

DBJ Discussion Paper Series, No. 1802

Liquidity Risk?

Jeffrey Pontiff

(Boston College)

Rohit Singla

(University of Michigan)

February 2019

Discussion Papers are a series of preliminary materials in their draft form. No quotations, reproductions or circulations should be made without the written consent of the authors in order to protect the tentative characters of these papers. Any opinions, findings, conclusions or recommendations expressed in these papers are those of the authors and do not reflect the views of the Institute.

Liquidity Risk?^ψ

Jeffrey Pontiff

and

Rohit Singla

February 5, 2019

Abstract

We revisit the role of liquidity risk. We successfully replicate Pastor and Stambaugh's (2003) gamma liquidity risk index and, within their time period, concur with their risk premium estimate. An out-of-their-time-period analysis finds post-time-period returns that are higher and pre-time-period returns that are lower than in-time-period returns. Modest variations to the index that are intended to improve power—such as value weighting, including zero volume days, including all stock price levels, and a modification intended to reduce estimation error—all cast doubt on whether the gamma premium is compensation for liquidity risk. We create five alternative liquidity risk indices from various popular liquidity proxies. Using time-series that start in either 1932 or 1968, none of the ten specifications produce statistically significant risk premia.

Keywords: Liquidity, Risk, Factor Model, Replication.

JEL Code: G00, G14, L3, C1

^ψPontiff (pontiff@bc.edu) is at Boston College. Singla (singlaro@umich.edu) is at the University of Michigan. Part of this work was completed by Pontiff while he was a Shimomura Fellow at the Development Bank of Japan. We thank George Aragon, Joel Hasbrouck, Jennifer Koski, Lubos Pastor, Ronnie Sadka, Robert Stambaugh, Kumar Venkataraman (referee), and Ivo Welch (editor) for valuable conversations. Lubos Pastor and Robert Stambaugh deserve special thanks for providing us with their original code. We thank Ming Lu for coding advice and Maximilan Papile for research assistance. SAS code that was used to generate this paper's results is available on Jeffrey Pontiff's Boston College website. Both co-authors are aware of concurrent work by Professor Robert Novy-Marx. We want our analysis to be independent and uninfluenced by Professor Novy-Marx's paper. As such, we have not read his paper and we are unaware of his results.

“Liquidity” describes the extent to which an asset can be quickly traded without too much of a price concession. Although the topic of liquidity is ubiquitous in the finance and economics literature, there is much disagreement on a precise measurement of it (see, for example, Goyenko, Holden, and Trzcinka, 2009)

Pastor and Stambaugh (2003) spawned a new literature by introducing the concept of “liquidity risk.” If liquidity contains a systematic component, security covariance with systematic liquidity risk might be priced. Investors demand higher expected returns for securities that deliver poor performance in in bad times. Thus, securities with lower returns in times of low liquidity will have higher equilibrium expected returns. Pastor and Stambaugh introduce a new measure of liquidity, gamma, which measures securities’ recovery from volume-related return shocks. They show that stocks with higher covariance to innovations of average, market-wide gamma have higher expected returns.

Using the Pastor-Stambaugh gamma index as a launching pad and focusing on simple measures of priced risk, our broad goal is to better understand the role of liquidity risk in asset pricing. Our investigation is comprised of two steps. First, we revisit the specifics of Pastor and Stambaugh’s liquidity index. Throughout our investigation we focus on direct implications that are likely to broadly steer liquidity risk-research. As the original, impactful liquidity risk paper, Pastor and Stambaugh (PS) provide a good starting place. The PS gamma index requires only daily CRSP data and the methodology can produce an index that starts 1920’s. Liquidity risk indices such as Sadka (2006), although appealing, require intraday data that is only available starting in the 1980’s. We also do not consider Acharya and Pedersen

(2005), which is more grounded in theory, since they produce a more complicated multi-measure expected return that has not been adopted by empirical research to the extent of Pastor and Stambaugh.

Our estimation precisely replicates the Pastor and Stambaugh gamma index (and we make this code available). Keeping with our focus, we evaluate their index based on direct tests of whether unconditional beta risk with respect to the index commands a return premium. We examine both the robustness of their index with their time period as well as the robustness of the index to construction decisions that are intended to increase statistical power. We find that gamma-risk continues to be associated with positive returns after the end of the Pastor and Stambaugh's original time period. On the other hand, our findings pose challenges for interpreting gamma-risk-related returns as compensation for liquidity risk. We consider four simple modifications of the gamma index that are expected to strengthen the detection of priced liquidity risk. Pastor and Stambaugh's estimation of gamma omits data on zero-volume days and does not include stocks with prices under \$5 or over \$1,000. Our first modification includes days with no trading volume. Our second modification includes all firms, regardless, of share price level. Pastor and Stambaugh's gamma index is an equal-weighted average of firm-level gammas. Our third modification constructs a value-weighted index. The original gamma estimation consumes a degree of freedom estimating an intercept. A wide range of theories imply that this estimate should be equal to zero. Our fourth modification imposes this restriction on the estimation. Although we expect all four modifications to strengthen the detection of priced liquidity risk, they all result in

lower return premium estimates and none of the modifications command statistical significance at the 5 percent level.

The second step of our investigation departs from the gamma index and considers liquidity risk more broadly. Guided by the literature, we construct four liquidity risk indices based on popular liquidity measures and we construct a fifth new liquidity risk measure that avoids theoretical pitfalls of common measures. The liquidity indices and the returns associated with these indices display remarkably low correlations. Ten out of ten specifications fail to generate a statistically significant liquidity risk premium. Our findings suggest that it is premature to conclude liquidity risk is priced or that we even have a reliable liquidity risk measure.

1. Replicating the Pastor-Stambaugh gamma index

Pastor and Stambaugh estimate the following daily return-generating process,

$$r_{i,d+1,t}^e = \theta_{i,t} + \phi_{i,t} r_{i,d,t} + \gamma_{i,t} \text{sign}(r_{i,d,t}^e) \cdot v_{i,d,t} + \epsilon_{i,d,t} \quad d = 1, \dots, D, \quad (1)$$

This regression is estimated at the firm-month level. d denotes the day in month t .

