates, such as the date that the tick-size regime change was announced.

This study's results are in contrast to the prevailing view that liquidity positively impacts firm value. I examine why my results differ from prior findings, and I show that the choice of liquidity measure primarily explains the differences in results. Overall, my findings are consistent with a small but growing literature that argues that increased liquidity has a small or, in certain cases, negative impact on the firm. This topic is of importance to a broad audience as it has implications for a number of stakeholders including firm management tasked with maximizing firm value and policymakers and regulators who shape liquidity through financial market laws and regulations.




Figure 1.1: Effective spreads of NYSE and AMEX stocks surrounding 1997 tick-size reduction
This figure presents equally weighted average daily effective spreads of NYSE- and AMEX-listed firms surrounding the 1997 tick-size reduction from eighths to sixteenths. The implementation date of the tick-size reduction is denoted as trade day 0. Firms are sorted into terciles based on the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Tercile 1 contains stocks with the largest spread reduction.




Figure 1.2: Effective spreads of NASDAQ stocks surrounding 1997 tick-size reduction

This figure presents equally weighted average daily effective spreads of NASDAQ-listed firms surrounding the 1997 tick-size reduction from eighths to sixteenths. The implementation date of the tick-size reduction is denoted as trade day 0. Firms are sorted into terciles based on the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Tercile 1 contains stocks with the largest spread reduction.




Figure 1.3: Effective spreads of NYSE and AMEX stocks surrounding decimalization

This figure presents equally weighted average daily effective spreads of NYSE- and AMEX-listed firms surrounding decimalization. The implementation date of the tick-size reduction is denoted as trade day 0. Firms are sorted into terciles based on the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Tercile 1 contains stocks with the largest spread reduction.




Figure 1.4: Effective spreads of NASDAQ stocks surrounding decimalization

This figure presents equally weighted average daily effective spreads of NASDAQ-listed firms surrounding decimalization. The implementation date of the tick-size reduction is denoted as trade day 0. Firms are sorted into terciles based on the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Tercile 1 contains stocks with the largest spread reduction.

Table 1.1: Descriptive statistics of key measures around the 1997 tick-size reduction

This table presents descriptive statistics of key measures around the 1997 tick-size reduction from eighths to sixteenths. Panel A presents descriptive statistics for NYSE- and AMEX-listed stocks, and Panel B presents descriptive statistics for NASDAQ-listed stocks. Change in effective spread ($\Delta ESPRD$) is the percentage change from the month before the 1997 tick-size reduction to the month in which the tick-size reduction was implemented. Return-on-assets (ROA) and investment ($INVST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over the three months preceding the tick-size regime change announcement date. Log market equity ($SIZE$) is constructed on the day preceding the announcement date. Log market-to-book (MB) is constructed using two different dates. Market value is computed on the day before the announcement date, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. See Section 1.2.2 for further details on sample construction.

	N	Mean	SD	P5	P25	P50	P75	P95
Panel A: NYSE and AMEX stocks								
$\Delta ESPRD$ (%)	1,723	-11.123	21.503	-41.502	-24.118	-12.798	-0.593	23.954
$ESPRD$	1,696	0.157	0.064	0.112	0.129	0.144	0.167	0.243
$SIZE$	1,714	6.309	1.800	3.304	5.087	6.311	7.568	9.348
MB	1,416	0.851	0.742	-0.253	0.419	0.807	1.243	2.143
ROA	1,544	0.127	0.122	0.008	0.081	0.133	0.184	0.290
MOM (%)	1,714	5.283	15.698	-20.564	-2.490	6.052	13.964	29.612
$INVST$	1,458	0.179	0.400	-0.137	0.007	0.084	0.220	0.771
$\Delta DEPTH$ (%)	1,723	-13.530	36.967	-57.906	-39.671	-21.670	2.808	66.038
Panel B: NASDAQ stocks								
$\Delta ESPRD$ (%)	1,468	-5.530	18.765	-33.079	-17.982	-6.608	4.994	25.226
$ESPRD$	1,440	0.244	0.136	0.091	0.151	0.215	0.311	0.498
$SIZE$	1,453	5.198	1.321	3.173	4.269	5.127	6.006	7.523
MB	1,043	1.063	0.791	-0.066	0.510	0.988	1.497	2.457
ROA	1,288	0.049	0.233	-0.399	0.009	0.107	0.180	0.298
MOM (%)	1,453	-6.138	25.556	-48.702	-21.563	-4.711	10.063	31.918
$INVST$	1,060	0.420	0.723	-0.195	0.046	0.201	0.487	1.873
$\Delta DEPTH$ (%)	1,468	3.196	18.303	-18.730	-3.190	0.518	5.685	35.406

Table 1.2: NYSE and AMEX stocks sorted by shock to spreads associated with the 1997 tick-size reduction

This table presents means and medians (in parentheses) of key measures of NYSE- and AMEX-listed firms sorted into terciles based on the percentage change in effective spread ($\Delta ESPRD$) from the month before the 1997 tick-size reduction to the month in which the tick-size reduction was implemented. Key firm characteristics are presented in Panel A and raw returns following key event dates are presented in Panel B. The second column presents the number of firms in each group. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over the three months preceding the tick-size regime change announcement date. Log market equity ($SIZE$) is constructed on the day preceding the announcement date. Log market-to-book (MB) is constructed using two different dates. Market value is computed on the day before the announcement date, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Panel B presents raw compounded returns over the one-, three-, or six-month window following tick size implementation and over the three-day or one-week window following tick-size announcement. Finally, raw compounded returns are presented over a window that combines three-day announcement returns with one-month implementation returns. See Section 1.2.2 for further details on sample construction.

Panel A: Firm Characteristics									
Group	N	$\Delta ESPRD$	$ESPRD$	$SIZE$	MB	ROA	MOM	$INVEST$	$\Delta DEPTH$
1	575	-31.928% (-29.117%)	0.148 (0.132)	6.787 (6.949)	0.834 (0.815)	0.131 (0.131)	3.365% (3.725%)	0.172 (0.075)	-12.681% (-23.341%)
2	574	-12.595% (-12.778%)	0.152 (0.144)	6.365 (6.346)	0.881 (0.837)	0.128 (0.138)	5.692% (6.155%)	0.186 (0.085)	-13.520% (-21.290%)
3	574	11.190% (6.113%)	0.171 (0.157)	5.773 (5.788)	0.839 (0.773)	0.122 (0.131)	6.800% (7.241%)	0.178 (0.091)	-14.391% (-19.744%)
Panel B: Raw stock returns									
Group	Following implementation			Following announcement			Combined		
	1 month	3 month	6 month	3 day	1 week	3 day & 1 mo.			
1	3.032% (1.724%)	12.892% (10.324%)	7.820% (8.811%)	1.030% (0.937%)	1.198% (1.064%)	4.140% (3.162%)			
2	4.023% (3.123%)	14.158% (11.275%)	12.005% (9.858%)	0.938% (0.957%)	1.004% (1.302%)	4.997% (4.295%)			
3	5.670% (4.525%)	17.029% (14.417%)	18.845% (16.605%)	0.868% (0.714%)	0.908% (0.943%)	6.587% (5.529%)			

Table 1.3: NASDAQ stocks sorted by shock to spreads associated with the 1997 tick-size reduction

This table presents means and medians (in parentheses) of key measures of NASDAQ-listed firms sorted into terciles based on the percentage change in effective spread ($\Delta ESPRD$) from the month before the 1997 tick-size reduction to the month in which the tick-size reduction was implemented. Key firm characteristics are presented in Panel A and raw returns following key event dates are presented in Panel B. The second column presents the number of firms in each group. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over the three months preceding the tick-size regime change announcement date. Log market equity ($SIZE$) is constructed on the day preceding the announcement date. Log market-to-book (MB) is constructed using two different dates. Market value is computed on the day before the announcement date, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Panel B presents raw compounded returns over the one-, three-, or six-month window following tick-size implementation and over the three-day or one-week window following tick-size announcement. Finally, raw compounded returns are presented over a window that combines three-day announcement returns with one-month implementation returns. See Section 1.2.2 for further details on sample construction.

Panel A: Firm Characteristics									
Group	N	$\Delta ESPRD$	$ESPRD$	$SIZE$	MB	ROA	MOM	$INVEST$	$\Delta DEPTH$
1	490	-24.653% (-23.123%)	0.250 (0.219)	5.360 (5.255)	1.053 (0.981)	0.039 (0.104)	-6.362% (-5.695%)	0.442 (0.229)	8.172% (1.850%)
2	489	-6.679% (-6.600%)	0.237 (0.206)	5.252 (5.169)	1.115 (1.042)	0.059 (0.110)	-7.943% (-5.406%)	0.431 (0.195)	0.911% (0.160%)
3	489	14.783% (11.763%)	0.247 (0.220)	4.981 (4.930)	1.014 (0.950)	0.049 (0.107)	-4.102% (-2.409%)	0.383 (0.186)	0.495% (0.086%)
Panel B: Raw stock returns									
Group	Following implementation			Following announcement			Combined		
	1 month	3 month	6 month	3 day	1 week	3 day & 1 mo.			
1	-0.011% (1.491%)	12.756% (10.619%)	10.974% (8.181%)	-0.107% (-0.152%)	-2.303% (-2.098%)	-0.142% (-0.009%)			
2	4.208% (3.297%)	16.231% (13.277%)	15.201% (11.950%)	0.039% (0.000%)	-1.939% (-1.862%)	4.210% (3.204%)			
3	4.224% (2.899%)	18.004% (14.291%)	16.297% (13.291%)	-0.558% (-0.538%)	-2.744% (-2.542%)	3.653% (1.803%)			

Table 1.4: Regression estimates of the impact of a shock to liquidity from the 1997 tick-size reduction on firm value

This table presents OLS regression estimates of the impact of a shock to liquidity from the 1997 tick-size reduction on firm value. I report estimated slope coefficients and robust t -statistics in parentheses. Panel A reports estimates for sample stocks listed on the NYSE and AMEX, while Panel B reports estimates for NASDAQ-listed stocks. The dependent variable (y_i) is raw compounded return over a one-, three-, or six-month window following tick size implementation, over a three-day or one-week window following tick size announcement, or over a period that combines the three-day window following announcement with the one-month window following implementation. The independent variables are listed in the first column. $\Delta ESPRD$ is the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over the three months preceding the tick-size regime change announcement date. Log market equity ($SIZE$) is constructed on the day preceding the announcement date. Log market-to-book (MB) is constructed using two different dates. Market value is computed on the day before the announcement date, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. See Section 1.2.2 for further details on sample construction.

Panel A: Regression estimates for NYSE and AMEX stocks										
	y_i = Return following implementation			y_i = Return following announcement			y_i = Combined return			
	1 month	3 month	6 month	3 day	1 week	3 day & 1 month				
$\Delta ESPRD$	0.043 (3.78)	0.076 (3.19)	0.207 (6.06)	0.136 (3.48)	-0.002 (-0.56)	0.004 (0.86)	-0.003 (-0.51)	0.010 (1.72)	0.040 (3.23)	0.035 (2.59)
$SIZE$	-0.224 (-1.18)	-2.019 (-4.16)	-1.985 (-3.14)	0.340 (4.66)	0.570 (6.03)	0.113 (0.52)				
MB	1.046 (2.36)	0.208 (0.18)	-1.573 (-0.98)	-0.057 (-0.27)	-0.483 (-1.75)	1.035 (2.04)				
ROA	-2.338 (-0.87)	2.121 (0.30)	-0.132 (-0.01)	-0.922 (-0.49)	1.839 (0.79)	-3.373 (-0.97)				
MOM	0.068 (2.62)	0.148 (2.19)	0.157 (1.87)	-0.019 (-1.92)	-0.041 (-2.80)	0.049 (1.60)				
$INVEST$	1.909 (2.25)	3.413 (2.00)	-0.561 (-0.24)	-0.189 (-0.61)	-0.561 (-1.44)	1.781 (1.72)				
$ESPRD$	0.030 (0.75)	0.105 (1.32)	0.328 (2.63)	0.015 (0.89)	0.020 (0.94)	0.046 (1.16)				
$\Delta DEPTH$	0.006 (0.65)	-0.007 (-0.37)	0.028 (1.13)	-0.006 (-2.01)	-0.010 (-2.34)	-0.001 (-0.16)				
Obs.	1,723	1,723	1,385	1,720	1,385	1,720	1,385	1,720	1,723	1,385
R^2 (%)	0.823	0.588	4.358	2.086	4.983	0.018	3.649	0.015	6.721	2.739

Panel B: Regression estimates for NASDAQ stocks									
	y_i = Return following implementation			y_i = Return following announcement			y_i = Combined return		
	1 month	3 month	6 month	3 day	1 week	3 day & 1 month	3 day	1 week	3 day & 1 month
$\Delta ESPRD$	0.108 (5.20)	0.125 (2.92)	0.153 (2.52)	0.156 (2.16)	-0.007 (-0.76)	-0.005 (-0.45)	-0.003 (-0.30)	-0.009 (-0.75)	0.101 (4.45)
$SIZE$	0.249 (0.74)	-1.029 (-1.30)	-0.629 (-0.53)	0.199 (1.25)	-0.090 (-0.52)	0.458 (1.32)	0.095 (3.55)	0.095 (3.55)	0.095 (3.55)
MB	-0.008 (-0.01)	-0.517 (-0.43)	-1.419 (-0.85)	-0.436 (-1.62)	-0.841 (-2.87)	-0.526 (-0.77)	-0.526 (-0.77)	-0.526 (-0.77)	-0.526 (-0.77)
ROA	2.937 (1.21)	20.544 (3.86)	17.321 (2.28)	3.031 (2.60)	4.216 (3.48)	6.029 (2.39)	6.029 (2.39)	6.029 (2.39)	6.029 (2.39)
MOM	0.062 (2.83)	0.164 (3.22)	0.328 (4.36)	-0.017 (-1.70)	-0.013 (-1.17)	0.045 (1.94)	0.045 (1.94)	0.045 (1.94)	0.045 (1.94)
$INVST$	0.146 (0.23)	-1.305 (-0.96)	-1.131 (-0.51)	0.090 (0.29)	-0.401 (-1.13)	0.298 (0.43)	0.298 (0.43)	0.298 (0.43)	0.298 (0.43)
$ESPRD$	0.081 (2.68)	0.135 (2.20)	0.269 (2.98)	0.026 (2.09)	0.038 (2.47)	0.110 (3.20)	0.110 (3.20)	0.110 (3.20)	0.110 (3.20)
$\Delta DEPTH$	-0.023 (-0.87)	-0.004 (-0.09)	-0.005 (-0.07)	-0.021 (-1.70)	-0.025 (-1.69)	-0.043 (-1.48)	-0.043 (-1.48)	-0.043 (-1.48)	-0.043 (-1.48)
Obs.	1,468	1,468	1,468	1,454	1,455	1,468	1,468	1,468	1,468
R^2 (%)	1.809	0.579	0.430	0.040	0.006	1.333	1.333	4.591	4.900

Table 1.5: Descriptive statistics of key measures around the shift to decimalization

This table presents descriptive statistics of key measures around the shift to decimalization. Panel A presents descriptive statistics for NYSE- and AMEX-listed stocks, and Panel B presents descriptive statistics for NASDAQ-listed stocks. Change in effective spread ($\Delta ESPRD$) is the percentage change from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over a three-month window ending one month before decimalization. Log market equity ($SIZE$) is constructed one month prior to decimalization. Log market-to-book (MB) is constructed using two different dates. Market value is computed one month prior to decimalization, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. See Section 1.2.2 for further details on sample construction.

	N	Mean	SD	P5	P25	P50	P75	P95
Panel A: NYSE and AMEX stocks								
$\Delta ESPRD$ (%)	1,565	-18.268	20.800	-47.722	-31.021	-20.211	-8.123	18.488
$ESPRD$	1,551	0.123	0.070	0.068	0.087	0.107	0.139	0.215
$SIZE$	1,565	6.431	2.097	2.856	4.946	6.555	7.859	9.995
MB	1,304	0.559	0.991	-1.126	-0.015	0.597	1.114	2.137
ROA	1,391	0.111	0.187	-0.040	0.075	0.124	0.184	0.296
MOM (%)	1,565	-8.816	40.266	-85.240	-24.065	1.123	16.511	36.520
$INVEST$	1,348	0.154	0.403	-0.151	-0.012	0.062	0.190	0.757
$\Delta DEPTH$ (%)	1,565	-29.334	30.386	-64.208	-50.448	-36.569	-15.820	31.159
Panel B: NASDAQ stocks								
$\Delta ESPRD$ (%)	2,352	-21.953	25.254	-56.020	-38.354	-25.384	-10.708	24.552
$ESPRD$	2,339	0.149	0.087	0.050	0.091	0.131	0.184	0.302
$SIZE$	2,351	4.900	1.666	2.232	3.746	4.854	5.953	7.715
MB	1,721	0.606	0.968	-0.948	-0.008	0.566	1.181	2.247
ROA	1,988	-0.003	0.277	-0.541	-0.068	0.052	0.151	0.283
MOM (%)	2,351	-15.963	41.405	-93.869	-40.222	-7.874	11.990	41.101
$INVEST$	1,758	0.446	0.980	-0.229	0.005	0.132	0.444	2.346
$\Delta DEPTH$ (%)	2,352	-18.769	30.388	-55.760	-37.546	-24.717	-7.521	42.048

Table 1.6: Univariate sorts by shock to spreads associated with decimalization

This table presents means and medians (in parentheses) of key measures sorted into terciles based on the percentage change in effective spread ($\Delta ESPRD$) from the month before decimalization to the month in which decimalization was implemented. Panel A presents summary statistics for NYSE- and AMEX-listed stocks, and Panel B presents summary statistics for NASDAQ-listed stocks. The second column presents the number of firms in each group. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over a three-month window ending one month before decimalization. Log market equity ($SIZE$) is constructed one month prior to decimalization. Log market-to-book (MB) is constructed using two different dates. Market value is computed one month prior to decimalization, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Returns are raw compounded returns over the one-, three-, or six-month window following decimalization. See Section 1.2.2 for further details on sample construction.

Group	N	$\Delta ESPRD$	Firm characteristics						Return following implementation			
			$ESPRD$	$SIZE$	MB	ROA	MOM	$INVEST$	$\Delta DEPTH$	1 month	3 month	6 month
			Panel A: NYSE and AMEX stocks									
1	522	-38.400% (-36.495%)	0.110 (0.094)	6.816 (7.004)	0.599 (0.693)	0.102 (0.121)	-9.332% (3.423%)	0.146 (0.062)	-31.449% (-38.104%)	-1.832% (-0.374%)	0.077% (1.233%)	3.208% (3.298%)
2	522	-20.369% (-20.211%)	0.119 (0.108)	6.683 (6.815)	0.662 (0.708)	0.127 (0.141)	-5.477% (1.991%)	0.173 (0.067)	-30.442% (-37.389%)	-0.551% (-0.578%)	2.434% (2.054%)	5.976% (3.499%)
3	521	4.010% (-0.228%)	0.139 (0.120)	5.792 (5.740)	0.417 (0.354)	0.104 (0.116)	-11.644% (-2.698%)	0.142 (0.055)	-26.105% (-32.441%)	-1.100% (-0.290%)	0.449% (0.661%)	7.106% (4.830%)
Panel B: NASDAQ stocks												
1	784	-45.940% (-44.017%)	0.116 (0.099)	4.600 (4.426)	0.434 (0.397)	-0.051 (0.028)	-24.489% (-16.813%)	0.532 (0.159)	-16.738% (-24.208%)	18.427% (12.685%)	15.261% (10.964%)	-6.407% (-8.933%)
2	784	-25.116% (-25.384%)	0.147 (0.131)	5.256 (5.233)	0.810 (0.768)	0.018 (0.068)	-15.159% (-9.127%)	0.442 (0.150)	-21.444% (-26.561%)	22.925% (17.979%)	25.470% (19.004%)	2.012% (-1.341%)
3	784	5.197% (-1.723%)	0.183 (0.158)	4.844 (4.900)	0.544 (0.518)	0.016 (0.068)	-8.233% (-0.014%)	0.377 (0.106)	-18.125% (-23.731%)	20.633% (13.473%)	26.848% (18.937%)	7.380% (2.782%)

Table 1.7: Regression estimates of the impact of a shock to liquidity from decimalization on firm value

This table presents OLS regression estimates of the impact of a shock to liquidity from decimalization on firm value. I report estimated slope coefficients and robust t -statistics in parentheses. The first six columns report estimates for sample stocks listed on the NYSE and AMEX, while the last six columns report estimates for NASDAQ-listed stocks. The dependent variable (y_i) is raw compounded return over a one-, three-, or six-month window following tick size implementation. The independent variables are listed in the first column. $\Delta ESPRD$ is the percentage change in effective spread from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. Return-on-assets (ROA) and investment ($INVEST$) are constructed as of the fiscal year-end preceding the tick-size reduction. Effective spread ($ESPRD$) and momentum (MOM) are constructed over a three-month window ending one month before decimalization. Log market equity ($SIZE$) is constructed one month prior to decimalization. Log market-to-book (MB) is constructed using two different dates. Market value is computed one month prior to decimalization, and book value is computed as of the prior fiscal year end. $\Delta DEPTH$ is the percentage change in average number of shares quoted at the best bid and ask from the month before the tick-size reduction to the month in which the tick-size reduction was implemented. See Section 1.2.2 for further details on sample construction.

NYSE and AMEX stocks				NASDAQ stocks							
y_i = Return following decimalization implementation											
	1 month	3 month	6 month	1 month	3 month	6 month					
$\Delta ESPRD$	0.023 (1.50)	0.026 (1.32)	0.029 (0.88)	0.101 (2.41)	0.030 (0.61)	-0.019 (-0.82)	0.045 (2.02)	0.155 (4.50)	0.135 (3.66)	0.214 (5.21)	0.152 (3.46)
$SIZE$		-0.354 (-1.48)	-0.093 (-0.22)	-1.264 (-2.15)	-0.484 (-1.16)	-1.738 (-2.50)	2.008 (2.47)	1.704 (0.49)	4.203 (0.77)	18.639 (3.03)	
MB		-1.641 (-3.43)	-3.942 (-4.28)	-7.121 (-5.40)	-0.243 (-10.79)	-0.067 (-2.08)	2.368 (1.86)	1.305 (0.94)	0.120 (3.45)	-3.458 (-3.17)	0.019 (0.16)
ROA		7.802 (2.95)	19.236 (4.26)	19.662 (2.94)	0.357 (0.46)	0.224 (0.20)	-0.277 (-2.96)	-0.057 (-1.73)	-0.055 (-1.42)	1.707 (3.46)	5.211 (3.71)
MOM		0.086 (6.18)	0.164 (6.95)	0.172 (4.77)	-0.243 (-10.79)	-0.067 (-2.08)	2.368 (1.86)	1.305 (0.94)	0.120 (3.45)	-3.458 (-3.17)	0.019 (0.16)
$INVEST$		-3.190 (-4.34)	-1.189 (-0.85)	-4.966 (-2.37)	0.357 (0.46)	0.224 (0.20)	-0.277 (-2.96)	-0.057 (-1.73)	-0.055 (-1.42)	1.707 (3.46)	5.211 (3.71)
$ESPRD$		-0.069 (-1.30)	-0.099 (-1.03)	-0.122 (-0.76)	-0.248 (-3.86)	-0.277 (-2.96)	-0.057 (-1.73)	-0.055 (-1.42)	1.707 (3.46)	5.211 (3.71)	
$\Delta DEPTH$		-0.020 (-1.45)	-0.043 (-1.90)	-0.014 (-0.38)	-0.099 (-5.10)	-0.057 (-1.73)	-0.055 (-1.42)	1.707 (3.46)	5.211 (3.71)		
Obs.	1,565	1,289	1,289	1,565	1,289	1,289	1,565	1,289	1,289	1,565	1,289
R^2 (%)	0.143	0.096	0.128	0.232	0.370	0.781	0.029	15.920	0.853	1.754	5.211

Table 1.8: Impact of shock to liquidity from decimalization on firm value using alternative measures

This table presents regression estimates of the impact of a shock to liquidity from decimalization on firm value, using alternative measures of liquidity and firm value. I report estimated slope coefficients and robust t -statistics in parentheses. Panel A presents estimates using annual change in liquidity and firm value, and Panel B presents estimates using monthly changes. RET is compounded raw stock return over the period following decimalization. Change in Tobin's Q (ΔQ), change in effective spread ($\Delta ESPRD$), and change in each of the two relative effective spread measures ($\Delta RESPRD_1$ and $\Delta RESPRD_2$) are measured from the period before to the period after decimalization. $\Delta RESPRD_1$ is scaled by concurrent bid-ask midpoint, while $\Delta RESPRD_2$ is scaled using prior-period stock price. These measures along with the control variables are described in detail in Sections 1.2.2 and 1.4.1.

	Panel A: Regression estimates using annual change in liquidity and firm value									
	NYSE and AMEX stocks					NASDAQ stocks				
	RET	RET	RET	ΔQ	ΔQ	RET	RET	RET	ΔQ	ΔQ
$\Delta ESPRD$	0.672 (5.96)			0.559 (7.71)		0.183 (2.99)			0.474 (6.60)	
$\Delta RESPRD_1$		-12.375 (-2.89)			-6.982 (-2.58)		-6.926 (-3.90)		-15.472 (-5.12)	
$\Delta RESPRD_2$			10.617 (3.11)		8.514 (4.03)		2.647 (1.64)			7.564 (2.94)
$SIZE$	0.058 (0.06)	-2.771 (-2.52)	-4.393 (-4.61)	-0.735 (-0.89)	-2.292 (-2.60)	-0.249 (-0.23)	-4.226 (-3.14)	-1.683 (-1.57)	-1.322 (-1.08)	3.405 (2.13)
MB	-3.890 (-2.01)	-3.554 (-1.77)	-4.133 (-2.07)	-6.414 (-2.98)	-8.122 (-3.70)	-2.389 (-1.34)	-4.358 (-2.56)	-3.158 (-1.81)	-1.955 (-0.96)	-4.592 (-2.29)
ROA	18.878 (1.35)	10.012 (0.77)	10.749 (0.76)	-17.016 (-1.22)	-5.950 (-0.45)	40.099 (5.62)	44.027 (6.07)	39.962 (5.57)	5.257 (0.40)	5.910 (0.45)
$INVT$	-2.149 (-0.56)	-4.619 (-1.22)	-4.154 (-1.11)	3.156 (0.98)	1.268 (0.41)	-0.612 (-0.27)	-3.203 (-1.42)	-0.911 (-0.41)	1.854 (0.74)	0.085 (0.03)
$ESPRD$	-0.243 (-0.99)	-0.177 (-0.08)	0.674 (2.85)	-0.228 (-1.31)	1.490 (0.98)	0.646 (3.63)	-4.071 (-2.54)	0.830 (5.13)	-0.248 (-1.14)	9.936 (4.36)
$\Delta DEPTH$	-0.197 (-3.40)	-0.097 (-1.80)	-0.264 (-4.56)	-0.219 (-5.57)	-0.210 (-5.28)	-0.309 (-7.57)	-0.181 (-4.26)	-0.338 (-8.50)	-0.212 (-4.26)	-0.259 (-5.16)
Obs.	1,197	1,197	1,197	990	990	1,417	1,417	1,417	1,113	1,113
R^2 (%)	12.37	13.55	9.71	24.32	18.75	16.68	19.23	16.40	14.24	17.04
					20.09					12.86

Panel B: Regression estimates using monthly change in liquidity and firm value						
	NYSE and AMEX stocks			NASDAQ stocks		
	RET	RET	RET	RET	RET	RET
$\Delta ESPRD$	0.026 (1.32)			0.045 (2.02)		
$\Delta RESPRD_1$		-4.162 (-5.53)			-3.132 (-2.88)	
$\Delta RESPRD_2$			0.339 (0.56)			0.910 (1.22)
$SIZE$	-0.354 (-1.48)	-1.025 (-3.86)	-0.454 (-1.87)	-0.484 (-1.16)	-0.588 (-1.04)	-0.693 (-1.57)
MB	-1.641 (-3.43)	-1.750 (-3.73)	-1.599 (-3.35)	2.008 (2.47)	1.842 (2.31)	1.920 (2.40)
ROA	7.802 (2.95)	6.399 (2.23)	7.502 (2.62)	1.704 (0.49)	1.710 (0.51)	1.520 (0.44)
MOM	0.086 (6.18)	0.094 (6.77)	0.085 (5.66)	-0.243 (-10.79)	-0.227 (-10.15)	-0.243 (-10.81)
$INVEST$	-3.190 (-4.34)	-3.092 (-4.15)	-3.200 (-4.33)	0.357 (0.46)	0.278 (0.36)	0.282 (0.36)
$ESPRD$	-0.069 (-1.30)	-2.093 (-5.66)	-0.063 (-1.21)	-0.248 (-3.86)	-0.806 (-1.44)	-0.225 (-3.62)
$\Delta DEPTH$	-0.020 (-1.45)	-0.026 (-2.06)	-0.021 (-1.52)	-0.099 (-5.10)	-0.116 (-6.24)	-0.100 (-5.17)
Obs.	1,289	1,289	1,289	1,707	1,707	1,707
R^2 (%)	9.10	12.33	8.97	15.92	16.48	15.84

Chapter 2

Micro(structure) before Macro? The Predictive Power of Aggregate Illiquidity for Stock Returns and Economic Activity²¹

²¹Chen, Y., Eaton, G., and Paye, B. Submitted to *Journal of Financial Economics*, 02/10/2016.

Abstract

This study constructs and analyzes various measures of trading costs in US equity markets covering the period 1926–2013. These measures contain statistically and economically significant predictive signals for stock market returns and real economic activity. We decompose illiquidity proxies into a component capturing aggregate volatility and a residual. The predictive content of these components differs in important ways. Specifically, we find strong evidence that the component of illiquidity uncorrelated with volatility forecasts stock market returns. In contrast, aggregate volatility accounts for most of the forecasting power of illiquidity for real economic activity.

2.1 Introduction

Liquidity conditions fluctuate over time in securities markets, and there is abundant evidence of commonality in liquidity.²² Fluctuations in aggregate liquidity potentially constitute a priced risk factor for stocks and other assets (Acharya and Pedersen (2005)), and numerous papers provide empirical evidence concerning the pricing of liquidity risk.

The implications of time-varying aggregate liquidity range beyond the cross-sectional pattern of security returns. Theoretical models link aggregate liquidity with time-variation in the equity premium. In the model of Acharya and Pedersen (2005), for example, persistent variation in liquidity implies that liquidity forecasts future stock returns. A related theoretical literature emphasizes connections between financial frictions and macroeconomic activity. Illiquidity, coupled with other financial frictions, can generate nonlinear amplification effects and exacerbate economic downturns (Brunnermeier and Pedersen (2009)).²³ This theoretical literature suggests that empirical measures of aggregate illiquidity should contain predictive signals for future stock returns and economic activity.

Empirically, Amihud (2002) and Jones (2002) find that certain aggregate liquidity measures forecast US stock market returns. However, relatively few subsequent papers emphasize such measures as return forecasting variables. For example, Welch and Goyal (2008) document the in- and out-of-sample forecasting power associated with a large set of return forecasting variables, but none of their variables explicitly measure aggregate liquidity. Bekaert, Harvey, and Lundblad (2007) find that a liquidity proxy based on the proportion of zero return days forecasts stock returns in emerging markets, but it is unclear whether similar results obtain for the US

²²Early studies documenting commonality in equity market liquidity include Chordia, Roll, and Subrahmanyam (2000), Hasbrouck and Seppi (2001) and Huberman and Halka (2001). Various studies extend evidence of commonality in liquidity beyond US equity markets. Chordia, Sarkar, and Subrahmanyam (2005) find evidence of commonality in aggregate spreads and depth across US equity and Treasury bond markets. Brockman, Chung, and Pérignon (2009) and Karolyi, Lee, and Van Dijk (2012) document commonality in liquidity across countries and exchanges. Subrahmanyam (2007) finds commonality between liquidity in equity markets and liquidity in real estate investment trusts (REITs) markets. Mancini, Ranaldo, and Wrampelmeyer (2013) identify common variation in liquidity across a number of asset classes and markets, including major foreign exchange currency pairs, US equities, and US government and corporate bond markets.

