Notes on Bonds: Illiquidity Feedback During the Financial Crisis[*][‡]

David Musto
Greg Nini
Krista Schwarz

September 28, 2016

Abstract: This paper traces the evolution of extreme illiquidity discounts among Treasury securities during the financial crisis; bonds fell more than six percent below more-liquid but otherwise identical notes. Using high-resolution data on market quality and trader identities and characteristics, we find that the discounts amplify through feedback loops, where cheaper, less-liquid securities flow to investors with longer horizons, thereby increasing their illiquidity and thus their appeal to these investors. The effect of the widened liquidity gap on transactions costs is further amplified by a surge in the price liquidity providers charged for access to their balance sheets in the crisis.

Notes on Bonds: Illiquidity Feedback During the Financial Crisis

Abstract: This paper traces the evolution of extreme illiquidity discounts among Treasury securities during the financial crisis; bonds fell more than six percent below more-liquid but otherwise identical notes. Using high-resolution data on market quality and trader identities and characteristics, we find that the discounts amplify through feedback loops, where cheaper, less-liquid securities flow to investors with longer horizons, thereby increasing their illiquidity and thus their appeal to these investors. The effect of the widened liquidity gap on transactions costs is further amplified by a surge in the price liquidity providers charged for access to their balance sheets in the crisis.

1. Introduction

The rapid contraction of the financial industry in 2007 - 2009 pulled many prices out of their usual orbits. Of the many resulting anomalies, perhaps none was as stark as the disparity in the Treasury market between bonds, i.e. Treasuries issued with 30 years to maturity, and notes, i.e. all other coupon-paying Treasuries. Bonds traded at a relative discount that reached six percent of face value, even with cash flows matched exactly. So extreme a violation of the law of one price indicates another strong force differentiating asset prices, and the notable remaining difference between off-the-run bonds and notes is that the notes are more liquid. In this paper we ask how this seemingly modest liquidity difference, which normally commands a very small price, came to command such a large price during the crisis. The answers to this question help us understand not only this extreme event but also the whole range of market disruptions characterized by dramatic differences in relative pricing of comparable securities, such as the subsequent Treasury "flash crash" in October 2014 and the March 2016 spike in Treasury "fails," and they also help evaluate the view that post-crisis developments have led to a lower-liquidity "new normal" in market quality.[1]

Our investigation adapts the analysis by Dow (2004) of the role of investor clientele in amplifying illiquidity discounts. In Dow (2004), investors with more private information about the security's payoffs select less liquid securities, and this selection lowers the liquidity further, which in turn increases the selection. So asymmetric information about payoffs can power a feedback loop that amplifies liquidity differences. We adapt this analysis to a setting where the power comes not from

---

[1] A Joint Staff Report by the U.S. Treasury, "The U.S. Treasury Market on October 15, 2014," investigates the 10-year Treasury note's price volatility, stating that in the post-crisis market environment, traditional liquidity measures "may need to be complemented by other measures in light of these changes to obtain a more meaningful picture of the state of market liquidity in the current structure." A Federal Reserve Bank of New York Staff Report, "What's behind the March Spike in Treasury Fails," investigates Treasury security settlement fails that spiked to a post-crisis high in March 2016, averaging $95 billion per day over the month, in both on- and off-the-run securities. The report states that, "When fails materialize, the risk of a self-fulfilling dynamic arises whereby some large holders of benchmark issues become reluctant to lend if they fear that the securities will not be returned at the end of the loan."

investors' private information about payoffs, but, as in Amihud and Mendelson (1986), from the investors' need for liquidity. That is, we ask whether the feedback could instead be that investors who expect to hold a security longer select less liquid securities, which lowers liquidity and further increases the selection. The market for off-the-run Treasuries, where this price divergence occurred, presents an ideal opportunity to find out.

The off-the-run Treasury market is ideal for three reasons. First, the precise identity between Treasury securities promising the same future cash flows isolates the effect of liquidity differences; their payments come from the same obligor with the same seniority, they are denominated in the same currency, they are subject to the same tax schedules, their trades settle the same way, they are equally valid as repo collateral, and are otherwise the same in every meaningful way. But, the ease with which the securities are traded is free to vary. Second, the liquidity indeed varies widely; despite the consistently high liquidity of on-the-run Treasuries, liquidity differences among *off*-the-run Treasuries are large. As we document below, bonds trade at twice the bid-ask spread of ten-year notes on average,[2] and the average of the time since the last buy or sell of a bond, on the platform that provided our data, is almost *three days*, over eight times longer than the same average for a ten-year note. And finally, for a large and varied group of Treasury traders, i.e. every U.S. insurance company, we see not only all their trades but also extensive data on their circumstances, particularly those relevant to investment horizons. So given the nature of the securities and the available data, the off-the-run Treasury market is ideal to examine whether less-liquid securities fall to a further disadvantage by flowing to longer-horizon investors.