We utilize Pastor and Stambaugh's notation. $r_{i,d+1,t}^e = r_{i,d,t} - r_{m,d,t}$, where $r_{i,d,t}$ and $r_{m,d,t}$ are the respective returns on stock i and the CRSP value-weighted index. $v_{i,d,t}$ is the dollar volume and $\epsilon_{i,d,t}$ is the residual of stock i on day d in month t .

Our estimation follows Pastor and Stambaugh's lead. Individual regressions are estimated for each stock-month. Thus, the three parameters estimated from equation (1) will tend to be noisy, since the typical month contains twenty trading

days. If a stock's first or last month's data in CRSP is partial, it is excluded. All stocks must have a CRSP share code of 10 or 11, and they must be listed on either the New York Stock (NYSE) or American Stock (AMEX) exchanges. We (like Pastor and Stambaugh) do not include NASDAQ stocks to estimate gammas during this period since there was variation in the conventions that NASDAQ used to compute volume. All stock-months must contain 16 or more observations. The gamma estimate for each month does not have a look-ahead bias, in that data used is from the current month and the last trading day of the previous month. Following the construction used to post data to Professor Stambaugh's and Professor Pastor's websites (and WRDS), we make an exception for September of 2001. During this month only 11 or more stock-month level observations are required. This accommodates exchange closing due to the New York City terrorist attacks on September 11.

Our estimation in this section uses two filters that were used in the original estimation, although Section 3 reconsiders these filters. First, we only use stocks with closing stock prices in the previous month that are greater than or equal to \$5 and less than or equal to \$1,000. Second, we only use stocks with non-zero trading volume. The non-zero trading volume restriction is not explained in the original paper, but this restriction is not unheard of in the liquidity literature. The Amihud (2002) measure also only uses observations from non-zero volume days.

Gamma critique. The crucial parameter that equation (1) estimates is stock-level monthly liquidity, $\gamma_{i,t}$. γ is usually negative. On days where a stock's return in excess of the market is negative (positive), higher volume is associated with a higher (lower) excess return in the following day. γ 's are intended to measure liquidity. The

more negative a stock's γ , the more illiquid the stock. Although Pastor and Stambaugh focus on Grossman and Miller (1988) and Campbell, Grossman, and Wang (1993), their framework fits broadly into the dealer inventory models that go back as far as Ho and Stoll (1981, 1983). Liquidity provider risk aversion is the core assumption behind these models. Market clearing implies a lower (higher) price for the risky asset, accompanied with higher (lower) expected returns. γ reflects the reaction of expected returns to volume-related price movements. Put another way, high expected returns induce the risk-averse market maker to provide liquidity. The decision to estimate γ using return bounce-back over the next day as the dependent variable is a judgment call: the parameter could have been estimated with rolling multi-day returns as the dependent variable.

Another class of liquidity models results from asymmetric information between the liquidity provider and informed traders (such as Kyle, 1985). Both inventory models and asymmetric information models predict contemporaneous relation between absolute returns and volume. This relation is at the core of the definition in the first sentence of this paper. Negative gammas are artifacts of inventory models but not asymmetric information models. For example, in Kyle (1985) and Glosten-Milgrom (1985)—gamma is precisely equal to zero. In the context of asymmetric information models, a non-zero gamma implies a trading strategy. In this case, negative gamma is a market inefficiency, and an investor can expect to profit by buying (selling) high volume stocks with negative (positive) returns

Like PS, we compute a series of cross-sectional averages of γ , $\hat{\gamma}_t$, by dividing the sum of γ 's in a given month by the number of cross-sectional observations. These cross-sectional averages are used to construct an innovation measure (following PS's equation 6).

$$\Delta\hat{\gamma}_t = \left(\frac{m_t}{m_1}\right) \left(\frac{1}{N_t}\right) \sum_{i=1}^{N_t} (\hat{\gamma}_{i,t} - \hat{\gamma}_{i,t-1}) \quad (2)$$

m_t is the market value of all stocks used in the index in the previous month, and m_1 is the value of all stocks used in the index in August of 1962. Gamma communicates the return-reversal cost of a \$1 trade in stock i . Scaling the difference by the total market capitalization makes it easier to compare gamma shocks in different periods. For example, this adjustment rids the differences from variation that is mechanically attributable to market movements.

We use PS's equation (7), to estimate the following time series process,

$$\Delta\hat{\gamma}_t = a + b\Delta\hat{\gamma}_{t-1} + c\left(\frac{m_{t-1}}{m_t}\right)\hat{\gamma}_{t-1} + \mu_t \quad (3)$$

This specification produces the series, μ_t , which is an estimate of the innovation to the liquidity series. The Pastor-Stambaugh gamma index, L_t , is the fitted value of μ_t , scaled by 100. We investigate various liquidity measures and we also investigate the extent to which the PS index proxies for liquidity risk. In the interest of clarity, we refer to the PS index as the "gamma index."

Using the data from August, 1962 through December 2017, we compare our estimates of the level of gamma (the cross-sectional average of $\hat{\gamma}_{i,t}$'s) and the gamma index, L_t , to the data available on Professor Lubos Pastor's website. Our estimates are virtually identical, with a correlation of one (up to 5 significant digits).

Next, we estimate individual stocks' sensitivities to gamma risk. Following Pastor and Stambaugh, we apply the same exclusions that were used to create the liquidity index. Since estimating sensitivity to gamma risk does not require volume data, we include NASDAQ stocks and we require that stocks have complete return data for the last 60 months. The liquidity index is re-estimated, as per equation (3), each year such that no future information is used. The following factor model is estimated at each year-end.

$$r_{i,t} = \beta_i^0 + \beta_i^L L_t + \beta_i^M MKT_t + \beta_i^S SMB_t + \beta_i^H HML_t + \epsilon_{i,t} \quad (4)$$

Equation (4) is the well-known Fama-French three factor model with the addition of gamma innovation risk. For each year-end, we use β_i^L to sort stocks into 10 value-weighted portfolios. If the difference in these portfolio returns is non-zero, this is evidence that liquidity risk is "priced." This exercise follows pages 673 to 677 of Pastor and Stambaugh.