²³Brunnermeier and Pedersen (2009) develop a model linking aggregate liquidity and real activity driven by liquidity providers' funding requirements. Feedback between market and funding liquidity can generate liquidity spirals during times of financial distress, which often precede economic contractions.

and other developed markets. The extent to which liquidity predicts macroeconomic outcomes has received even less attention in the literature, although Næs, Skjeltorp, and Ødegaard (2011) report positive evidence for several measures.²⁴

This study provides new, relatively comprehensive evidence concerning the predictive content of aggregate liquidity for stock returns and real economic activity. We construct a variety of measures of market frictions at monthly and quarterly frequencies using daily CRSP data for stocks listed on the New York Stock Exchange (NYSE).²⁵ These time series cover an 88 year period from 1926–2013 and include multiple proxies for the effective spread, several alternative measures of price impact, and a measure of aggregate price delay. Since our measures relate positively to transaction costs, we refer to them as “illiquidity” measures.

Consistent with earlier evidence in Jones (2002) for bid-ask spreads on Dow Jones stocks, aggregate illiquidity contains a cyclical component. Liquidity tends to erode during economic recessions, and markets are considerably more illiquid during both the Great Depression and the recent financial crisis. Alternative aggregate spread measures comove closely from 1926 through the early 1990s. In the late 1990s, most proxies indicate a decrease in the spread that accelerates from 2001 until shortly before the financial crisis, when spreads increase. The average effective spread for all NYSE stocks, computed using high-frequency (TAQ) data, exhibits a similar structural shift. Sharp effects for several proxies implicate tick-size reductions occurring in 1997 and 2001 as contributing to the reduction in spreads. These results for aggregate measures are consistent with earlier literature regarding the liquidity effects of tick-size reductions.²⁶ To focus

²⁴Beber, Brandt, and Kavajecz (2011) and Kayacetin and Kaul (2009) analyze the predictive content in aggregate equity order flow for the state of the economy. These studies focus on relatively short time periods due to limitations on data availability.

²⁵Our convention of basing measures on NYSE stocks follows a number of papers in the literature, including Amihud (2002), Goyenko, Holden, and Trzcinka (2009), and Næs, Skjeltorp, and Ødegaard (2011). As a robustness check, we also construct our proxies using firms on both NYSE and AMEX. Our main findings continue to obtain for these alternative measures.

²⁶Studies documenting lower effective spreads following tick-size reductions include Bacidore, Battalio, and Jennings (2003), Bessembinder (2003), Chakravarty, Wood, and Van Ness (2004), Jones and Lipson (2001), and Van Ness, Van Ness, and Pruitt (2000). The aggregate effective spread and several low-frequency proxies continue to fall for several years following decimalization in 2001. A possible explanation for this pattern is the introduction of ‘OpenBook’ to the NYSE, which allowed subscribing market participants real time access to the limit order book. Boehmer, Saar, and Yu (2005) report evidence of reductions in the effective spread associated with this move toward ex ante limit order book transparency.

on cyclical fluctuations in market liquidity, we construct break-adjusted versions of proxies that decrease following tick-size reductions.

Most aggregate illiquidity measures are positively correlated with stock market volatility. In some cases, the correlation is quite strong. The connection between measures of market frictions and measures of volatility runs deep: we show that several common proxies for the effective spread are mechanically log-linear in the standard deviation of returns. In similar fashion, price impact measures of the form proposed by Amihud (2002) are approximately log-linear in return volatility. Motivated by these observations, we decompose illiquidity measures into a component reflecting aggregate volatility and a residual. Importantly, we show that a simple measure of illiquidity based on the frequency of zero return days (Lesmond, Ogden, and Trzcinka (1999)) can be interpreted as a spread measure that is already volatility-adjusted. After extracting a volatility component, other common spread measures are highly correlated with the zeros measure in the aggregate.

The remainder of the study analyzes the predictive content associated with various aggregate illiquidity measures, with a focus on assessing the predictive content of the volatility component embedded in such measures and the residual component. Specifically, we test whether illiquidity measures and their components forecast stock market returns and various measures of economic activity, including output as measured by the growth rate in industrial production (or GDP growth and its components at the quarterly frequency) and the growth rate in unemployment.

There is little evidence that raw illiquidity measures forecast stock returns. However, weak forecasting performance is potentially attributable to unmodeled level shifts in illiquidity. Indeed, we obtain stronger evidence of stock return forecasting power upon adjusting illiquidity measures for structural breaks following tick-size reductions. This finding resembles results in Lettau and Van Nieuwerburgh (2008), who show that adjusting for level shifts in financial ratios produces stronger and more stable evidence of return forecasting power. The zeros measure of transaction costs proposed by Lesmond, Ogden, and Trzcinka (1999) significantly forecasts returns and performs best among break-adjusted measures. Since the zeros proxy can be interpreted as a volatility-adjusted spread measure, this suggests that stronger evidence of stock

return forecasting power will obtain for volatility-adjusted illiquidity measures. Consistent with this intuition, we find that many volatility-adjusted measures significantly forecast returns.

The predictive power associated with several volatility-adjusted measures is economically significant. One-month ahead predictive regressions for excess returns yield in-sample R^2 -values of approximately 1.2–1.5% for these measures. Because illiquidity is persistent, predictability at short horizons mechanically generates stronger predictability at longer horizons. At the annual horizon, R^2 -values for these variables are in the 7-12% range. These values are on par with (in-sample) R^2 -values at the annual horizon for stock return forecasts based on the dividend yield (Fama and French (1988)).

The intuition behind the importance of volatility-adjustment in stock return forecasting regressions is as follows. The return forecasting coefficient associated with unadjusted illiquidity measures is a weighted average of the forecasting coefficient for the volatility component and that for the volatility-adjusted (residual) component. In the data, the stock return forecasting coefficient associated with lagged aggregate volatility is negative.²⁷ As a result, estimates of stock return forecasting coefficients for raw measures are often close to zero and insignificant, whereas slope estimates for volatility-adjusted measures are positive and significant. While we show that illiquidity measures adjusted for volatility forecast stock returns, it is important to emphasize that our results do not necessarily imply a negative risk-return relation. For example, Ghysels, Santa-Clara, and Valkanov (2005) show that alternative volatility models that incorporate a longer window length in forecasting variance yield a significant positive relation.

Additional results address a variety of potential concerns. Since illiquidity measures exhibit long memory and shocks to the illiquidity are contemporaneously (negatively) correlated with return shocks, econometric concerns arise related to the finite sample bias and associated inference problems studied by Stambaugh (1999). We show that key results regarding return predictability continue to obtain under alternative bootstrap inference methods that accom-

²⁷This is consistent with French, Schwert, and Stambaugh (1987), who find an insignificant and sometimes negative risk-return relation using a volatility proxy based on the sum of squared daily return in the prior month (a measure highly correlated with our aggregate volatility measure).

moderate these features of the data. We continue to find evidence that many volatility-adjusted illiquidity measures forecast stock returns after including other common financial predictors in the model (short interest rates, spreads, etc.). In addition, these measures outperform both the historical average and other common return forecasting variables out-of-sample.

Turning to measures of economic activity, we find strong evidence that aggregate liquidity predicts economic activity, controlling for past activity. Consistent with theory and intuition, more illiquid markets predict lower future output growth and higher future unemployment. As with stock return forecasts, it is important to adjust for the structural shifts. Most series that exhibit a downward break following tick-size reductions only predict economic activity after break adjustment. The predictive power of illiquidity measures is stronger for quarterly data and for 6- and 12-month ahead forecasts based on overlapping monthly observations. There is an important caveat to these results, however. Upon decomposing illiquidity measures into a volatility component and a residual, we find that the bulk of predictive power for real economic activity derives from the volatility component embedded in illiquidity measures.²⁸ Conditioning on a wider information set including lagged returns, volatility, and various bond spreads yields stronger evidence of marginal predictive power associated with volatility-adjusted illiquidity measures. Still, the associated incremental forecasting power is relatively small. As with stock return forecasting results, key findings continue to obtain under bootstrap inference methods and in an out-of-sample forecast comparison.

Our key results survive various robustness checks. In particular, we show that the predictive power of illiquidity measures is not spuriously generated by the procedure we use to adjust measures for structural shifts around tick size reductions. Similarly, our results are not driven by a form of “look-ahead bias” related to projecting illiquidity measures on volatility, because we obtain similar results using an analytical decomposition that does not rely on estimation.

Jones (2002) represents an important predecessor to our study. Jones (2002) documents the cyclical behavior of aggregate liquidity and provides early evidence regarding the predictive

²⁸Our results concerning the predictive power of the volatility component of illiquidity measures are consistent with Campbell, Lettau, Malkiel, and Xu (2001), who find that volatility helps predict GDP growth. In related work, Giglio, Kelly, and Pruitt (2015) find that systemic risk measures are able to predict macroeconomic downturns.

power of liquidity measures for US stock returns. We show that many illiquidity measures are subject to structural shifts, with an important shift occurring around tick-size reductions near the turn of the century, and demonstrate the importance of adjusting for this shift from a forecasting perspective. We also illustrate the importance of adjusting for volatility in uncovering the relation between illiquidity and the equity premium. In so doing, we extend evidence in Bekaert, Harvey, and Lundblad (2007) to US data and develop an interpretation of the zeros proxy as a volatility-adjusted illiquidity measure. Finally, our evidence that illiquidity measures forecast economic activity is broadly consistent with Næs, Skjeltorp, and Ødegaard (2011), although we argue that a considerable portion of the associated forecasting power is attributable to the volatility component embedded in such measures.

2.2 Data

This section describes our data sources and the variables used in the empirical analysis. There are three categories of variables considered: 1) illiquidity and other market friction measures; 2) macroeconomic variables; and 3) other financial variables.

2.2.1 Illiquidity and other market friction measures

We construct a variety of firm-level measures of illiquidity and frictions at monthly and quarterly horizons for NYSE-listed stocks. Below we list and briefly describe these measures. The Appendix provides additional discussion regarding the various measures and implementation details.

Roll measure: The Roll (1984) liquidity measure (*ROLL*) is an estimate of the effective bid-ask spread based on the serial covariance of successive price movements.²⁹

²⁹The Roll liquidity proxy is only defined when the first-order serial covariance of daily returns for the corresponding period is negative. Empirically, the first-order sample covariance can be positive. In such cases, we set the Roll proxy to zero. The Appendix briefly discusses alternative conventions. Our main results are robust to the particular approach employed.

Corwin and Schultz measure: The Corwin and Schultz (2012) liquidity measure (*CS*) uses high and low daily price quotes to estimate bid-ask spreads. The key idea behind the proxy is that, although the high-low ratio reflects both variance and the spread, the variance component is proportional to the return interval, while the spread component is not.

Fong, Holden and Trzcinka measure: The Fong, Holden, and Trzcinka (2014) liquidity measure (*FHT*) incorporates two key elements of transaction costs: return volatility and the proportion of zero return. The measure represents a simplification of a more general variation proposed by Lesmond, Ogden, and Trzcinka (1999).

Effective tick measure: The effective tick measure (*TICK*) proposed by Holden (2009) and Goyenko, Holden, and Trzcinka (2009) is based on observable price clustering.

Zeros measure: This liquidity measure (*ZEROS*), proposed by Lesmond, Ogden, and Trzcinka (1999), is based on the proportion of trading days with zero returns. A variation on this measure is defined as the fraction of zero return days among trading days with nonzero trading volume. Key results in the study are robust to the version of zeros measure used.

Amihud illiquidity measure: Amihud (2002) proposes a measure of price impact (*AMI*) based on the ratio of the daily absolute return to daily trading volume. The measure is computed as the average of this ratio over a certain time interval (e.g., month, quarter), including only days with positive trading volume. The measure captures the magnitude of daily price response per dollar of trading volume. To control for the effects of inflation on dollar trading volume, we construct *AMI* in real terms. However, the inflation adjustment does not make a meaningful difference and our results continue to hold if we construct *AMI* in nominal terms. Brennan, Huh, and Subrahmanyam (2013) construct a version of the Amihud measure based on the ratio of absolute daily returns to daily turnover, calculated as daily share volume divided by total shares outstanding (*AMITO*), in order to separate turnover from firm size effects. Finally, following Næs, Skjeltorp, and Ødegaard (2011) we compute a first-differenced version of the (log) Amihud measure, defined as $\Delta ami_t \equiv \ln(AMI)_t - \ln(AMI)_{t-1}$.

Pastor and Stambaugh measure: The Pástor and Stambaugh (2003) measure (PS) is a measure of price impact based on the extent of volume-related return reversal. The original measure proposed by Pástor and Stambaugh (2003) is one of liquidity (larger values indicate more liquidity). Our PS measure multiplies the original measure by negative one to convert to a measure of illiquidity.

Hou and Moskowitz measure: The Hou and Moskowitz (2005) measure (HM) attempts to capture the delay in which stock prices respond to information. The measure is computed based on regressions of firm returns on current and lagged market returns. Loosely, the measure is larger when lagged market returns explain more variation in firm returns (relative to current market returns). Hou and Moskowitz (2005) interpret their measure as capturing a variety of market frictions including effects related to incomplete information, short-sale constraints, asymmetric information, taxes, and liquidity. Hou and Moskowitz (2005) find that the cross-sectional premium associated with their delay measure relates to measures of investor recognition, rather than traditional measures of liquidity.

Effective spread (high-frequency): A common measure of the cost dimension of market liquidity using high-frequency (intradaily) trading data is the *percent effective spread*, defined as

$$ESPD = 2 \frac{|P_k - M_k|}{M_k}, \quad (2.1)$$

where P_k represents the price of the k -th trade, and M_k represents the midpoint of the best bid and offer (BBO) prevailing at the time of the k -th trade. At the firm or asset level, this measure is aggregated over some unit of time (trading day, month, quarter) by computing the dollar-volume-weighted average of the percent effective spread for the time interval. The effective spread represents an estimate of the round trip cost of trading over the corresponding interval. We obtained this measure for all NYSE firms over the period 1993–2007 as a benchmark for evaluating the performance of the various low-frequency spread proxies.

For each firm-level measure of market illiquidity, we construct an aggregate measure by computing the equally-weighted average across securities for each month or quarter. This procedure follows prior studies, including Jones (2002), Amihud (2002), Pástor and Stambaugh (2003), and Næs, Skjeltorp, and Ødegaard (2011). Results in the study are based on NYSE firms; however, we have also computed analogous measures for the universe of NYSE and AMEX firms, and our main findings continue to obtain.

Most of the aggregate illiquidity measures are strictly positive. For these measures, it is convenient to apply a logarithmic transformation, which can be motivated by both statistical and economic considerations. From a statistical perspective, raw aggregate liquidity measures are positively skewed and leptokurtotic, and the logarithmic transformation dampens these deviations from the Gaussian distribution. From an economic perspective, we show in Section 2.3.2 that the logarithm of many illiquidity proxies can be decomposed into the sum of a component capturing aggregate volatility and a residual. Lowercase notation indicates log-transformed variables. For example, $roll_t \equiv \ln(ROLL_t)$.

2.2.2 Macroeconomic and financial variables

We collect macroeconomic variables at both the monthly and quarterly frequencies. To measure output at the monthly frequency, we use real, seasonally-adjusted industrial production (IP). These data cover the period 1926.1–2013.12. As an alternative measure of business conditions, we examine the monthly, seasonally adjusted US civilian unemployment rate (UE), available over the period 1948.2–2013.12. At the quarterly frequency, we follow Næs, Skjeltorp, and Ødegaard (2011) and examine output, unemployment, investment, and consumption. The output measure is real, seasonally adjusted gross domestic product (GDP). Unemployment (UE) is the seasonally adjusted unemployment rate. Our measure of investment (INV) is real, gross private domestic investment (seasonally adjusted). Consumption ($CONS$) is real, seasonally-adjusted personal consumption expenditures. These data are available over the period 1947.1–2013.4, with the exception of UE which is available beginning in 1948. The macroeconomic variables either trend over time, or exhibit considerable persistence and potential structural shifts. For

this reason, we frequently focus on growth rates. For example, the variable Δip denotes the growth rate of industrial production, computed as log-differences.

In addition to forecasting macroeconomic outcomes, we examine the ability of illiquidity measures to forecast excess stock market returns. Our primary stock return measure (ret) is the log excess return on the NYSE index with the 3-month T-bill rate serving as the reference return. As a robustness check we also consider an excess return based on the value-weighted CRSP universe and obtain similar results. The empirical analysis incorporates a number of financial variables often used to forecast macroeconomic outcomes. Among these, the most crucial is aggregate volatility, since a central focus of our study involves contrasting the volatility-related and non-volatility-related forecasting content in various illiquidity measures. Our primary measure of aggregate volatility is the cross-sectional average (for NYSE firms) of the natural logarithm of the sample standard deviation returns within the corresponding month or quarter, denoted vol . The motivation behind this particular measure of aggregate volatility will be made clear in Section 2.3.2 of the study. The term spread ($TERM$) is the difference between the 10-year Treasury yield and the yield on the 3-month Treasury bill. The default spread (DEF) is the yield on Moody's Baa-rated corporate bonds and the yield on Moody's Aaa-rated corporate bonds. The commercial paper to Treasury spread ($CPSP$) is the spread between the 3-month commercial paper rate and the rate on 3-month Treasury bills (see Appendix for details). All financial variables are obtained over the period 1926–2013 at both the monthly and quarterly frequencies.

2.3 Aggregate Illiquidity

This section analyzes the dynamics of aggregate illiquidity. Sections 2.3.1 characterizes the time-series behavior of the various aggregate illiquidity measures. Section 2.3.2 discusses approaches to decomposing illiquidity measures into a component related to aggregate volatility and a residual. Section 2.3.3 discusses structural shifts in many illiquidity measures around the turn of the century related to tick-size reductions.

2.3.1 Time series properties of aggregate spread measures

Many of the low-frequency illiquidity measures explicitly aim to measure the round trip cost of trading (bid-ask spread). Figure 2.1 shows that the associated aggregate time series comove closely. The affinity among the various spread proxies is particularly high during a sixty year period stretching from the 1940s through the 1990s. Consistent with Jones (2002), the aggregate spread series exhibit cyclical behavior. Spreads tend to increase during economic recessions and financial crises. The aggregate spread increases dramatically around the Great Depression, although the extent of the increase differs somewhat across alternative proxies.

The spread proxies diverge in recent years. To better illustrate this divergence, Figure 2.2 provides time series plots for the sub-period 1993-2013. In addition to the low-frequency spread proxies, the figure plots the aggregate effective spread computed using intraday TAQ data over the period 1993–2007. Around the turn of the century, the spread measures bifurcate, with the *FHT* and *TICK* measures suggesting an economically significant, structural shift downward in the aggregate effective spread culminating shortly before the onset of the financial crisis in 2007–2008. The effective spread computed using TAQ data conforms closely to the *FHT* and *TICK* proxies throughout the 1993–2007 period. In contrast, *ROLL* and *CS* are upward biased measures of the aggregate effective spread: these proxies are roughly 3 times the aggregate effective spread from 2003-2006. In addition to shifting downward, the volatility of the effective spread and many of the spread proxies decreases after the turn of the century. Log-transformed series exhibit less of an apparent volatility shift, suggesting a proportional reduction in the spread during this period.

The Internet Appendix provides time series plots for the other illiquidity proxies we consider. Similar to the spread measures, the other illiquidity proxies tend to fluctuate inversely with the business cycle, rising during recessions and financial crises. Second, many illiquidity measures appear subject to a downward level shift around the turn of the century. Structural shifts are particularly clear for the zeros proxy. This measure exhibits two sharp downward structural breaks, one following June 1997 and the other following January 2001: dates that coincide with minimum tick-size reductions for the NYSE. The stark pattern for the zeros proxy suggests

that downward structural shifts around the turn of the century for spreads and other illiquidity measures may be driven by tick-size reductions.

Table 2.1 presents descriptive statistics for illiquidity measures and other selected variables at the monthly frequency over the sample period 1948(2)–1997(6). This sample period is chosen to minimize the potential of spurious correlations and artificially inflated persistence parameters attributable to structural shifts. Most illiquidity measures are positively skewed and leptokurtotic, even after applying a logarithmic transformation. They are also persistent. First-order sample autocorrelations range from 0.33 (the Hou-Moskowitz measure) to 0.98 (the effective tick measure). Autocorrelations for the illiquidity measures die out slowly. The table shows the sample autocorrelation at lag 36 (3 years), along with the autocorrelation at lag 36 implied by the first-order sample autocorrelation under an AR(1) process. The sample autocorrelation at a 36-month lag is considerably larger than that implied by an AR(1) process for many illiquidity measures. This suggests that aggregate illiquidity measures exhibit *long memory* properties. The autocovariances of a long memory process decay at a slower rate relative to the geometric decay rate associated with autocovariances for standard short memory processes, such as stationary ARMA processes.³⁰ Table 2.1 reports estimates of the long-memory parameter d using the local Whittle estimator proposed by Shimotsu and Phillips (2005). The d estimates are statistically greater than zero for all illiquidity measures, and most estimates fall in the range 0.3–0.9. Augmented Dickey-Fuller tests of the null of a unit root ($d = 1$) are rejected for most measures. Consequently, the aggregate illiquidity measures appear to exhibit an order of integration between zero and one. Consistent with numerous

³⁰The autocorrelation γ_k associated with a long memory time series follows the asymptotic relation

$$\gamma_k \approx ck^{2d-1} \text{ as } k \rightarrow \infty, \quad (2.2)$$

where $c > 0$ is a constant. Granger and Joyeux (1980) and Hosking (1981) show that a stationary long memory time series can be modeled by extending the standard ARIMA class of models to permit *fractional integration*. Such a process may be written as

$$(1 - L)^d(x_t - \mu_x) = w_t. \quad (2.3)$$

where $\mu_x = E(x_t)$, w_t is a short-memory, zero mean time series, $-0.5 < d < 0.5$ and (setting $\mu_x = 0$), $(1 - L)^d x_t$ has the interpretation $1 - dx_{t-1} + \frac{d(d-1)}{2!}x_{t-2} - \dots$, with d the so-called fractional differencing parameter. The case $d = 0$ corresponds to the absence of long memory (a short memory series). Values of d above 0.5 are also possible; however, in such cases the resulting process is no longer covariance stationary. The case $d = 1$ corresponds to an integrated process for x_t .

previous studies, estimates in Table 2.1 imply that volatility is also a long memory process.³¹ In contrast, output growth, unemployment growth, and stock market returns yield d estimates that are close to zero.

Panel B of Table 2.1 displays sample correlations among the illiquidity proxies. The spread proxies tend to be highly correlated over the 1948–1997 period. Most other illiquidity measures, including Amihud measures and the *zeros* measure, are positively correlated with the spread measures and with each other. Interestingly, the *hm* price delay measure is *negatively* correlated with spread and price impact measures, suggesting that this measure emphasizes other sources of frictions.

The illiquidity proxies are generally negatively correlated with the growth rate in industrial production and positively correlated with unemployment growth, as intuition suggests. Stock market returns are negatively correlated with most illiquidity measures: a positive market return tends to coincide with an improvement in aggregate liquidity. Stock market volatility is positively correlated with most illiquidity measures. The price delay proxy *hm* correlates positively with contemporaneous stock returns, and negatively with stock market volatility. Hou and Moskowitz (2005) find that stocks with large price delay effects tend to be small, with limited institutional ownership and analyst coverage. To the extent that the aggregate *hm* measure captures attentional effects, it is plausible that average investor attention increases in volatile, turbulent market conditions associated with recessions.

The last two rows of Table 2.1 reports an estimate of the contemporaneous correlation between *shocks* to illiquidity measures and shocks to either excess stock returns or growth in industrial production. To obtain these estimates, we first apply a univariate time series filter to each variable, and then compute the sample correlation based on the resulting residuals. The filtering procedure applied to illiquidity measures fractionally differences the series based on the estimate $\hat{d}$ reported in Table 2.1, and then applies an autoregressive filter with one lag to capture remaining short memory dynamics. Since there is little persistence in stock returns, return shocks are simply deviations from the mean. We apply an autoregressive filter with 4 lags

³¹Examples of studies documenting long memory in stock market volatility include Ding, Granger, and Engle (1993), Baillie, Bollerslev, and Mikkelsen (1996) and Andersen and Bollerslev (1997).

to growth in industrial production. Shocks to most illiquidity measures are negatively correlated with contemporaneous shocks to excess returns. This is consistent with theoretical predictions (see, e.g., Acharya and Pedersen (2005)). The sign of the correlation estimate is positive rather than negative for the *hm* measure. The magnitude of most correlation estimates falls in the range 0.25–0.5. There appears to be little correlation between shocks to illiquidity measures and shocks to growth in industrial production. We reference these correlation estimates later in discussing the econometric properties of predictive regressions involving illiquidity measures.

2.3.2 Extracting a volatility component from illiquidity measures

Most measures of transaction costs are positively correlated with stock market volatility (see Table 2.1). Below we show that several standard illiquidity measures are *mechanically* related to volatility, and that the associated relation is log-linear. This provides a convenient basis for decomposing aggregate illiquidity measures into a volatility component and a residual. The decomposition we make is empirical in nature. That is, we decompose illiquidity proxies into an empirical measure of return volatility (the log of the standard deviation of returns) and a residual. However, motivation for our decomposition also stems from theory. For example, the celebrated price impact parameter λ in the model of Kyle (1985) is loglinear in the volatility of fundamental asset value.³²

As a first example, consider the spread proxy proposed by Fong, Holden, and Trzcinka (2014). For the i -th firm in period t , the measure is defined as:

$$FHT_{i,t} = 2\hat{\sigma}_{i,t}\Phi^{-1}\left(\frac{1 + ZEROS_{i,t}}{2}\right), \quad (2.4)$$

where $\hat{\sigma}_{i,t}$ is the standard deviation of returns over time interval t (e.g., a month or quarter), $ZEROS_{i,t}$ is the fraction of zero return days out of the total trading days during time interval t , and $\Phi^{-1}(\cdot)$ is the inverse cumulative normal distribution. Assuming $ZEROS_{i,t} > 0$ and taking

³²The decompositions that follow involve the (sample) standard deviation of daily returns. Although this empirical measure reflects fundamental volatility, it can also incorporate microstructure effects due to, e.g., bid-ask bounce. Our aggregate measure of volatility (*vol*) correlates highly with log realized volatility for the S&P 500 Index, which suggests that in the aggregate the measure primarily reflects fundamental volatility.

logarithms yields:

$$fht_{i,t} = \ln(2) + \ln(\hat{\sigma}_{i,t}) + g(ZEROS_{i,t}), \quad (2.5)$$

where $g(ZEROS_{i,t}) \equiv \ln\left(\Phi^{-1}\left(\frac{1+ZEROS_{i,t}}{2}\right)\right)$. Averaging across firms gives

$$\overline{fht}_t = \ln(2) + vol_t + \overline{g(ZEROS_{i,t})}, \quad (2.6)$$

where overbars indicate cross-sectional averages and $vol_t \equiv \overline{\ln(\hat{\sigma}_{i,t})}$, the cross-sectional average of the logarithm of firm volatility in period t . This shows that, ignoring the constant $\ln(2)$, the aggregate quantity $\overline{fht}_t$ is the sum of two components. The first is the cross-sectional average of the log of estimated volatility at the firm level (aggregate volatility), and the second is the cross-sectional average of a nonlinear transformation of the *ZEROS* illiquidity measure.

In practice, a technical issue arises due to the fact that the quantity $g(ZEROS_{i,t})$ is undefined when $ZEROS_{i,t} = 0$. This case occurs in the data, particularly in recent years following decimalization. Consequently, we consider the closely-related aggregate decomposition:

$$\widetilde{fht}_t \equiv fht_t - vol_t, \quad (2.7)$$

in which $fht_t = \ln(\overline{FHT}_{i,t})$ is substituted for $\overline{fht}_t$. The two quantities differ by a Jensen's inequality correction term, and the substitution ensures that $\widetilde{fht}$ is well-defined.

The decomposition of Eq. (2.6) suggests that the *zeros* measure *itself* can be interpreted as a volatility-adjusted measure of the aggregate effective spread. Indeed, the correlation between $\widetilde{fht}$ and *zeros* in monthly data over the period 1926–2013 is approximately 0.94. Bekaert, Harvey, and Lundblad (2007) examine the relation between the zeros illiquidity measure and volatility in a set of emerging markets. They find the correlation between the two to be small in magnitude on average across countries and conclude that the zeros measure captures “unique aspects of liquidity not entirely driven by the presence or absence of news in a particular period.” This suggests that the decomposition of Eq. (2.6) should be roughly orthogonal, i.e.,

that $\widetilde{fht}$ defined in Eq. (2.7) should be roughly uncorrelated with aggregate volatility. Below, we confirm this conjecture in the data.

Similar algebra leads to a volatility-adjusted Roll (1984) spread proxy. The Roll (1984) proxy is defined as:

$$ROLL_{i,t} = 2\sqrt{-\widehat{\text{Cov}}(R_{d,i,t}, R_{d-1,i,t})}, \quad (2.8)$$

where $R_{d,i,t}$ ($R_{d-1,i,t}$) is the return for the i -th firm on trading day d (day $d-1$) of time interval t and $\widehat{\text{Cov}}(\cdot)$ represents the sample covariance. Taking logs and averaging across firms provides a spread decomposition similar to Eq. (2.6):

$$\overline{roll}_t = \ln(2) + vol_t + (1/2)\overline{\ln(-\hat{\rho}_1(R_{t,i}))}, \quad (2.9)$$

where $\hat{\rho}_1$ denotes the sample autocorrelation at lag one, and it is implicitly assumed that $\hat{\rho}_1(R_{t,i}) < 0$. In the data, it is sometimes the case that $\hat{\rho}_1(R_{t,i}) > 0$, in which case the second term on the right-hand side of Eq. (2.9) is undefined. For this reason, we define a volatility-adjusted aggregate measure as $\widetilde{roll}_t \equiv roll_t - vol_t$. Similar to the construction of $\widetilde{fht}$, this involves substituting $roll_t = \ln(\overline{ROLL}_{i,t})$ for $\overline{roll}_t$.

These explicit decomposition results motivate the following generalized decomposition for an arbitrary illiquidity measure ILQ_t :

$$\widetilde{ILQ}_t = ILQ_t - vol_t. \quad (2.10)$$

This is useful for proxies such as *tick* and *cs*, for which it is more difficult to explicitly derive loglinear decompositions.³³ Since these proxies also measure the effective spread, one can conjecture that they are also log-linear in volatility and form $\widetilde{tick}$ and $\widetilde{cs}$ based on Eq. (2.10). Two remarks are in order. First, nothing in the derivation of decompositions of the form of Eq. (2.10) implies that the associated decomposition is orthogonal, in the sense that vol_t and

³³In the Internet Appendix, we show that, under certain assumptions, it is possible to analytically derive a decomposition similar to Eqs. (2.6) and (2.9) for Amihud-type measures.

the residual $\widetilde{ILQ}_t$ are uncorrelated. Second, it is not clear ex ante that a decomposition of the form of Eq. (2.10) is sensible for certain illiquidity measures such as *PS* and *hm*.

An alternative, data-driven decomposition of an arbitrary illiquidity measure ILQ_t can be based on the linear regression:

$$ILQ_t = \alpha + \beta vol_t + \epsilon_t, \quad (2.11)$$

where $\beta \equiv \text{Cov}(ILQ_t, vol_t) / \text{Var}(vol_t)$. The regression residual ϵ_t in Eq. (2.11) is uncorrelated with vol_t by virtue of the definition of β . Let $\hat{\alpha}$ and $\hat{\beta}$ denote estimates of the parameters α and β in Eq. (2.11). We define a volatility-adjusted illiquidity measure as the residual: $ILQ_t^\perp \equiv \hat{\epsilon}_t$, where $\hat{\epsilon} = ILQ_t - \hat{\alpha} - \hat{\beta} vol_t$. The notation emphasizes that this form of volatility-adjustment orthogonalizes the illiquidity measure with respect to aggregate log volatility. In the special case that $\beta = 1$ in Eq. (2.11), the analytical decomposition of Eq. (2.10) becomes an *orthogonal* decomposition of the spread, and is equivalent to the linear projection approach up to an intercept and in a population sense. Consequently, it is of interest to test the restriction $\beta = 1$ in the data.