Our Treasury market data source is key to this study, because the usual sources (e.g. eSpeed, Brokertec) reveal almost nothing about the off-the-run market. This is because the market segment they represent, the interdealer market, is not where off-the-run transactions take place. For example,

---

[2] Throughout the paper, we refer to securities originally issued with *n* years to maturity as *n*-year securities.

one could not infer from interdealer data sources, and probably would not even suspect, that days go by between trades of an off-the-run bond. The major venue for off-the-run trading is not the interdealer market but rather the customer-to-dealer market, as shown on the TradeWeb platform,[3] an online request-for-quote platform backed by the largest Treasury dealers. This is the source of our database, which gives key metrics of market quality, including bids and asks for each CUSIP, as well as the time since each CUSIP last traded. Thus, we not only see the transactions costs charged by liquidity providers to demanders, we also have a good proxy for how long the liquidity providers can expect to wait before a later trade takes any new supply off of their books.

Figure 1 shows the evolution of the anomaly we study. Each line in Panel A represents the price of a ten-year note minus the price of a replicating portfolio that matches cash flows exactly, comprised of a bond and the bond's principal STRIP.[4] For example, the red line represents three securities that mature on 2/15/15: a 4% note issued in 2005, an 11.25% bond issued in 1985, and the principal STRIP from this bond. The line shows the price of the note minus the sum of $(4/11.25)$ times the price of the bond and $(1-4/11.25)$ times the price of the STRIP, which is the net revenue from buying the note and shorting a portfolio with identical cash flows. Beginning in August 2007, the price of the note rises relative to the price of the replicating portfolio, and this price difference continues over the next two years to follow the familiar contours of the crisis.[5] The pattern repeats across all pairs of notes and bonds extant at the time. Panel B shows the average price difference across all of the maturity-matched bonds and notes in our sample.

---

[3] An independent market research report from February 2014 cites TradeWeb as the leading electronic U.S. Treasury trading platform. Greenwich Associates (2014) estimates that 25 percent of total fixed income trading volume is executed electronically, and that 45 percent of institutional investors use electronic platforms for at least part of their fixed income transactions.

[4] STRIPS stands for Separate Trading of Registered Interest and Principal of Securities. These are single cash flow securities formed from the individual coupon and principal components of Treasury coupon securities.

[5] A short time-series animation of the cross-section of actual Treasury yields illustrates the magnitude and systematic nature of the price divergences: http://finance.wharton.upenn.edu/~kschwarz/movie.html.

The implicit arbitrage is significantly profitable even net of repo-market financing costs. When we assemble the components of the arbitrage from a database of concurrent repo transactions, we find that funding the short side of the trade offsets only a small fraction of the gross profit, so the explicit costs to establish the trade are small change relative to the price divergence.

We analyze the divergence in three stages, the first of which identifies the cross-sectional determinants of illiquidity discounts. We start by subtracting from each security's price the price implied by a smoothed yield curve, to then relate the cross section of relative prices to the underlying liquidity drivers, such as issue size, and the manifestations of liquidity, such as bid-ask spreads.

We find a strong negative relation between measures of a security's liquidity and the security's relative price. Securities that are older, have smaller outstanding quantities, have been stripped less, have lower trading volume, and higher bid-ask spreads, all of which distinguish bonds from notes, trade at larger discounts. These relations become substantially stronger when market-wide liquidity declines, as happened during the crisis. Controlling for the aggregate state of liquidity over time, these security-specific liquidity variables account for most of the widening of the bond/note price gap.

The second stage of the analysis identifies the investor clienteles attracted to the more- or less-liquid securities. The key database for this identification is the universe of transactions in Treasury securities by all U.S. insurance companies, which we combine with extensive information on the insurers' financial conditions. This database allows us to follow each insurer's activity over twelve years, including the date, direction and identity of each security purchased or sold, while tracking the insurer's changing circumstances, such as its leverage. Collectively, insurers hold about three percent of the stock of U.S. Treasuries,[6] so they are not the largest aggregate clientele for these securities, but their circumstances vary in ways that we can observe, and that can influence their relative demand

---

[6] The Federal Reserve's Financial Accounts data, Table L.209, shows the combined holdings of Life Insurers and Property and Casualty Insurers, as of 12/31/2008, to be $171 Billion, compared with the $6.1 Trillion of publicly held Treasury securities (other than savings bonds) then outstanding (see http://www.federalreserve.gov/releases/z1/current/ z1r-4.pdf).

for more or less liquid assets. Thus, we can identify the investor characteristics and circumstances driving demand for the more liquid securities at their premium prices.