Table 1 compares our estimate of the return to this portfolio during the same time period as Pastor and Stambaugh's paper and with the data available on Professor Lubos Pastor's website. We estimate a liquidity premium of 3.91% per year, with a t-statistic of 2.00. Both parameters are to identical to estimates calculated from the website. These set of results are very similar to the finding reported in Table 8 of Pastor and Stambaugh, which reports a liquidity premium of 4.15% with a t-statistic of 2.08. This minor difference is almost certainly attributable to year-to-year corrections to the CRSP data.

(Table 1)

Pastor and Stambaugh devote considerable effort to construct predicted liquidity betas. They estimate time-series and cross-sectional variation in liquidity betas as a function of historical liquidity betas and various other right hand side variables such as stock momentum, volume, and return volatility. We have not pursued this avenue for two reasons. First, studies influenced by Pastor and Stambaugh almost never follow this approach. We read the 40 most cited papers (according to Google Scholar) that, in turn, cite Pastor Stambaugh. None of these papers use the predicted beta approach. Second, some of the conditioning variables used by Pastor and Stambaugh have already been shown to predict cross-sectional returns. As such, our analysis to focuses on the simple, unconditional liquidity betas used by the literature.

2. Out-of-Time-Period Stability

We compare the impact of gamma during the time period of Pastor and Stambaugh's, before their data starts, and after their time period ends. Our before-time-period investigation is possible since CRSP daily is currently available before 1962, whereas when Pastor and Stambaugh was written, this data was unavailable.

To conduct this analysis, we re-estimate the gamma innovation measure every year, only using historic data. Thus, in the 1960's estimation of equation (3) reflects data going back to the 1920s. To the extent, that the equation (3) parameters are stable, this is expected to produce a better estimate of gamma risk than the estimate in Section 1.

Figure 1 plots the aggregate gamma index for the entire series. Pastor and Stambaugh considered many specification variations. They settled on the equation (1) specification, in part, based on the resulting index's low level during the stock market crash of 1987. Therefore, pre- and post-time-period evidence is particularly valuable. Figure 1 shows that the gamma index is much less volatile, after their time period, and in particular, during the pre-1962 period. Despite equation 3's attempt to whiten the gamma process, the series has a first-order autocorrelation of 0.18. The financial drama of the great crash of 1929 and the great depression fail to register liquidity levels that could be construed as being low during the 60's through 90's. Post 1999, the financial crisis of 2007 to 2008 received incredible attention from the press and academia as being a period of market failure. The good news for the gamma index is that low index levels during this period are noticeable. The low levels are not confined to one month, rather they span several months.

The most negative level of the post-1999 gamma index is October 2002. This period is challenging, in that nothing in the media during this time seems consistent with a liquidity explanation.¹

(Figure 1 goes here)

An analysis outside of Pastor-Stambaugh's original time period is presented in Table 2. Our estimate of the price of gamma-risk during the Pastor-Stambaugh time period has a negligible drop of 9 basis points. This occurs since the liquidity index varies slightly from the Table 1 index since the estimation now starts in 1926. Before the start of the Pastor and Stambaugh time period, we estimate the price of

¹ We thank Brian Weller for insightful feedback about the behavior of the gamma index during this period.

gamma risk as -2.72 percent and statistically insignificant from zero. After the end of their time period, we estimate a 5.56 percent gamma risk premium with a t-statistic of 1.84. Using the entire time-series, we estimate an annual gamma risk premium of 1.45 percent with a t-statistic of 1.19. Thus, using data for the longest possible time series, we are unable to reject the null that gamma risk is not priced. The difference between the early data and the more-recent data is dramatic. For example, focusing on the 1968 to 2017 time period, the liquidity premium yields a hefty 4.45 percent annual returns and clearly rejects the null (t-statistic=2.67). Kamara et al. (2016) also document a time trend of increasing liquidity premium.

(Table 2)

McLean and Pontiff (2016) examine 97 variables documented to predict returns in academic papers. Using indicator variables, they estimate statistical bias and predictability decay from publication-informed trading. Their sample did not include the risk premium on the Pastor and Stambaugh liquidity measure. McLean and Pontiff estimate a decay of 26% after the original sample ends and a total decay of 58% post publication. Pastor and Stambaugh's study is unusual, in that Table 2 shows that the estimate of the return premium on gamma innovations increase post-sample by 46% $[(5.56-3.82)/3.82]$.

Readers will have different interpretations of Table 2. A reader who thinks that compensation for gamma-risk should be relatively stable over time will focus on the longest time-series and conclude that the price of gamma risk is insignificant from zero and likely small. A reader who believes that market participants' concerns about liquidity risk should lead to financial innovation that, in turn, mitigates this

risk (such as Alchian, 1950), is likely to be skeptical about whether gamma risk is liquidity risk, since the data before Pastor-Stambaugh's time period displays a negative price of gamma-risk. A reader who believes that financial innovation (or perhaps government policies) has enabled liquidity risk to be an increasingly systematic problem over time (for example, Taleb, 2012), will gravitate towards the more recent data and conclude that gamma risk is priced liquidity risk.

PS show that their gamma index is negatively correlated with market return volatility. They report a correlation between their gamma index and the within-month daily standard deviation of the value-weighted market (presumably restricted to AMEX and NYSE firms) of -0.57. During PS's original time period, our index has a correlation -0.56 with within-month daily standard deviations of the CRSP value-weighted index. Outside of PS's original time period, we find a less pronounced, albeit still negative, correlation of -0.38.

PS consider the extent to which correlations between stock market returns and bond returns vary in normal periods versus periods that the gamma index denotes as being illiquidity periods. They are interested in whether during apparently illiquid periods there is a "flight to quality." In low liquidity months they expect correlations of stock returns with low-risk bond returns to be negative as money flows between these investment classes.