Table 2.2 shows results for analytical and projection-based decompositions of illiquidity measures. For each measure, the left-hand columns of Table 2.2 report statistics for regressions of ILQ_t on vol_t . Specifically, the table shows $\hat{\beta}$, the OLS estimate of β , as well as the associated standard error and *t*-statistics for tests of the null of $\beta = 0$ and the alternative null of $\beta = 1$. The null $\beta = 0$ reflects a situation in which the corresponding illiquidity measure is uncorrelated with aggregate (log) volatility. The alternative null $\beta = 1$ corresponds to the value imposed by the analytical decomposition formula. The fifth column of the table shows the R^2 -value associated with the OLS regression, expressed as a percentage.

OLS estimates of $\hat{\beta}$ are positive and statistically different from zero for most proxies. One exception is the *zeros* measure, for which the null that $\beta = 0$ cannot be rejected. Recall that the *zeros* measure is closely related to the volatility-adjusted measure $\widetilde{fht}$. The fact that the null of $\beta = 0$ cannot be rejected for *zeros* indicates that $\widetilde{fht}$ is essentially uncorrelated with aggregate volatility, i.e., the components *vol* and $\widetilde{fht}$ effect a roughly orthogonal decomposition of the

fht illiquidity measure. A second exception is the *hm* measure, for which $\hat{\beta}$ is negative and significant. This is not surprising in light of previous evidence that *hm* is negatively correlated with traditional illiquidity measures. Although volatility is positively correlated with most illiquidity measures, there is heterogeneity across measures with regard to the importance of the volatility component, as measured by the regression R^2 -value. For example, aggregate volatility explains a majority of the variation in the *roll* and *cs* measures: the regression R^2 -values for these measures are around 75% and 65%, respectively. In contrast, the R^2 -values associated with the alternative spread proxies *fht* and *tick* are roughly 15%. OLS estimates of β are close to one for the spread and price impact measures, particularly *fht* and *tick*. Indeed, the null hypothesis that $\beta = 1$ cannot be rejected for these measures. The β estimate for the *amito* price-impact measure is also statistical indistinguishable from one. Estimates of β are larger than one for *roll*, *cs*, and *ami*, and the null that $\beta = 1$ is rejected.

The apparent long memory behavior of both volatility and aggregate illiquidity raises potential concerns regarding the appropriateness of OLS as an estimator of β . As a robustness check, we also estimate β using a narrow-band least squares estimator, denoted $\hat{\beta}^{NB}$, proposed by Robinson (1994). This is a frequency-domain least squares estimator such that estimation is carried out over a degenerating band of frequencies around the origin. Robinson and Marinucci (2003) and Christensen and Nielsen (2006) provide asymptotic distributional results. To implement the estimator, it is necessary to choose a bandwidth parameter m , which we set to the integer part of $T^{0.6}$, where T is the sample size. The narrow-band least squares estimates of β are often slightly larger than the OLS estimates; however, differences are not pronounced.

The final two columns of Table 2.2 show that, for most proxies, the analytical and projection-based approaches yield very similar volatility-adjusted measures. The statistic $\hat{\rho}(ILQ^\perp, \widetilde{ILQ})$ denotes the sample correlation between the volatility-adjusted proxy $ILQ_t^\perp$ and the analytical volatility-adjusted proxy $\widetilde{ILQ}_t$. With the exception of *PS* and *hm*, the correlation between the two alternative volatility-adjusted measures is close to one. $R_{\widetilde{ILQ}}^2$ represents the implicit R^2 -value associated with the analytical decomposition, i.e., the quantity $1 - \frac{\text{Var}(\widetilde{ILQ}_t)}{\text{Var}(ILQ_t)}$, expressed as a percentage. This value need not be positive, because it is possible that the analytical

adjustment produces a fit inferior to the sample average. Empirically, the R^2 -value associated with the simple analytical decomposition is close to (but necessarily smaller than) the regression R^2 -value for most measures. Not surprisingly, the analytical decomposition does not work well for the measures PS and hm . The β estimate for PS is close to zero rather than one, while $\hat{\beta}$ for hm is negative. The null hypothesis that $\beta = 1$ is strongly rejected for both measures.

Figure 2.3 plots monthly time series for various volatility-adjusted spread proxies over the period 1926–2013 along with the unadjusted *zeros* illiquidity measure. The volatility-adjusted measures are based on the analytical decomposition of Eq. (2.10), but results are similar for measures based on the regression approach. There are three main takeaways from the plot. First, the series $\widetilde{fht}$ and $\widetilde{tick}$ are extremely similar to *zeros*. The former is consistent with the analytical decomposition of Eq. (2.6), but the plot shows that, despite the lack of an explicit algebraic decomposition, $\widetilde{tick}$ is also very highly correlated with *zeros*. Second, $\widetilde{fht}$ and $\widetilde{tick}$ exhibit sharp downward shifts upon tick-size reductions in 1997 and 2001 similar to *zeros*. This shows that the downward structural breaks associated with the raw measures *fht* and *tick* are driven by the residual component after adjusting for volatility. (Aggregate volatility does not display similar structural shifts around tick-size reductions.) Finally, although the series $\widetilde{roll}$ and $\widetilde{cs}$ exhibit similar short- to medium-run dynamics to $\widetilde{fht}$ and $\widetilde{tick}$, their long-run behavior differs. Specifically, $\widetilde{roll}$ and $\widetilde{cs}$ are higher during the Great Depression and do not respond to tick-size reductions in 1997 and 2001. The measure $\widetilde{roll}$ is considerably noisier than other measures, although its lower frequency dynamics are similar to $\widetilde{cs}$.

To avoid presenting redundant results, the subsequent analysis focuses on the projection-based measures $ILQ_t^\perp$. These measures are more sensible for the PS and hm proxies. Results in Table 2.2 imply that there is little difference between using regression-based or analytical decompositions for the remaining spread and price impact measures.

2.3.3 Break-adjusted liquidity measures

Tick-size reductions for the NYSE in 1997 and 2001 lead to considerable reductions in the aggregate effective spread relative to previous decades. The low frequency spread proxies *fht*

and *tick* also shift downward at this time, and Figure 2.3 shows that the structural shift is driven by the volatility-adjusted component of these measures. Our research focuses on cyclical fluctuations in aggregate illiquidity and connections with the business cycle and stock returns. Consequently, we attempt to control for the effects of structural shift(s) driven by tick-size reductions on the illiquidity proxies.

We break adjust the series as follows. First, we compute deviations from the mean over the sample period prior to the onset of tick-size reductions (i.e., prior to July, 1997). Between July, 1997 and the onset of decimalization (February, 2001), we break-adjust the series by computing deviations from 1997–2001 average levels. Because several of the illiquidity measures appear to gradually trend downward over the 2001–2005 period, we adjust for the structural shift over this period by computing residuals from a (linear) trend regression over the period. Finally, break-adjusted series post-2005 are computed as deviations from the mean over this period. We refer to these series as ‘break-adjusted’ and indicate this adjustment notationally via the superscript ‘B’. For example, fht_t^B indicates the break-adjusted aggregate *fht* time series. Break-adjustments are applied both to raw illiquidity proxies and volatility-adjusted versions. For example, the series $fht_t^{\perp,B}$ indicates the volatility-adjusted spread measure based on *fht*, with an adjustment for structural breaks due to tick-size reductions. The Internet Appendix contains a figure illustrating the effects of the break-adjustment procedure. As expected, break-adjustment restores the close affinity between spread measures such as *fht* and *roll* over the period following tick-size reductions.

2.4 The Predictive Power of Aggregate Illiquidity

This section addresses the question of whether aggregate liquidity measures contain information that is helpful for forecasting stock returns and various measures of economic activity.

2.4.1 Model Specification and Inference Issues

Our initial tests are based on predictive regressions that incorporate lagged illiquidity measures. Much of our analysis focuses on monthly data. The output measure at the monthly frequency is the growth rate in industrial production (Δip). Both output growth and unemployment growth exhibit some persistence. For example, the first-order sample autocorrelation of Δip in post-war data is approximately 0.4. The question, therefore, is not whether output growth is predictable – past output growth is clearly informative regarding future output growth. Rather, the question is whether lagged illiquidity provides incremental predictive power conditional on past output growth. In other words, our empirical tests focus on whether illiquidity *Granger causes* the economic variables of interest.

The predictive regressions that form the basis of our tests take the form

$$y_{t+1} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \gamma' z_t + \epsilon_{t+1}, \quad (2.12)$$

where y_{t+1} is the economic variable of interest (e.g., output growth), x_t indicates a particular aggregate illiquidity measure, and z_t represents a vector-valued set of additional macroeconomic and financial variables included in the forecasting model. The model includes J lags of the dependent variable to accommodate serial correlation. Based on information criteria, we set $J = 4$ for output and unemployment growth, and $J = 0$ for stock returns. Under the null hypothesis of no Granger causality, $\beta = 0$. This hypothesis may be tested in a standard regression framework. Our initial tests focus on regressions that do not include additional macroeconomic variables such as bond spreads or stock return volatility ($z_t = 0$ in Eq. (2.12)).

The high degree of persistence associated with various illiquidity measures complicates statistical inference in the linear predictive regression model of Eq. (2.12). Persistent regressors can cause econometric problems in predictive regressions, particularly when there is correlation between innovations in the target variable and predictor (Stambaugh (1999)). OLS estimates of slope coefficients in such regressions can be biased and/or inconsistent. In addition, the t -statistic testing the null hypothesis of no Granger causality converges to a nonstandard

asymptotic distribution when the regressor follows a local to unity process (see, e.g., Campbell and Yogo (2006) and Jansson and Moreira (2006)). The illiquidity measures of interest in this study are better characterized as long memory processes as opposed to local to unity processes. However, it turns out that long memory regressors generate similar econometric problems in predictive regressions (Maynard and Phillips (2001), Maynard, Smallwood, and Wohar (2013)).

To assess the severity of inference problems associated with predictive regressions involving aggregate illiquidity measures, we conduct a Monte Carlo simulation analysis. The simulation analysis focuses on the finite sample bias of slope coefficient estimates and the size associated with Granger causality tests. We simulate data under the null of no predictability according to the following system of equations:

$$\begin{aligned} y_t &= \mu_y + \epsilon_{1,t}, \\ x_t &= \mu_x + (1 - L)^{-d} u_t, \quad u_t = (1 - \phi L)^{-1} \epsilon_{2,t-1} \\ \epsilon_t &= (\epsilon_{1,t}, \epsilon_{2,t})' \sim \text{i.i.d. } N(0, \Sigma), \end{aligned} \tag{2.13}$$

where $|\phi| < 1$ and Σ is a symmetric, positive definite matrix with (potentially) nonzero off-diagonal elements. The contemporaneous correlation between shocks is $\rho = \Sigma_{12}/\sqrt{\Sigma_{11}\Sigma_{22}}$. The forecasting variable x potentially exhibits long-memory behavior via the parameter d . The parameter ϕ permits additional stationary short-memory dynamics. Setting $d = 0$ results in the restricted first-order VAR frequently studied in the context of predictive regressions for stock returns. The simulation analysis considers alternative ρ values ranging from -0.95 to 0.95, and values of the fractional differencing parameter d ranging from 0 to 1. Remaining parameters in the system (2.13) are calibrated based on monthly data for excess stock returns (y_t) and the *fht* illiquidity measure (x_t) over the sample period 1948–1997. In the data, calibrations of the parameter ϕ are close to zero, and consequently our main results set $\phi = 0$ for simplicity. (Unreported results with ϕ equal to in-sample estimates are similar.) We chose a sample size (T) of 480 (40 years of monthly data) and report Monte Carlo results based on 2,500 simulations.

Table 2.3 presents results. Panel A of Table 2.3 reports the simulation-based estimate of the bias associated with the OLS estimate of $\hat{\beta}$ in the predictive regression. Panel B reports the bias as a percentage of the simulation-based root mean square error (RMSE). That helps convey the economic significance of the bias associated with β . Finally, Panel C shows the simulated size of a two-sided t -test of the null that $\beta = 0$ based on the OLS estimate of β at the 5% level. Consistent with theory, the Monte Carlo results indicate essentially no bias when there is no contemporaneous correlation between shocks to y_t and x_t . As expected, the OLS estimate of β is negatively (positively) biased when $\rho > 0$ ($\rho < 0$). The magnitude of the bias increases as d increases from zero to 0.5. As d increases from 0.5 to 1.0 (the unit root case), the magnitude of the bias decreases. This is due to the fact that the simulation design holds fixed the variance of *shocks* across cases considered, but not the variance of x_t . In fact, for $d > 0.5$ the variance of x_t explodes similar to the unit root case. Upon normalizing by the RMSE (Panel B), the severity of the bias steadily increases as d rises from zero to unity. The bias is economically significant when there is long memory and when substantial correlation exists between contemporaneous shocks to y_t and x_t .

Granger causality tests are oversized for nonzero values of ρ , particularly in combination with larger values of d (Panel C). The case $d = 1$ and $\rho = -0.95$ produces rejections of the null approximately 27% of the time. This is consistent with Stambaugh (1999) and related papers focusing on the dividend yield and similar financial ratios. Although the frequency of rejections falls toward 5% as the value of d decreases from one to zero, the t -tests continue to be noticeably oversized whenever $d > 0.4$ and $|\rho| > 0.2$.

The simulation results in Table 2.3 indicate that biased coefficients and oversized Granger-causality tests are potential concerns, particularly for stock return forecasting regressions.³⁴ As an alternative to conventional inference methods, we employ a bootstrap procedure analogous to that frequently used in exchange rate and stock return forecasting regressions involving near unit-root predictors (Mark (1995), Killian (1999)). The bootstrap is implemented as

³⁴Empirical estimates of ρ reported in Table 2.1 for various illiquidity measures tend to cluster around a value of approximately -0.3 when the target variable is stock returns, but are close to zero when the target variable is output growth.

follows. First, we estimate the equation system (2.13), which imposes the null hypothesis of no predictability, using the full sample of observations. This yields a time series of estimated shocks $\hat{\epsilon}_t$. We then generate bootstrapped time series based on the system (2.13) by drawing with replacement from the estimated residuals $\hat{\epsilon}_t$. This procedure preserves the dependence structure of the predictor variable, including potential long-memory behavior, as well as the contemporaneous correlation between shocks $\epsilon_{1,t}$ and $\epsilon_{2,t}$. Additional implementation details appear in the Internet Appendix. We report bootstrap-based p -values for one- and two-sided tests of the null that $\beta = 0$ in Eq. (2.12). In addition, we report a bias-corrected estimate of β , defined as $\hat{\beta}^* \equiv \hat{\beta} - (1/B) \sum_{i=1}^B \hat{\beta}_i^*$, where $\hat{\beta}$ is the OLS estimate of β and $\hat{\beta}_i^*$ indicates bootstrap estimates of β under the null of no predictability.

2.4.2 Predictive Regressions: Main Empirical Results

Table 2.4 presents results for one month-ahead univariate stock return forecasting regressions. The table shows estimated slope coefficients, robust standard errors with associated one- and two-sided p -values, and the regression R^2 -value. For the one-sided test, the null hypothesis is $\beta \leq 0$. The right-hand portion of the table shows bias-corrected slope estimates and bootstrap-based p -values based on the bootstrap procedure discussed in Section 2.4.1.

Most raw illiquidity measures fail to successfully forecast excess market returns (Panel A). Break adjustment strengthens the case for stock return predictability (Panel B), as several break-adjusted illiquidity measures are statistically significant predictors of returns. The best-performing predictor is the break-adjusted zeros measure, which is highly significant under both standard and bootstrap-based inference procedures. Removing the volatility component further strengthens the evidence that illiquidity predicts stock returns (Panel C). The measures $fht^{\perp,B}$, $tick^{\perp,B}$, and $ami^{\perp,B}$ are all highly significant, irrespective of whether inference is based on conventional or bootstrap methods. In addition, the R^2 -values associated with these variables are in excess of 1.4% and compare favorably to in-sample R^2 -values for a large body of popular stock return forecasting variables considered by Welch and Goyal (2008). Two other measures, $cs^{\perp,B}$ and $amito^{\perp,B}$, exhibit positive slope coefficients and p -values that are

significant at the 10% level based on one-sided tests. The slope coefficients associated with all of these variables are positive, so that an increase in volatility-adjusted illiquidity is associated with higher expected returns, in accordance with theory.³⁵

The improved stock return-forecasting performance of volatility-adjusted measures is driven by the fact that the estimated return-forecasting coefficient associated with aggregate volatility (*vol*) is negative. The slope coefficients for unadjusted illiquidity measures are essentially weighted averages of return-forecasting coefficients for aggregate volatility and the corresponding volatility-adjusted illiquidity measure. The negative return-forecasting coefficient for volatility pulls the return-forecasting coefficients for raw illiquidity measures closer to zero, and removing this component reveals a stronger predictive relation. Bootstrap-based inference results are generally consistent with standard inference methods. A feature that mitigates the “Stambaugh bias” in our setting is the relatively modest correlation between contemporaneous shocks to stock returns and illiquidity measures. In addition, volatility-adjusted measures behave similarly to the *zeros* measure, and shocks to this measure are nearly uncorrelated with stock returns (see Table 2.1). This helps explain why bias-corrections to slope estimates are more pronounced for unadjusted and break-adjusted measures in Panels A and B, relative to the volatility- and break-adjusted measures in Panel C.

Table 2.5 presents one month-ahead forecasting results for two macroeconomic targets for the period 1948–2013: the growth rate in industrial production Δip and the change in unemployment rate Δue .³⁶ The table presents results analogous to those for stock return forecasting regressions in Table 2.4 with two modifications. First, we omit *p*-values for one-sided tests to conserve space. Second, because both Δip and Δue exhibit serial correlation, we include four lags in the benchmark specification ($J = 4$ in Eq. (2.12)). Consequently, we report ΔR^2 , defined as the increase in the R^2 -value for the predictive regression relative to this benchmark univariate autoregressive model, expressed as a percentage.

³⁵The slope coefficient associated with the *hm* measure is negative, rather than positive, which is not surprising in light of the fact that this measure is procyclical rather than countercyclical.

³⁶Unemployment growth is not available prior to 1947, and we include 4 lags in the predictive regression; hence the sample start point of 1948. We also estimated predictive regressions for industrial production and stock returns incorporating earlier data. Industrial production is erratic during the period from the Great Depression through World War II. Results are discussed briefly later in the study.

The results in Table 2.5 indicate that aggregate illiquidity Granger causes economic activity. Consistent with economic theory and intuition, higher illiquidity levels forecast lower future output growth (*ip*) and higher future unemployment (*ue*). Among the raw illiquidity measures (Panel A) the spread measures *roll* and *cs*, several of the Amihud measures, and the *PS* measure significantly predict growth in industrial production, unemployment, or both, controlling for lags of the respective economic activity measure. The incremental R^2 is nontrivial, ranging from around 0.5% to 2.0% for specifications with significant predictors. The break-adjusted price delay measure $\widetilde{hm}$ forecasts real activity; however, the direction of the relation differs from that of the illiquidity variables. This reflects the previously documented fact that aggregate price delay tends to decrease, rather than increase, in recessions.

Although several raw illiquidity measures (e.g., *fht* and *tick*) do not appear to predict future economic activity, Panel B of Table 2.5 shows that break-adjusted versions of these measures do forecast real activity. Indeed, the vast majority of break-adjusted measures are significant predictors of growth in both *ip* and *ue*. Interestingly, the *zeros* measure does *not* appear to forecast economic activity, even after break-adjustment. Section 2.2 shows that the *zeros* measure can be interpreted as a spread measure adjusted for volatility. Consequently, the results in Panel B of Table 2.5 suggest that at least a portion of the forecasting power of illiquidity measures for real activity arises from the volatility component embedded in these measures. Panel C provides explicit evidence on this point by reporting forecasting results for aggregate volatility and, separately, volatility-adjusted illiquidity measures. Aggregate volatility significantly forecasts growth in industrial production and growth in unemployment. The associated slope coefficients take the expected sign: high volatility predicts reduced economic activity, controlling for past activity. Separating out the volatility component weakens the predictive power associated with most illiquidity measures. Most slope coefficients associated with volatility-adjusted measures in Panel C are statistically insignificant, and ΔR^2 -values generally fall relative to analogous values reported in Panel B. There is some evidence of predictability associated with volatility-adjusted measures, however. For example, the volatility-adjusted *amito* measure continues to forecast growth in industrial production and unemployment.

There is little difference in inference based on the bootstrap method versus standard methods. In addition, bootstrap-based bias correction produces little change in slope coefficients relative to OLS estimates. The small magnitude of bias and similarity of inference results derive from the fact that shocks to macroeconomic growth measures are close to uncorrelated with contemporaneous shocks to illiquidity measures.

The persistent nature of illiquidity measures implies that forecasting coefficients and R^2 -values should increase with the forecasting horizon. Table 2.6 presents results for longer horizon predictive regressions for industrial production, unemployment, and stock market returns. The regression specification takes the form:

$$\sum_{h=1}^H y_{t+h} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \epsilon_{t+H}. \quad (2.14)$$

When the forecasting target is excess market returns, the dependent variable is the excess market return over the next H -months. For industrial production (unemployment), the dependent variable in specification (2.14) corresponds to growth in industrial production (unemployment) over the following H months. The table presents estimated slope coefficients and ΔR^2 -values for forecast horizons of 6 and 12 months. To conserve space, we only present results for break-adjusted measures (Panel A) and break- and volatility-adjusted measures (Panel B).

The Granger causality test results for longer-horizon regressions are similar to those at the one-month horizon.³⁷ The economic significance of the predictive power afforded by the illiquidity measures increases with the forecast horizon. This is particularly true for stock return forecasts, where slope coefficients and R^2 -values increase noticeably for 6- and 12-month return forecasts relative to one-month-ahead results. R^2 -values for significant return-forecasting measures (*zeros*, *tick*, etc.) fall in the 3–6% range at the 6-month horizon and in the 7–12%

³⁷Because the regression involves overlapping observations, we base inference on Newey-West standard errors employing $H + 2$ lags. These inference results are likely imperfect due to the long memory characteristic of illiquidity measures. We prefer to rely on previously presented one-step ahead test results for Granger causality questions. Our interest in longer horizon regressions lies in characterizing the economic significance of predictive relations.

range at a 12-month horizon. These horizon effects are similar to long-run stock return forecasts for other persistent predictors in the literature, most notably the dividend yield.

The economic significance of the predictive content in illiquidity for output growth (Δip) and unemployment growth (Δue) also increases at longer horizons; however, ΔR^2 -values do not increase much between 6- and 12-month forecast horizons. As with the short-horizon case, the volatility component embedded in illiquidity measures accounts for most of the predictive power for macroeconomic outcomes. Many of the volatility-adjusted illiquidity measures in Panel B are insignificant even over the longer forecasting horizons, and ΔR^2 -values tend to be relatively small in the limited cases where such variables are statistically significant. Some volatility-adjusted measures do perform significantly better than the benchmark at longer horizons. For example, the (break adjusted) price impact measure $\widetilde{amito}$ performs particularly well for unemployment growth ($\Delta R^2 \approx 3.5\%$).

It is possible that apparent forecasting power associated with illiquidity measures simply reflects the fact that these measures are correlated with other financial variables frequently employed in forecasting models. Indeed, earlier results indicate that the volatility component embedded in illiquidity measures accounts for at least a portion of these variables' ability to predict economic activity. Table 2.7 presents results for selected measures when additional financial forecasting variables z_t appear in the predictive regression of Eq. (2.12). Results in Panel A add aggregate volatility vol_t and two lags of market (excess) returns to the model. Results in Panel B include these variables, as well as the short interest rate, the term spread, the default spread, and the commercial paper-to-Treasury spread.

We highlight two key insights from the results in Table 2.7. First, many volatility-adjusted illiquidity measures remain statistically significant predictors of stock returns upon including the short rate, default spread, and term spread in the forecasting model. This is notable, since the short rate and bond spreads are popular return forecasting instruments in the literature. Controlling for these variables slightly reduces the incremental predictive power associated with volatility-adjusted illiquidity, but the ΔR^2 -values associated with the illiquidity measures remain economically significant. Second, a number of volatility-adjusted illiquidity

measures remain significant or borderline-significant predictors of economic activity after including additional financial predictors. This is consistent with the general spirit of results in Næs, Skjeltorp, and Ødegaard (2011), although the associated ΔR^2 -values tend to be small. Interestingly, including interest rate information in the form of the short rate and bond spreads (Panel B) actually increases the predictive power of volatility-adjusted illiquidity measures for real activity relative to including only volatility and lagged returns (Panel A).

A drawback of relying on monthly data is that many standard macroeconomic variables, including GDP, are only available at the quarterly frequency. To further explore relations between aggregate liquidity and macroeconomic quantities, we run predictive regressions for measures of quarterly economic activity using (lagged) quarterly liquidity proxies. We examine five target variables: 1) excess stock market returns *ret*, 2) the growth rate in GDP Δgdp , 3) the growth rate of investment Δinv , 4) the growth rate of consumption $\Delta cons$, and 5) the growth rate in unemployment Δue . The latter four variables are studied by Næs, Skjeltorp, and Ødegaard (2011). At the quarterly frequency we include 2 lags of the dependent variable for measures of real economic activity, and no lags for excess return forecasting regressions.

Results in Table 2.8 are qualitatively similar to those at the monthly frequency. Few raw illiquidity measures forecast stock returns; however, the majority of volatility- and break-adjusted illiquidity measures (Panel B) significantly forecast returns. The associated return predictability is economically significant, with R^2 -values mostly in the 1–3% range for successful predictors. Turning to measures of economic activity, raw illiquidity measures less subject to structural breaks (*roll*, *cs*, Δami , and *PS*) forecast economic activity. A number of other measures, including *fht* and *ami*, become significant following break adjustment (results omitted to conserve space). Slope coefficients generally take the expected sign, i.e., indicate that higher illiquidity corresponds to weaker future economic activity.³⁸ Panel B shows results upon decomposing illiquidity measures into a volatility component and a (break-adjusted) residual

³⁸ Among raw measures, the differenced Amihud measure Δami forecasts economic activity particularly well at the quarterly frequency. This measure was analyzed by Næs, Skjeltorp, and Ødegaard (2011), and our results are consistent with theirs. The ΔR^2 -value for this variable falls considerably when additional financial variables are included in the specification, although the variable remains significant for GDP growth and investment growth. (See the Internet Appendix for explicit results.)

component. Similar to results at the monthly frequency, aggregate volatility forecasts real activity for quarterly data, while most volatility-adjusted measures are insignificant. The Internet Appendix provides additional results that incorporate other financial predictors. Granger causality test results for the illiquidity measures remain qualitatively similar.

2.4.3 Out-of-sample forecasting performance

Stock and Watson (2003) find that forecasts based on individual variables derived from asset prices are unstable. Many recent papers in the stock return predictability literature assess predictive ability using an out-of-sample research design. There remains contention in the literature regarding whether in- or out-of-sample model comparison is preferred. Without taking a stand on this issue, this section provides an out-of-sample analysis to complement earlier in-sample results.³⁹

The out-of-sample analysis proceeds as follows. Forecasting models are estimated using a rolling procedure with an initial estimation sample of 30 years of data. Main results are based on monthly data with an initial estimation sample of 1948–1977, and associated out-of-sample forecasts cover the period 1978–2013. Out-of-sample forecasts based on models incorporating illiquidity measures are compared to forecasts from benchmark models that omit such variables. The choice of benchmark depends on the degree of persistence in the forecasting target. The benchmark model for output growth (Δip) and unemployment growth (Δue) includes four lags. Following most literature concerning stock return predictability, the benchmark forecast for stock market returns is the historical average. Results are qualitatively similar for various modifications of the out-of-sample design (see Section 2.4.4 and the Internet Appendix). Although our analysis follows a standard ‘pseudo out-of-sample’ design, it is important to note that we include results for break-adjusted measures. Results for such measures provide ro-

³⁹For current perspectives regarding in-sample versus out-of-sample inference, see Diebold (2015) and published commentary on this article in the same issue.

bustness checks regarding in-sample findings, but do not reflect feasible real-time forecasting methods.⁴⁰

Let y_{t+1} represent the forecasting target and $\hat{y}_{t+1}$ denote a one-step ahead forecast of y_{t+1} from a predictive regression. The corresponding sample mean square prediction error (MSPE) for the predictive regression is

$$\hat{\sigma}^2 \equiv P^{-1} \sum (y_{t+1} - \hat{y}_{t+1})^2, \quad (2.15)$$

where P denotes the number of out-of-sample forecasts. Similarly, let $\hat{y}_{t+1}^{BMK}$ denote the one-step ahead forecast based on the benchmark model, with corresponding sample MSPE equal to $\hat{\sigma}_0^2$. We report the out-of-sample R^2 statistic (Campbell and Thompson (2008) and Welch and Goyal (2008)). This statistic is defined as

$$R_{OOS}^2 = 1 - \frac{\hat{\sigma}^2}{\hat{\sigma}_{BMK}^2}, \quad (2.16)$$

expressed as a percentage. The R_{OOS}^2 statistic serves as the out-of-sample analog of the ΔR^2 statistic reported for in-sample regressions.

Although it is informative, the R_{OOS}^2 -value represents a particular sample realization. Hence, it is of interest to test whether differences in forecast performance are statistically significant. Statistical inference associated with out-of-sample forecast performance is nuanced when forecasts rely on estimated parameters (West (1996)). We conduct two alternative tests for differences in out-of-sample forecast performance. The first tests the null hypothesis of no Granger causality ($\beta = 0$ in Eq. (2.12)) in an out-of-sample design. A variety of approaches to inference exist in the literature (see Clark and McCracken (2011) for a comprehensive review). We opt to apply a conceptually simple approach proposed by Clark and West (2007). The test is based

⁴⁰It can be argued that the break dates associated with tick-size reductions were known ex ante. In additional, unreported results, we find that a backward-looking procedure that stochastically detrends measures beginning in 1997 performs similarly to our main break-adjusted measures.

on the quantity

$$CW = \hat{\sigma}_{BMK}^2 - \hat{\sigma}^2 + \underbrace{P^{-1} \sum (\hat{y}_{t+1} - \hat{y}_{t+1}^{BMK})^2}_{\text{Adjustment}}. \quad (2.17)$$

The final term in (D.0.1) captures an adjustment for additional noise associated with the larger model's forecast under the null hypothesis. The test statistic associated with CW is not asymptotically normal; however simulation evidence in Clark and West (2007) indicates that standard normal critical values lead to tests with actual sizes close to nominal size. The test is one-sided, because, under the alternative, $\sigma^2 < \sigma_{BMK}^2$.

We also apply the test for superior predictive ability proposed by Giacomini and White (2006). The Giacomini and White (2006) test focuses on differences in finite sample predictive ability, i.e., predictive ability incorporating the impact of parameter estimation on forecast performance. The distinction is important, because a variable may Granger cause stock returns, but fail to forecast returns more accurately than the benchmark forecast when parameters are estimated using finite samples, due to the bias-variance tradeoff. The Giacomini and White (2006) test statistic for equal (unconditional) predictive ability is:

$$\text{Giacomini-White} = \frac{\hat{\sigma}_0^2 - \hat{\sigma}^2}{\hat{\sigma}_P / \sqrt{P}}, \quad (2.18)$$

where $\hat{\sigma}_P$ is a heteroskedasticity and autocorrelation consistent (HAC) estimator of the asymptotic variance $\sigma_P^2 = \text{var} \left[\sqrt{P}(\hat{\sigma}_0^2 - \hat{\sigma}^2) \right]$. In contrast with the Clark-West test, the Giacomini-White test is two-sided.

Table 2.9 presents results for the out-of-sample analysis for one-month-ahead forecasts using monthly data. For each variable, the table shows the ΔR_{OOS}^2 summarizing out-of-sample performance, as well as the t -statistic associated with the Clark and West (2007) and Giacomini and White (2006) tests. To conserve space, the table first presents results for a subset of raw illiquidity measures less affected by structural shifts around tick-size reductions (Panel A), whereas Panel B presents results for volatility- and break-adjusted illiquidity measures. For comparison, Panel C presents results for aggregate volatility and other financial variables.

Consistent with in-sample results, out-of-sample Granger causality tests fail to reject the null of no predictability for most raw illiquidity measures, and most yield a negative ΔR_{OOS}^2 value. The volatility- and break-adjusted illiquidity measures in Panel B fare better. Six of the nine break- and volatility-adjusted measures outperform the benchmark out-of-sample. In contrast, and consistent with the spirit of results in Welch and Goyal (2008), aggregate volatility and the other financial predictors generally fail to outperform the historical average (Panel C). In addition, the Clark-West test rejects the null of no predictability for the break adjusted zeros measure $zeros^B$, as well as for $fht^{\perp,B}$, $tick^{\perp,B}$, and $ami^{\perp,B}$. Similarly, the Giacomini-White test rejects the null of equal predictive ability for these variables.