Our principal result is that the insurers with more appetite for liquidity tilt increasingly toward the more liquid and expensive securities as their premiums increase. Those paying up for the liquid securities include those investors that generally trade more, and those with financial circumstances, such as high leverage, that indicate a need for more liquid investments. This result bears out the prediction of Amihud and Mendelson (1986) that less-liquid securities attract investors with longer horizons (distinct from the relationship between expected returns and liquidity that has been well-documented in the literature), and because the selection increases with the premium price, it also bears out the view that the selection powered a feedback loop, amplifying the liquidity difference.

The third stage asks how the liquidity difference came to command such a high price in the crisis. We answer this question by decomposing the cross section of bid-ask spreads into two prices charged by the liquidity provider: the price charged to finance the position until the next trade, and the price charged to bear the position's interest rate risk over that time. We can do this because we have a large cross section of securities that differ in interest rate risk and also an excellent proxy for the expected time to the security's next trade, i.e. the time since its last trade. We achieve the decomposition by regressing the cross section of spreads on the securities' durations and the times since their last trades. Based on the estimated coefficients, the regression shows a sharper spike in the price charged for financing a position over a longer expected holding period than the price charged for bearing interest rate risk. That is, liquidity providers showed a particular aversion at the peak of the crisis to expanding their balance sheets to bridge the arrival of trades, and a less pronounced increase in aversion to bearing the securities' price risk over this period. These prices fall and stabilize after the crisis, but while the price charged per unit of financing has remained stable, the quantity of financing risk has increased, i.e. more time elapses between trades, and this has widened average

spreads and the difference in spreads between more and less liquid Treasury securities. This effect persists until the end of our sample period.

Putting the evidence together, the Treasury market experienced a large price divergence that presented a substantial arbitrage opportunity, net of funding costs, to willing traders. Differences in liquidity that normally command small prices came to command very high prices, which resulted in notable differences in bid-ask spreads within the market for off-the-run Treasury securities. Contributing to this amplification was a feedback loop driven by the selection of longer-horizon traders into less-liquid securities that intensified as the liquidity premium increased. The resulting differences in liquidity premia led to large differences in prices for identical securities.

2. Background and Data

In this section we first review the existing literature that relates liquidity to asset prices and discuss the findings on price discrepancies in the Treasury market. Next, we introduce the data that we use for our analysis. We describe our dataset of secondary-market Treasury security prices, which we complement with additional security-level data to form proxies for relative liquidity. We also describe the secondary market trading activity and portfolio holdings of Treasuries for U.S. insurance companies, which we use to gain insight into the trading activity of a large end-buyer of Treasuries.

2.1. Literature Review

Amihud and Mendelson (1986) provide a formal treatment of the relation between bid-ask spreads and security prices. Since the bid-ask spread represents a cost of trading a security, it affects the ultimate return realized by investors. Amihud and Mendelson (1986) show that expected asset returns are increasing in the bid-ask spread, as investors demand a discount to buy securities that are more costly to trade. This effect persists when investors have different holding period horizons, but

it is somewhat muted. Clienteles with a long horizon rationally purchase securities with higher bid-ask spreads, since the longer horizon makes the higher expected return attractive relative to the higher trading cost. Conversely, clienteles with a short horizon purchase securities with a low bid-ask spread, since the shorter horizon makes the lower trading cost more attractive. This selection of investors results in an increasing and concave relation of expected returns to bid-ask spreads, which Amihud and Mendelson (1986) confirm empirically using stocks traded on the NYSE.

Dow (2004) endogenizes bid-ask spreads in a model with clienteles that vary based on private information and shows that there are strategic complementarities between investor security demand and the bid-ask spreads posted by market markers. Investors with less private information congregate into securities with lower spreads, which are justified due to the lack of adverse selection risk. Conversely, investors with more private information choose securities with higher spreads to capitalize more fully on their private information, which justifies the higher spreads and crowds out less informed traders. Dow (2004) shows that multiple equilibria are possible, with similar securities having either low or high liquidity. We conjecture that a similar feedback mechanism could operate through heterogeneity in investors' trading horizon.

Amihud and Mendelson (1991) empirically examine the relationship between bid-ask spreads and yields on Treasury securities. They compare the pricing of Treasury bills, which have an original maturity less than one year, and Treasury notes that have only six months remaining until maturity and thus are past their penultimate coupon payment. Both are pure discount Treasury securities, so a straightforward arbitrage relation exists. They find that the notes have wider bid-ask spreads and higher yields, so they allow profits for traders who do not try to unwind their positions before maturity. Thus the authors conclude that this is an illiquidity premium available to patient capital.[7]

---

[7] Taxes also explained some of the variation in prices at that time (Kamara, 1994), but the tax code has been changed to remove this differential treatment of notes and bills.