Table 3 compares our in-time-period and out-of-time-period results with those that PS report in their Table 1². Two caveats are in order. First, our analysis,

² Table 1, panel A, in PS refers to their main time period as starting in January 1962. This is almost certainly a typo, since availability of CRSP data would have prevented them from starting before

like PS's is asking a lot from the data. The number of observations in low gamma subsamples is sparse—a little more than a dozen observations. Also, our results, although in the spirit of Pastor and Stambaugh, should not be considered a replication. Our index is constructed using an autocorrelation adjustment that is estimated starting in 1926—whereas PS start in 1962. We use CRSP bond data, whereas PS use Ibbotson Associates bond data. Our market return index is the CRSP value-weighted index, whereas their value-weighted index is confined to NYSE and AMEX listings.

(Table 3 goes here)

Like Pastor and Stambaugh, in Table 3 we call a month a “low liquidity” month, if the gamma index is more than two standard deviations lower than the index mean during the original PS time period. Using CRSP data, we measure $-\Delta R_{f,t}$ as the next period's return minus this period's return on 30 day T-bills. Pastor and Stambaugh's version of this variable is last month's treasury yield minus this month's. Again, using CRSP we measure $R_{GB,t}$ as the return on 20 year Treasury bonds, whereas Pastor and Stambaugh use the Ibbotson Associate return on long-term government bonds. We are unable to retrieve corporate bond returns, so they are not included in our Table 3.

Panel A of Table 3 considers the same time period as PS. The estimates are similar, but not exact. Like PS, we find that in low gamma months the correlation between $-\Delta R_{f,t}$ and market returns is negative and in other months it is positive. Using Fisher's r-to-z transformation, our p-value on this difference is 0.19, whereas

August 1962. The number of observations that we report (starting in August 1962) corresponds to what PS report.

PS report a bootstrap p-value of 0.09. We suspect that the difference in p-values levels is attributable to the fact that our p-values assume a two-tailed test and PS likely report values for a one-tail test. We find differences in correlations between 20 year bond returns and market returns in normal and low-gamma months that are similar to PS, although not quite as pronounced. Our two-sided p-value on the difference is 0.10, roughly twice the 0.045 reported by PS.

Column (3) considers the correlation between the changes in volume and market returns. PS report that their change is the average percentage equally weighted change for NYSE-AMEX stocks. We report the median percentage change of this variable on for all CRSP stocks with share codes of 10 or 11³. Like PS, we find negative correlations between change of volume and returns in low-gamma months, and positive correlations in other months. This difference is statistically significant at all standard levels.

Panel B repeats the Panel A statistics for data outside PS' original time period. We continue to use the mean and standard deviation during their time period to sort gammas into normal and low months. The differences in correlation between market returns and changes in Treasury bill prices are very similar to the PS results during the same time period. This is consistent with PS' appeal to "flight to quality," although the 0.24 p-value is unable to reject the null. On the other hand, the difference in correlations switches signs for 20 year government bonds. Outside of the PS time period, in low-gamma months, stock market returns are positively correlated with long-term government bond returns, and in other months they are

³We focus on medians since percentage changes in volume are susceptible to extreme outliers, since high volume can follow periods with no volume or very low volume.

negatively correlated. We are unable to use all the months outside of the PS time period for this test, since long-term government bond data is not always available. The volume evidence also varies. Both low-gamma and normal months exhibit positive correlations with changes in volume.

What do these results tell us about PS' liquidity risk measure? Regarding capital flows from stocks to Treasury bills, although the inside/outside period point estimates are very close, both fail to reject the null. Regarding capital flows from stocks to long-term government bonds, the outside period result flips. This result challenges a "flight to quality" interpretation. The inside/outside volume differences appear similar, suggesting consistency. This being noted, PS' original investigation of volume is descriptive and not a test of a specific story.

3. Modifying the gamma index

In this section, we consider four simple modifications in the construction of the Pastor-Stambaugh gamma index and the extent to which these modifications influence our estimate of the gamma risk premium. We focus on modifications that are expected to improve estimates of average returns based on the presumption that these returns are compensation for liquidity risk. These 4 modifications to the gamma index (and the modification in Section 4) were selected before we knew how the estimated price of risk would be affected, and all modifications that we estimated are presented in this paper.⁴

⁴Although we did not know how the modifications would affect the price of risk, for two modifications we knew something about their correlation with the base-line PS gamma index. First, PS report that the correlation between gamma index innovations and a value-weighted version is

Including observations on zero-volume days. The estimation of equation (1) does not include data from zero volume days. While not discussed in Pastor and Stambaugh, this decision is consistent with other liquidity measures such as that of Amihud (2002). Since August 1962, zero volume days account for 7.8 percent of CRSP daily data (with non-missing returns, non-missing prices, and a share code of 10 or 11). Lesmond, Holden, and Trzcinka (1999) argue that zero volume days are a critical consideration for measuring liquidity. On zero volume days, we follow CRSP in calculating returns with bid-ask averages instead of trade prices. We expect that inclusion of zero volume days will improve equation (1)'s estimation of liquidity since doing so allows the estimation to use more data over a wider liquidity range.

Inclusion of all price levels. As noted by Demsetz (1968) Stock price levels proxy for transaction costs and liquidity. Pastor and Stambaugh exclude stocks with prices that are either less than 5 dollars or greater than 1,000 dollars. Although PS do not provide an explanation for this restriction, this restriction is common in studies that use stock price data. Authors such as Amihud (2002) attribute the restriction to concerns about relative bid-ask spreads being too big, especially during periods of minimum tick sizes, and some authors (such as Boni and Womack, 2006) attribute the restriction to difficulties that traders have shorting stocks with prices under \$5. The under \$5 restriction eliminates 24.5% of the CRSP daily data (with non-missing returns, non-missing prices, and a share code of 10 or 11). Including this data should improve the index for two reasons. First, additional data

0.77. Our own specification produces a correlation of 0.95. Second, when we first attempted to replicate the gamma index we did not exclude zero volume days, and we realized that this produced an index with very little correlation to the original index. We learned about this exclusion from conversations with professors Pastor and Stambaugh.

should decrease estimation error. Second, stocks from these price levels are likely to be very sensitive to changes in market-wide liquidity.