Clark-West test results indicate that the *roll* and *cs* raw illiquidity measures forecast output growth, and all three raw measures forecast unemployment growth. The corresponding out-of-sample R^2 -values are positive, indicating that forecasts employing the illiquidity measures outperform the benchmark model; however, the null of equal predictive ability in the Giacomini-White test cannot be rejected. Most volatility- and break- adjusted measures do not successfully outperform the univariate benchmark out-of-sample for growth in industrial production or unemployment (Panel B). This is largely due to removing the volatility component from the measures. Panel C shows out-of-sample results for volatility, lagged returns, various bond spreads, and the short interest rate for comparison to illiquidity measures. Aggregate volatility Granger causes real activity based on the Clark-West test, and the out-of-sample R^2 -value is positive, although insignificant based on the Giacomini and White test.

Figure 2.4 plots the cumulative out-of-sample MSE difference between one month-ahead stock return forecasts based on selected illiquidity measures, and the historical average benchmark. This form of plot, popularized by Welch and Goyal (2008), illustrates the time series behavior of forecast performance relative to the benchmark. An upward slope indicates periods in which forecasts based on the corresponding illiquidity measure outperform the benchmark, whereas a downward slope indicates periods in which the benchmark is more accurate. Forecasts for models incorporating illiquidity measures are based on Eq. (2.12), with no lags included ($J = 0$). The figure shows results for three of the better performing measures based on results

in Table 2.9: the break-adjusted zeros measure $zeros^B$ and the volatility- and break-adjusted measures $tick^{\perp,B}$ and $ami^{\perp,B}$. To provide additional insight and demonstrate robustness of key results, we base Figure 2.4 on an analysis in which the initial estimation sample begins in 1926. This yields a longer time series of out-of-sample performance differences covering 1956:2 – 2013:12.⁴¹

Stock return forecasts based on illiquidity measures outperform the benchmark over the longer out-of-sample analysis period. The plot nearly always slopes upward or is roughly flat. In addition, there are several periods during which the illiquidity measures produce much more accurate forecasts relative to the benchmark. Most notable is the period from the onset of the financial crisis onward. Although this period yields the greatest gains relative to the benchmark, illiquidity measures perform well during several other periods, including around October 1987 and during the early 1970s. The time series of performance differences is similar under an alternative recursive design in which the estimation window expands over time. Overall, the results indicate that successful illiquidity measures produce stock return forecasts that generally perform at least as well as the historical average, and significantly outperform the benchmark in periods characterized by sharp liquidity contractions.

2.4.4 Robustness checks

We conduct a number of robustness checks for our main results. A potential concern is that the break-adjustment process applied to many illiquidity measures induces spurious forecasting power. We address this concern in two ways. First, we apply an alternative break-adjustment procedure that is backward-looking, and consequently feasible in real time. This approach adjusts for structural change driven by tick size reductions by computing deviations from the mean prior to 1997, and residuals from a two-year historical moving average starting in 1997. Our main results regarding the predictive power of break- and volatility-adjusted illiquidity series are robust to this alternative approach. As a second test, we run forecasting regressions using raw illiquidity measures for sample periods that end prior to tick-size reductions begin-

⁴¹Table 2.9 begins estimation in 1948 so that all three series considered (stock returns, output growth, and unemployment growth) share common sample periods.

ning in 1997. Key results regarding the predictive content of illiquidity measures continue to obtain over this sample period, illustrating that our findings are not driven by break-adjustment procedures. In particular, the (unadjusted) *zeros* measure remains a significant predictor of returns at both the one-month and one-year ahead horizons, and the ΔR^2 -value at the one-year horizon is approximately 9%. Most volatility-adjusted measures also significantly forecast returns over this sample period. In fact, we find stronger evidence of stock return forecasting power for the volatility-adjusted measures $roll^\perp$ and $cs^\perp$ over this sample period, suggesting that the predictive power of these measures breaks down somewhat in recent years.

A second potential concern is that key forecasting results might be driven by a ‘look-ahead’ bias related to the need to estimate the parameter β in the projection of illiquidity measures on aggregate volatility. However, measures based on the analytical decomposition formula, which simply subtracts aggregate log volatility, produce similar forecasting results to those based on the regression approach.

Our aggregate illiquidity measures are constructed using NYSE-listed firms. We alternatively compute measures using the NYSE/AMEX universe of firms and obtain similar forecasting results. Basing out-of-sample forecasts on a recursive, rather than rolling, scheme does not alter our key findings.

2.5 Concluding Remarks

This study constructs and analyzes long time series of measures of illiquidity. Illiquidity measures contain a component related to the US business cycle. Aggregate measures of trading costs contain important predictive signals for stock market returns and real economic activity. We decompose illiquidity proxies into a component capturing aggregate volatility and a residual component. The two components contain differential predictive content. Specifically, we find strong and robust evidence that the component of illiquidity uncorrelated with volatility forecasts stock market returns. On the other hand, the aggregate volatility component accounts for most of the predictive content of illiquidity for real economic activity.

Several directions for future research seem promising. First, the long time series of illiquidity measures we construct potentially constitute the basis for new empirical work on liquidity and asset prices. An important strand of literature considers liquidity to be a priced risk factor (e.g., Pástor and Stambaugh (2003), Acharya and Pedersen (2005)). Existing empirical literature tests whether securities' exposure to shocks to aggregate liquidity (measured by proxies) are priced. This study shows that aggregate liquidity measures contain a component reflecting aggregate volatility, and that predictive content associated with this component differs from the residual component. These results motivate testing for cross-sectional variation in expected returns as a function of exposure to shocks to the two components separately. Second, the availability of high-frequency, intraday data in recent decades affords the potential to compute improved measures of liquidity, such as the percentage effective spread, realized spread, and various high-frequency measures of price impact (see, e.g., Hasbrouck (2009) and Goyenko, Holden, and Trzcinka (2009)). An approach similar to that developed in this study might be used to decompose these high-frequency measures into a component reflecting volatility and a residual, with an eye toward analyzing the empirical properties of the two components.




Figure 2.1: Alternative effective spread proxies: 1926-2013

This figure shows monthly time series for alternative low-frequency proxies for the aggregate effective spread over the period 1926–2013. The top panel shows the spread proxies in levels, while the bottom panel plots log versions of the series. *ROLL* is the proxy of Roll (1984). *CS* is the proxy proposed by Corwin and Schultz (2012). *FHT* is the proxy proposed by Fong, Holden, and Trzcinka (2014). *TICK* is the effective tick measure proposed by Holden (2009). Log versions of series are denoted using lowercase letters (e.g., *roll*).




Figure 2.2: Alternative effective spread proxies: 1993-2013

This figure shows monthly time series for alternative low-frequency proxies for the aggregate effective spread over the period 1993–2013. $ROLL$ is the proxy of Roll (1984). CS is the proxy proposed by Corwin and Schultz (2012). FHT is the proxy proposed by Fong, Holden, and Trzcinka (2014). $TICK$ is the effective tick measure proposed by Holden (2009). $ESPD_t$ is the effective spread computed using the TAQ data over the period 1993–2007. Log versions of series are denoted using lowercase letters (e.g., $roll$).




Figure 2.3: Spread proxies adjusted for aggregate volatility

This figure depicts monthly time series for various (log) spread proxies adjusted for aggregate volatility, along with the raw illiquidity measure $zeros$ for comparison. All series are standardized. The sample period is 1926–2013.




Figure 2.4: Out-of-sample Stock Return Forecasting Performance

This figure plots the out-of-sample cumulative difference in mean square error, relative to the historical average benchmark, for one month-ahead stock return forecasts for selected illiquidity measures. Forecasts for models incorporating illiquidity measures are based on Eq. (2.12), with no lags included ($J = 0$). The figure shows results for three of the better performing measures based on results in Table 2.9: the break-adjusted zeros measure $zeros^B$, and the volatility- and break-adjusted measures $tick^{\perp,B}$ and $ami^{\perp,B}$. The initial out-of-sample forecast is based on parameters estimated using a 30-year historical sample beginning in 1926:2. Consequently, the initial out-of-sample forecast is for 1956:2. Panel A shows results under a rolling estimation scheme, while Panel B shows results under a recursive estimation scheme.

Table 2.1: Descriptive Statistics: Raw Illiquidity Measures and Selected Macroeconomic and Financial Variables

This table presents summary statistics for aggregate illiquidity measures and selected macroeconomic and financial variables. Panel A presents the sample mean, standard deviation, skewness and kurtosis for each variable, along with several statistics characterizing the degree of persistence in the time series. Sample means and standard deviations for the variables Δip , Δue and ret are monthly percentages. $\hat{\rho}(1)$, and $\hat{\rho}(36)$ are the sample autocorrelations at lags 1 and 36, respectively. $\hat{\rho}(1)^{36}$ is the correlation at lag 36 implied by the sample autocorrelation $AC(1)$, assuming a first order autoregressive (AR(1)) process for the corresponding variable. $\hat{d}$ is the exact local Whittle estimate of the long-memory parameter d proposed by Shimotsu and Phillips (2005). Asymptotic standard errors are provided for this statistic. ADF is the augmented Dickey-Fuller test statistic for the null hypothesis of a unit root, including an intercept, with autoregressive order determined by the AIC criterion with a maximum lag length of 18. The associated simulated p -value is also provided. The top portion of Panel B provides selected pairwise sample correlation estimates. The bottom portion of Panel B shows the estimated correlation between shocks to the corresponding variable and shocks to either excess returns or growth in industrial production. See Section 2.2 for variable definitions. The sample period is monthly for 1948.2-1997.6.

	Illiquidity measures								Other variables				
	<i>roll</i>	<i>cs</i>	<i>fht</i>	<i>tick</i>	<i>zeros</i>	<i>ami</i>	<i>amito</i>	<i>PS</i>	<i>hm</i>	<i>ret</i>	Δip	Δue	$\overline{vol}$
Panel A: Univariate descriptive statistics													
Mean	-4.66	-4.91	-4.63	-4.54	-1.67	0.68	3.27	0.01	-0.45	0.58	0.29	0.07	-3.93
St. Dev.	0.23	0.24	0.28	0.30	0.17	0.71	0.32	0.02	0.16	4.06	1.05	4.19	0.21
Skew.	0.78	0.62	0.34	0.61	-0.51	0.14	0.55	2.65	-1.52	-0.65	0.30	0.46	0.70
Kurt.	4.08	3.10	3.06	3.64	3.22	3.55	3.89	17.77	6.75	5.89	7.49	6.96	4.89
$\hat{\rho}_1$	0.66	0.96	0.97	0.98	0.88	0.96	0.88	0.37	0.33	0.05	0.40	0.14	0.83
$\hat{\rho}_{36}$	0.22	0.45	0.54	0.48	0.34	0.38	0.12	0.03	0.07	0.01	-0.02	-0.09	0.32
$\hat{\rho}_1^{36}$	0.00	0.21	0.30	0.41	0.01	0.26	0.01	0.00	0.00	0.00	0.00	0.00	0.00
$\hat{d}$	0.66	0.83	0.87	0.89	0.76	0.85	0.63	0.35	0.36	-0.03	0.02	0.12	0.66
$SE(\hat{d})$	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10	0.10
ADF	-3.00	-3.67	-2.72	-2.66	-2.98	-3.47	-3.42	-4.55	-7.17	-9.23	-8.39	-7.65	-3.52
p -value	0.04	0.01	0.07	0.08	0.04	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.01
Panel B: Selected Correlations													
$roll$	1.00	0.80	0.76	0.77	0.25	0.70	0.40	0.40	-0.34	-0.12	-0.21	0.16	0.61
ami	0.70	0.87	0.83	0.90	0.66	1.00	0.50	0.15	-0.21	-0.01	-0.21	0.19	0.31
ret	-0.12	-0.02	-0.01	-0.02	0.10	-0.01	-0.21	-0.32	0.39	1.00	0.00	-0.03	-0.14
Δip	-0.21	-0.23	-0.19	-0.20	-0.10	-0.21	-0.24	-0.12	0.07	0.00	1.00	-0.51	-0.16
Δue	0.16	0.20	0.15	0.16	0.09	0.19	0.25	0.13	-0.09	-0.03	-0.51	1.00	0.12
$\overline{vol}$	0.61	0.55	0.64	0.57	-0.30	0.31	0.31	0.46	-0.37	-0.14	-0.16	0.12	1.00
$\hat{\rho}(\epsilon_{RET}, \epsilon_X)$	-0.27	-0.37	-0.33	-0.48	0.04	-0.41	-0.46	-0.33	0.48				
$\hat{\rho}(\epsilon_{\Delta ip}, \epsilon_X)$	-0.07	-0.03	-0.07	-0.05	-0.09	-0.09	-0.08	-0.03	-0.01				

Table 2.2: Extracting a Volatility Component from Illiquidity Measures

The table presents results for projection-based and analytical decompositions of illiquidity measures into a volatility component and a residual. Let ILQ_t denote a generic illiquidity measure. Projection-based decompositions are based on regressions of the form

$$ILQ_t = \alpha + \beta \overline{vol}_t + \epsilon_{t+1},$$

where $\overline{vol}_t$ denotes (log) average volatility. For each illiquidity measure, the Table reports $\hat{\beta}$, the OLS estimate of β , as well as the associated standard error. t -statistics are provided for the typical null of $\beta = 0$ and for the null $\beta = 1$, which is imposed under the analytical decomposition. The fifth column of the table shows the R^2 -value associated with the OLS regression, expressed as a percentage. $\hat{\beta}^{NB}$ denotes the narrow-band least squares estimate of β with a bandwidth of $T^{0.6}$. $\hat{\rho}(ILQ_t^\perp, \widetilde{ILQ})$ denotes the sample correlation between the volatility-adjusted proxy $ILQ_t^\perp$, defined as the residuals from the regression of ILQ_t on $\overline{vol}_t$, and the analytical volatility-adjusted proxy $\widetilde{ILQ}_t$, defined as $ILQ_t - \overline{vol}_t$. $R_{\widetilde{ILQ}}^2$ denotes the R^2 -value associated with the analytical decomposition, i.e., the quantity $1 - \frac{\text{Var}(\widetilde{ILQ}_t)}{\text{Var}(ILQ_t)}$, expressed as a percentage. This value need not be positive, because the analytical adjustment may produce a fit inferior to the sample average. The corresponding values for PS and hm are in fact negative and large (in magnitude). To conserve space, we indicate this as " $\ll 0$." The sample period is 1926–2013. See Section 2.2 for variable definitions.

ILQ	OLS estimation of β					$\hat{\beta}^{NB}$	$\hat{\rho}(ILQ_t^\perp, \widetilde{ILQ})$	$R_{\widetilde{ILQ}}^2$
	$\hat{\beta}$	SE	$t(\beta = 0)$	$t(\beta = 1)$	R^2			
<i>roll</i>	1.19	0.07	18.11	2.95	75.55	1.26	0.96	73.55
<i>cs</i>	1.62	0.17	9.74	3.72	64.65	1.80	0.89	55.20
<i>fht</i>	1.10	0.30	3.63	0.33	16.57	1.18	1.00	16.43
<i>tick</i>	0.95	0.26	3.57	-0.21	14.52	1.01	1.00	14.48
<i>zeros</i>	-0.05	0.30	-0.16	-3.49	0.04	0.01	0.92	-18.95
<i>ami</i>	1.43	0.23	6.15	1.86	30.19	1.55	0.98	27.43
<i>amito</i>	0.71	0.19	3.78	-1.56	17.80	0.76	0.98	14.76
<i>PS</i>	0.02	0.01	3.17	-154.80	7.31	0.02	0.07	$\ll 0$
<i>hm</i>	-0.26	0.07	-3.57	-17.10	14.78	-0.19	0.45	$\ll 0$

Table 2.3: Finite-Sample Bias in Predictive Regressions with Long Memory Variables

This table presents Monte Carlo evidence regarding the properties of the OLS estimator of the slope coefficient β in the predictive regression $y_{t+1} = \alpha + \beta x_t + \epsilon_{t+1}$ under the null hypothesis of no predictability ($\beta = 0$), when the forecasting variable x_t follows the process $(1-L)^d x_t = \nu_t$, where d represents the long memory parameter. Panel A presents the Monte Carlo estimate of the bias associated with $\hat{\beta}$, the OLS estimate of the slope coefficient in the predictive regression (multiplied by 100 for readability). Panel B shows the magnitude of the bias, expressed as a percentage of the root mean square error (RMSE) associated with the estimator, where the RMSE is computed based using the Monte Carlo results. Panel C addresses the size of the standard (two-sided) t -test of the null hypothesis that $\beta = 0$ based on the OLS estimator. Specifically, Panel C reports the fraction of simulations in which the null hypothesis is rejected at the 5% level. Results are based on 2,500 simulations.

d/ρ	-0.95	-0.9	-0.8	-0.4	-0.2	0	0.2	0.4	0.8	0.9	0.95
Panel A: Bias $((1/S) \sum_{i=1} \hat{\beta}_i)$											
0.0	0.00	0.06	0.05	0.05	-0.05	-0.12	-0.01	-0.01	-0.15	-0.12	-0.12
0.2	0.45	0.42	0.31	0.22	0.07	-0.01	0.05	-0.20	-0.44	-0.36	-0.51
0.4	0.88	0.85	0.70	0.41	0.21	0.03	-0.26	-0.44	-0.87	-0.87	-0.94
0.5	1.24	1.11	1.03	0.56	0.23	0.02	-0.22	-0.47	-1.06	-1.26	-1.24
0.6	1.36	1.26	1.13	0.60	0.28	-0.02	-0.25	-0.57	-1.16	-1.27	-1.37
0.8	1.14	1.10	0.95	0.46	0.26	-0.01	-0.23	-0.46	-0.94	-1.09	-1.18
1.0	0.63	0.61	0.54	0.27	0.14	-0.01	-0.15	-0.27	-0.50	-0.61	-0.65
Panel B: Absolute value of bias as % of RMSE											
0.0	0.17	2.20	1.91	1.69	1.71	4.17	0.46	0.38	5.52	4.61	4.49
0.2	16.77	15.18	11.48	8.22	2.48	0.29	1.83	7.24	16.07	13.22	18.54
0.4	37.27	35.89	29.40	17.28	9.43	1.15	11.42	19.06	36.19	36.50	38.85
0.5	54.37	50.39	46.89	27.61	11.68	0.85	11.20	23.62	48.51	54.05	54.47
0.6	66.23	62.76	58.20	35.75	17.25	1.26	15.39	34.20	58.79	62.90	65.52
0.8	76.51	74.75	70.53	45.83	27.95	1.68	24.65	45.90	70.49	74.20	76.59
1.0	76.58	74.64	72.75	52.35	32.17	1.76	32.89	53.09	71.40	75.46	75.31
Panel C: Empirical size of t-test (%)											
0.0	5.92	5.00	5.24	5.44	4.96	5.40	4.64	5.32	5.44	4.36	5.28
0.2	5.20	5.48	5.20	4.92	5.28	5.48	5.36	5.24	5.92	5.88	5.32
0.4	5.40	5.84	6.04	6.00	4.68	5.36	5.40	5.28	6.32	5.84	5.96
0.5	7.64	6.36	7.04	6.28	5.24	5.08	5.32	5.84	6.72	8.24	7.68
0.6	9.60	9.16	8.44	6.40	6.32	5.12	6.04	6.04	8.40	9.16	10.64
0.8	18.08	17.04	14.00	7.56	6.64	5.32	6.12	7.56	13.92	17.16	18.40
1.0	27.36	25.52	19.76	9.24	5.80	5.04	6.32	9.20	20.20	25.92	28.24

Table 2.4: Forecasting regressions for stock returns: Monthly frequency

The table presents results for one-month-ahead predictive regressions for excess returns on the NYSE index using various measures of aggregate illiquidity. The table reports results from regressions of the form

$$y_{t+1} = \alpha + \beta x_t + \epsilon_{t+1},$$

where y_{t+1} is the dependent variable and x_t represents the specified forecasting variable. For each forecasting variable, the table presents OLS and bootstrap estimates. The presented OLS estimates include the slope coefficient $\hat{\beta}$, the associated Newey-West standard error, a one-sided p -value ($p\text{-val}_{1s}$), a two-sided p -value ($p\text{-val}_{2s}$), and an R^2 value. The presented bootstrap estimates include the slope coefficient $\hat{\beta}^*$ and the associated one-sided p -value ($p\text{-val}_{1s}$) and two-sided p -value ($p\text{-val}_{2s}$). The sample period is 1948–2013. See Section 2.2 for variable definitions.

	OLS					Bootstrap		
	$\hat{\beta}$	SE	$p\text{-val}_{1s}$	$p\text{-val}_{2s}$	R^2	$\hat{\beta}^*$	$p\text{-val}_{1s}$	$p\text{-val}_{2s}$
Panel A: Raw illiquidity measures								
<i>roll</i>	-0.32	0.86	0.65	0.71	0.05	-0.40	0.82	0.47
<i>cs</i>	0.27	0.86	0.38	0.75	0.03	0.11	0.32	0.38
<i>fht</i>	0.17	0.18	0.17	0.35	0.13	0.10	0.28	0.38
<i>tick</i>	0.20	0.19	0.15	0.30	0.16	0.10	0.23	0.25
<i>zeros</i>	0.17	0.20	0.20	0.39	0.14	0.20	0.16	0.40
<i>ami</i>	0.40	0.18	0.01	0.02	0.61	0.32	0.03	0.03
<i>amito</i>	0.15	0.33	0.33	0.65	0.04	0.00	0.47	0.57
Δami	-0.78	0.76	0.85	0.30	0.14	-0.76	0.85	0.29
<i>PS</i>	-6.45	6.67	0.83	0.33	0.21	-6.97	0.92	0.20
<i>hm</i>	-0.43	0.86	0.69	0.61	0.07	-0.27	0.65	0.55
Panel B: Selected break-adjusted illiquidity measures								
<i>roll</i> ^B	-0.05	1.04	0.52	0.96	0.00	0.06	0.47	0.95
<i>cs</i> ^B	-0.02	0.72	0.51	0.97	0.00	-0.19	0.72	0.95
<i>fht</i> ^B	0.97	0.66	0.07	0.15	0.46	0.82	0.03	0.04
<i>tick</i> ^B	0.99	0.61	0.05	0.10	0.51	0.75	0.04	0.04
<i>zeros</i> ^B	2.81	0.93	0.00	0.00	1.27	2.93	0.00	0.00
<i>ami</i> ^B	0.59	0.24	0.01	0.01	0.90	0.50	0.01	0.01
<i>amito</i> ^B	0.37	0.65	0.29	0.57	0.07	0.19	0.29	0.34
<i>hm</i> ^B	-0.91	1.00	0.82	0.37	0.16	-0.85	0.87	0.24
Panel C: Decomposition of illiquidity into volatility and residual components								
<i>vol</i>	-0.69	0.87	0.79	0.43	0.18	-0.85	0.96	0.18
<i>roll</i> ^{⊥,B}	0.34	0.77	0.33	0.66	0.03	0.33	0.34	0.65
<i>cs</i> ^{⊥,B}	0.80	0.58	0.09	0.17	0.29	0.88	0.05	0.15
<i>fht</i> ^{⊥,B}	2.39	0.70	0.00	0.00	1.59	2.47	0.00	0.00
<i>tick</i> ^{⊥,B}	2.26	0.66	0.00	0.00	1.59	2.28	0.00	0.00
<i>ami</i> ^{⊥,B}	0.82	0.24	0.00	0.00	1.45	0.79	0.00	0.00
<i>amito</i> ^{⊥,B}	0.70	0.70	0.16	0.31	0.23	0.59	0.08	0.11
<i>PS</i> ^{⊥,B}	-4.91	6.87	0.76	0.47	0.10	-4.85	0.82	0.37
<i>hm</i> ^{⊥,B}	-1.43	0.95	0.93	0.13	0.32	-1.32	0.92	0.13

Table 2.5: Forecasting regressions for macroeconomic activity: Monthly frequency

The table presents results for one-month-ahead predictive regressions for growth in industrial production (Δip) and growth in unemployment (Δue) using various measures of aggregate illiquidity. The table reports results from regressions of the form

$$y_{t+1} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \epsilon_{t+1},$$

where y_{t+1} is the dependent variable and x_t represents the specified forecasting variable. For each forecasting variable, the table presents OLS and bootstrap estimates. The presented OLS estimates include the slope coefficient $\hat{\beta}$, the associated Newey-West standard error, a two-sided p -value, and ΔR^2 , defined as the increase in R^2 -value relative to a similar regression that excludes x_t (the above model with $\beta = 0$). The presented bootstrap estimates include the slope coefficient $\hat{\beta}$ and the associated two-sided p -value. Four lags of the dependent variable are included in the predictive regression ($J = 4$). The sample period is 1948–2013. See Section 2.2 for variable definitions.

	$y_{t+1} = \Delta ip$						$y_{t+1} = \Delta ue$					
	OLS				Bootstrap		OLS				Bootstrap	
	$\hat{\beta}$	SE	p -val	ΔR^2	$\hat{\beta}^*$	p -val	$\hat{\beta}$	SE	p -val	ΔR^2	$\hat{\beta}^*$	p -val
Panel A: Raw illiquidity measures												
<i>roll</i>	-0.43	0.12	0.00	1.53	-0.45	0.00	1.49	0.46	0.00	1.21	1.51	0.00
<i>cs</i>	-0.56	0.14	0.00	2.00	-0.56	0.00	1.94	0.48	0.00	1.54	1.97	0.00
<i>fht</i>	0.01	0.03	0.65	0.01	0.01	0.75	0.08	0.10	0.43	0.04	0.09	0.52
<i>tick</i>	0.00	0.03	0.96	0.00	0.00	0.96	0.12	0.11	0.28	0.07	0.14	0.36
<i>zeros</i>	0.04	0.03	0.13	0.18	0.04	0.35	-0.02	0.12	0.90	0.00	-0.01	0.93
<i>ami</i>	-0.07	0.03	0.03	0.36	-0.07	0.01	0.40	0.14	0.00	0.67	0.41	0.01
<i>amito</i>	-0.08	0.06	0.17	0.17	-0.08	0.24	0.57	0.21	0.01	0.62	0.58	0.04
Δami	-0.35	0.15	0.02	0.52	-0.34	0.03	1.49	0.68	0.03	0.60	1.51	0.04
<i>PS</i>	-2.33	0.95	0.01	0.49	-2.33	0.04	6.74	3.87	0.08	0.27	6.61	0.16
<i>hm</i>	0.35	0.12	0.00	0.81	0.35	0.03	-0.94	0.49	0.06	0.39	-0.92	0.15
Panel B: Selected break-adjusted illiquidity measures												
<i>roll</i> ^B	0.59	0.19	0.00	1.22	0.59	0.00	-2.01	0.77	0.01	0.92	-2.01	0.01
<i>cs</i> ^B	-0.47	0.11	0.00	1.91	-0.47	0.00	1.26	0.40	0.00	0.93	1.29	0.00
<i>fht</i> ^B	-0.33	0.13	0.01	0.87	-0.34	0.00	0.95	0.44	0.03	0.50	1.02	0.03
<i>tick</i> ^B	-0.38	0.12	0.00	1.23	-0.38	0.00	1.11	0.41	0.01	0.71	1.17	0.00
<i>zeros</i> ^B	0.09	0.21	0.66	0.02	0.07	0.56	0.27	0.81	0.74	0.01	0.31	0.68
<i>ami</i> ^B	-0.16	0.05	0.00	1.08	-0.16	0.00	0.56	0.17	0.00	0.89	0.57	0.00
<i>amito</i> ^B	-0.45	0.10	0.00	1.80	-0.45	0.00	1.89	0.42	0.00	2.06	1.92	0.00
<i>hm</i> ^B	0.48	0.18	0.01	0.86	0.48	0.01	-1.59	0.76	0.04	0.60	-1.56	0.03
Panel C: Decomposition of illiquidity into volatility and residual components												
<i>vol</i>	-0.48	0.17	0.00	1.48	-0.48	0.00	1.28	0.67	0.06	0.71	1.30	0.02
<i>roll</i> ^{⊥,B}	-0.02	0.16	0.89	0.00	-0.03	0.89	0.55	0.62	0.37	0.10	0.56	0.41
<i>cs</i> ^{⊥,B}	0.09	0.15	0.56	0.07	0.09	0.42	-0.09	0.61	0.89	0.00	-0.08	0.85
<i>fht</i> ^{⊥,B}	0.11	0.15	0.45	0.07	0.10	0.37	-0.06	0.59	0.92	0.00	-0.04	0.93
<i>tick</i> ^{⊥,B}	-0.08	0.14	0.57	0.04	-0.09	0.54	0.48	0.54	0.37	0.08	0.54	0.32
<i>ami</i> ^{⊥,B}	-0.07	0.05	0.18	0.18	-0.07	0.12	0.35	0.21	0.09	0.31	0.37	0.05
<i>amito</i> ^{⊥,B}	-0.24	0.10	0.02	0.47	-0.24	0.01	1.41	0.38	0.00	1.05	1.44	0.00
<i>PS</i> ^{⊥,B}	-1.67	0.97	0.09	0.22	-1.64	0.14	5.39	3.89	0.17	0.15	5.27	0.26
<i>hm</i> ^{⊥,B}	0.34	0.19	0.07	0.34	0.34	0.07	-1.31	0.82	0.11	0.32	-1.24	0.12

Table 2.6: Long-horizon forecasting regressions

The table presents results for 6- and 12-month ahead predictive regressions for growth in industrial production (Δip), growth in unemployment (Δue), and excess returns on the NYSE index (ret) using various measures of aggregate illiquidity:

$$\sum_{h=1}^H y_{t+h} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \epsilon_{t+1},$$

For each forecasting variable, the table presents the estimated slope coefficient $\hat{\beta}$, and the ΔR^2 , defined as the increase in R^2 -value relative to a similar regression that excludes x_t (the above model with $\beta = 0$). Statistical significance is determined using Newey-West standard errors. ***, **, and * represent statistical significance at the .01, .05, and .10 levels, respectively. When the dependent variable is growth in industrial production or unemployment, four lags of y_t are included in the predictive regression ($J = 4$). The predictive regression for excess returns does not include lagged returns. The sample period is 1948–2013. See Section 2.2 for variable definitions.