A related comparison is between on-the-run, i.e. "new" securities, and just-off-the-run, i.e. "old", securities. New securities tend to trade at premium prices, and this is generally attributed to their greater liquidity (e.g. Fleming (2003), Goldreich, Hanke and Nath (2005), Barclay, Hendershott and Kotz (2006), and Pasquariello and Vega (2009)). There is a predictable convergence that patient capital can wait for, in that the new security becomes old, and thus likely to trade in line with the existing old securities, when the next security is floated. However, Krishnamurthy (2002) finds that the cost of borrowing the new bond to short it tends to offset much of the profitability of this trade. Fontaine and Garcia (2012) extend the analysis of how a security's age affects its liquidity by examining other pairs of Treasury securities with the same maturity but different age. They assume that age-driven price differences reflect liquidity, and from these pairs construct an index that is correlated with various measures of funding liquidity. Hu, Pan and Wang (2013) also construct an aggregate measure of liquidity from the Treasury market, which they term "noise," as the square root of the mean squared deviation of individual Treasury yields from a smoothed curve. They interpret the measure as capturing liquidity shocks, which among other things show a local peak around the collapse of Bear Stearns and the global peak after the Lehman bankruptcy apparent in Figure 1.

Matched-maturity comparisons are also possible between other securities that share the same creditworthiness. One such comparison is that of Refcorp bonds, which arose from the Savings and Loan crisis, with Treasuries, both of which are backed by the full faith and credit of the U.S. Federal Government. Longstaff (2004) shows that Treasury securities often trade at a large premium over same-maturity Refcorp issues, and attributes this premium to current and expected liquidity. In the euro-area, Schwarz (2016) shows that the yield spread between comparable German federal government and KfW agency securities, which have an identical guarantee, is a real-time, tradable proxy for market liquidity and liquidity risk. Also, inflation-protected Treasury securities

economically equate to regular Treasury securities paired with inflation swaps, but as Fleckenstein, Longstaff and Lustig (2014) show, their prices moved far apart in the crisis.

There is also a literature examining the relationship between liquidity and yields of corporate bonds, which includes Bao, Pan and Wang (2011) and Dick-Nielsen, Feldhutter, and Lando (2010). It is more challenging to precisely pin down the liquidity component in yields of corporate bonds due to the influence of credit risk on bond prices. Liquidity tends to correlate with creditworthiness in the corporate market, making the effect of liquidity difficult to distinguish completely from the effect of creditworthiness.

To summarize, there is generally a liquidity premium in securities markets, including in the Treasury market. Newer securities are more liquid and expensive than older, and the newest securities, the on-the-run, are the most expensive of all, due presumably to the intense focus of traders on these particular securities. Our analysis avoids confusion with this on-the-run effect by focusing exclusively on the off-the-run Treasuries.

2.2. Data

Treasury Transactions on the TradeWeb Platform

We construct a dataset of daily observations for all nominal Treasury securities outstanding, including STRIPS. We begin with bid and ask price quotes from TradeWeb, a large electronic trading platform that specializes in customer-to-dealer trades of fixed income securities. We have this data from May 3, 2004 through September 30, 2011. We use TradeWeb, rather than the usual interdealer platforms (eSpeed, BrokerTec and GovPX) because the customer-to-dealer market is the primary venue for off-the-run Treasury market transactions. These data capture the trades of a variety of market participants (e.g. hedge funds, pension funds, insurers) with differing motives for trading. The

database flags the dates that each security is on-the-run, and we exclude those security-date pairs from our analysis.

The TradeWeb data report outstanding quotes averaged across market makers at four moments each day: 8:05 AM, 3:00 PM, 4:00 PM and 4:45 PM (U.S. Eastern time). From June 2008 onward, the data also report for each security the time since the last buy and the last sell. We take the average of these two times and label the result *TTT* (i.e. time-to-trade). For our daily observations we choose the 3:00 PM snapshot for two reasons. First, the time since the last trade is likely to be more informative as an indicator of a security's liquidity near the end of the trading day when most trades have occurred. And second, the 3:00 PM snapshot is likely more representative of general intraday liquidity as it is still within the window of high intradaily trading volume and moderate bid-ask spreads, before an end-of-day deterioration in market conditions (Fleming and Remolona (1997)).[8]

We complement the TradeWeb data with a database of specials repo transactions from a large interdealer broker. Our repo sample contains all transactions intermediated by the broker from April 1, 2004 through March 1, 2009. For each repo we see the identity of the specific security that served as collateral for the repo, the repo rate, the trade date, the settlement date, and the term. Not all securities are repoed on each day; on average, twenty percent of the maturity-matched bonds and notes have observable repo rates. To this repo database we add the daily general collateral (GC) rate reported by Bloomberg.