Value-weighted Index. Pastor and Stambaugh's liquidity index is equal weighted. They consider a value-weighted index, but they do not estimate the price of liquidity risk relative to a value-weighted index. Their decision to focus on an equal-weighted index is based on the equal-weighted index's low levels during times of purported low liquidity—October 1987 and September 1998. Almost all theories of market equilibrium produce value-weighted pricing implications. For example, representative agent models imply that the agent holds the value-weighted market. Similarly, market clearing with multiple agents implies a value-weighted equality. As such, we consider a value-weighted version of the Pastor and Stambaugh index.

$\theta_{i,t}$ restricted to zero. Estimation of equation (1) produces parameter estimates that are very noisy. Since August 1962, the average month has 21.0 trading days. Thus, in an average month (for stocks with a full array of data), 21 observations are used to estimate three parameters: $\theta_{i,t}$, $\phi_{i,t}$, and $\gamma_{i,t}$. For some stock-months (even if we disregard September 2011), 16 observations are used to estimate these parameters. One way to estimate more efficient liquidity parameters, $\gamma_{i,t}$, is to impose theoretically sensible restrictions on one of the other two coefficients. We consider specifications where the intercept, $\theta_{i,t}$, is restricted to be zero. This restriction implies that a stock with a zero return in given day, is expected to have a return in the following day that equates to the return of the value-weighted market. This restriction is in line with previous research. Using monthly

data, Simin (2008) shows that forecasts of individual stock returns that are equated to the market return have lower mean-square-errors than estimates from asset pricing models.

Using data from August 1962 to December 2017, Table 4 reports correlations between the Pastor-Stambaugh index and the four modified indices (panel A), correlations between long-short decile portfolio returns for this set of indices, and return premiums for long-short portfolio based on these indices. Panel A shows the Pastor-Stambaugh index, the value-weighted modification, and the intercept restricted version, all have correlations with one another in excess of 0.90. The correlation between the PS index and value-weighted modified index is 0.95, which is substantially higher than the 0.77 correlation that Pastor and Stambaugh document between these indices using data that ends in 1999.

(Table 4 goes here)

Including all stocks regardless of price and including zero volume days creates indices that are largely orthogonal to the remaining indices. For example, the highest correlation, 0.27, is between the PS index and the modification that includes zero volume days. The lowest correlation, 0.15, is between the modification that includes all stocks regardless of price with the intercept restricted modification,

Panel B looks at correlations between long-short extreme decile portfolio returns, that are formed based on slope coefficients with these indices. The population of stocks used to form these populations is the same for all indices, except the indices that includes all stocks regardless of price. For this specification, the long-short portfolio, like the index, includes stocks regardless of price level. The

correlations in panel B, with a few exceptions, tend to shrink compared with panel A.

Panel C compares estimates of annual return premia from the long-short deciles portfolios. All modifications result in lower estimates of the liquidity risk premium than the premium estimated from the PS index. All modifications produce positive estimates of the risk premium, although the only modification with statistically significant estimate is the intercept-restricted modification (t-statistic of 1.89). The news from this panel is mixed. On one hand, a modification intended to improve the power of the gamma estimation, continues to command a statistically significant risk premium, albeit with a lower point estimate. On the other hand, modifications that are expected to do a better job measuring liquidity risk (such as including all prices and zero volume days), and a specification expected to do a better job capturing priced risk (value-weighted index), are unable to produce statistically significant estimates of risk premia. An interpretation of these results is that the Pastor-Stambaugh gammas convey information about expected returns, but this information is either unrelated to liquidity or unrelated to priced risk. The specifications that include all prices and zero volume days, challenge the notion that the gamma index is liquidity risk, while the value-weighted results cast doubt on whether gamma risk is priced.

4. Liquidity risk beyond Pastor-Stambaugh

The literature has developed many proxies for liquidity. In this section, we explore three liquidity proxies from the literature, a version of the gamma index

with two modifications, and a fifth proxy that we develop ourselves that is intended to avoid some of the pitfalls of other measures.

In selecting liquidity measures, we focus on measures that only require CRSP data. Measures that use TAQ data, such as Sadka (2006) are appealing, but data availability prevents the creation of a long-time series of liquidity shocks. We seek an even-playing field to compare measures. As such, the indices are calculated by taking the market-capitalization-weighted average of each stock's monthly change in the three liquidity measures from the month before. Market capitalization is calculated with last month's price and shares outstanding. This decision to weight by market capitalization recognizes that illiquidity among large market capitalization stocks is more likely to be a cause of economy-wide risk than small stocks. Stocks are only excluded from each index if their CRSP share code is not equal to 10 or 11, if we are unable to calculate their liquidity measure in the current or past month, if we are unable to calculate their market capitalization in the previous month, or if they are not traded on the NYSE or AMEX exchanges.

Similar to the methodology used in the previous section to estimate the price or risk of the PS gamma index, betas on the indices' innovations are calculated in annual factor models that include the Fama-French three factors. Stocks are included from all exchanges as long as they have a full 60 months of data. Betas are used to allocate stocks in decile portfolios in the following year. The price of risk is the annual return of the value-weighted portfolio that holds the top decile and shorts the value-weighted portfolio that holds the bottom decile. Our first four liquidity measures are: the proportion of zero returns (Lesmond, Ogden, and

Trczinka, 1999), the Amihud liquidity measure (Amihud, 2002), relative bid-ask spreads, and a hybrid index that avoids common pitfalls of other measures. For comparison, we also include a variation of the PS gamma index that is constructed using the methodology in Section 1, except that two modifications are included simultaneously--value-weighting and inclusion all stocks regardless of price. For the non-gamma indices, high values are associated with illiquidity, whereas for the modified gamma index high values are associated with liquidity. As such, we multiply index innovation for non-gamma indices by negative one, such that the betas of all indices can be interpreted as liquidity betas.

Proportion of zero returns. Lesmond, Ogden, and Trczinka (1999) develop a transaction cost proxy that is the proportion zero return days. Their insight is that trade occurs when the value of information exceeds transaction costs. Variation in transaction costs causes illiquid securities to have a higher a proportion of zero return days. Following Lesmond et al., the zero return index is based on the monthly percentage of zero return days at the stock level.