	$\sum_{h=1}^6 ret_{t+h}$			$\sum_{h=1}^{12} ret_{t+h}$			$\sum_{h=1}^6 \Delta ip_{t+h}$			$\sum_{h=1}^{12} \Delta ip_{t+h}$			$\sum_{h=1}^6 \Delta ue_{t+h}$			$\sum_{h=1}^{12} \Delta ue_{t+h}$		
	$\hat{\beta}$	ΔR^2		$\hat{\beta}$	ΔR^2		$\hat{\beta}$	ΔR^2		$\hat{\beta}$	ΔR^2		$\hat{\beta}$	ΔR^2		$\hat{\beta}$	ΔR^2	
Panel A: Selected break-adjusted illiquidity measures																		
$roll^B$	-6.57**	1.16		-7.18	0.68		4.03***	4.04		6.07***	3.66		-11.69***	2.62		-19.49**	2.68	
cs^B	1.51	0.17		2.46	0.22		-3.10***	5.94		-5.32***	7.07		7.18***	2.56		12.38*	2.85	
fh^B	6.09**	2.57		10.91***	4.13		-1.93**	2.16		-2.95*	2.04		3.20	0.48		3.41	0.20	
$tick^B$	6.60***	3.21		11.75***	5.10		-2.08**	2.67		-3.03**	2.29		4.41*	0.95		5.48	0.55	
$zeros^B$	13.14***	3.87		26.62***	7.81		0.95	0.19		2.25	0.42		-1.27	0.03		-7.12	0.30	
ami^B	3.89***	5.50		7.45***	10.10		-0.80**	1.98		-0.98	1.19		2.40*	1.39		2.92	0.77	
$amito^B$	3.23	0.79		5.71	1.24		-2.63***	4.50		-3.82**	3.85		10.71***	5.63		18.67***	6.42	
hm^B	-11.53***	3.70		-16.14***	3.58		3.00***	2.33		3.62	1.36		-7.99*	1.26		-10.11	0.75	
Panel B: Decomposition of illiquidity into volatility and residual components																		
$\overline{vol}$	-1.02	0.06		-3.49	0.33		-3.37***	5.25		-5.76***	6.22		6.88*	1.73		13.02	2.32	
$roll^{\perp,B}$	4.84	0.87		9.58*	1.70		-0.03	0.00		0.28	0.01		2.24	0.14		3.52	0.12	
$cs^{\perp,B}$	2.90	0.53		7.32	1.68		0.91	0.49		1.63	0.63		-0.23	0.00		-1.89	0.06	
$fh^{\perp,B}$	10.35***	4.17		20.70***	8.25		1.30	0.61		2.63	1.01		-3.06	0.26		-9.39	0.91	
$tick^{\perp,B}$	10.88***	5.17		20.88***	9.49		0.24	0.02		1.07	0.19		0.41	0.01		-3.17	0.12	
$ami^{\perp,B}$	4.56***	6.32		9.07***	12.52		-0.15	0.06		0.16	0.03		1.28	0.35		0.75	0.05	
$amito^{\perp,B}$	3.61	0.84		7.13	1.64		-1.16	0.79		-1.31	0.41		8.42***	3.15		14.60**	3.54	
$PS^{\perp,B}$	0.68	0.00		4.44	0.01		-18.01***	1.84		-22.83**	1.20		56.05***	1.37		102.32***	1.71	
$hm^{\perp,B}$	-14.31***	4.52		-20.77***	4.71		1.92*	0.76		1.42	0.17		-6.47	0.66		-6.26	0.23	

Table 2.7: Monthly forecasts including macroeconomic and financial variables

The table presents one-month-ahead predictive regressions for excess returns on the NYSE index (ret), growth in industrial production (Δip), and growth in unemployment (Δue) using various measures of aggregate illiquidity:

$$y_{t+1} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \gamma' z_t + \epsilon_{t+1},$$

where y_{t+1} is the dependent variable and x_t indicates a volatility-adjusted aggregate illiquidity measure of interest. The vector z_t represents an additional set of macroeconomic and financial forecasting variables included in the model. For each forecasting variable, the table presents the estimated slope coefficient $\hat{\beta}$, the associated t -statistic based on Newey-West standard errors, and ΔR^2 , defined as the increase in R^2 -value relative to a similar regression that excludes x_t (the above model with $\beta = 0$). We set $J = 4$ for Δip and Δue , and $J = 0$ for ret . The sample period is 1948–2013. See Section 2.2 for variable definitions.

	$y_{t+1} = ret$			$y_{t+1} = \Delta ip$			$y_{t+1} = \Delta ue$		
	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2
Panel A: $z_t = (\overline{vol}_t, ret_t, ret_{t-1})$									
$zeros^B$	2.55	2.90	0.94	-0.19	-0.97	0.10	1.20	1.63	0.24
$roll^{\perp,B}$	0.14	0.16	0.00	-0.23	-1.51	0.22	1.12	2.01	0.36
$cs^{\perp,B}$	0.57	0.88	0.10	-0.24	-1.62	0.33	0.93	1.66	0.32
$fht^{\perp,B}$	2.21	3.27	1.23	-0.10	-0.69	0.04	0.62	1.17	0.12
$tick^{\perp,B}$	2.10	3.39	1.33	-0.19	-1.40	0.20	0.84	1.65	0.25
$ami^{\perp,B}$	0.76	3.31	1.23	-0.09	-1.95	0.33	0.45	2.24	0.48
$amito^{\perp,B}$	0.64	0.95	0.17	-0.24	-2.19	0.42	1.56	3.93	1.11
$PS^{\perp,B}$	-1.88	-0.29	0.01	-0.66	-0.72	0.03	1.49	0.39	0.01
$hm^{\perp,B}$	-2.31	-2.44	0.73	0.09	0.43	0.02	-0.44	-0.50	0.03
Panel B: $z_t = \text{Full set of controls}$									
$zeros^B$	2.15	2.31	0.63	-0.41	-1.95	0.41	2.17	2.84	0.74
$roll^{\perp,B}$	-0.53	-0.59	0.06	-0.39	-2.53	0.58	1.65	3.00	0.68
$cs^{\perp,B}$	-0.35	-0.50	0.03	-0.58	-3.40	1.40	2.13	3.45	1.23
$fht^{\perp,B}$	1.91	2.59	0.82	-0.34	-2.15	0.46	1.67	2.96	0.73
$tick^{\perp,B}$	1.81	2.66	0.89	-0.40	-2.70	0.78	1.75	3.33	0.96
$ami^{\perp,B}$	0.65	2.54	0.79	-0.16	-3.18	0.88	0.71	3.61	1.09
$amito^{\perp,B}$	0.66	0.89	0.14	-0.20	-1.71	0.24	1.36	2.99	0.70
$PS^{\perp,B}$	-0.13	-0.02	0.00	0.08	0.09	0.00	-1.89	-0.51	0.02
$hm^{\perp,B}$	-1.93	-1.85	0.47	0.01	0.04	0.00	-0.01	-0.01	0.00

Table 2.8: Forecasting regressions: Quarterly frequency

The table presents results for one-quarter-ahead predictive regressions for growth in GDP (Δgdp), growth in unemployment (Δue), growth in investment (Δinv), growth in consumption ($\Delta cons$), and excess returns on the NYSE index (ret) using various measures of aggregate illiquidity. The table reports results from regressions of the form

$$y_{t+1} = \alpha + \sum_{j=1}^J \rho_j y_{t+1-j} + \beta x_t + \epsilon_{t+1},$$

where y_{t+1} is the dependent variable and x_t represents the specified forecasting variable. For each forecasting variable, the table presents the estimated slope coefficient $\hat{\beta}$, the associated t -statistic based on Newey-West standard errors, and ΔR^2 , defined as the increase in R^2 -value relative to a similar regression that excludes x_t (the above model with $\beta = 0$). When the dependent variable is Δgdp , Δue , Δinv , or $\Delta cons$; two lags of the dependent variable are included in the predictive regression ($J = 2$). The predictive regression for excess returns does not include lagged returns. The sample period is 1948–2013. See Section 2.2 for variable definitions.

	$y_{t+1} = ret$			$y_{t+1} = \Delta gdp$			$y_{t+1} = \Delta ue$			$y_{t+1} = \Delta inv$			$y_{t+1} = \Delta cons$		
	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2	$\hat{\beta}$	t -stat	ΔR^2
Panel A: Raw illiquidity measures															
<i>roll</i>	3.08	1.35	1.63	-0.61	-2.97	3.78	2.94	2.65	1.71	-3.26	-2.97	4.54	-0.48	-2.94	3.05
<i>cs</i>	2.78	1.08	0.82	-0.72	-3.11	3.06	4.23	3.30	2.05	-4.00	-2.92	3.99	-0.51	-2.38	2.05
<i>fh</i>	0.45	0.73	0.25	0.07	1.34	0.36	0.08	0.28	0.01	0.01	0.02	0.00	0.08	1.85	0.69
<i>tick</i>	0.54	0.85	0.33	0.06	1.11	0.26	0.18	0.54	0.04	-0.05	-0.17	0.01	0.07	1.59	0.55
<i>zeros</i>	0.33	0.50	0.15	0.11	1.95	1.01	-0.12	-0.36	0.03	0.20	0.63	0.15	0.11	2.45	1.40
<i>ami</i>	1.35	2.35	1.88	-0.05	-0.71	0.15	0.67	1.55	0.54	-0.58	-1.50	0.90	0.00	-0.06	0.00
<i>amito</i>	0.59	0.50	0.16	-0.01	-0.10	0.00	1.05	1.57	0.59	-0.67	-1.13	0.54	0.07	0.81	0.19
Δami	-0.97	-0.48	0.11	-0.99	-5.17	7.38	4.18	3.66	2.50	-6.12	-6.44	12.00	-0.57	-3.28	3.34
<i>PS</i>	56.30	2.90	2.62	-5.46	-2.29	1.57	28.39	2.34	0.80	-25.38	-2.36	1.39	-4.51	-2.04	1.46
<i>hm</i>	-4.23	-2.17	2.33	0.61	2.83	3.11	-2.12	-1.65	0.73	2.85	2.04	2.80	0.56	3.21	3.34
Panel B: Decomposition of illiquidity into volatility and residual components															
$\overline{vol}$	0.29	0.11	0.01	-0.69	-2.81	2.93	3.05	1.89	1.12	-2.95	-2.18	2.29	-0.52	-2.35	2.23
<i>zeros</i> ^B	7.11	2.40	2.06	0.11	0.29	0.03	-0.83	-0.35	0.03	0.02	0.01	0.00	0.17	0.55	0.11
<i>roll</i> ^{⊥,B}	3.19	1.87	1.32	-0.14	-0.61	0.17	0.71	0.51	0.08	-1.16	-0.94	0.47	-0.09	-0.45	0.09
<i>cs</i> ^{⊥,B}	0.97	0.52	0.12	0.17	0.73	0.23	-0.50	-0.32	0.04	0.56	0.43	0.10	0.18	0.84	0.34
<i>fh</i> ^{⊥,B}	5.13	2.15	1.83	0.23	0.79	0.23	-1.41	-0.85	0.17	0.84	0.58	0.13	0.21	0.88	0.28
<i>tick</i> ^{⊥,B}	4.04	2.30	1.73	0.00	0.00	0.00	0.22	0.16	0.01	-0.33	-0.27	0.03	0.03	0.14	0.01
<i>ami</i> ^{⊥,B}	2.19	3.20	2.84	-0.06	-0.59	0.14	0.40	0.63	0.12	-0.51	-0.95	0.41	-0.02	-0.21	0.02
<i>amito</i> ^{⊥,B}	1.74	0.76	0.35	-0.32	-1.45	0.74	2.81	2.27	1.08	-2.23	-1.85	1.49	-0.08	-0.43	0.07
<i>PS</i> ^{⊥,B}	70.76	3.12	3.35	-4.22	-1.49	0.78	25.18	1.66	0.51	-21.18	-1.63	0.79	-3.89	-1.53	0.90
<i>hm</i> ^{⊥,B}	-9.90	-5.29	5.03	0.64	1.93	1.40	-3.46	-1.34	0.76	4.56	2.14	2.85	0.42	1.41	0.78

Table 2.9: Out-of-sample forecasting performance: monthly frequency

This table presents out-of-sample forecasts for monthly data. CW is the t -statistic for the Clark-West test. ΔR_{OOS}^2 is the out-of-sample R^2 -value relative to the benchmark. GW is the t -statistic for the Giacomini and White (2006) test for equal predictive ability. The forecasting target is alternatively the growth rate in industrial production Δip , the growth rate in unemployment Δue , or excess stock market returns ret . The benchmark model for output and unemployment growth rates is an AR(4) specification. The benchmark model for excess returns is the historical average. Forecasts are based on parameter estimates using a rolling window of 30 years. The out-of-sample forecasts cover the period 1978–2013. Panel A presents results for raw illiquidity measures, Panel B presents results for break- and volatility-adjusted measures, and Panel C shows results for volatility and various other financial predictors.

	$y_{t+1} = ret$			$y_{t+1} = \Delta ip$			$y_{t+1} = \Delta ue$		
	ΔR_{OOS}^2	CW	GW	ΔR_{OOS}^2	CW	GW	ΔR_{OOS}^2	CW	GW
Panel A: Selected raw liquidity measures									
<i>roll</i>	-0.32	0.51	-0.26	0.71	2.71	0.28	0.17	2.26	0.06
<i>cs</i>	-0.46	0.46	-0.44	1.42	3.46	0.66	1.14	2.67	0.52
<i>Δami</i>	-0.05	0.48	-0.09	0.11	0.74	0.11	0.68	1.88	0.44
Panel B: Break- and volatility-adjusted illiquidity measures									
<i>zeros^B</i>	1.76	2.53	1.89	-0.35	-0.33	-0.80	-0.35	-0.37	-0.88
<i>roll^{⊥,B}</i>	0.04	0.73	0.06	-0.13	-0.18	-0.49	-0.61	-0.87	-1.36
<i>cs^{⊥,B}</i>	-0.27	0.01	-0.54	-0.75	0.51	-1.11	-0.64	-0.18	-0.92
<i>fht^{⊥,B}</i>	2.43	2.98	2.17	0.10	0.80	0.36	-0.19	-0.59	-0.83
<i>tick^{⊥,B}</i>	2.14	2.98	2.00	-0.11	0.45	-0.22	-0.36	-0.25	-0.90
<i>ami^{⊥,B}</i>	1.44	2.78	1.79	-0.02	0.94	-0.03	-0.20	0.78	-0.24
<i>amito^{⊥,B}</i>	0.23	1.08	0.68	-3.98	-0.41	-2.25	-1.35	1.65	-0.72
<i>PS^{⊥,B}</i>	-0.75	-1.30	-1.61	-1.49	-0.37	-1.68	-1.11	0.31	-0.93
<i>hm^{⊥,B}</i>	-0.73	-1.26	-1.44	-0.34	0.67	-0.34	0.14	1.52	0.14
Panel C: Volatility and other financial variables									
<i>vol</i>	-0.48	-0.02	-0.46	1.34	3.03	0.59	1.89	2.68	0.74
<i>TERM</i>	-0.33	1.31	-0.22	0.09	2.40	0.04	-0.34	2.25	-0.13
<i>DEF</i>	-1.37	-0.61	-1.49	2.45	2.35	1.21	1.05	1.84	0.83
<i>CPSP</i>	-1.01	-0.58	-1.09	5.48	4.30	1.81	3.51	3.44	1.19
<i>TBL</i>	-2.62	0.27	-1.61	-0.55	1.79	-0.21	-1.71	0.98	-0.79
<i>ret</i>	0.14	0.74	0.19	0.24	0.78	0.49	0.59	1.77	0.41
<i>ret_{t-1}</i>	-0.22	-0.13	-0.54	0.85	1.69	0.37	0.57	1.79	0.43

Chapter 3

Payout Yields and Stock Return

Predictability: How Important is the Measure of Cash Flow?⁴²

⁴²Eaton, G. and Paye, B. Accepted by *Journal of Financial and Quantitative Analysis*. Reprinted here with permission of the publisher.

Abstract

We compare the stock return forecasting performance of alternative payout yields. The net payout yield produces more accurate forecasts relative to alternatives, including the traditional dividend yield. This remains true even after excluding several years during the Great Depression when issuance was unusually high. The measure of cash flow used to form the yield matters economically. Long-term investors' hedging demand for stock is considerably reduced when net payout, rather than dividends, serves as the cash flow measure. An agent relying on an incorrect payout measure is willing to pay an economically significant 'management fee' to switch to the optimal policy.

3.1 Introduction

The dividend yield is among the most common variables used to forecast stock returns. The intuition behind this variable is simple: when prices are high relative to dividends, we expect low future returns, and vice versa. The extent of stock return predictability associated with the dividend yield remains controversial; however, due to econometric issues that complicate inference in predictive regressions.⁴³

Dividends represent only one component of cash flows between firms and investors.⁴⁴ Boudoukh, Michaely, Richardson, and Roberts (2007) examine two alternative yields: the total payout yield, which incorporates dividends and share repurchases, and the net payout yield, which incorporates dividends, repurchases and stock issuances. Both outperform the traditional dividend yield in predicting stock returns.⁴⁵ Results are particularly strong for the net payout yield: Boudoukh, Michaely, Richardson, and Roberts (2007) report an R^2 -value in excess of 20% for a predictive regression based on this yield using annual data. Some controversy surrounds this finding, however. Referring to the result, Cochrane (2008) states “The stunning 23.4% R^2 -value comes from one large outlier in the early 1930s.” (p. 1567) Overall, there does not appear to be consensus regarding the form of payout yield that most effectively captures time-variation in the equity premium.

This study provides new evidence on the predictive power of alternative payout yields. We make two main contributions. First, we show that the net payout yield produces more accurate stock return forecasts relative to alternatives even after excluding extraordinary issuance years during the Great Depression. Second, we demonstrate that the payout measure used to form

⁴³Early studies reporting evidence that the dividend yield forecasts stock returns include Rozeff (1984), Shiller, Fischer, and Friedman (1984), Keim and Stambaugh (1986), Campbell and Shiller (1988b), and Fama and French (1988). Shiller (1981), Campbell and Shiller (1988a), Campbell and Ammer (1993), and Cochrane (1991, 1992, 1994, 2008) discuss the connection between return predictability and variation in the price-dividend ratio. Goetzmann and Jorion (1993), Nelson and Kim (1993), Stambaugh (1999), and Campbell and Yogo (2006) are examples of papers that focus on inference issues in predictive regressions.

⁴⁴Fama and French (2001) find a decrease in the fraction of firms paying dividends between the 1970s and the late 1990s. Grullon and Michaely (2002) and Skinner (2008) associate this decrease with firms’ tendency to substitute share repurchases for dividends following the SEC’s enactment of safe harbor provisions. Brav, Graham, Harvey, and Michaely (2005) provide survey evidence consistent with the substitution hypothesis. Allen and Michaely (2003) and Farre-Mensa, Michaely, and Schmalz (2014) review this area.

⁴⁵See also Ackert and Smith (1993), Robertson and Wright (2006) and Kim and Park (2013).

the yield matters economically: investors’ portfolio policies differ in important ways depending on whether the state variable governing time-varying investment opportunities is the net payout yield or the traditional dividend yield. Calibration results indicate that an agent relying on an incorrect definition of payout is willing to pay an economically significant ‘management fee’ in order to switch to the correct payout measure.

Following Boudoukh, Michaely, Richardson, and Roberts (2007), we construct aggregate annual measures of dividends, stock repurchases and issuances over the period 1926–2014 and define a variety of associated payout yields. Predictive regressions using the net payout yield are sensitive to the inclusion of data from 1929 and 1930, two years when stock issuance was extraordinarily high relative to prices. Upon excluding these years, the R^2 for the regression falls from over 20% to around 10% – a value closer to the corresponding R^2 -values for regressions based on the dividend yield and total payout yield. We show that issuance is the particular component of net payout that exhibits high sensitivity to the inclusion of 1929–1930 in the sample.

Despite the sensitivity of results for the net payout yield to the sample period, we find strong evidence that both the total payout yield and net payout yield forecast returns, consistent with Boudoukh, Michaely, Richardson, and Roberts (2007).⁴⁶ Additional empirical tests provide further support for the net payout yield. The net payout yield implicitly imposes restrictions on the slope coefficients associated with payout components in predictive regressions. After excluding extraordinary issuance years around the Great Depression, we cannot reject the null hypothesis that these restrictions are valid. The net payout yield also produces more accurate out-of-sample forecasts relative to various alternative yields, including approaches that separate net payout into components.

The dividend yield, total payout yield, and net payout yield often co-move closely, and it is tempting to think that the choice of yield is not particularly important. To the contrary, we show that the cash flow measure used to construct the yield is economically significant. First, the cash flow measure affects the implied dynamics of the equity premium. Predictive

⁴⁶Butler, Cornaggia, Grullon, and Weston (2011) provide complementary evidence at the firm level regarding the relation between the level of net financing and future stock returns.

regressions based on the dividend yield imply extremely low-frequency variation in discount rates. In contrast, the net payout yield generates persistence properties similar to those of common business cycle proxies. Second, the payout measure significantly impacts investors' long-horizon portfolio decisions. To show this, we conduct calibration analyses for optimal long-horizon portfolio decisions under time-varying expected returns.

The calibration results reveal significant differences across the three yields with respect to investors' long-horizon portfolio policies. Most notably, there is a large difference in the extent of hedging demand, in the sense of Merton (1969, 1971, 1973). Shocks to both the dividend yield and total payout yield are strongly negatively correlated with contemporaneous shocks to returns. This negative correlation, coupled with the high persistence of these yields, generates an important hedging demand for stock (Campbell and Viceira (1999)). In contrast, contemporaneous shocks to the net payout yield and returns are only weakly correlated. Consequently, the hedging demand for stock is greatly reduced.

A final analysis measures the welfare implications associated with portfolio policies under alternative measures of payout. We simulate portfolio performance for two investors: one who makes allocations based on the net payout yield – assumed to be the correct state variable – and another who invests suboptimally based on the dividend yield. An agent who follows the suboptimal dividend yield policy is willing to pay an economically significant management fee to switch to the optimal policy. The maximal fee ranges from approximately 40 basis point to greater than 4% annually, depending upon risk aversion and the investment horizon. The fee increases in the investment horizon, and is economically large even for relatively risk averse investors. This is because agents basing decisions on the dividend yield allocate a nontrivial fraction of wealth to stock as hedging demand when, in reality, there is little (true) hedging demand for stock.

Surprisingly, in light of Boudoukh, Michaely, Richardson, and Roberts (2007), most recent stock return forecasting studies eschew the net payout yield in favor of the dividend yield.⁴⁷

⁴⁷These remarks are based on a survey of papers published between 2008–2014 in four leading finance journals: *Journal of Finance*, *Journal of Financial and Quantitative Analysis*, *Journal of Financial Economics*, and *Review of Financial Studies*. Keyword criteria identified 67 published papers over this period that include some form of predictive regression for stock market returns. Among these, 62 (92.5%) include the dividend yield as a

There are several potential explanations for the continued dominance of the dividend yield. Researchers might continue to use this yield in deference to earlier published work. Alternative payout yields are also more difficult to construct than the dividend yield. We conjecture that many researchers view the dividend, total payout, and net payout yields as close substitutes, and consequently view reliance on the dividend yield as innocuous. This study emphasizes that the choice of payout measure used to form the yield has important consequences for both the implied dynamics of the equity premium and investors' optimal long-horizon portfolio decisions.

3.2 Alternative payout yields and cash flow measures

3.2.1 Data sources and data construction

Stock return data are sourced from the University of Chicago's Center for Research in Security Prices (CRSP) covering the period 1926–2014. We measure aggregate stock returns using the value-weighted return, including dividends, for the CRSP universe. Our primary measure of the annual dividend yield is imputed from CRSP market returns by taking the capitalized series of returns over the uncapped series, i.e., $DY_{t+1} = D_{t+1}/P_{t+1} = (R_{t+1}/Rx_{t+1}) - 1$, where R_{t+1} and Rx_{t+1} denote gross annual CRSP index returns with and without dividends, respectively.⁴⁸

Dividend growth is computed as

$$\frac{D_{t+1}}{D_t} = \frac{DY_{t+1}}{DY_t} (Rx_{t+1}) = \frac{D_{t+1}}{P_{t+1}} \frac{P_t}{D_t} \frac{P_{t+1}}{P_t}. \quad (3.1)$$

Stock repurchases and issuances represent additional cash flows between firms and shareholders that are central to the questions addressed in this study. We construct annual time series

forecasting instrument. Seven papers (10.5%) consider the total payout yield, six papers (8.9%) consider the net payout yield, and eight papers (11.9%) include the net issuance yield. Details regarding the survey are available from the authors upon request.

⁴⁸Imputing the dividend yield from annual returns implicitly assumes that dividends received within the year are reinvested at the return R until the end of the year. Ang (2012), Chen (2009), and Engsted and Pedersen (2010) show that the extent of dividend growth predictability is sensitive to whether or not dividends are assumed to be reinvested. As a robustness check, we compute an alternative dividend yield assuming that dividends received during the year are consumed rather than reinvested. The alternative measure generates stronger evidence of dividend growth predictability, consistent with previous studies. Key results concerning stock return predictability and long-horizon portfolio choice are robust to the convention for measuring dividends. Consequently, reported results employ the measure DY .

of aggregate stock repurchases and issuances following the approach of Boudoukh, Michaely, Richardson, and Roberts (2007) and use these to compute alternative yield measures including the total payout yield, the net payout yield, the issuance yield, and the net issuance yield. We also calculate annual growth rates for total payout, net payout, and stock issuances. All growth rates are computed for the period 1927–2014 except for net payout growth, for which the sample period is 1931–2014 due to negative net payouts in several years prior to 1931. The Appendix provides details regarding the construction of yields and cash flow growth rates.

Engsted and Pedersen (2010) show that results for long-horizon dividend growth predictability in US data depend on whether payout is measured in real or nominal terms. They attribute the difference in results to the fact that the dividend yield negatively predicts long-term inflation in US post-war data. We analyze variables measured in real terms, and adjust payout growth rates for inflation as described in the Appendix.

Below, we briefly describe the variables used in the empirical analysis:

Market returns: The real excess return (R^e) on the value-weighted CRSP portfolio is the real index return, including dividends, less the real return on a 90-day Treasury bill sourced from CRSP. r denotes the real log return and r^e is the log excess return.

Dividend yield: The dividend yield (DY) for the value-weighted CRSP portfolio is imputed using annual index returns including and excluding dividends, respectively. The log dividend yield is denoted dy .

Total payout yield: The total payout yield (TPY) is the sum of the dividend yield and the repurchase yield, where the latter is defined as aggregate repurchases divided by year-end market capitalization. The log total payout yield is denoted tpy .

Net payout yield: The net payout yield (NPY) is the dividend yield minus the net issuance yield, where the net issuance yield is the dollar value of issuances minus repurchases divided by year-end market capitalization. We construct two versions of the log net payout yield. First, following Boudoukh, Michaely, Richardson, and Roberts (2007), we define $\overline{npy} \equiv \ln(0.1 + NPY)$ to obtain a log net payout yield for the full sample period

1926–2012 that includes two years (1929 and 1930) in which $NPY < 0$. Alternatively, we define $npy \equiv \ln(NPY)$, where the series is available for the period 1931–2014.

Issuance yield: The issuance yield (IY) is defined as aggregate issuances divided by year-end market capitalization. The log issuance yield is denoted iy .⁴⁹

Net issuance yield: The net issuance yield (NIY) is the dollar value of issuances minus repurchases divided by year-end market capitalization. Because the net issuance yield is positive for some years in our sample and negative for others, we construct a log net issuance yield as $\overline{niy} \equiv \ln(0.1 + NIY)$.

Dividend growth: Dividend growth (ΔD) is the real annual percentage change in the dividend level. Log dividend growth is $\Delta d \equiv \ln(D_{t+1}/D_t)$.

Total payout growth: Total payout growth (ΔTP) is the real annual percentage change in the total payout level. Log total payout growth is $\Delta tp \equiv \ln(TP_{t+1}/TP_t)$.

Net payout growth: Net payout growth (ΔNP) is the real annual percentage change in the net payout level. Log net payout growth is $\Delta np \equiv \ln(NP_{t+1}/NP_t)$.

Issuance growth: Issuance growth (ΔI) is the real annual percentage change in the issuance level. Log issuance growth is $\Delta i \equiv \ln(I_{t+1}/I_t)$.

3.2.2 Time series properties of alternative payout yields

Figure 3.1 presents time series plots of various payout yields. The top panel shows the traditional dividend yield along with the total payout yield and net payout yield. The bottom panel shows the issuance and net issuance yields. The divergence of the total payout yield from the dividend yield over recent decades is consistent with well-documented shifts in the pattern of corporate

⁴⁹Baker and Wurgler (2000) find that a variable related to issuance activity – the equity share of new issues – forecasts aggregate returns, and interpret their results as indicative of managers’ ability to time market returns via their issuing activity. Butler, Grullon, and Weston (2005a) find that the predictive ability of the equity share of new issues vanishes once the Great Depression and Oil Crisis of the early 1970s are removed from the sample. Although the issuance yield we study is positively correlated with the equity share of new issues, the correlation is not overly large (around 0.4). Moreover, the issuance yield does not exhibit unusual behavior around the Oil Crisis of the 1970s.

payout (see, e.g., Farre-Mensa, Michaely, and Schmalz (2014)). The net payout yield plunges around the Great Depression. The bottom panel of Figure 3.1 confirms that this plunge is largely attributable to a spike in issuance activity in the years 1929 and 1930. Issuance does not exhibit a similar spike around the recent financial crisis of 2007–2009, suggesting that extreme issuance activity is not a systematic feature of financial crises.

The persistence properties of payout yields prove important in our empirical analysis. It is well-known that the dividend yield is very persistent. The first-order sample autocorrelation for the annual dividend yield over the period 1926–2014 is $\hat{\rho}_{DY} \approx 0.90$. Boudoukh, Michaely, Richardson, and Roberts (2007) show that alternative payout yields are substantially less persistent. For our sample, $\hat{\rho}_{TPY} \approx 0.77$ and $\hat{\rho}_{NPY} \approx 0.67$.

3.3 Stock return forecasts and the measure of payout

3.3.1 Inference issues in predictive regressions

A predictive regression for (excess) stock returns takes the form:

$$R_{t+1}^e = \alpha + \beta' X_t + \epsilon_{t+1}, \quad (3.2)$$

where R_{t+1}^e denotes the real excess return on the market and X_t represents a vector of variables used to forecast returns. The null hypothesis of interest is $\beta = 0$ (no predictability). Under standard assumptions, OLS based on Eq. (3.2) is consistent for β ; however, the OLS estimate can be substantially biased in finite samples (Stambaugh (1999)). For example, suppose that X_t is scalar-valued and follows a first-order autoregressive process:

$$X_{t+1} = \theta + \phi X_t + \eta_{t+1}, \quad (3.3)$$

with $|\phi| < 1$. Stambaugh (1999) shows that the bias of the OLS estimate of β in Eq. (3.2) is:

$$E[\hat{\beta} - \beta] \approx -\frac{\sigma_{\epsilon\eta}}{\sigma_\eta^2} \left(\frac{1 + 3\phi}{T} \right), \quad (3.4)$$

where $\hat{\beta}$ is the OLS estimate of β , $\sigma_{\epsilon\eta} = \text{Cov}(\epsilon, \eta)$, and $\sigma_{\eta}^2 = \text{Var}(\eta)$. Eq. (3.4) implies that $\hat{\beta}$ is biased *upward* when $\phi > 0$ and $\sigma_{\epsilon\eta} < 0$, as is the case for the financial ratios considered in this study. Consequently, standard *t*-tests of the null of no predictability reject too often.

The “Stambaugh bias” is particularly important for financial ratios, because these ratios tend to be highly persistent and exhibit strong (negative) contemporaneous correlation with returns. To address the problem of small sample bias, we apply corrected slope coefficients and associated standard errors proposed in Amihud and Hurvich (2004) and Amihud, Hurvich, and Wang (2009). This approach involves augmented regressions applicable to a multiple-predictor generalization of the model analyzed by Stambaugh (1999). The traditional predictive regression is augmented with proxies for the residual series in a Gaussian VAR (1) model for the predictors. We refer the reader to Amihud and Hurvich (2004) and Amihud, Hurvich, and Wang (2009) for details.

3.3.2 In-sample evidence

Table 3.1 presents results for one-step-ahead predictive regressions of excess returns on alternative yields at an annual frequency. We consider five alternative yields to forecast stock returns: the dividend yield, the total payout yield, the net payout yield, the issuance yield, and the net issuance yield. The regression model is specified either in levels or in logarithms. Panel A presents results for the sample period 1927–2014, which represents an extension of the sample period analyzed in Boudoukh, Michaely, Richardson, and Roberts (2007).⁵⁰ Stock issuance activity was extraordinarily high in the years 1929 and 1930. To investigate the influence of this period, Panel B presents results for the alternative sample period 1931–2014. Finally, Panel C presents results for the post-Treasury Accord period 1952–2014.

Results for the sample period 1927–2014 are similar to those reported by Boudoukh, Michaely, Richardson, and Roberts (2007). Focusing on the regression results in levels (left-hand side of Table 3.1), the OLS slope coefficient on the dividend yield is 3.28. As Cochrane (2011) empha-

⁵⁰Boudoukh, Michaely, Richardson, and Roberts (2007) focus on log returns and log yields, and analyze the sample period 1927–2003, along with the subperiod 1984–2003. We obtain results very similar to theirs for identical sample periods.

sizes, this estimate is economically large: a one percentage point increase in the dividend yield implies an increase in the expected market excess return of over three percentage points. When statistical inference is based on conventional standard errors, the dividend yield is a statistically significant predictor of stock returns (t -statistic of 2.58). The associated regression R -squared is 5.85%. Consistent with previous literature, correcting estimates for the Stambaugh bias weakens the evidence for predictability. The Amihud and Hurvich (2004) bias-corrected slope coefficient falls to approximately 2.4 and the corrected t -statistic falls to 1.69.