Additional Liquidity Variables

In addition to *TTT* we calculate three additional measures of a security's current liquidity. From the TradeWeb data we have the percentage bid-ask spread, *Bid-Ask*, which is $\frac{Ask-Bid}{(Bid+Ask)/2}$ as of 3:00 PM. For volume and relative trading volume, there is no publicly available data for trading of

---

[8] None of our results are sensitive to this choice.

individual Treasury securities, so we use the insurance data described below. For each security on each day, we construct *Volume* for each security on each day as the sum of buys and sells of the security by insurers, scaled by the sum of all Treasury security trades by insurers that day.

We gather data on security characteristics from the Treasury Department's Monthly Statement of the Public Debt. We collect each security's initial size and the fluctuation over time due to repurchases and re-openings. We also see fluctuations in the effective issue size due to stripping and reconstitution. These data are all end-of-month, which we match with the corresponding month of our daily variables. We construct *ln(Out)* as the log of the principal outstanding and *Share Stripped* as the share of the security's principal outstanding held in stripped form. From the issuance and maturity dates, we calculate *ln(Age)* as the log of the time since issuance and *ln(TTM)* as the log of the time remaining until maturity. The last two variables are computed daily.

Table 1 reports summary statistics for the liquidity variables, along with some additional information about the securities, and it highlights significant variation across securities with different original-issue maturities. 30-year bonds, in particular, are notably less liquid by all of the measures. The average quoted percentage bid-ask spread is 3 basis points on average across all securities, but nearly 7 basis points for bonds alone. Bonds also trade far less frequently than other Treasuries; their average *TTT* is almost 3 days, compared to 8 hours for a 10-year note. Bonds are also much smaller, older, and more stripped, on average, than other securities. Consistent with prevailing interest rates at their time of issue, bonds also have higher coupons and shorter duration, on average, than notes.

Trading and Holdings Data for Insurance Companies

Our final dataset reports the transactions in and holdings of Treasury securities by U.S. insurance companies, from January 1, 1998 through December 31, 2011. U.S. insurance companies are required to report every purchase and sale of a Treasury security, indicating the date, size, and

direction of each transaction, and to report their quarter-end holdings of each security. These data are packaged and resold by eMaxx.[9] We limit our sample to the 2,321 insurers that show holdings and transactions in coupon Treasury securities, including STRIPS, and that also can be matched with accounting data from SNL Financial over the same sample period as our Treasury price data, from 2006 through 2011.[10]

For each insurer we calculate two statistics summarizing its trading activity over the sample period. To calculate how long an insurer tends to hold a position, we calculate *Holding Horizon* as the volume-weighted average number of days between the purchase and sale of each Treasury security that the insurer trades.[11] The average *Holding Horizon*, reported in Panel A of Table 2, is about 2.5 years. For a measure of portfolio turnover we compute *Churn*, which is an insurer's average trading volume per month in Treasury securities divided by its average holding of Treasury securities.[12] The mean value of *Churn* is 0.08 in Panel A of Table 2 indicates that the typical insurer trades about 8% of the portfolio each month. Since *Holding Horizon* tends to decrease and *Churn* to increase as an insurer trades more, they are negatively correlated across insurers; Table 2, Panel B, reports a correlation of -0.34.

From the SNL Financial accounting data we construct several measures of insurers' financial strength and liquidity. *Capital/Assets* is a measure of the leverage of the insurer, where *Capital* is an

---

[9] The original data from eMaxx is at the level of insurer-date-CUSIP-investment advisor, and we aggregate the data to insurer-date-CUSIP by summing transactions that have identical insurer, date, and CUSIP. eMaxx is a Reuters subsidiary that obtains the source data from the statutory filings of regulated insurance companies.

[10] We construct our insurer trading style variables with the 1998 through 2011 sample period, but for all of our analysis, we use the 2006 through 2011 sample period.

[11] Holding horizon measures the average number of days that insurer j holds a security i in its portfolio. $Horizon_j = \frac{\sum_{i=1}^{I}\left(Q_{i,j}Days_{i,j}\right)}{\sum_{i=1}^{I}Q_{i,j}}$ where $Q_{i,j}$ is the quantity that is both bought and subsequently sold of security i for insurer j, and $Days_{i,j}$ is the holding horizon in days for security i and insurer j. We drop observations for which there is not a matching purchase and sale in our sample.