Amihud Liquidity Measure. Amihud (2002) proposes a measure of price impact that is based on the ratio of absolute return and trading volume. This measure is calculated as,

$$A_{i,m} = \frac{1}{D_{i,m}} \sum_{t=1}^{D_{i,m}} \frac{|r_{i,t}|}{Volume_{i,t}}. \quad (xx)$$

$A_{i,m}$ is stock i 's Amihud measure for month m , $D_{i,m}$ is the number of days in month m with both return and dollar volume data, $r_{i,t}$ is stock i 's return on day t , and $Volume_{i,t}$ is stock i 's dollar trading volume on day t . Like the Pastor and Stambaugh measure, the Amihud measure does not use information from zero volume days. Thus, a

disadvantage of the Amihud measure is that it assigns the same liquidity measure to stocks with the same average ratio of absolute return to volume, despite the fact some stocks might trade much more frequently than the other. An advantage of the Amihud measure is that it is constructed with contemporaneous returns and volume. As mentioned earlier, a necessary condition for all theories of illiquidity is a contemporaneous relation between absolute returns and volume.

Bid-ask spread index. Amihud and Mendelsohn (1996) develop a transaction cost asset pricing model. Clientele groups with longer holding periods buy securities with higher relative bid-ask spreads. In equilibrium, securities with higher levels of relative bid-ask spreads have higher expected returns. Using NYSE bid-ask spread data, Amihud and Mendelsohn find empirical support for their model.

Bid-ask spread data is not generally available. Corwin and Schultz (2012) propose a procedure that estimates levels of bid-ask spreads based on high and low daily prices. They show that their estimates have a correlation of 0.9 with actual bid-ask spreads. Using Corwin and Schultz's method, we estimate bid-ask spreads for all CRSP stocks, and create a relative spread by dividing the estimated spread of a stock in a given month by the average price of that stock in that month.

Hybrid index. A hybrid index is constructed based on a liquidity measure that avoids the pitfalls of other measures. The hybrid measure is defined as,

$$H_{i,m} = \left(\frac{Max_m}{D_{i,m}} \right) \left(\frac{1}{D_{i,m}} \right) \sum_{t=1}^{D_{i,m}} \frac{|r_{i,t} - r_{m,t}|}{Turnover_{i,t}}$$

where $D_{i,m}$ is the number of days in month m for which stock i has turnover and return data. Max_m is the total number of trading days in month m . $r_{i,t}$ and $r_{m,t}$ are the

respective returns of the stock i and the market for day t . For estimation, we use the CRSP equal-weighted index return with dividends as our market proxy. $Turnover_{i,t}$ is the number of shares stock i that are traded in day t as a fraction of shares outstanding. The second and third expressions are similar to the Amihud measure, in that they attempt to measure the contemporaneous volume-related price impact. The numerator of the last expression recognizes that return movements that correspond to broad market movements are unlikely to provide information about illiquidity. The first expression magnifies the illiquidity measure for stocks that do not trade every day. Thus, a stock that only trades on half of the days, will be assigned an H measure that is twice that of stock which trades every day. This adjustment measures liquidity under the assumption that a trader is unable to trade on a zero volume day, and proportionately increases trading on non-zero volume days.

The decision to use turnover instead of dollar volume in the denominator of the Hybrid index avoids a mechanical relation between the liquidity measure and market capitalization. This characterization dovetails with the decision to value-weight the measures to create the index.

Results. Table 5 and 6 present results for the investigation of alternative liquidity risk indices. Table 5 focuses on the most recent time period and Table 6 reports on the full time period. Correlations between innovations and the indices are low. These correlations are even lower than those calculated on a different set of liquidity measures (detrended using a second-order auto-regressive process) by Koraczyck and Sadka (2008). Post 1962 (Table 5, panel A) the highest correlation,

0.39, is between the Amihud index and the Hybrid index. The fact that this is the highest correlation among the group is not too surprising since the construction of both involves the division of an absolute return measure by a trading intensity measure. The second highest correlation is 0.05. The lowest correlation, -0.12, is between the zero return index and the bid-ask spread index. The full time period, Table 6 results demonstrate higher correlations—the strongest being 0.95, again between the Amihud and Hybrid index. The second highest correlation is 0.37. The message is clear—these indices share little commonality. Correlations between long-short portfolio returns based on these indices tell the same story.

The low correlations in tables 5 and 6 reinforce a re-occurring challenge in the liquidity literature. Different liquidity measures often seem at odds with one another (for example, Goyenko et al., 2009). Although there is no unifying explanation for these differences, a contributing factor is the non-observability of uncompleted trades. Market participants make decisions to not trade when illiquidity is high. In days of high illiquidity, all trades may be prohibitively expensive. In turn, this results in no volume and unobservable returns based on transaction prices—which are inputs to most liquidity measures.⁵

Panel C of Table 5 and Panel F of Table 6, present average annualized risk premiums. None of the ten estimates is significant at the 5% level. Using recent data, the best performer is the zero return index, with an annualized risk premium of 3.44 percent (t-statistic of 1.74) and using the full time series, the best performer is the hybrid index with a risk premium of 2.25 percent (t-statistic of 1.52).

⁵ We thank the referee, Kumar Venkataraman, for pointing this out.

5. Conclusions

Pastor and Stambaugh correctly estimate a return premium on their gamma index that is positive and statistically significant. Although the result does not hold before their time period, it is stronger after their time period. This is unusual given that the typical finding of return predictability decays by nearly 60% once a paper is published (McLean and Pontiff, 2016).

Is gamma risk priced illiquidity risk? On one hand, the aforementioned predictability continues after the end of Pastor and Stambaugh's original sample. Despite this, we fail to find other supporting evidence that gamma risk return predictability is compensation for liquidity risk. Four modifications to the gamma index that are designed to strengthen the measurement of priced liquidity risk all yield lower returns that fail to reject the null of no premium at the 5% level. Outside of Pastor and Stambaugh's time period, we fail to find that low gamma-periods exhibit a "flight-to-quality." Departing from the gamma index and considering liquidity risk more broadly, we fail to find any evidence that indices created from common liquidly variables capture priced risk.