The evidence for predictability is stronger for the total payout and net payout yields. OLS estimates of slope coefficients for these variables are 4.87 and 5.50, respectively. Corrections for the Stambaugh bias are more modest, particularly for the net payout yield. This is not surprising, because the net payout yield is much less persistent than the dividend yield. The corrected estimates provide strong statistical support for predictability, particularly for the net payout yield. Consistent with results reported by Boudoukh, Michaely, Richardson, and Roberts (2007), the R^2 -value for the net payout yield forecasting regression is nearly 20%, which is more than double that for the total payout yield and over three times the R^2 for the dividend yield regression. Forecasting regressions for the issuance and net issuance yields suggest that issuance activity predicts returns: both variables are statistically significant and the associated slope coefficients are economically significant. In short, results in Panel A suggest that issuance activity forecasts stock returns, and incorporating issuance activity via the net payout yield dramatically improves stock return forecasts relative to the traditional dividend yield.

Results in Panel B show that some of the apparent predictive power of issuance activity is driven by ‘outlier’ issuance behavior in 1929 and 1930. When the sample period begins in 1932, the issuance yield and net issuance yield are no longer statistically significant and the regression R^2 -value drops dramatically for both variables. Forecasting results for the net payout yield are also altered by omitting 1929–1930 from the sample. The R^2 -value for the regression drops from approximately 20% to roughly 10%, and the bias-corrected estimate of the slope coefficient falls to 4.08. In contrast, regression results for the dividend yield and total payout yield are not materially affected by this slight alteration of the sample period. These results

illustrate that the 1929–1931 period is only unusual with respect to issuance activity and stock returns.⁵¹

Results in Panel C, covering the sample period 1952–2014, are qualitatively similar to those in Panel B with respect to the predictive ability for the issuance and net issuance yields. Slope coefficients and R^2 -values for the dividend, total payout, and net payout yields fall relative to the 1932–2012 sample period. Bias-corrected estimates for the dividend and total payout yields are no longer statistically significant. These results are consistent with previous studies documenting a breakdown in predictive power for payout yields and many other return forecasting variables over recent decades (see, e.g., Goyal and Welch (2008)). Interestingly, the net payout yield remains statistically significant at the 5% level based on the Amihud and Hurvich (2004) bias-corrected t -statistic.

The right-hand portion of Table 3.1 contains results for regressions in log form, i.e., where the dependent variable is the log excess market return and the forecasting variable is the natural logarithm of the corresponding financial ratio. For comparability with Boudoukh, Michaely, Richardson, and Roberts (2007), the log net payout yield is defined as $\overline{np\bar{y}}$ (see Section 3.2). There is little qualitative difference between the results of predictive regressions in log form from those in levels. Results in Panels B and C show that the R^2 -value for the log net payout yield falls when the years 1929–1930 are excluded from the sample, and slope coefficients on the log issuance yield and net issuance yield are insignificant for these sample periods.⁵²

Basing a predictive regression on the net payout yield implicitly constrains the manner in which payout components impact future returns. It is possible to test these restrictions using multivariate predictive regressions that separate the net payout yield into components. We

⁵¹In 1929 and 1930 the issuance yield spikes to over 6%. Market returns in the subsequent years (1930 and 1931) are approximately -29% and -43%, respectively. The dividend yield was 3.8% and 4.8% in 1929 and 1930, respectively. Neither value is unusual by historical standards. Our sample period also includes data from the financial crisis of 2007–2008 and its aftermath. There is no major increase in stock issuance in this period, and predictive regressions based on the issuance or net issuance yield that exclude 1929–1930, but include the recent financial crisis, deliver insignificant slope coefficients and low R^2 -values.

⁵²The R^2 for the log issuance yield regression over the 1927–2014 period (Panel A) is significantly lower than that based on the level of the issuance yield. The intuitive explanation for this fact is that the log transformation dampens the extent to which 1929–1930 is an outlier in terms of issuance activity.

consider two alternative decompositions of the net payout yield:

$$R_{t+1}^e = \alpha + \beta_1 DY_t + \beta_2 (-NIY_t) + \epsilon_{t+1}, \text{ or} \quad (3.5)$$

$$R_{t+1}^e = \alpha + \beta_1 TPY_t + \beta_2 (-IY_t) + \epsilon_{t+1}. \quad (3.6)$$

These specifications include the opposite of the issuance yield or net issuance yield to facilitate comparison of the associated slope coefficient with that of the total payout or dividend yield.⁵³ In each case, the restriction $\beta_1 = \beta_2$ corresponds to the null hypothesis that the restrictions imposed by the net payout yield are valid. Although inference can be based on a standard Wald test, it is important to account for finite-sample bias. Because the dividend yield and total payout yield are more persistent than the issuance and net issuance yields, the Stambaugh bias is likely to be stronger for β_1 than for β_2 . In this case, OLS estimates of the difference $\beta_1 - \beta_2$ are biased. Under the null hypothesis that $\beta_1 = \beta_2$, it is likely that $E[\hat{\beta}_1 - \hat{\beta}_2] > 0$, where $\hat{\beta}_1$ and $\hat{\beta}_2$ refer to OLS estimates of β_1 and β_2 , respectively. To account for this possibility, we adopt the bias-corrected estimation and testing approach for multivariate predictive regressions proposed by Amihud, Hurvich, and Wang (2009).

Results appear in Table 3.2. For the 1927–2014 sample period, the Wald test of the restriction $\beta_1 = \beta_2$ is rejected for bias-corrected estimates based on Eq. (3.6). The rejection is due to the fact that, after bias-adjustment, the estimated slope coefficient on $-IY$ ($\hat{\beta}_2 \approx 7.3$) is much larger than the estimated slope coefficient for TPY ($\hat{\beta}_1 \approx 2.0$). However, this apparent rejection of the restriction imposed by the net payout yield should be interpreted with caution, because the estimated slope coefficient on $-IY$ is inflated due to the inclusion of the years 1929–1930 in the sample.

Panels B and C consider alternative sample periods of 1932–2014 and 1952–2014, respectively. For these sample periods, the Wald test is unable to reject the null hypothesis that $\beta_1 = \beta_2$. This conclusion is robust to whether inference is based on standard OLS estimates

⁵³It is conceptually possible to decompose the net payout yield into three components reflecting dividends, repurchases, and issuances, respectively. However, stock repurchases are negligible prior to the early 1980s, and the repurchase yield exhibits (upward) trending behavior for much of the subsequent two decades. For these reasons, we focus on the specifications of Eqs. (3.5)–(3.6).

or bias-corrected estimates. The point estimates of the slope coefficients on the two yield components are often quite close. For example, over the 1932–2014 sample period (Panel B), the bias-corrected estimates of the coefficients on DY and $-NIY$ are 3.61 and 3.36, respectively. These results show that omitting the anomalous issuance years of 1929–1930 actually *supports* the case for the net payout yield. Otherwise, it would seem prudent to separate payout components in return forecasting regressions.

An alternative set of tests based on Eqs. (3.5)–(3.6) adopts the null hypothesis that $\beta_2 = 0$. This amounts to testing that the correct measure of payout is dividends (Eq.(3.5)) or total payout (Eq.(3.6)). The null hypothesis that $\beta_2 = 0$ is typically rejected at conventional levels for bias-adjusted estimates. For example, bias-corrected t -statistics for $\hat{\beta}_2$ for the 1932–2014 sample period (Panel B) are approximately two for both decompositions.

3.3.3 Out-of-sample analysis

Out-of-sample research designs split the data into two subsets, the first constituting an estimation sample and the second constituting a hold-out sample. Forecasts for the hold-out sample are constructed using past information, as though the forecasts were constructed in real time. Out-of-sample research designs are often motivated by concerns about structural change and to provide a measure of robustness to potential data mining. In an influential paper, Goyal and Welch (2008) document that many return forecasting models fail to outperform the historical average out-of-sample.

We compute out-of-sample forecasts of returns using alternative payout yield measures. The benchmark forecast is the historical average. To compute forecasts, we consider both rolling and recursive estimation schemes, with an initial estimation window of 30 years. Let $\hat{R}_{t+1}^e$ denote the one-step ahead return forecast based on a particular predictive regression. The corresponding sample mean square prediction error (MSPE) is

$$\hat{\sigma}^2 \equiv P^{-1} \sum (R_{t+1}^e - \hat{R}_{t+1}^e)^2, \quad (3.7)$$

where P denotes the number of out-of-sample forecasts. Similarly, let $\bar{R}_{t+1}$ denote the one-step ahead forecast based on the historical average excess return, with corresponding sample MSPE equal to $\hat{\sigma}_0^2$. To convey the economic significance of differences in forecast performance, we report the out-of-sample R^2 statistic considered by Campbell and Thompson (2008) and Goyal and Welch (2008), defined as

$$R_{OOS}^2 = 1 - \frac{\hat{\sigma}^2}{\hat{\sigma}_0^2}. \quad (3.8)$$

Statistical inference in out-of-sample research designs is nuanced, due to the fact that forecasts rely on estimated parameters. We conduct two alternative tests for differences in out-of-sample forecast performance. The first is the test proposed by Clark and West (2007) for equal MSPE. The second is the test for superior predictive ability proposed by Giacomini and White (2006). The null hypothesis in these alternative tests differs in a subtle, but important way. In essence, the Clark and West (2007) test is an out-of-sample test for Granger causality, while the Giacomini and White (2006) test focuses on differences in predictive ability, factoring in the impact of parameter estimation on forecast performance. The latter test is more stringent: a variable may Granger cause stock returns (β is nonzero in a predictive regression), but fail to outperform the historical average in practice, due to the bias-variance tradeoff. See the Appendix and Clark and McCracken (2011) for additional discussion and explicit expressions for test statistics.

Table 3.3 presents out-of-sample results. We highlight three points. First, the net payout yield achieves the highest out-of-sample R^2 -value among the candidate yields and outperforms the historical average under both the rolling and recursive designs. Second, forecasts based on the issuance yield and net issuance yield perform relatively poorly out-of-sample. Third, return forecasts based on the net payout yield, which imposes the restriction of equal slope coefficients on the components of net payout, perform better than forecasts from multivariate models that separate components of payout. This pattern of results is consistent with earlier, in-sample results supporting the use of the net payout yield. Out-of-sample MSPE differences are subject to considerable sampling error and the Giacomini and White (2006) test often fails

to reject the null of equal predictive ability for alternative yields. There are some exceptions. For example, the issuance yield and net issuance yield significantly *underperform* the historical average out-of-sample under the rolling scheme.⁵⁴

The Clark and West (2007) test rejects the null of no predictability for the dividend yield, total payout yield, and net payout yield. In some cases (e.g., the dividend yield), these rejections occur in spite of the fact that the corresponding variable fails to outperform the historical average out-of-sample. The intuition behind this result is that, under the null of no predictability, the forecast based on the yield is expected to underperform the benchmark due to additional estimation error. In the data, forecasts for these models perform sufficiently well relative to the benchmark (although not necessarily outperforming it) such that we reject the null. The Clark and West (2007) fails to reject the null of no predictability for the issuance yield or net issuance yield. These findings complement in-sample results from univariate regressions once the 1929–1930 period is removed from the sample.

3.3.4 Joint tests incorporating returns and cash flow growth

The Campbell and Shiller (1988a) linearized present value identity implies that variation in the dividend yield must reflect news about future dividend growth or future returns. Cochrane (2008) advocates tests for the null of no stock return predictability that jointly exploit information in slope estimates for returns and dividend growth.

Let cf_t represent the logarithm of an arbitrary cash flow measure, such as total payout or stock issuances. The following system generalizes the VAR linking stock return and cash flow

⁵⁴Asymptotic results in Giacomini and White (2006) require the estimation sample to remain bounded, which is inconsistent with an estimation window that expands forever. We report these test results for comparison with results using a rolling window; however, they should be treated with caution.

growth predictability considered in, e.g., Cochrane (2008):

$$r_{t+1} = a_r + \beta_r yld_t + \epsilon_{t+1}^r \quad (3.9)$$

$$\Delta cf_{t+1} = a_{cf} + \beta_{cf} yld_t + \epsilon_{t+1}^{cf} \quad (3.10)$$

$$yld_{t+1} = a_{yld} + \phi yld_t + \epsilon_{t+1}^{yld}, \quad (3.11)$$

$$cf_{t+1} - d_{t+1} = a_{cf,d} + \beta_{cf,d} yld_t + \epsilon_{t+1}^{cf,d} \quad (3.12)$$

where the payout yield yld_t is defined as cash flow divided by market capitalization.⁵⁵ When dividends represent the cash flow measure, $cf_t = d_t$, $yld_t = dy_t$, and Eq. (3.12) can be dropped.

The following approximate identity links slope coefficients in the system (3.9)–(3.12):

$$\beta_r \approx 1 - \rho\phi + \beta_{cf} - (1 - \rho)\beta_{cf,d}, \quad (3.13)$$

where the parameter $\rho \equiv 1/(1 + \overline{D/P})$, and $\overline{D/P}$ represents the dividend yield around which one linearizes.⁵⁶ For the case of dividends, the identity (3.13) reduces to $\beta_r \approx 1 - \rho\phi + \beta_d$. Assuming $|\phi| < 1$, $\beta_r = 0$ implies $\beta_d < 0$. The dividend yield can only fail to predict stock returns if higher yields forecast lower future dividend growth. Estimates of $(1 - \rho)\beta_{cf,d}$ are small for the alternative cash flow measures (≈ 0.01). Consequently, the null hypothesis of no stock return predictability continues to imply $\beta_{cf} < 0$.

We conduct joint tests for stock return or cash flow predictability following Cochrane (2008) and Engsted and Pedersen (2010). To illustrate the approach, consider the null of no stock return predictability ($\beta_r^N = 0$). Under this hypothesis, Eq. (3.13) implies that $\beta_{cf}^N = \rho\phi +$

⁵⁵The cash flow measure must be strictly positive to ensure a well-defined log payout yield yld_t and log cash flow growth Δcf_t . Dividends, total payout, and issuances satisfy this requirement; however, net payout can be negative. Empirically, negative values for NPY only occur in 1929 and 1930, coinciding with extraordinary stock issuance during these years. Consequently, we analyze sample periods starting in 1931.

⁵⁶Starting from an aggregate version of the Campbell-Shiller log-linearization of returns:

$$\begin{aligned} r_{t+1} &\approx \rho(mv_{t+1} - d_{t+1}) + \Delta d_{t+1} - (mv_t - d_t) \\ &= -\rho(yld_{t+1}) + \Delta cf_{t+1} + yld_t - (1 - \rho)(cf_{t+1} - d_{t+1}), \end{aligned} \quad (3.14)$$

where mv is log market value, and the second line follows from the first after some algebra. Projecting both sides of (3.14) on yld_t gives (3.13). It is implicitly assumed that the log ratio of cash flow to dividends $cf_t - d_t$ is stationary. An analogous identity links the error terms: $\epsilon_{t+1}^r \approx \epsilon_{t+1}^{cf} - \rho\epsilon_{t+1}^{yld} - (1 - \rho)\epsilon_{t+1}^{cf,d}$.

$(1 - \rho)\beta_{cf,d} - 1$. The parameter ρ is calibrated based on the historical average of the dividend yield. Setting ϕ and $\beta_{cf,d}$ to corresponding empirical estimates gives the implied value of β_{cf}^N . We then simulate data from the system (3.9)–(3.12) under the null. Samples are generated in accordance with the Campbell-Shiller identity, following Cochrane (2008).

The Monte Carlo analysis computes estimates of slope coefficients in forecasting regressions for 10,000 simulated samples under the null of no predictability. We report the simulation-based estimate of the probability of obtaining OLS estimates ‘more extreme’ than the actual estimates in the data. A more extreme coefficient is defined as a coefficient that is further than the actual estimate from the corresponding coefficient under the null, in a one-sided sense. For example, consider a test of no stock return predictability using the dividend yield. In the data, it is the case that $\hat{\beta}_r > \beta_r^N = 0$ and $\hat{\beta}_{cf} > \beta_{cf}^N$. We report the simulated probability $P(\hat{\beta}_r^* > \hat{\beta}_r)$ and $P(\hat{\beta}_{cf}^* > \hat{\beta}_{cf})$, where $\hat{\beta}_r^*$ and $\hat{\beta}_{cf}^*$ denote coefficients from simulated data. The notation $P(\hat{\beta}_r^* \leq \hat{\beta}_r)$ and $P(\hat{\beta}_{cf}^* \leq \hat{\beta}_{cf})$ indicates the probability of a ‘more extreme’ estimate under the null defined in this way. $P(\text{Both})$ indicates the simulated probability that both estimates are more extreme than those observed in the data.

Table 3.4 presents results for the joint tests. Panels A and B report results for the null of no stock return predictability. Parameters under the null are based either on OLS estimates (Panel A) or bias-corrected estimates (Panel B). The joint test strongly rejects the null of no stock return predictability for the dividend, total payout, and net payout yields. Rejections for the dividend and total payout yields are due to the lack of predictability for cash-flow growth (Cochrane (2008)). In contrast, the rejection for net payout is primarily due to the large return forecasting coefficient, rather than the cash flow growth-forecasting coefficient. Finally, incorporating issuance growth strengthens the evidence that the issuance yield forecasts returns relative to standard predictive regressions (Table 3.1). The simulated p -values for the joint tests are 7-9%, indicating borderline evidence against the null. This does not contradict earlier evidence suggesting that net payout is the appropriate cash flow measure for a return-forecasting yield, because issuance is a component of net payout.

Results for joint tests of predictability in cash flow growth (Panels C and D of Table 3.4) show that the evidence of predictability depends highly on how cash flow is measured. When dividends or total payout serve as the cash flow measure, the joint tests fail to reject the null of no predictability. In contrast, there is strong evidence of predictable variation in the growth rate of net payout. The fact that there is little evidence of predictability for growth in total payout, but strong evidence of predictability for growth in net payout, suggests that there is a strong predictable component to issuance growth. Indeed, the joint tests strongly reject the null of no predictability in issuance growth.

In unreported results, we decompose variation in each payout yield into components capturing changes in expected returns, changes in expected cash flow growth, and changes in the expected ‘spread’ between (log) cash flow and dividends. Most variation in the dividend and total payout yields reflects changes in expected returns, while roughly equal portions of variation in the net payout yield reflect changes in expected returns and net payout growth. These results corroborate evidence in ? and Larrain and Yogo (2008).

3.4 Economic significance of the payout measure

3.4.1 The implied persistence of the equity premium

Fama and French (1989) argue that the dividend yield captures variation in expected returns related to persistent aspects of business conditions. Over recent decades; however, the dividend yield has become increasingly persistent and variation in the dividend yield does not closely track real economic activity.

When market excess returns depend linearly on a lagged payout yield, the equity premium inherits the same persistence pattern as the forecasting yield. The highly persistent nature of the dividend yield implies very low-frequency cyclical variation in discount rates. The implied half-life of return responses to a shock to the dividend yield is approximately 7–9 years, depending on the sample period and whether the dividend yield is in level or log form. The total and net payout yields imply less persistent time series variation in the equity premium. For example,

the half-life of return responses to a shock to the net payout yield is approximately 2–3 years, which is reasonably consistent with macroeconomic shocks in the context of the business cycle.

3.4.2 Long-horizon portfolio choice

To characterize the economic significance of the choice of payout measure from a different perspective, we consider implications for investors' long-horizon portfolio decisions. Our analysis applies optimal portfolio policy rules derived by Jurek and Viceira (2011). Let $r_{1,t+1}$ denote the potentially time-varying log return on a benchmark asset of interest and $x_{t+1} = r_{t+1} - r_{1,t+1}\iota$ denote an $N \times 1$ vector of log returns in excess of the benchmark reflecting other investment options. (The symbol ι represents an $N \times 1$ vector of ones.) Finally, let s_{t+1} denote an $M \times 1$ vector of state variables capturing time-variation in expected returns for the benchmark and other assets. Jurek and Viceira (2011) assume that asset returns and state variables jointly follow a first-order VAR process:

$$z_{t+1} = \Phi_0 + \Phi_1 z_t + v_{t+1}, \quad (3.15)$$

where z_{t+1} is the $(N + 1 + M) \times 1$ vector defined as $z_{t+1} = (r_{1,t+1}, x'_{t+1}, s'_{t+1})'$. Zero-mean shocks v_{t+1} are homoskedastic and normally distributed with covariance matrix Σ_v .

An investor with power utility maximizes the expected utility of terminal wealth K periods in the future. More specifically, the investor chooses a sequence of portfolio weights $\alpha_{t+K-\tau}$ from time t to time $t + K - 1$ such that

$$\left\{ \alpha_{t+K-\tau}^{(\tau)} \right\}_{\tau=1}^{\tau=K} = \arg \max E_t \left[\frac{1}{1-\gamma} W_{t+K}^{1-\gamma} \right], \quad (3.16)$$

when $\gamma \neq 1$ and

$$\left\{ \alpha_{t+K-\tau}^{(\tau)} \right\}_{\tau=1}^{\tau=K} = \arg \max E_t [\log(W_{t+K})]$$

when $\gamma = 1$. Following Jurek and Viceira (2011), we index the sequence of portfolio weight vectors by the time at which they are chosen (the subscript), as well as the time remaining to

the horizon (the superscript τ). Wealth evolves according to $W_{t+1} = W_t(1 + R_{p,t+1})$, where $1 + R_{p,t+1}$, the gross return on wealth, is given by

$$1 + R_{p,t+1} = \alpha'_t(R_{t+1} - R_{1,t+1}) + (1 + R_{1,t+1}).$$

Jurek and Viceira (2011) derive an analytic solution for the optimal dynamic rebalancing strategy via a backwards recursion argument. The solution strategy follows Campbell and Viceira (1999) and Campbell, Chan, and Viceira (2003) by employing an approximation to the log return on the wealth portfolio. The quality of the approximation is increasing in the frequency of portfolio rebalancing, and is exact in continuous time.⁵⁷ Under this approximation, optimal total portfolio demand is an affine function of the vector of state variables z_t :

$$\alpha_{t+K-\tau}^{(\tau)} = A_0^{(\tau)} + A_1^{(\tau)} z_{t+K-\tau}, \quad (3.17)$$

where $A_0^{(\tau)}$ and $A_1^{(\tau)}$ are horizon-dependent coefficients. Total portfolio demand $\alpha_{t+K-\tau}^{(\tau)}$ can be decomposed into two components: a ‘myopic’ component that does not depend on the investment horizon, and a second, horizon-dependent component reflecting hedging demand in the sense of Merton (1969, 1971, 1973). We refer readers to Jurek and Viceira (2011) for explicit expressions characterizing the portfolio rules. Hedging demand under the portfolio rules depends upon the correlation between asset returns and the vector of state variables, which reflects the ability of assets to hedge changes in future investment opportunities.

To assess how alternative yield definitions influence long-horizon portfolio decisions, we perform a calibration analysis based on the results of Jurek and Viceira (2011). The vector z_t in the calibration is:

$$z_t = (r_{1,t}, r_t^e, YLD_t)', \quad (3.18)$$

⁵⁷Campbell and Viceira (1999) derive optimal portfolio and consumption rules for an infinitely-lived individual with Epstein-Zin-Weil preferences who faces a constant risk-free rate and a single risky asset with time-varying expected returns. Campbell, Chan, and Viceira (2003) extend results to multiple risky assets and state variables. Campbell, Chan, and Viceira (2003) use a numerical procedure to solve for optimal consumption and portfolio choice, whereas Jurek and Viceira (2011) obtain an (approximate) analytical solution in the simpler setting without intermediate consumption. An earlier version of this study provides additional calibration results in the Campbell and Viceira (1999) setting. Key results regarding hedging demand continue to obtain.

where $r_{1,t}$ is the real log return on a 90-day Treasury bill, r_t^e is the log market excess return and YLD_t is the payout yield of interest (e.g., the dividend yield). Consequently, there are two investment options, the stock market and a short-term Treasury bill, and s_t consists of a single state variable: the payout yield.⁵⁸ The value-weighted CRSP portfolio serves as the market proxy and returns are in excess of the 90-day Treasury bill return.

We estimate the VAR system (3.15) using annual data over the period 1932–2014 to calibrate VAR parameters. Table 3.5 presents estimation results. The first six columns present estimated slope coefficients $\hat{\Phi}_1$ and the associated t -statistics in parentheses. Under the unrestricted VAR model, the predictive regressions for excess stock returns (r_t^e) include the lagged real T-bill rate and lagged excess returns, as well as the lagged payout yield. This has little impact on the forecasting power of lagged payout yields. Each of the three payout yields are economically and statistically significant predictors of market excess returns, and slope coefficients are similar to those in Panel B of Table 3.1.

Table 3.5 also presents calibration results associated with the covariance matrix of shocks $\hat{\Sigma}_v$. The table shows the standard deviation of shocks on the diagonal, and off-diagonal elements represent pairwise correlations between corresponding shocks. Shocks to the dividend yield are strongly correlated with contemporaneous shocks to market returns (correlation of -59.97%). Total payout yield shocks and stock return shocks are similarly correlated (-59.98%). In contrast, contemporaneous shocks to the net payout yield and returns are relatively weakly correlated (-33.46%). This proves to be critical in the calibration results that follow. Intuitively, since shocks to the net payout yield are only weakly correlated with contemporaneous shocks to returns, hedging is likely to represent a less important component of demand for risky stock when the net payout yield serves as the state variable.

⁵⁸We considered a more general VAR with two additional state variables: the default spread and term spread. These additional variables do little to alter key results. The state variable YLD_t is defined as the *level* of the yield rather than the log yield. Alternatively defining the state variable as the log yield has little impact on the pattern of long-horizon allocations to stock. We elect to use the level of the yield as the state variable because this affords the simple decomposition $NPY = DY - NIY$, which facilitates our subsequent welfare analysis (see Section 3.4.3). Finally, key results are robust to calibrations based on more recent sample periods, e.g., 1952–2014, and to employing bias-corrected VAR slope coefficients rather than OLS estimates.

Figure 3.2 shows the average allocation to the stock market as a function of the investment horizon in years for alternative calibrations in which the included payout yield is either the dividend yield, total payout yield, or net payout yield. The top panel shows results for a relatively low level of risk aversion ($\gamma = 2$) and the bottom panel shows results for a higher risk aversion level ($\gamma = 10$). Results at a one-year horizon reflect the average myopic demand for risky stock. This is higher under the alternative payout yields than under the dividend yield because the slope coefficient on these yields in the predictive regression for excess stock returns is larger than that for the dividend yield. Demand for stock rises more rapidly as a function of investment horizon for the dividend yield relative to the alternative payout yields, because hedging demand is stronger for the dividend yield.

As the investment horizon increases, the average allocation to stock under the total payout yield increases more rapidly than that under the net payout yield. This reflects the fact that hedging demand is stronger under the total payout yield. For less risk averse investors (top panel), the average allocation to the stock is nearly the same under the dividend and net payout yield at long horizons. This is due to the fact that the larger hedging demand under the dividend yield offsets the lower myopic component of demand at long horizons. At a higher level of risk aversion, hedging demand is more important, and consequently the long-horizon allocation to stock under the net payout yield is much lower than under the dividend and total payout yields. Figure 3.2 shows that there is a certain investment horizon at which average demand for stock under the net payout and dividend yields is equal. This occurs at roughly a 22-year horizon for less risk-averse investors (top panel) and at a much shorter 5-year horizon for more risk-averse investors (bottom panel). Intuitively, the crossing point reflects the horizon at which larger hedging demand under the dividend yield counterbalances the larger myopic demand under the net payout yield.

3.4.3 Welfare implications

A final empirical analysis characterizes welfare effects associated with using the correct or incorrect form of payout. We consider a scenario in which dividends and net issuance only

influence asset returns via net payout. There are two investor types: one who employs the correct state variable and optimal strategy, and another who omits net issuance and relies on the dividend yield.⁵⁹ We simulate a large number of paths for asset returns and state variables and track realized wealth and utility for each investor type. This allows us to compare welfare outcomes and determine the fee that an investor basing decisions on dividends would pay to switch to the optimal strategy based on net payout.

The simulation analysis sets the true state vector to $z_t = (r_{1,t}, r_t^e, DY_t, NIY_t)'$, where DY_t and NIY_t are the dividend yield and net issuance yield, respectively. The dividend and net issuance yields are separately included so that simulated paths for the two yields preserve the correlation structure observed in the data. We calibrate the VAR parameters using annual data for 1932–2012, with one adjustment. Specifically, we reset the slope coefficients on lagged DY_t and NIY_t such that predictable variation in market risk premia and interest rates are effectively driven by the net payout yield. For example, we set the slope coefficients on lagged DY_t and NIY_t in the market excess returns equation to β^{NPY} and $-\beta^{NPY}$, respectively, where β^{NPY} represents the slope coefficient in a regression of market excess returns on lagged NPY_t .

For each simulated path, we track asset allocations and wealth outcomes for two investors: one who applies the optimal dynamic portfolio strategy and a second who incorrectly relies on $\tilde{z}_t = (r_{1,t}, r_t^e, DY_t)'$ as the vector of state variables, and makes dynamic allocations accordingly. The investment horizon is K years, with $K \in \{1, 5, 10, 20, 30\}$. Initial wealth is normalized to one. The simulations yield one million wealth outcomes W_{t+K} and $\tilde{W}_{t+K}$ for policies based on z_t and $\tilde{z}_t$, respectively, and the associated realized utilities. We compute a certainty equivalent wealth associated with each policy as $W_{CE} = U^{-1}(\overline{EU})$, where $\overline{EU}$ is the simulation-based expected utility of final wealth.

Table 3.6 presents results. Panel A reports the annual growth rate r_{CE}^g implied by the certainty equivalent wealth associated with the optimal dynamic portfolio strategy, computed as $(1 + r_{CE}^g)^K = W_{CE}$. Panel B shows the analogous certainty equivalent growth rate $\tilde{r}_{CE}^g$ for the suboptimal policy based on the dividend yield. Panel C presents the percentage reduction

⁵⁹The assumed suboptimality of the dividend yield is a convention. One could instead assume that the dividend yield is optimal and a policy that (needlessly) incorporates net issuance is suboptimal.

in the certainty equivalents of wealth from basing dynamic allocations on the dividend yield rather than the net payout yield. Panel D presents the maximal annual management fee ϕ an agent would be willing to pay to switch from the dynamic strategy based on the dividend yield to the optimal strategy. This is computed as the solution to $(1 + r_{CE}^g)(1 - \phi) = 1 + \tilde{r}_{CE}^g$.⁶⁰

Log utility investors allocate aggressively to stock. The growth rate corresponding to the certainty equivalent wealth level for the optimal portfolio policy is around 18.5% and is essentially constant across investment horizon, due to the myopic nature of portfolio demand under log utility. For a fixed investment horizon, the certainty equivalent growth rate declines as risk aversion increases. For a fixed level of risk aversion, the growth rate generally increases with the investment horizon for $\gamma > 1$, although the magnitude of increases are modest. The certainty equivalent growth rates associated with the suboptimal dividend-based policy (Panel B) are, naturally, lower than those for the optimal policy. They also decrease in magnitude with the level of risk aversion. For fixed $\gamma > 1$, the certainty equivalent growth rates for the suboptimal policy decrease considerably at long horizons. Consequently, the difference between the certainty equivalent growth rates associated with the optimal and suboptimal policies increases at long horizons.

The certainly equivalent growth rates in Table 3.6 are relatively large at low risk aversion levels. These results are partially attributable to the equity premium puzzle. They also reflect the economic value of market timing induced by return predictability, which is documented by Campbell and Viceira (1999) in a similar portfolio choice setting. As in Campbell and Viceira (1999) and Jurek and Viceira (2011), our analysis implicitly assumes that parameters describing the investment opportunity set are known to investors. The concluding section of the study discusses potential extensions to incorporate parameter uncertainty.

For myopic investors ($K = 1$), the percentage loss in certainty equivalent wealth associated with relying on the dividend-based portfolio policy ranges from approximately 4% to slightly

⁶⁰The analytical portfolio rules of Jurek and Viceira (2011) are based on an approximation of the log portfolio return. Our simulation analysis assumes that portfolio returns obey the approximation exactly in order to avoid technical issues associated with the existence of expected utility under the power utility specification. Results for $\gamma = 20$ are omitted because we found that simulation results at this risk aversion level remain unstable even with one million simulated paths.

under 0.4% as risk aversion ranges from $\gamma = 1$ to $\gamma = 10$. The percentage loss in certainty equivalent wealth rises rapidly as the horizon increases. For an investment horizon of 20 years, the percentage reduction in certainty equivalent wealth of relying on the dividend-based policy is around 57% for a log utility investor and approximately 15% for an investor with $\gamma = 10$. The horizon-dependence of these results stem from two features of the data: 1) differences in the certainty equivalent growth rate accumulate over time; and 2) these growth rate differences increase with horizon for $\gamma > 1$.