[12] $Churn_j = \frac{\sum_{t=1}^{T}\sum_{i=1}^{I}\left(Buys_{i,j,t} + Sells_{i,j,t}\right)}{\sum_{t=1}^{T}\sum_{i=1}^{I}Holding_{i,j,t}}$, where $Holding_{i,j,t}$ is the level of holdings of CUSIP i for insurer j on day t.

accounting measure of policyholder surplus, and *Assets* is the book value of total assets.[13] *Annuity* is an indicator variable equal to 1 if the insurer sells more annuities than life insurance or property/casualty insurance, and equal to 0 otherwise. As discussed in McMenamin et al. (2012), annuities provide customers with some ability to withdraw their savings, which creates liquidity risk for insurers. *RBC* is an indicator variable equal to 1 if the ratio of actual capital to risk-based capital, measuring the insurer's capital adequacy, is greater than the median for all insurers in our sample, and 0 otherwise.[14] Finally, *Net Income* is an indicator variable equal to 1 if the insurer's net income is positive, and 0 otherwise.

3. Explaining Relative Price Differences in the Cross-Section

3.1. The Divergence and the Arbitrage

In this section we document the divergence and gauge the cost of trading on it, which is determined by the relevant repo rates. In the days after the Lehman bankruptcy, there are additional considerations arising from the high propensity of delivery fails.

On 11/20/08, Treasury coupon securities closed the day at the yields displayed in Figure 2. The figure highlights nine bonds (shown with blue asterisks) that share a maturity date with a note (black dots). All of the matching notes are 10-year Treasuries. Because their maturity dates match, they pay principal and coupons on the same dates. But, the bonds date from an era of higher rates, and thus have larger coupons, which imply lower yields in an upward sloping yield curve. However, the bond yields on this date are instead much higher. So the bonds appear to be significantly cheap relative to the notes, and the coupon difference makes the differential even more surprising.[15]

---

[13] The vast majority of insurers are private, so it is impractical to use market values.

[14] Merrill, Nadauld, Sherlund, and Stulz (2012) show that insurers with below-median levels of regulatory capital were more likely to sell securities at fire sale prices during the recent financial crisis.

[15] The bonds with more than ten years remaining to maturity follow a curve largely consistent with the nine highlighted bonds, but the absence of matching notes prevents a precise comparison of their prices.

For an exact comparison we construct an exact match by using the bond's principal STRIP. The cash flows of a note with coupon C_N are replicated by a portfolio with (C_N/C_B) of a bond with coupon C_B and $(1-C_N/C_B)$ of the bond's principal STRIP.[16] We subtract the price of this replicating portfolio from that of the note (using the midpoint of bid and ask prices), yielding the time series plots in Figure 1 which show the large divergence.

Did this divergence offer a profitable trading opportunity? To trade on the divergence, one must short the note, buy the matching portfolio of the bond and its STRIP, and wait for convergence. Shorting the note is generally accomplished through the repo market by lending cash and accepting the note as collateral. The interest rate earned on the loan depends on the demand to borrow the specific security. For generic securities that are not in high demand, the repo rate is the GC rate. For securities in high demand – those deemed "special" – the repo rate will fall below the GC rate. The bond purchase can also be financed through the repo market, where the GC rate is the highest rate one would generally pay. The net cost of financing the arbitrage trade is the difference between the GC rate paid to finance the purchase of the bond and the potentially special rate paid to short the note.

We gauge this carrying cost by reference to the rates in the repo database. On each day, we take the repo rate for any of the notes in Figure 1 to be the average rate across the repos of any of those notes on that day (the database does not show a repo transaction for each security on each day).[17] If there are no repos of those notes that day, we take the repo rate for any of the notes to be zero, which is the lowest rate for any repo in the database.[18] So our estimate assumes that repo rates are

---

[16] We use bond principal strips throughout, as note principal strips are relatively scarce (see Table 1).
[17] Our data shed light on financing costs, but there are necessarily some limitations on what they can illuminate due to the short maturity of the typical repo transaction and the much longer term over which an investor might potentially wish to hold the trade. In our database, 82 percent of the transactions are overnight, so we estimate the realized financing costs of a hypothetical investor by stringing together a series of overnight rates. Both Duffie (1996) and Jordan and Jordan (1997) use exclusively overnight repo data, and Keane (1996) states that term repos longer than 90 days are in practice not available.
[18] There are 834 zero-rate overnight Treasury transactions in our repo database, but none with a negative rate.

similar across the notes, which are likely to be special due to the arbitrage opportunity (Duffie (1996)), and also assumes the worst-case scenario when no rates are observed. We further assume the worst-case scenario for financing the long leg by assuming the trader pays the GC rate, i.e. the rate that imputes no scarcity rents because the money borrower is indifferent to which Treasuries are provided as collateral.