References

- Acharya, Viral V. and Lasse H. Pedersen, 2005, Asset pricing with liquidity risk, *Journal of Financial Economics* 77, 375-410.
- Alchian, Armen 1950, Uncertainty, evolution, and economic theory, *Journal of Political Economy* 58, 211-221.
- Amihud, Yakov, 2002, Illiquidity and stock returns: Cross-section and time series effects, *Journal of Financial Markets* 5, 34-56.
- Amihud, Yakov and Haim Mendelson, 1986, Asset pricing and the bid-ask spread, *Journal of Financial Economics* 17, 223-249.
- Boni, Leslie, and Kent L. Womack, 2006, Analysts, industries and price momentum, *Journal of Financial and Quantitative Analysis* 41, 85-109.
- Campbell, John Y., Sanford J. Grossman, and Jiang Wang, 1993, Trading volume and serial correlation in stock returns, *Quarterly Journal of Economics* 108, 905-39.
- Corwin, Shane A. and Paul Schultz, 2012, A Simple Way to Estimate Bid-Ask Spreads from Daily High and Low Prices, *Journal of Finance* 67, 719-760.
- Glosten, Lawrence R. and Paul R. Milgrom, 1985, Bid, ask and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 14, 71-100.
- Goyenko, Ruslan Y., Craig W. Holden, and Charles A. Trzcinka, 2009, Do liquidity measures measure liquidity?, *Journal of Financial Economics* 92, 153-181.
- Grossman, Sanford J., and Merton H. Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617-637.
- Ho, Thomas, and Hans Stoll, 1981, Optimal dealer pricing under transactions and return uncertainty, *Journal of Financial Economics* 9, 47-73.
- Ho, Thomas, and Hans Stoll, 1983, The dynamics of dealer markets under competition, *Journal of Finance* 38, 1053-1074.
- Federal Reserve Bank of Boston, 2004, Stock market report—2003 review, <https://www.bostonfed.org/-/media/Documents/stock-market-report/smrye2003.pdf>
- Kamara, Avraham, Robert Korajczyk, Xiaoxia Lou, and Ronnie Sadka, 2016, Horizon Pricing, *Journal of Financial and Quantitative Analysis* 51, 1769-1793.

- Korajczyk, Robert and Ronnie Sadka, 2008, Pricing the commonality across alternative measures of liquidity, *Journal of Financial Economics* 87, 42-72.
- Kyle, Albert S., 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315-1335.
- Lesmond, D., Ogden, J., Trzcinka, C., 1999, A new estimate of transaction costs. *Review of Financial Studies* 12, 1113-1141.
- McLean, R. David and Jeffrey Pontiff, 2016, Does academic research destroy stock return predictability?, *Journal of Finance* 71, 5-32.
- Miller, Edward, 1977, Risk, uncertainty, and divergence of opinion, *Journal of Finance* 32, 1151-1168.
- Sadka, Ronnie, 2006, Momentum and post-earnings-announcement drift anomalies: The role of liquidity risk, *Journal of Financial Economics* 80, 309-349.
- Simin, Timothy, 2008 "The (Poor) Predictive Performance of Asset Pricing Models", *Journal of Financial and Quantitative Analysis* 43, 391-2968.
- Taleb, Nassim Nicholas, 2012, *Antifragile: things that gain from disorder*, Random House (New York, New York).

**Figure 1:
Aggregate Liquidity Measure**




Table 1: Replication of PS Decile Return Spread

Description: This table compares annualized alphas for a long-short portfolio that is long a value-weighted portfolio of stocks in the highest decile of gamma betas and is short a value-weighted portfolio of stocks in the lowest decile of gamma betas. t-statistics are in parentheses. The time period is 1968:1 to 1999:12.

Interpretation: Our code perfectly replicates the results produced on Lubos Pastor and Robert Stambaugh's website. Our code produces estimates that are very close to those produced in their paper.

Annual alpha of long-short liquidity beta portfolio		
(1)	(2)	(3)
Replication results	Estimate from Table 8 of Pastor Stambaugh	Website estimate
3.91	4.15	3.91
(2.00)	(2.08)	(2.00)

Table 2: Outside of PS time period comparison

Description: This table compares annualized alphas for a long-short portfolio that is long a value-weighted portfolio of stocks in the highest decile of gamma betas and is short a value-weighted portfolio of stocks in the lowest decile of gamma betas. t-statistics are in parentheses.

Interpretation: Before the start of Pastor and Stambaugh’s original time period, gamma-risk is associated with a negative statistically insignificant risk premium. In contrast to the typical return-predictability findings (McLean and Pontiff, 2016), after the end of Pastor and Stambaugh’s original time period, gamma risk commands a higher return premium.

Note: The column (2) risk premium estimate differs from the Table 1, column (2) risk premium estimate because the gamma index this table is whitened using information that starts in the 1932. Table 1, like the Pastor and Stambaugh paper, whitens the series using data that starts in 1968.

	(1) Pre-PS-time period	(2) PS time period	(3) Post-PS time period	(4) Entire CRSP	(5) PS plus post- PS time period
Start	1932:1	1968 :1	2000:1	1932: 1	1968:1
End	1967:12	1999:12	2017:12	2017:12	2017:1
	-2.72 (-1.36)	3.82 (1.93)	5.56 (1.84)	1.45 (1.13)	4.45 (2.67)

Table 3: Comparison of correlations of PS Index innovations during and outside the original PS period

Description: The table reports correlations between monthly returns on the CRSP aggregate value-weighted index, $R_{S,t}$, and (i) the current month's return on 30 day Treasury bills minus the next month's return on 30 day Treasury bills, $-\Delta R_{f,t}$; (ii) the return on 20 year Treasury bonds; and (iii) the median percentage change in monthly dollar volume.