Finally, the maximum annual management fee ϕ that an agent would be willing to pay to switch from the dividend-based policy to the optimal policy (Panel D of Table 3.6) ranges from around 40 basis points to over 4%, depending on investment horizon and risk aversion level. The fee is higher for lower risk aversion levels and (typically) at longer horizons. It is particularly high at long horizons ($K = 30$ years), even for larger risk aversion levels. This is due to the fact that, when $\gamma > 1$, the wedge between certainty equivalent growth rates associated with the optimal and suboptimal policies increases at long horizons. Overall, the results indicate that there are important welfare effects associated with correctly identifying the particular payout variable that governs variation in expected returns.

3.5 Concluding Remarks

Predictive regressions for stock returns based on alternative payout yields are sensitive to the sample period examined, due to extraordinary issuance activity around the Great Depression. Despite this point, we find that alternative payout yields outperform the traditional dividend yield. However, differences in forecasting performance, as measured by in- and out-of-sample R^2 -values, are relatively modest. We show that it is not preferable to separate payout components in predictive regressions, because slope coefficients associated with (net) issuance are estimated very imprecisely. When stock repurchases and issuances are incorporated into a net payout yield, variation in the yield captures both changes in expected returns and changes in expected (net) cash flow growth. The particular choice of yield matters economically for dy-

dynamic portfolio decisions, particularly via the extent of hedging demand associated with risky stock, which is much lower under a net measure of payout.

We note two interesting directions for future research. First, optimal long-horizon portfolio and consumption policies depend on model parameters that are unknown to investors. In addition, our results suggest a substantial degree of model uncertainty with respect to the appropriate payout yield to include in predictive regressions for stock returns. Bayesian techniques, including Bayesian model averaging methods, provide a way to formally incorporate uncertainty regarding parameter values and models into return forecasts and long-horizon portfolio choice. It would be interesting to examine the implications of using alternative payout yields in Bayesian implementations of predictive regressions such as Barberis (2000), Johannes, Kortweg, and Polson (2014), and Wachter and Warusawitharana (2009, 2013). Second, van Binsbergen and Koijen (2010) model expected returns and expected dividend growth as latent variables that follow an exogenous process. In the context of a present-value model, they exploit the entire history of price-dividend ratios and dividend growth to estimate expected returns and expected dividend growth. A similar approach might be applied to improve return and cash flow growth forecasts using alternative cash flow measures.




Figure 3.1: Time series plots for alternative yields.

The first panel plots the dividend yield, total payout yield and net payout yield. The dividend yield is dividend over price. The total payout yield is the dividend yield plus the repurchases yield. The net payout yield is the dividend yield less the net issuance yield, where the net issuance yield is the issuance yield less the repurchase yield. The second panel plots the issuance yield and net issuance yield. See Section 3.2 and the Appendix for details regarding variable construction. All series are constructed annually covering the period 1926–2014.




Figure 3.2: Average allocation to stock as a function of investment horizon. The figure shows the average allocation to stock, expressed as a percentage, as a function of the investment horizon in years for optimal portfolios held by an investor in the Jurek and Viceira (2011) portfolio choice framework. Results are presented for three alternative definitions of the payout yield: the dividend yield (DY), total payout yield (TPY), and net payout yield (NPY). The dynamics of asset returns and state variables are governed by a first-order VAR. Results are based on VAR parameter estimates over the sample period 1932–2014.

Table 3.1: Predictive regressions for stock returns using alternative yields

This table presents univariate predictive regressions for stock returns using alternative payout yields. The dependent variable is either the real excess return (R^e) on the value-weighted CRSP index or the corresponding log excess return (r^e). (DY) is the aggregate dividend-to-price ratio. TPY , NPY , IY and NIY denote the total payout yield, net payout yield, issuance yield, and net issuance yield, respectively. Lowercase symbols represent natural logarithms of corresponding uppercase symbols. The log net payout yield npy and log net issuance yield niy are constructed as the natural logarithm of 0.1 plus the corresponding ratio to ensure that the logarithm is well-defined. See Section 3.2 for details concerning data construction. $\hat{\beta}$ is the OLS estimate of the coefficient on the corresponding yield in a univariate predictive regression. The table reports the heteroskedasticity-corrected standard error for this estimate ($SE(\hat{\beta})$) and the corresponding t -statistic. $\hat{\beta}^c$ represents the bias-corrected coefficient proposed by Amihud and Hurvich (2004). $SE(\hat{\beta}^c)$ denotes the associated standard error, and t -stat (AH) is the corresponding t -statistic. $\rho(\%)$ is the cross correlation between the predictive regression residuals and the residuals from an estimated AR(1) model for the predictor. R^2 is the standard R^2 -value for the regression. Panel A presents results for the sample period 1927–2014, Panel B presents results for 1932–2014, and Panel C presents results for 1952–2014.

	Dependent variable = R^e					Dependent variable = r^e				
	DY	TPY	NPY	IY	NIY	dy	tpy	npy	iy	niy
Panel A: Sample period 1927–2014										
$\hat{\beta}$	3.28	4.87	5.50	-7.27	-3.69	0.10	0.21	0.70	-0.07	-0.34
$SE(\hat{\beta})$	1.27	1.35	1.27	2.10	2.26	0.05	0.06	0.14	0.04	0.29
t -stat	2.58	3.61	4.34	-3.45	-1.63	2.07	3.31	4.93	-1.72	-1.19
$\hat{\beta}^c$	2.44	4.26	5.31	-7.33	-3.58	0.06	0.17	0.68	-0.07	-0.33
$SE(\hat{\beta}^c)$	1.45	1.72	1.20	2.04	1.49	0.05	0.08	0.15	0.03	0.16
t -stat (AH)	1.69	2.48	4.41	-3.60	-2.40	1.12	2.25	4.67	-2.46	-2.03
$\rho(\%)$	-63.97	-63.39	-36.60	-7.61	12.67	-66.10	-65.27	-31.98	-10.34	13.64
R^2	5.85	8.85	19.89	13.05	6.74	4.19	7.89	21.48	6.48	4.99
Panel B: Sample period 1932–2014										
$\hat{\beta}$	3.35	4.48	4.08	-3.13	-0.59	0.11	0.19	0.50	-0.03	-0.01
$SE(\hat{\beta})$	1.21	1.35	1.26	3.16	1.83	0.04	0.06	0.16	0.03	0.20
t -stat	2.76	3.32	3.25	-0.99	-0.32	2.49	3.11	3.15	-1.19	-0.07
$\hat{\beta}^c$	2.57	3.90	3.82	-3.42	-0.52	0.07	0.16	0.47	-0.03	-0.01
$SE(\hat{\beta}^c)$	1.34	1.59	1.38	3.22	1.88	0.05	0.07	0.18	0.03	0.18
t -stat (AH)	1.92	2.45	2.76	-1.06	-0.28	1.51	2.32	2.60	-1.20	-0.03
$\rho(\%)$	-59.71	-61.54	-36.51	-21.81	6.75	-63.40	-63.89	-32.91	-16.89	6.03
R^2	7.46	9.24	9.95	1.18	0.13	6.43	8.91	8.86	1.60	0.01
Panel C: Sample period 1952–2014										
$\hat{\beta}$	3.21	3.97	4.35	0.33	-0.48	0.10	0.17	0.53	0.01	-0.01
$SE(\hat{\beta})$	1.80	1.83	1.95	3.83	1.89	0.05	0.08	0.25	0.05	0.21
t -stat	1.78	2.17	2.23	0.09	-0.25	1.83	2.13	2.07	0.14	-0.05
$\hat{\beta}^c$	1.74	3.17	4.03	-0.10	-0.28	0.04	0.13	0.49	0.00	0.01
$SE(\hat{\beta}^c)$	1.92	1.93	2.01	4.35	1.87	0.06	0.08	0.26	0.07	0.18
t -stat (AH)	0.90	1.64	2.00	-0.02	-0.15	0.75	1.59	1.90	0.01	0.04
$\rho(\%)$	-58.69	-58.61	-32.91	-20.71	13.79	-61.73	-62.94	-32.42	-20.73	12.77
R^2	4.64	6.78	7.31	0.01	0.11	4.63	7.08	6.67	0.02	0.01

Table 3.2: Multivariate regressions: decomposing the net payout yield

This table presents results for in-sample regressions that employ multiple yield measures. The yield measures decompose the net payout yield into two components. The dependent variable is real excess returns on the VW CRSP portfolio. Model 1 decomposes the net payout yield as $NPY = DY + (-NIY)$, with DY representing the dividend yield and NIY representing the net issuance yield. Model 2 uses an alternate decomposition of the net payout yield as $NPY = TPY + (-IY)$, with TPY representing the total payout yield and IY representing the issuance yield. For each model, the table presents coefficient estimates, standard errors, and t-statistics for the corresponding variable using two alternative methods: ordinary least squares (OLS) with heteroskedasticity-corrected standard errors and the bias-corrected estimation and testing approach for multivariate predictive regressions proposed by Amihud, Hurvich, and Wang (2009). ‘Wald’ refers to a Wald statistic testing the null hypothesis that $\beta_1 = \beta_2$. The corresponding p -value is also provided. Panel A presents results for 1927–2014. Panel B presents results for 1932–2014 and Panel C presents results for 1952–2014.

		Model 1: $DY + (-NIY)$					Model 2: $TPY + (-IY)$				
	Method	DY	$-NIY$	Wald	p -val	R^2 (%)	TPY	$-IY$	Wald	p -val	R^2 (%)
Panel A: 1927–2014											
Coefficient	OLS	5.30	5.74	0.05	0.82	19.96	4.53	6.94	0.92	0.34	20.70
Standard Error		1.32	1.82				1.34	2.04			
t-statistic		4.03	3.16				3.37	3.39			
Coefficient	AHW	4.50	5.71	0.73	0.39		1.97	7.32	6.58	0.01	
Standard Error		1.19	1.29				1.30	1.56			
t-statistic		3.77	4.43				1.52	4.70			
Panel B: 1932–2014											
Coefficient	OLS	4.36	3.19	0.40	0.53	10.38	4.43	2.82	0.27	0.60	10.19
Standard Error		1.32	1.91				1.32	2.93			
t-statistic		3.32	1.67				3.36	0.96			
Coefficient	AHW	3.61	3.36	0.02	0.87		1.72	4.71	1.12	0.29	
Standard Error		1.22	1.65				1.46	2.39			
t-statistic		2.95	2.03				1.17	1.97			
Panel C: 1952–2014											
Coefficient	OLS	5.34	3.51	0.74	0.39	8.55	4.43	3.01	0.17	0.68	7.49
Standard Error		2.08	2.16				1.98	3.80			
t-statistic		2.57	1.63				2.24	0.79			
Coefficient	AHW	4.37	3.70	0.13	0.71		2.55	5.54	0.64	0.42	
Standard Error		1.93	1.93				2.32	3.47			
t-statistic		2.26	1.91				1.10	1.60			

Table 3.3: Out-of-sample evidence: stock return forecasts

This table presents results for out-of-sample forecasts of returns using alternative payout yield measures. $OOS R^2$ is the out-of-sample R^2 -value relative to the historical average, as defined in Campbell and Thompson (2008). The statistics GW t -stat and GW p -value represent the t -statistic and associated p -value for the Giacomini and White (2006) test for equal predictive ability. For univariate models the historical average return serves as the benchmark in these tests. For multivariate models, the benchmark model is the univariate model using the first predictor. For example, for the model employing both the dividend yield and net issuance yield (DY, NIY), the benchmark model is the univariate dividend yield (DY) forecasting model. The statistics CW t -stat and CW p -value correspond to the out-of-sample test for equal MSPE proposed by Clark and West (2007). Panel A presents results where forecasts are based on estimates computed using a rolling window consisting of the 30 preceding years. Panel B presents results for a recursive estimation window.

	Univariate models					Multivariate models	
	DY	TPY	NPY	IY	NIY	DY, NIY	TPY, IY
Panel A: Rolling estimation window of 30 years							
OOS R^2	-0.64	0.42	1.57	-6.75	-11.22	-5.52	-2.66
GW t -stat	-0.07	0.05	0.18	-1.78	-1.64	-0.59	-0.77
GW p -value	0.94	0.96	0.86	0.08	0.10	0.55	0.44
CW t -stat	1.32	1.63	1.54	-1.33	-1.05	-0.00	-0.21
CW p -value	0.09	0.05	0.06	0.91	0.85	0.50	0.58
Panel B: Recursive estimation scheme							
OOS R^2	-6.06	1.03	4.48	-5.68	-21.33	-0.47	3.70
GW t -stat	-0.47	0.10	0.40	-0.91	-1.51	0.42	0.49
GW p -value	0.64	0.92	0.69	0.36	0.13	0.68	0.62
CW t -stat	1.47	1.79	1.88	0.06	-0.09	1.96	1.68
CW p -value	0.07	0.04	0.03	0.47	0.54	0.02	0.05

Table 3.4: Joint tests for predictability in stock returns or cash flow growth

This table presents results for joint tests of the null of either no stock return predictability or no cash flow growth predictability similar to those proposed by Cochrane (2008). The left-hand portion of the table shows parameter values for the null hypothesis, while the right-hand portion of the table presents simulated p -values. Panels A and B present results for the null of no stock return predictability ($\beta_r^N = 0$), where the slope parameters ϕ^N and β_{cf-d}^N are set to estimates from the VAR system described by Eqs. (3.9)–(3.12). The parameter β_{cf}^N is set to the implied value under the identity assuming no stock return predictability, i.e., $\beta_{cf}^N = \rho\phi^N + (1 - \rho)\beta_{cf,d}^N - 1$. Results in Panel A are based on OLS estimates for the parameters ϕ and β_{cf-d} , while results in Panel B are based on bias-corrected estimates for these parameters. Panels C and D present results for similar tests of the null of no predictability in cash flow growth. The slope parameters ϕ^N and β_{cf-d}^N under the null are set to estimates from the VAR system and the parameter β_r^N is set to the implied value under the identity assuming $\beta_{cf}^N = 0$. Simulated p -values correspond to the probability, based on 10,000 simulations, of obtaining estimates more extreme than the actual estimates in the data, denoted $P(\hat{\beta}_r^* \leq \hat{\beta}_r)$ and $P(\hat{\beta}_{cf}^* \leq \hat{\beta}_{cf})$ for return and cash flow growth forecasting coefficients, respectively. $P(\text{Both})$ reports the simulation-based probability that both estimates are more extreme than those observed in the data. See section 3.3.4 for details regarding the convention we use in assigning a simulated coefficient as more extreme. The sample period is 1932–2014.

	Null hypothesis				Simulated p -values		
	β_r^N	β_{cf}^N	ϕ^N	β_{cf-d}^N	$P(\hat{\beta}_r^* \leq \hat{\beta}_r)$	$P(\hat{\beta}_{cf}^* \leq \hat{\beta}_{cf})$	$P(\text{Both})$
Panel A: Stock returns; OLS estimates							
Dividends	0.00	-0.10	0.94	0.00	0.16	0.01	0.01
Total payout	0.00	-0.22	0.82	-0.12	0.04	0.00	0.00
Net payout	0.00	-0.32	0.69	0.22	0.01	0.04	0.00
Issuances	0.00	-0.39	0.61	0.77	0.12	0.57	0.09
Panel B: Stock returns; Bias-corrected estimates							
Dividends	0.00	-0.05	0.99	0.00	0.14	0.06	0.04
Total payout	0.00	-0.18	0.86	-0.11	0.03	0.00	0.00
Net payout	0.00	-0.29	0.73	0.24	0.01	0.12	0.00
Issuances	0.00	-0.35	0.65	0.81	0.12	0.42	0.07
Panel C: Cash flow growth; OLS estimates							
Dividends	0.10	0.00	0.94	0.00	0.64	0.41	0.36
Total payout	0.22	0.00	0.82	-0.12	0.23	0.35	0.15
Net payout	0.32	0.00	0.69	0.22	0.00	0.02	0.00
Issuances	0.39	0.00	0.61	0.77	0.00	0.00	0.00
Panel D: Cash flow growth; Bias-corrected estimates							
Dividends	0.05	0.00	0.99	0.00	0.34	0.38	0.21
Total payout	0.18	0.00	0.86	-0.11	0.57	0.34	0.12
Net payout	0.29	0.00	0.73	0.24	0.00	0.02	0.00
Issuances	0.35	0.00	0.65	0.81	0.00	0.00	0.00

Table 3.5: VAR calibration for long-horizon portfolio choice analysis

This table presents OLS estimates of the VAR(1) system that forms the basis of our calibration exercise:

$$z_{t+1} = \Phi_0 + \Phi_1 z_t + v_{t+1}.$$

The vector $z_t = (r_{1,t}, r_t^e, YLD_t)'$, where $r_{1,t}$ is the real log return on a 90-day Treasury bill, r_t^e is the log market excess return, where the value-weighted CRSP portfolio serves as the market proxy and returns are in excess of the 90-day Treasury bill return, and YLD_t is the payout yield of interest. Results are presented for the dividend yield (Panel A), the total payout yield (Panel B), and the net payout yield (Panel C). We omit estimation results for $\hat{\Phi}_0$, the intercepts in the VAR model, for brevity. The first six columns of the table present estimated slope coefficients $\hat{\Phi}_1$, along with associated t -statistics in parentheses. The t -statistics are computed using Newey-West standard errors with two lags. The next three columns present calibration results associated with $\hat{\Sigma}_v$, the covariance matrix of shocks to the system. Specifically, the table presents estimates of the standard deviation for each shock (on the diagonal) and estimates of the pairwise correlation between shocks, expressed as a percentage (off-diagonal). The right-most columns of the table present the R^2 -value for each equation in the system and the F -stat and p -value associated with the null hypothesis that all slope coefficients are equal to zero for the corresponding equation in the system. The sample period is 1932–2014.

VAR estimated coefficients and t -statistics					Reg. error s.d./corr. %			Model statistics				
$r_{1,t}$					r_t^e	YLD_t	$r_{1,t}$	r_t^e	YLD_t	R^2 %	F -stat	p -value
Panel A: $YLD_t = DY_t$												
$r_{1,t}$	0.52	(5.27)	-0.01	(-0.63)	-0.26	(-1.02)	2.90	7.13	-11.93	34.79	13.87	0.00
r_t^e	0.30	(0.80)	-0.10	(-1.11)	3.71	(2.76)	7.13	17.39	-59.97	9.87	2.85	0.04
YLD_t	-0.06	(-2.18)	-0.00	(-0.70)	0.85	(13.98)	-11.93	-59.97	0.59	84.29	139.49	0.00
Panel B: $YLD_t = TPY_t$												
$r_{1,t}$	0.53	(5.55)	-0.02	(-0.78)	-0.23	(-0.86)	2.91	6.15	-2.58	34.30	13.58	0.00
r_t^e	0.17	(0.46)	-0.05	(-0.63)	4.46	(2.88)	6.15	17.36	-59.98	10.15	2.94	0.04
YLD_t	-0.06	(-2.08)	0.00	(0.66)	0.74	(11.09)	-2.58	-59.98	0.75	63.52	45.27	0.00
Panel C: $YLD_t = NPY_t$												
$r_{1,t}$	0.51	(5.88)	-0.01	(-0.71)	-0.49	(-2.08)	2.84	11.38	10.14	37.42	15.55	0.00
r_t^e	0.26	(0.81)	-0.09	(-1.08)	4.13	(3.10)	11.38	17.26	-33.46	11.16	3.27	0.03
YLD_t	-0.10	(-3.71)	-0.01	(-0.65)	0.65	(8.95)	10.14	-33.46	0.89	59.94	38.91	0.00

Table 3.6: Welfare simulation results by investment horizon and risk aversion level

This table presents results for welfare simulations. We simulate one million time series paths for the vector $z_t = (r_{1,t}, r_t^e, DY_t, NIY_t)'$, where DY_t and NIY_t are the dividend yield and net issuance yield, respectively. The vector z_t follows a VAR(1) process with parameters calibrated using annual data for 1932–2014. The slope coefficients on lagged DY_t and NIY_t for the Treasury-bill and market excess return equations are set to β^{NPY} and $-\beta^{NPY}$, respectively, where β^{NPY} is the OLS slope coefficient for the lagged net payout yield in a predictive regressions for returns. This modification implies that predictable variation in risk premia and interest rates are driven by the net payout yield. For each simulated path, we track asset allocations and wealth outcomes for two types of investor: one who applies the optimal dynamic portfolio strategy within the Jurek and Viceira (2011) framework, and a second who incorrectly relies on $\tilde{z}_t = (r_{1,t}, r_t^e, DY_t)'$ as the vector of state variables, and makes dynamic allocations accordingly. The investment horizon is K years, with $K \in \{1, 5, 10, 20, 30\}$. Initial wealth is normalized to one. The simulations yield one million wealth outcomes W_{t+K} and $\tilde{W}_{t+K}$ for policies based on z_t and $\tilde{z}_t$, respectively, and the associated realized utilities, e.g., $\frac{1}{1-\gamma} W_{t+K}^{1-\gamma}$ (the case $\gamma = 1$ corresponds to log utility). We then compute a certainty equivalent wealth level associated with each policy as $W_{CE} = U^{-1}(\overline{EU})$, where $\overline{EU}$ is the simulation-based expected utility of final wealth. Panel A reports the annual growth rate r_{CE}^g implied by the certainty equivalent wealth associated with the optimal dynamic portfolio strategy, computed as $(1 + r_{CE}^g)^K = W_{CE}$. Panel B shows the analogous certainty equivalent growth rate $\tilde{r}_{CE}^g$ for the suboptimal policy based on the dividend yield. Panel C presents the welfare loss, expressed as a percentage reduction in the certainty equivalents of wealth, associated with basing dynamic allocations on the simulated dividend yield rather than the net payout yield. Panel D presents the maximal annual management fee ϕ that an agent would be willing to pay in order to switch from the dynamic strategy based on the dividend yield to the optimal dynamic strategy based on the net payout yield. This is computed as the solution to $(1 + r_{CE}^g)(1 - \phi) = 1 + \tilde{r}_{CE}^g$, where r_{CE}^g and $\tilde{r}_{CE}^g$ are the growth rates implied by the certainty equivalent wealth levels under the optimal policy and the alternative policy based on the dividend yield, respectively.

γ	$K = 1$	$K = 5$	$K = 10$	$K = 20$	$K = 30$
Panel A: Certainty equivalent growth rate r_{CE}^g for optimal policy (annual %)					
1	18.41	18.43	18.43	18.44	18.44
2	8.44	9.35	9.67	9.96	10.10
4	3.98	4.57	4.86	5.14	5.30
10	1.17	0.92	0.94	1.00	1.20
Panel B: Certainty equivalent growth rate $\tilde{r}_{CE}^g$ for dividend-based policy (annual %)					
1	13.44	13.47	13.47	13.48	13.49
2	6.26	6.83	6.92	6.82	6.71
4	2.96	3.29	3.33	3.09	3.14
10	0.78	0.41	0.34	0.20	0.60
Panel C: Welfare loss associated with suboptimal dividend yield strategy (%)					
1	-4.19	-19.28	-34.84	-57.44	-72.26
2	-2.02	-11.00	-22.45	-43.89	-60.96
4	-0.99	-5.98	-13.70	-32.54	-46.30
10	-0.39	-2.48	-5.77	-14.74	-16.31
Panel D: Maximum management fee ϕ to switch to optimal strategy (annual %)					
1	4.19	4.19	4.19	4.18	4.18
2	2.02	2.30	2.51	2.85	3.09
4	0.99	1.23	1.46	1.95	2.05
10	0.39	0.50	0.59	0.79	0.59

Bibliography

- Acharya, Viral V., and Lasse Heje Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375–410.
- Ackert, Lucy F., and Brian F. Smith, 1993, Stock price volatility, ordinary dividends, and other cash flows to shareholders, *Journal of Finance* 48, 1147–1160.
- Admati, Anat R., and Paul Pfleiderer, 2009, The “Wall Street Walk” and shareholder activism: Exit as a form of voice, *Review of Financial Studies* 22, 2645–2685.
- Allen, Franklin, and Roni Michaely, 2003, Payout policy, *Handbook of the Economics of Finance* 1, 337–429.
- Amihud, Yakov, 2002, Illiquidity and stock returns: cross-section and time-series effects, *Journal of Financial Markets* 5, 31–56.
- , and Clifford M. Hurvich, 2004, Predictive regressions: A reduced-bias estimation method, *Journal of Financial and Quantitative Analysis* 39, 813–841.
- , and Yi Wang, 2009, Multiple-predictor regressions: Hypothesis testing, *Review of Financial Studies* 22, 413–434.
- Amihud, Yakov, and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223–249.
- Andersen, Torben G., and Tim Bollerslev, 1997, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4, 115–158.

- Ang, Andrew, 2012, Predicting dividends in log-linear present value models, *Pacific-Basin Finance Journal* 20, 151–171.
- Asparouhova, Elena, Hendrik Bessembinder, and Ivalina Kalcheva, 2010, Liquidity biases in asset pricing tests, *Journal of Financial Economics* 96, 215–237.
- Bacidore, Jeffrey, Robert H. Battalio, and Robert H. Jennings, 2003, Order submission strategies, liquidity supply, and trading in pennies on the new york stock exchange, *Journal of Financial Markets* 6, 337–362.
- Back, Kerry, Tao Li, and Alexander Ljungqvist, 2015, Liquidity and governance, *Working paper*.
- Baillie, Richard T., Tim Bollerslev, and Hans Ole Mikkelsen, 1996, Fractionally integrated generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 74, 3–30.
- Baker, Malcolm, and Jeffrey Wurgler, 2000, The equity share in new issues and aggregate stock returns, *The Journal of Finance* 55, 2219–2257.
- Barberis, Nicholas, 2000, Investing for the long run when returns are predictable, *The Journal of Finance* 55, 225–264.
- Barclay, Michael J., Eugene Kandel, and Leslie M. Marx, 1998, The effects of transaction costs on stock prices and trading volume, *Journal of Financial Intermediation* 7, 130–150.
- Beber, Alessandro, Michael W. Brandt, and Kenneth A. Kavajecz, 2011, What does equity sector orderflow tell us about the economy?, *Review of Financial Studies* 24, 3688–3730.
- Bekaert, Geert, Campbell R. Harvey, and Christian Lundblad, 2007, Liquidity and expected returns: Lessons from emerging markets, *Review of Financial Studies* 20, 1783–1831.
- Ben-Rephael, Azi, Ohad Kadan, and Avi Wohl, 2015, The diminishing liquidity premium, *Journal of Financial and Quantitative Analysis*, *Forthcoming*.
- Bessembinder, Hendrik, 2003, Trade execution costs and market quality after decimalization, *Journal of Financial and Quantitative Analysis* 38, 747–777.

- Bharath, Sreedhar T., Sudarshan Jayaraman, and Venky Nagar, 2013, Exit as governance: An empirical analysis, *The Journal of Finance* 68, 2515–2547.
- Bhide, Amar, 1993, The hidden costs of stock market liquidity, *Journal of Financial Economics* 34, 31–51.
- Boehmer, Ekkehart, Gideon Saar, and Lei Yu, 2005, Lifting the veil: An analysis of pre-trade transparency at the NYSE, *The Journal of Finance* 60, 783–815.
- Boehmer, Ekkehart, and Juan (Julie) Wu, 2013, Short selling and the price discovery process, *Review of Financial Studies* 26, 287–322.
- Boudoukh, Jacob, Roni Michaely, Matthew Richardson, and Michael R. Roberts, 2007, On the importance of measuring payout yield: Implications for empirical asset pricing, *Journal of Finance* 62, 877–915.
- Brav, Alon, John R. Graham, Campbell R. Harvey, and Roni Michaely, 2005, Payout policy in the 21st century, *Journal of Financial Economics* 77, 483–527.
- Brennan, Michael, Sahn-Wook Huh, and Avanidhar Subrahmanyam, 2013, An analysis of the Amihud illiquidity premium, *Review of Asset Pricing Studies* pp. 133–176.
- Brennan, Michael J., and Avanidhar Subrahmanyam, 1996, Market microstructure and asset pricing: On the compensation for illiquidity in stock returns, *Journal of Financial Economics* 41, 441–464.
- Brockman, Paul, Dennis Y. Chung, and Christophe Pérignon, 2009, Commonality in liquidity: A global perspective, *Journal of Financial and Quantitative Analysis* 44, 851–882.
- Brunnermeier, Markus K., and Lasse Heje Pedersen, 2009, Market liquidity and funding liquidity, *Review of Financial studies* 22, 2201–2238.
- Butler, Alexander W., Jess Cornaggia, Gustavo Grullon, and James P. Weston, 2011, Corporate financing decisions, managerial market timing, and real investment, *The Journal of Financial Economics* 101, 666–683.

- Butler, Alexander W., Gustavo Grullon, and James P. Weston, 2005a, Can managers forecast aggregate market returns?, *The Journal of Finance* 60, 963–986.
- , 2005b, Stock market liquidity and the cost of issuing equity, *Journal of Financial and Quantitative Analysis* 40, 331–348.
- Campbell, John Y., and John Ammer, 1993, What moves the stock and bond markets? a variance decomposition for long-term asset returns, *The Journal of Finance* 48, 3–37.
- Campbell, John Y., Yeung Lewis Chan, and Luis M. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.
- Campbell, John Y., Martin Lettau, Burton G. Malkiel, and Yexiao Xu, 2001, Have individual stocks become more volatile? An empirical exploration of idiosyncratic risk, *Journal of Finance* 56.
- Campbell, John Y., and Robert J. Shiller, 1988a, The dividend-price ratio and expectations of future dividends and discount factors, *Review of Financial Studies* 1, 195–228.
- , 1988b, Stock prices, earnings, and expected dividends, *The Journal of Finance* 43, 661–676.
- Campbell, John Y., and Samuel B. Thompson, 2008, Predicting excess stock returns out of sample: Can anything beat the historical average?, *Review of Financial Studies* 21, 1509–1531.
- Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *The Quarterly Journal of Economics* 114, 433–495.
- Campbell, John Y., and Motohiro Yogo, 2006, Efficient tests of stock return predictability, *Journal of Financial Economics* 81, 27–60.
- Chakravarty, Sugato, Robert A. Wood, and Robert A. Van Ness, 2004, Decimals and liquidity: A study of the NYSE, *Journal of Financial Research* 27, 75–94.

- Chen, Long, 2009, On the reversal of return and dividend growth predictability: A tale of two periods, *Journal of Financial Economics* 92, 128–151.
- Chordia, Tarun, Richard Roll, and Avanidhar Subrahmanyam, 2000, Commonality in liquidity, *Journal of Financial Economics* 56, 3–28.
- , 2001, Market liquidity and trading activity, *Journal of Finance* 56, 501–530.
- Chordia, Tarun, Asani Sarkar, and Avanidhar Subrahmanyam, 2005, An empirical analysis of stock and bond market liquidity, *Review of Financial Studies* 18, 85–129.
- Christensen, Bent Jesper, and Morten Ørregaard Nielsen, 2006, Asymptotic normality of narrow-band least squares in the stationary fractional cointegration model and volatility forecasting, *Journal of Econometrics* 133, 343–371.
- Chung, Kee H., and Hao Zhang, 2014, A simple approximation of intraday spreads using daily data, *Journal of Financial Markets* 17, 94–120.
- Clark, Todd E., and Michael W. McCracken, 2011, Testing for unconditional predictive ability, in *Oxford Handbook on Economic Forecasting* . pp. 415–440 (Oxford University Press).
- Clark, Todd E., and Kenneth D. West, 2007, Approximately normal tests for equal predictive accuracy in nested models, *Journal of Econometrics* 138, 291–311.
- Cochrane, John H., 1991, Volatility tests and efficient markets, *Journal of Monetary Economics* 27, 463–485.
- , 1992, Explaining the variance of price-dividend ratios, *Review of Financial Studies* 5, 243–280.
- , 1994, Permanent and transitory components of GNP and stock prices, *The Quarterly Journal of Economics* 109, 241–265.
- Cochrane, John H, 2008, The dog that did not bark: A defense of return predictability, *Review of Financial Studies* 21, 1533–1575.