The financing rates are plotted in Panel A of Figure 3, which shows the GC rate as the solid black line and note rate as the dotted blue line, throughout our repo sample period. The key quantity, the trade's net funding cost, is the difference between these rates, which is generally small; it is between 0 and 23 basis points on 75 percent of the days in our sample period. It spikes on occasion, but not in late 2008, which is the period of greatest divergence between prices in the cash market. At that time, the latitude for high net funding costs narrowed, as the overnight GC rate declined in line with the policy-driven overnight Fed Funds rate, and the special repo rate did not become negative. The carry cost of the trade was therefore reduced to almost zero. Panel B of Figure 3 plots the cumulative funding cost, from the beginning to the end of the sample period, as implied by the difference in rates from Panel A. The realized financing cost of the trade, from February 22, 2008, the date the average divergence across bond/note pairs first rose above 1 percent, to the end of our sample, long after the divergence subsided, is just 0.21 percent. A trade established at the very peak of the average divergence, on December 16, 2008, and unwound at the end of our sample, would have cost only 0.02 percent. From this we conclude that the financing costs were small relative to the apparent arbitrage opportunity.

The low cost of financing the arbitrage trade at the time of greatest apparent opportunity raises questions, most immediately whether this is really the cost that a hypothetical trader would have paid. Krishnamurthy (2002) shows that the estimated financing costs borne by well-positioned traders are in the neighborhood of the gross profit implied by the on-the-run Treasury price divergence. In a time

of extreme market stress, it stands to reason that some traders would not be able to access financing at all, even in collateralized markets. Thus, the costs in Figure 3 are best characterized as the costs facing a large trader.

Another question raised is why the specials rates do not fall below zero when there is likely strong demand to borrow particular securities. Repo rates are not bounded below by law or any other convention, and indeed, negative rates have occurred after the sample period. The answer likely lies in a trader's alternative to fail to deliver a shorted Treasury, i.e. to simply not deliver the security to the buyer, who would then not pay for it until it is delivered. As Evans et al. (2009) note, a delivery fail is economically equivalent to a repo rate of 0, plus any expected cost of getting bought in, i.e. the expected cost arising from the trader's prime broker buying the security in order to deliver it, thereby ending the short and generating transactions costs.[19] So if these expected costs are sufficiently small, a short-seller would benefit from failing, rather than accepting a negative repo rate. That the alternative to fail was available and attractive is apparent in the incidence of both Treasury repos and Treasury delivery fails (as reported by the Federal Reserve Bank of New York), shown in Panel C of Figure 3. In late 2008, the number of weekly repo transactions in our database reveals a steep drop, suggesting a combination of low demand to arbitrage (e.g. Gorton and Metrick (2012), Mitchell and Pulvino (2012)) and reluctance to arbitrage specifically through the repo market. The steep concurrent rise of fails indicates that at least some of this drop reflects substitution from repos to failing, and the option to fail placed a floor under repo rates that limited the cost of financing the arbitrage.

To summarize, the carry cost we estimate from repo transactions is low relative to the price divergence. The low carry cost in the months after Lehman is likely best representative of the best-situated traders, so there could have been traders who wanted to profit from the divergence but could

---

[19] A new market convention adopted in May 2009 made it much more costly to fail to deliver a Treasury security by imposing a penalty for failing.

not exploit it. The low cost was supported by the option at that time to fail to deliver, which protected traders from negative repo rates. Traders could have expected costs to run higher, and for the divergence to expand or persist more than it did, but to the extent the data can reveal, the relative prices of notes and bonds presented a feasible arbitrage opportunity for patient capital. The next subsection addresses the role of relative liquidity in relative pricing by analyzing the relation between liquidity and prices for the full set of off-the-run Treasury securities.

3.2. Relative Pricing of All Treasury Securities

In this section, we relate the pricing of Treasury securities to their liquidity. We consider both the standard measures of liquidity, such as bid-ask spreads, and the potential fundamental drivers of liquidity, such as issue size and age. The goal is to establish both the general cross-sectional relation and also the change over time through the sample period, which runs from 2006 to 2011, so from before to after the crisis. For this estimation we expand our sample from the maturity-matched pairs to the universe of off-the-run coupon Treasury securities.