Interpretation: This estimation produces two-tailed p-values that are roughly twice that reported by PS. PS does not report whether their test is two- or one-tailed, so it is likely that their test is one-tailed. There are not many low gamma months and as such, our results and PS results suffer from low statistical power. Although the correlation differences between low gamma month and other month are similar during and outside of the PS time period for both the change in the risk free return and the median change of monthly dollar value, the correlation differences with the return on 20 year Treasury bonds switch sign.

Panel A: Within PS time period correlation of monthly CRSP value-weighted index return with other variables. 1962:8 to 1999:12

	Correlation of $R_{S,t}$ with			Number of		
	(1) $-\Delta R_{f,t}$	(2) $R_{GB,t}$	(3) Vol_t	Observations		
All months	-0.01	0.30	0.47	449	449	449
Low-gamma months	-0.35	-0.13	-0.21	14	14	14
Other months	0.03	0.35	0.52	435	435	435
p-value (difference)	0.19	0.10	0.01			

Panel B: Outside PS time period correlation of monthly stock returns with other variables in months with large liquidity drops. 1927:1–1962:7 and 2000:1–2017:12

	Correlation of $R_{S,t}$ with			Number of		
	(1) $-\Delta R_{f,t}$	(2) $R_{GB,t}$	(3) Vol_t	Observations		
All months	-0.02	-0.22	0.45	654	463	654
Low-gamma months	-0.30	0.19	0.26	16	15	16
Other months	0.02	-0.24	0.46	638	448	638
p-value (difference)	0.24	0.14	0.41			

Table 4: Performance of modified PS gamma indices

Description: This table examines the PS gamma index and four modifications of the gamma index. The four modifications are including zero volume days in the gamma estimation, including all prices regardless of level, value-weighting gammas to create the index instead of equal-weighting, and estimating the gammas in an equation that restricts the intercept to be equal to zero. Correlations between the innovations, correlations between correlations between long-short extreme deciles portfolios, and the mean returns on these portfolios are reported. t-statistics are in parentheses.

Interpretation: Panel A and B show that the zero volume and all price modifications create results with low correlation with PS. Panel C shows that these produce liquidity portfolio returns that are lower than the PS estimate and that fail to reject the null.

Panel A. Correlations between PS variant liquidity innovations: 1962:8 to 2017:12

	Value weighted Index	All prices included	Zero volume included	Baseline PS index
Zero volume included				0.27
All prices included			0.11	0.17
Value weighted index		0.17	0.26	0.95
$\theta_{i,t}$ restricted to zero	0.92	0.15	0.24	0.97

Panel B. Correlations between PS variant liquidity returns: 1968:1 to 2017:12

	Value weighted Index	All prices included	Zero volume included	Baseline PS index
Zero volume included				0.26
All prices included			0.02	0.10
Value weighted index		0.15	0.25	0.82
$\theta_{i,t}$ restricted to zero	0.77	0.18	0.25	0.87

Panel C. Average annualized returns of PS variant traded liquidity factors: 1968:1 to 2017:12

$\theta_{i,t}$ restricted to zero	Value weighted Index	All prices included	Zero volume included	Baseline PS index
3.13 (1.89)	1.02 (0.59)	0.78 (0.15)	0.82 (0.50)	4.45 (2.67)

Table 5: Performance of alternative liquidity risk indices—recent data

Description: This table examines five indices constructed from liquidity measures. Correlations between the innovations, correlations between correlations between long-short extreme deciles portfolios, and the mean returns on these portfolios are reported. t-statistics are in parentheses

Interpretation: Panel A and B show that there tends to be low correlation between the alternative indices and the returns of alternative indices. Panel C reports that none of the indices produce a liquidity portfolio return that is statistically significant from zero.

Panel A. Correlations between alternative liquidity innovations: 1962:8 to 2017:12

	Bid-ask spread index	Amihud index	Zero return index	Hybrid index
Zero return index				0.05
Amihud index			0.04	0.39
Bid-ask spread index		0.05	-0.12	0.04
PS dual modified index	0.04	0.03	-0.01	-0.01

Panel B. Correlations between alternative liquidity returns: 1968:1 to 2017:12

	Bid-ask spread index	Amihud index	Zero return index	Hybrid index
Zero return index				-0.08
Amihud index			-0.09	0.30
Bid-ask spread index		0.12	-0.20	0.13
PS dual modified index	-0.16	-0.06	0.27	0.01

Panel C. Average annualized returns of alternative traded liquidity factors: 1968:1 to 2017:12

PS dual modified index	Bid-ask spread index	Amihud index	Zero return index	Hybrid index
-1.26 (-0.71)	-1.76 (-0.89)	0.47 (0.24)	3.44 (1.74)	1.57 (0.90)

Table 6: Performance of alternative liquidity risk indices—entire time period

Description: This table examines five indices constructed from liquidity measures. Correlations between the innovations, correlations between correlations between long-short extreme deciles portfolios, and the mean returns on these portfolios are reported. t-statistics are in parentheses

Interpretation: Panel A and B show that there tends to be low correlation between the alternative indices and the returns of alternative indices. Panel C reports that none of the indices produce a liquidity portfolio return that is statistically significant from zero.

Panel A. Correlations between alternative liquidity innovations: 1927:1 to 2017:12

	Bid-ask spread index	Amihud index	Non-zero return index	Hybrid index
Zero return index				0.08
Amihud index			0.04	0.95
Bid-ask spread index		0.27	0.05	0.37
PS dual modified index	0.02	0.01	-0.00	0.01

Panel B. Correlations between alternative liquidity returns: 1932:1 to 2017:12

	Bid-ask spread index	Amihud index	Zero return index	Hybrid index
Zero return index				0.18
Amihud index			0.12	0.51
Bid-ask spread index		0.25	0.12	0.35
PS dual modified index	0.07	0.03	0.32	0.10

Panel C. Average annualized returns of alternative traded liquidity factors: 1932:1 to 2017:12

PS dual modified index	Bid-ask spread index	Amihud index	Zero return index	Hybrid index
-0.65 (-0.45)	-1.60 (-1.01)	1.86 (1.21)	1.63 (1.08)	2.25 (1.52)