- Cochrane, John H., 2011, Presidential address: Discount rates, *The Journal of Finance* 66, 1047–1108.
- Constantinides, George M., 1986, Capital market equilibrium with transaction costs, *The Journal of Political Economy* 94, 842–862.
- Corwin, Shane A., and Paul Schultz, 2012, A simple way to estimate bid-ask spreads from daily high and low prices, *The Journal of Finance* 67, 719–760.
- Diebold, Francis X., 2015, Comparing predictive accuracy, twenty years later: A personal perspective on the use and abuse of diebold–mariano tests, *Journal of Business & Economic Statistics* 33, 1–24.
- Ding, Zhuanxin, Clive W.J. Granger, and Robert F. Engle, 1993, A long memory property of stock market returns and a new model, *Journal of Empirical Finance* 1, 83–106.
- Edmans, Alex, 2009, Blockholder trading, market efficiency, and managerial myopia, *The Journal of Finance* 64, 2481–2513.
- , Vivian W. Fang, and Emanuel Zur, 2013, The effect of liquidity on governance, *Review of Financial Studies* 26, 1443–1482.
- Edmans, Alex, and Gustavo Manso, 2011, Governance through trading and intervention: A theory of multiple blockholders, *Review of Financial Studies* 24, 2395–2428.
- Engsted, Tom, and Thomas Quistgaard Pedersen, 2010, The dividend-price ratio does predict dividend growth: International evidence, *Journal of Empirical Finance* 17, 585–605.
- Fama, Eugene F., and Kenneth R. French, 1988, Dividend yields and expected stock returns, *Journal of Financial Economics* 22, 3–25.
- , 1989, Business conditions and expected returns on stocks and bonds, *Journal of Financial Economics* 25, 23–49.
- , 1992, The cross-section of expected stock returns, *the Journal of Finance* 47, 427–465.

- , 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- , 2001, Disappearing dividends: Changing firm characteristics or lower propensity to pay?, *Journal of Financial Economics* 60, 3–43.
- , 2015, A five-factor asset pricing model, *Journal of Financial Economics* 116, 1–22.
- Fang, Vivian W., Thomas H. Noe, and Sheri Tice, 2009, Stock market liquidity and firm value, *Journal of Financial Economics* 94, 150–169.
- Fang, Vivian W., Xuan Tian, and Sheri Tice, 2014, Does stock liquidity enhance or impede firm innovation?, *The Journal of Finance* 69, 2085–2125.
- Farre-Mensa, Joan, Roni Michaely, and Martin Schmalz, 2014, Payout policy, in Andrew W. Lo, and Robert C. Merton, ed.: *Annual Review of Financial Economics, Volume 6* (Annual Reviews: Palo Alto, CA).
- Fong, Kingsley, Craig Holden, and Charles Trzcinka, 2014, What are the best liquidity proxies for global research, *Working paper*.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.
- Gerken, William C., 2014, Blockholder ownership and corporate control: The role of liquidity, *The Quarterly Journal of Finance* 4.
- Ghysels, Eric, Pedro Santa-Clara, and Rossen Valkanov, 2005, There is a risk-return trade-off after all, *Journal of Financial Economics* 76, 509–548.
- Giacomini, Raffaella, and Halbert White, 2006, Tests of conditional predictive ability, *Econometrica* 74, 1545–1578.
- Giglio, Stefano, Bryan T. Kelly, and Seth Pruitt, 2015, Systemic risk and the macroeconomy: An empirical evaluation, *Journal of Financial Economics, Forthcoming*.

- Goetzmann, William N., and Philippe Jorion, 1993, Testing the predictive power of dividend yields, *The Journal of Finance* 48, 663–679.
- Goldstein, Michael A., and Kenneth A. Kavajecz, 2000, Eighths, sixteenths, and market depth: Changes in tick size and liquidity provision on the NYSE, *Journal of Financial Economics* 56, 125–149.
- Goyal, Amit, and Ivo Welch, 2008, A comprehensive look at the empirical performance of the equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka, 2009, Do liquidity measures measure liquidity?, *Journal of Financial Economics* 92, 153 – 181.
- Granger, Clive W.J., and Roselyne Joyeux, 1980, An introduction to long-memory time series models and fractional differencing, *Journal of Time Series Analysis* 1, 15–29.
- Grullon, Gustavo, and Roni Michaely, 2002, Dividends, share repurchases, and the substitution hypothesis, *The Journal of Finance* 57, 1649–1684.
- Grullon, Gustavo, Bradley Paye, Shane Underwood, and James P. Weston, 2011, Has the propensity to pay out declined?, *Journal of Financial and Quantitative Analysis* 46, 1–24.
- Hasbrouck, Joel, 2009, Trading costs and returns for US equities: Estimating effective costs from daily data, *The Journal of Finance* 64, 1445–1477.
- , and Duane J. Seppi, 2001, Common factors in prices, order flows, and liquidity, *Journal of Financial Economics* 59, 383–411.
- Holden, Craig W., 2009, New low-frequency spread measures, *Journal of Financial Markets* 12, 778–813.
- Holmström, Bengt, and Jean Tirole, 1993, Market liquidity and performance monitoring, *Journal of Political Economy* 101, 678–709.
- Hosking, Jonathan R.M., 1981, Fractional differencing, *Biometrika* 68, 165–176.

- Hou, Kewei, and Tobias J. Moskowitz, 2005, Market frictions, price delay, and the cross-section of expected returns, *Review of Financial Studies* 18, 981–1020.
- Hou, Kewei, Chen Xue, and Lu Zhang, 2014, Digesting anomalies: An investment approach, *Review of Financial Studies* 28, 650–705.
- Huberman, Gur, and Dominika Halka, 2001, Systematic liquidity, *Journal of Financial Research* 24, 161–178.
- Jansson, Michael, and Marcelo J. Moreira, 2006, Optimal inference in regression models with nearly integrated regressors, *Econometrica* 74, 681–714.
- Jegadeesh, Narasimhan, and Sheridan Titman, 1993, Returns to buying winners and selling losers: Implications for stock market efficiency, *The Journal of Finance* 48, 65–91.
- Johannes, Michael, Arthur Kortweg, and Nicholoas Polson, 2014, Sequential learning, predictability, and optimal portfolio returns, *The Journal of Finance* 69, 611–644.
- Jones, Charles M., 2002, A century of stock market liquidity and trading costs, *Working paper*.
- , and Marc L. Lipson, 2001, Sixteenths: Direct evidence on institutional execution costs, *Journal of Financial Economics* 59, 253–278.
- Jurek, Jakub W, and Luis M Viceira, 2011, Optimal value and growth tilts in long-horizon portfolios, *Review of Finance* 15, 29–74.
- Kahn, Charles, and Andrew Winton, 1998, Ownership structure, speculation, and shareholder intervention, *The Journal of Finance* 53, 99–129.
- Karolyi, G. Andrew, Kuan-Hui Lee, and Mathijs A. Van Dijk, 2012, Understanding commonality in liquidity around the world, *Journal of Financial Economics* 105, 82–112.
- Kayacetin, Nuri Volkan, and Aditya Kaul, 2009, Forecasting economic fundamentals and stock returns with equity market order flows, *Working paper*.

- Keim, Donald B., and Robert F. Stambaugh, 1986, Predicting returns in the stock and bond markets, *Journal of Financial Economics* 17, 357–390.
- Killian, Lutz, 1999, Exchange rates and monetary fundamentals: What do we learn from long-horizon regressions?, *Journal of Applied Econometrics* 14, 491–510.
- Kim, Chang-Jin, and Cheolbeom Park, 2013, Disappearing dividends: Implications for the dividend-price ratio and return predictability, *Journal of Money, Credit and Banking* 45, 933–952.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1335.
- , and Jean-Luc Vila, 1991, Noise trading and takeovers, *The RAND Journal of Economics* 22, 54–71.
- Larrain, Borja, and Motohiro Yogo, 2008, Does firm value move too much to be justified by subsequent changes in cash flow?, *Journal of Financial Economics* 87, 200–226.
- Lesmond, David A., Joseph P. Ogden, and Charles A. Trzcinka, 1999, A new estimate of transaction costs, *Review of Financial Studies* 12, 1113–1141.
- Lettau, Martin, and Stijn Van Nieuwerburgh, 2008, Reconciling the return predictability evidence, *Review of Financial Studies* 21, 1607–1652.
- Lynch, Anthony W., and Sinan Tan, 2011, Explaining the magnitude of liquidity premia: The roles of return predictability, wealth shocks, and state-dependent transaction costs, *The Journal of Finance* 66, 1329–1368.
- Mancini, Lorian, Angelo Ranaldo, and Jan Wrampelmeyer, 2013, Liquidity in the foreign exchange market: Measurement, commonality, and risk premiums, *The Journal of Finance* 68, 1805–1841.
- Mark, Nelson C., 1995, Exchange rates and fundamentals: Evidence on long-horizon predictability, *American Economic Review* 85, 201–218.

- Maug, Ernst, 1998, Large shareholders as monitors: Is there a trade-off between liquidity and control?, *The Journal of Finance* 53, 65–98.
- Maynard, Alex, and Peter C.B. Phillips, 2001, Rethinking an old empirical puzzle: Econometric evidence on the forward discount anomaly, *Journal of Applied Econometrics* 16, 671–708.
- Maynard, Alex, Aaron Smallwood, and Mark E. Wohar, 2013, Long memory regressors and predictive testing: A two-stage rebalancing approach, *Econometric Reviews* 32, 318–360.
- Merton, Robert C., 1969, Lifetime portfolio choice under uncertainty: The continuous time case, *Review of Economics and Statistics* 51, 247–257.
- Merton, Robert C, 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* pp. 867–887.
- Næs, Randi, Johannes A. Skjeltorp, and Bernt Arne Ødegaard, 2011, Stock market liquidity and the business cycle, *The Journal of Finance* 66, 139–176.
- Nelson, Charles R., and Myung J. Kim, 1993, Predictable stock returns: The role of small sample bias, *The Journal of Finance* 48, 641–661.
- Norli, Øyvind, Charlotte Ostergaard, and Ibolya Schindele, 2015, Liquidity and shareholder activism, *Review of Financial Studies* 28, 486–520.
- Pástor, Ľuboš, and Robert F. Stambaugh, 2003, Liquidity risk and expected stock returns, *Journal of Political Economy* 111, 642–685.
- Roberts, Michael R., and Toni M. Whited, 2013, Endogeneity in empirical corporate finance, *Handbook of the Economics of Finance* 2, 493–572.
- Robertson, Donald, and Stephen Wright, 2006, Dividends, total cash flow to shareholders, and predictive return regressions, *Review of Economics and Statistics* 88, 91–99.

- Robinson, Peter M., 1994, Semiparametric analysis of long-memory time series, *The Annals of Statistics* 22, 515–539.
- , and Domenico Marinucci, 2003, Semiparametric frequency domain analysis of fractional cointegration, in *Time series with long memory* . pp. 334–373 (Oxford University Press).
- Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *The Journal of Finance* 39, 1127–1139.
- Rozeff, Michael, 1984, Dividend yields are equity risk premiums, *Journal of Portfolio Management* 11, 68–75.
- Shiller, Robert J., 1981, Do stock prices move too much to be justified by subsequent changes in dividends?, *American Economic Review* 71, 421–436.
- , Stanley Fischer, and Benjamin M. Friedman, 1984, Stock prices and social dynamics, *Brookings Papers on Economic Activity* 1984, 457–510.
- Shimotsu, Katsumi, and Peter C.B. Phillips, 2005, Exact local whittle estimation of fractional integration, *The Annals of Statistics* 33, 1890–1933.
- Skinner, Douglas J., 2008, The evolving relation between earnings, dividends, and stock repurchases, *Journal of Financial Economics* 87, 582–609.
- Stambaugh, Robert F., 1999, Predictive regressions, *Journal of Financial Economics* 54, 375–421.
- Stephens, Clifford P., and Michael S. Weisbach, 1998, Actual share reacquisitions in open-market repurchase programs, *The Journal of Finance* 53, 313–333.
- Stock, James H., and Mark W. Watson, 2003, Forecasting output and inflation: The role of asset prices, *Journal of Economic Literature* 41, 788–829.
- Subrahmanyam, Avaniidhar, 2007, Liquidity, return and order-flow linkages between REITs and the stock market, *Real Estate Economics* 35, 383–408.

- van Binsbergen, Jules H., and Ralph S. J. Koijen, 2010, Predictive regressions: A present value approach, *Journal of Finance* 65, 1439 – 1471.
- Van Ness, Bonnie F., Robert A. Van Ness, and Stephen W. Pruitt, 2000, The impact of the reduction in tick increments in major U.S. markets on spreads, depth, and volatility, *Review of Quantitative Finance and Accounting* 15, 153–167.
- Vayanos, Dimitri, 1998, Transaction costs and asset prices: A dynamic equilibrium model, *Review of Financial Studies* 11, 1–58.
- Wachter, Jessica A., and Missaka Warusawitharana, 2009, Predictable returns and asset allocation: Should a skeptical investor time the market?, *Journal of Econometrics* 148, 162–178.
- , 2013, What is the chance that the equity premium varies over time? Evidence from regressions on the dividend-price ratio, *Working Paper*.
- Welch, Ivo, and Amit Goyal, 2008, A comprehensive look at the empirical performance of the equity premium prediction, *Review of Financial Studies* 21, 1455–1508.
- West, Kenneth D., 1996, Asymptotic inference about predictive ability, *Econometrica* 64, 1067–1084.

Appendix A

A Timeline of Key Dates for Tick-Size Reductions in U.S. Equity Markets

March 13, 1997 Congressman Michael Oxley introduces a bill to the U.S. House of Representatives directing the S.E.C. to require decimal pricing for all equity securities.

March 13, 1997 AMEX's board of directors approves switch in tick size from eighths to sixteenths.

March 25, 1997 NASDAQ's board of directors approves switch in tick size from eighths to sixteenths.

May 7, 1997 AMEX begins pricing securities in increments of sixteenths.

June 2, 1997 NASDAQ begins pricing securities in increments of sixteenths.

June 5, 1997 NYSE announces plans to switch to decimal pricing as soon as the essential systems are in place. As an interim step, the NYSE's board approves a switch from an eighths to an sixteenths regime, effective by the end of the month.

June 24, 1997 The NYSE begins pricing securities in increments of sixteenths.

December 24, 1997 The NYSE announces that it will delay conversion to decimal trading until 2000 in order to first deal with year 2000 (Y2K) issues.

January 28, 2000 The S.E.C. orders all major stock markets to begin phasing in decimal pricing by the end of the year.

June 8, 2000 The S.E.C. orders all major stock markets to fully implement decimal pricing by April 9, 2001. NASDAQ should begin phasing in decimal pricing no later than March 12, 2001. The other major stock markets should begin phasing in decimal pricing by September 5, 2000.

July 5, 2000 The NYSE announces the seven issues that comprise it's initial decimal pricing pilot group.

July 7, 2000 The AMEX announces the six issues that comprise it's initial decimal pricing pilot group.

August 16, 2000 The NYSE announces the 57 issues that comprise it's phase two decimal pricing pilot group.

August 28, 2000 Both NYSE's and AMEX's phase one pilot groups begin decimal trading.

August 30, 2000 The AMEX announces the 46 issues that comprise it's phase two decimal pricing pilot group.

September 25, 2000 Both NYSE's and AMEX's phase two pilot groups begin decimal trading.

November 6, 2000 The NYSE announces the 94 issues that comprise it's phase three decimal pricing pilot group.

November 9, 2000 The NYSE announces plans to fully convert to decimal pricing by the end of January, 2001.

December 4, 2000 The NYSE's phase three pilot group begins decimal trading.

December 27, 2000 The AMEX joins the NYSE in planning to fully implement decimal pricing on January 29, 2001.

January 29, 2001 Both the NYSE and AMEX fully convert to decimal pricing.

February 7, 2001 The NASDAQ announces the fifteen issues that comprise it's phase one decimal pricing pilot group.

March 12, 2001 The NASDAQ's phase one pilot group begins decimal trading.

March 24, 2001 The NASDAQ announces the stock issues that comprise it's phase two decimal pricing pilot group.

March 26, 2001 The NASDAQ's phase two pilot group begins decimal trading.

April 9, 2001 The NASDAQ fully converts to decimal pricing.

Appendix B

Construction of Aggregate Illiquidity and Macroeconomic Variables

B.1 Aggregate liquidity and price delay measures

Low-frequency market liquidity and price delay measures are constructed using daily data from the University of Chicago’s Center for Research in Security Prices (CRSP) over the period 1926-2013. We restrict the sample to common shares (CRSP share code of 10, 11, or 12) that trade on the NYSE. We construct the various measures for each stock i over each time interval t (monthly or quarterly intervals) and then compute the equally-weighted mean across securities for each time interval.⁶¹

Roll measure: The Roll (1984) liquidity measure (*ROLL*) is an estimate of the relative effective bid-ask spread based on the serial covariance of successive price movements. We compute the measure based on Eq. (2.8) in Section 2.3.2 of the study. The *ROLL* proxy is undefined whenever the first-order serial covariance of returns is positive. Researchers

⁶¹In addition to the proxies listed below, we also computed a bid-ask spread measure as the difference between the closing ask and bid quotes reported by CRSP, scaled by the average of the ask and bid quotes. Chung and Zhang (2014) find that this CRSP-based spread is highly correlated with the spread based on high-frequency (TAQ) data across stocks, and our series co-moves closely with other spread proxies. This low frequency proxy is not continuously available, due to the fact that CRSP ask and bid quotes for NYSE stocks are unavailable for certain portions of our sample period. Consequently, our empirical analysis focuses on other spread proxies.

have adopted one of three conventions to address this issue: (1) set the Roll proxy to zero, (2) treat the observation as missing, or (3) reverse the sign of the sample autocovariance, and multiply the resulting spread proxy by negative one, which yields a negative spread estimate. The latter approach suggests a negative trading cost, which seems undesirable. On the other hand, the first two approaches impart an upward bias to the cross-sectional average of the spread. Results in the main analysis set the Roll proxy to zero when the first-order sample covariance is positive. Our main conclusions regarding Granger causality continue to obtain under the alternative approaches.

Corwin and Schultz measure: The Corwin and Schultz (2012) liquidity measure (*CS*) uses high and low daily price quotes to estimate bid-ask spreads. Daily high prices are almost always buyer-initiated orders while daily low prices are almost always seller-initiated orders, which suggests that the ratio of high to low prices reflects both the stock’s variance and its spread. The variance component grows proportionately with the time interval but the spread component does not, allowing one to derive a spread estimator of high-low ratios over one and two day windows. See Corwin and Schultz (2012) for details on variable construction. We closely follow Corwin and Schultz’s implementation procedures including the suggestion to adjust the measure to 0 if the spread estimate is negative.

Fong, Holden, and Trzcinka measure: We compute the Fong, Holden, and Trzcinka (2014) liquidity measure (*FHT*) at the firm level based on Eq. (2.4) in Section 2.3.2. The *FHT* measure represents a simplification of the Lesmond, Ogden, and Trzcinka (1999) (LOT) measure. Fong, Holden, and Trzcinka (2014) find that the *FHT* measure outperforms both the LOT measure and a variation of the LOT measure, the LOT Y-split measure, proposed in Goyenko, Holden, and Trzcinka (2009).

Zeros measure: This liquidity measure (*ZEROS*), proposed by Lesmond, Ogden, and Trzcinka (1999), is the proportion of zero-return trading days in time interval t . *ZEROS* serves as a liquidity proxy because zero-return days often occur when transaction costs keep marginal investors from trading. The frequency of zero returns are increasing in transaction costs.

Effective tick measure: Holden (2009) and Goyenko, Holden, and Trzcinka (2009) propose the effective tick measure (*TICK*), a relative effective spread measure based on observable price clustering. *TICK* is constructed under the assumption that price clustering is completely determined by spread size. This implies that the simple frequency in which closing prices occur in particular price clusters can be used to estimate spread probabilities. Using these estimated probabilities, firm i 's effective tick measure is a weighted average of its estimated spreads over period t . See Holden (2009) and Craig Holden's website⁶² for further details on variable construction.

We construct *TICK* using daily prices from positive-volume days only. An alternative measure, *TICK2*, is computed in the same manner as *TICK* except that *TICK2* uses daily prices from all trade days. We rely on *TICK* in our empirical analysis, but our main conclusions continue to hold if we alternatively use *TICK2*. We construct the effective tick measure using a 1/8 fractional price grid from 1926 up until 1997 when the NYSE implemented a 1/16 tick size. We use a 1/16 fractional price grid from 1997-2000. In 2001, the NYSE implemented decimal pricing and we use a decimal price grid from 2001 until the end of our sample period in 2013.

Amihud illiquidity measure: Amihud (2002) proposes a price impact measure (*AMI*) that captures the sensitivity of stock price movements to trading volume level. *AMI* is computed using daily data and aggregated over time interval t and is constructed as follows

$$AMI_{i,t} = \frac{1}{D_{i,t}} \times \sum_{d=1}^{D_{i,t}} \frac{|R_{i,d,t}|}{DVOL_{i,d,t}},$$

where $D_{i,t}$ is the number of trading days for stock i during interval t , $|R_{i,d,t}|$ is the absolute value of return for stock i on day d during time interval i , and $DVOL_{i,d,t}$ is the respective dollar trading volume. Brennan, Huh, and Subrahmanyam (2013) construct a version of the Amihud measure based on the ratio of absolute daily returns to daily turnover, calculated as daily share volume divided by total shares outstanding (*AMITO*).

⁶²<http://www.kelley.iu.edu/cholden/examples.pdf>

To correct for inflation effects on dollar trading volume, we compute *AMI* and *AMIC* in real terms. Specifically, each month, we multiply dollar volume by a CPI adjustment factor where the base CPI period is January of 1926, and U.S. CPI data is collected from the U.S. Bureau of Labor and Statistics. Finally, we multiply *AMI* and *AMIC* by 10^6 , and require that daily dollar trading volume in *AMI* and *AMIC* exceeds \$1000 and daily volume exceeds 100 shares for *AMITO*.

Pastor and Stambaugh measure: Pástor and Stambaugh (2003) propose a price impact measure (*PS*) based on the following regression:

$$RET_{i,d+1,t}^e = \theta_{i,t} + \phi_{i,t}RET_{i,d,t} + \gamma_{i,t}sign(RET_{i,d,t}^e)DVOL_{i,d,t} + \epsilon_{i,d+1,t},$$

where $RET_{i,d,t}$ is the return for stock i on day d in time interval t , $RET_{i,d,t}^e$ is the return in excess of the CRSP value-weighted market index, and $DVOL_{i,d,t}$ is dollar volume. The $\gamma_{i,t}$ estimate is the liquidity proxy. The intuition is as follows: volume signed by the stock return in excess of the market return is a proxy for order flow and should be followed by a return reversal in the future if the stock is less than perfectly liquid. The greater the reversal, the lower the liquidity, and thus the more negative the $\gamma_{i,t}$ estimate.

Following Pástor and Stambaugh (2003), we adjust for stock market growth by multiplying the γ_t estimate by m_t/m_0 where m_t is the CRSP market index value at time t and m_0 is the market index value at August, 1962 (the date used by Pastor and Stambaugh). We exclude observations if less than 15 days of price data in a month are available, if the stock price is below \$5 or above \$1000, or if dollar volume is less than or equal to \$1000.

Hou and Moskowitz measure: Hou and Moskowitz (2005) propose an annual price delay measure that captures the speed of adjustment of prices to market news. We create an analogous measure at quarterly and monthly intervals. Specifically, we regress individual stock returns on contemporaneous and 5 days of lagged market returns over each quarterly

or monthly time interval:

$$RET_{i,d,t} = \alpha_{i,t} + \beta_{i,t}RET_{m,d,t} + \sum_{n=1}^5 \delta_{i,t}^{(-n)} RET_{m,d-n,t} + \epsilon_{i,d,t},$$

where $RET_{i,d,t}$ is the return for stock i on day d during time interval t , $RET_{m,d,t}$ is the corresponding return on the NYSE value-weighted index, and n represents the number of lags. The above model represents a modification of Hou and Moskowitz's approach as we compute quarterly or monthly price delay measures using daily, rather than weekly, returns. Additionally, we include 5 rather than 4 lags, though our main conclusions continue to hold if we only include 4 lags. This modified framework closely follows the approach taken in Boehmer and Wu (2013).

The regression estimates are then used to construct the price delay measure. Price delay is the fraction of variation of individual stock returns explained by lagged market returns:

$$HM = 1 - \frac{R^2_{(\text{restricted model})}}{R^2_{(\text{unrestricted model})}},$$

where $R^2_{(\text{unrestricted model})}$ is the R^2 from the full regression model and $R^2_{(\text{restricted model})}$ is the R^2 from the regression model after restricting all of the lagged $\delta_{i,t}$ coefficients to equal 0. If stock i promptly responds to market news then $\beta_{i,t}$ should be different from zero but the lagged $\delta_{i,t}$ coefficients should all be insignificant. However, if stock i responds to news with a delay then some or all of the lagged $\delta_{i,t}$ coefficients will be different from zero.

B.2 Macroeconomic and financial variables

We collect macroeconomic variables at both monthly and quarterly frequencies.

Output: We collect two output measures. The first, real gross domestic product (GDP), is available at a quarterly frequency from 1947-2013. The second output measure, industrial production growth (IP), is available at a quarterly frequency from 1926-2013. Both measures are seasonally-adjusted.

Unemployment: We collect the seasonally-adjusted overall unemployment rate (UE) at both quarterly and monthly frequencies from 1948-2013.

Investment: Our investment measure (INV) is real, seasonally-adjusted gross private domestic investment, available at a quarterly frequency from 1947-2013.

Consumption: The consumption measure ($CONS$) is real, seasonally-adjusted personal consumption expenditures, available at a quarterly frequency from 1947-2013.

We collect all macroeconomic variables, except for the unemployment rate, from the St. Louis Federal Reserve (FRED2) website. The unemployment rate is collected at the U.S. Bureau of Labor and Statistics. Our empirical analysis also incorporates a number of financial variables. Each variable is constructed at monthly and quarterly frequencies from 1926-2013.

Term spread: $TERM$ is the difference between long-term US government bond yields and the 3-month T-bill yield. Data are sourced from Welch and Goyal (2008), available at Amit Goyal's website.⁶³

Credit spread: $CRED$ is the yield on BAA-rated corporate bonds less the yield on long-term US bonds. Data are sourced from Welch and Goyal (2008).

Default spread: DEF is defined as the yield on BAA-rated corporate bonds minus the yield on AAA-rated corporate bonds. Data are sourced from Welch and Goyal (2008).

Commercial paper spread: $CPSP$ is the 3-month commercial paper yield less the 3-month T-bill yield. Both series are obtained from the Federal Reserve Statistical Release H.15 beginning in 1971. Prior to 1971, the commercial paper rate is based on the New York commercial paper rate from the NBER macro-history database.

Excess market return: Log market return (ret) is the log of one plus NYSE value-weighted return with distributions minus the log of one plus the 3-month T-bill rate. Stock return

⁶³<http://www.hec.unil.ch/agoyal/>

data are sourced from CRSP and the T-bill rates are sourced from Welch and Goyal (2008).

Volatility: Aggregate volatility (vol) is the cross-sectional average of the natural logarithm of the sample standard deviation of returns, where firm-level standard deviations are computed using daily returns over the corresponding month or quarter. This measure is constructed using U.S. common stocks listed on the NYSE.

Appendix C

Construction of Payout Variables

The dividend yield (DY) is imputed from value-weighted CRSP market returns for the AMEX, NASDAQ, and NYSE universe with and without dividends (VWRETD and VWRETX, respectively) as

$$DY = (1 + VWRETD)/(1 + VWRETX) - 1.$$

Aggregate repurchases (RP) are computed using Compustat data for the period 1971–2012. Repurchases are calculated as the dollar value of the purchase of common and preferred stock (PRSTKC on Compustat) adjusted for any reduction in outstanding preferred stock (PSTKRV).⁶⁴ The repurchase yield (RPY) is $RPY = RP/MV$, where MV is year-end CRSP market capitalization. This yield is set to zero from 1926–1970, since Compustat data for repurchases are unavailable prior to 1971. Although repurchases prior to 1971 are unlikely to be literally zero, there is minimal repurchasing activity prior to the early 1980s (Grullon and Michaely (2002)). The total payout yield (TPY) is computed as $TPY = DY + RPY$.

We use Compustat to collect issuance data beginning in 1971, since these data are unavailable in prior years. Issuances are the dollar value of the sale of common and preferred stock (SSTK) adjusted for any increase in outstanding preferred stock (PSTKRV). The equity is-

⁶⁴Boudoukh, Michaely, Richardson, and Roberts (2007) construct an additional, alternative measure based on the change in treasury stock. The treasury stock measure attempts to exclude repurchases potentially earmarked for compensation or payment-in-kind. We rely on the more inclusive measure; however, changing to the treasury stock measure does not alter our main empirical findings.

suance yield (IY) is $IY = I/MV$ and the net equity issuance yield (NIY) is $NIY = IY - RPY$. Prior to 1971, we compute the net issuance yield using monthly CRSP data on the change in shares outstanding. Specifically, net issuance for the i -th firm in month t is computed as $(shrout_t \times cfacshr_t - shrout_{t-1} \times cfacshr_{t-1}) \times (prc_t/cfacpr_t + prc_{t-1}/cfacpr_{t-1})/2$. This figure is aggregated across months and firms to form annual net issuance and the net issuance yield. This convention follows previous studies including Stephens and Weisbach (1998), Boudoukh, Michaely, Richardson, and Roberts (2007) and Grullon, Paye, Underwood, and Weston (2011). Finally, we compute the net payout yield (NPY) as $NPY = DY - NIY$. The correlation of our net payout yield data with the net payout yield data used in Boudoukh, Michaely, Richardson, and Roberts (2007) over the common sample period of 1926–2003 is in excess of 0.98.⁶⁵

To compute payout growth rates, we first compute an annual payout level for each yield by multiplying the yield by the CRSP value-weighted index price level excluding distributions.⁶⁶ We then calculate growth as the percentage change in year-over-year levels, and log growth as the log difference in payout level. Growth rates are computed for 1927–2014, except for net payout growth, for which the sample starts in 1931 due to negative net payouts in 1929 and 1930. We adjust growth rates to real terms using CPI data from the U.S. Department of Labor. Specifically, we multiply the one-period ahead payout level by a CPI factor, where the CPI factor is lagged end-of-the-year CPI over current end-of-the-year CPI.

⁶⁵We thank Michael Roberts for providing these data on his website.

⁶⁶We do not use aggregate market capitalization from CRSP because our sample contains a break in the number of firms covered in 1971. The break in the number of sample firms is due to the transition from CRSP data to Compustat data in 1971. Since aggregate market capitalization is a function of the number of sample firms, it too suffers from a level break in 1971. The CRSP value-weighted index price level does not suffer from a level break, since it is invariant to the number of sample firms.

Appendix D

Out-of-Sample Forecasting Tests

The null hypothesis in the Clark and West (2007) test is that $\sigma^2 = \sigma_0^2$, where σ^2 and σ_0^2 refer to population MSPEs for the predictive regression and benchmark model, respectively. The Clark and West (2007) test statistic takes the form:

$$\text{Clark-West} = \hat{\sigma}_0^2 - \hat{\sigma}^2 + \underbrace{P^{-1} \sum (R_{t+1}^e - \bar{R}_{t+1})^2}_{\text{Adjustment}}.$$

The final term in (D.0.1) captures an adjustment for additional noise associated with the larger model's forecast under the null hypothesis. Clark and West (2007) show that standard normal critical values lead to tests with actual sizes close to nominal size for reasonably large samples. The test is one-sided, since, under the alternative, $\sigma^2 < \sigma_0^2$.

Giacomini and White (2006) consider the problem of comparing two alternative “forecasting methods.” A forecasting method is a broader concept that encompasses not only a model specification, but also the explicit procedure used to obtain forecasts. The null hypothesis for the Giacomini and White (2006) test is that the two models possess equal (conditional) forecasting ability. Our application involves unconditional forecast comparisons. The corresponding Giacomini and White (2006) test statistic is:

$$\text{Giacomini-White} = \frac{\hat{\sigma}_0^2 - \hat{\sigma}^2}{\hat{\sigma}_P / \sqrt{P}},$$

where $\hat{\sigma}_P$ is a heteroskedasticity and autocorrelation consistent (HAC) estimator of the asymptotic variance $\sigma_P^2 = \text{var} \left[\sqrt{P}(\hat{\sigma}_0^2 - \hat{\sigma}^2) \right]$. In contrast with the Clark-West test, the Giacomini-White test is two-sided.