To expand the analysis beyond same-maturity pairs, we need to first calculate a baseline price for each security on each day. We do this by first fitting a smooth yield curve to all the securities, using the six-parameter model of instantaneous forward rates developed by Svensson (1994),[20] then computing for each security the price implied by this curve, and finally subtracting this implied price from each security's actual price. We call this difference the security's "fitting error" for that day, denoted FE_{it} for security i on day t. The mean fitting error is by construction close to zero (not precisely zero, since prices are not linear in yields), and the root mean square of the fitting errors on each day is similar to the aggregate noise measure in Hu, Pan and Wang (2013). The fitting errors line up well with the price differences of the exactly-matched pairs in Figure 1; the average fitting

---

[20] See Gürkaynak, Sack, and Wright (2006) for a discussion of the methodology:
http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html

error across 30-year bonds becomes significantly negative during the crisis while the average across 10-year notes becomes significantly positive, and the difference peaks at around 6 percent on the same dates.

We relate the fitting error to liquidity with pooled cross-sectional, time-series regressions. Each independent variable differs across securities in ways that could affect its relative liquidity. We include *Bid-Ask* and *Volume* to capture the variation of realized liquidity across securities and time. Existing research identifies both age (Amihud and Mendelson (1991), Fontaine and Garcia (2012)) and principal amount (Longstaff, Mithal and Neis (2005)) as potential drivers of liquidity, so we include *ln(Age)*, time-to-maturity as *ln(TTM)*, issue size outstanding as *ln(Out)*, and *Share Stripped* in the analysis. While existing research does not, to our knowledge, associate coupon with liquidity,[21] the difference in coupons between bonds and notes is quite large, so we include *Coupon* as a control variable. Panel B of Table 1 shows that age, coupon, and share stripped are all highly correlated, which makes it difficult to distinguish their separate effects in a multiple regression. Accordingly, we run simple regressions on each. Finally, we include an indicator for the maturity-matched bond/note pairs. This *Net Long Matched Pair* indicator is set to 1 for the notes in the pairs and -1 for the bonds, so that the estimated coefficient captures the average difference in prices between the maturity-matched securities.

The regression model takes the general form:

$$FE_{it} = \alpha + \beta Liq_{it} + \varepsilon_{it}$$ (1)

where Liq_{it} is a set of liquidity variables specific to security i on day t. Because *FE* increases with the security's relative price, a positive loading on a variable indicates a liquidity premium associated with

---

[21] Conceivably, high-coupon securities could be bid up by agents investing on behalf of consumers who don't understand the tradeoff between higher current yield and bigger subsequent capital losses (see, e.g., Donnelly (1988)).

that variable. We report the results, in Panel A of Table 3, from a simple regression for each independent variable and three multiple regressions with different sets of the independent variables.

The simple regressions find premia consistent with the literature; securities are more expensive when they are newer, larger and less stripped, and when they have lower spreads or trade more frequently. The price is also higher for securities with lower coupons, which also distinguishes notes from bonds, but the correlation between coupons and age and size makes this hard to interpret. In the multiple regressions, most of the premia persist but the effects of coupon and age are not stable, potentially due to their high correlation; older bonds issued in the 1980s have relatively higher coupons than notes that have the same time remaining to maturity. The note/bond price difference detected by the matched-pair indicator alone is cut by more than half with these dimensions of liquidity controlled for.

The model in Panel A fits static coefficients and thus captures the *general* relation of the liquidity variables to prices in the cross section. However, we are interested in understanding how differences in liquidity contributed to the widening of price differences during the crisis. Panel A of Figure 4 suggests that the impact of fundamental differences strengthened during the crisis. The figure plots the average fitting error by a security's size quartile. The figure illustrates the relationship captured by the regression results; on average, securities in the smaller quartile have lower prices. However, the difference in prices in the entire cross-section of securities varies notably over time and shows a pattern quite similar to that of the bond/note pairs alone.

To capture time-series variation in the contribution of liquidity to prices, we allow the coefficients in equation (1) to vary over time with an index of market-wide liquidity. A positive loading on this interaction means that the variable's contribution is bigger when aggregate liquidity is worse. Our index, labeled $AggLiq_t$ for day t, is the average bid-ask spread across the universe of securities that day, shown in Panel B of Figure 4. The average spread rises from around 2 basis points

pre crisis to much higher levels around Lehman before drifting back down, though not all the way to where they started. This pattern is stronger among the less liquid bonds, whose spreads exceed 20 basis points at the widest. These Treasury-market results are consistent with the equity-market results in Acharya and Pederesen (2005), where the most illiquid securities show higher sensitivity to changes in market liquidity.

The adapted model allows the equation explaining relative prices to vary across days by allowing the coefficients to vary through interaction with the index, and allowing the intercepts to vary through daily fixed effects. Thus, we fit the model

$$FE_{it} = a_t D_t + bLiq_{it} + \gamma Liq_{it} * AggLiq_t + \varepsilon_{it}$$

where D is an indicator for day t, and Liq is a set of liquidity variables. We repeat each regression from Panel A with this expanded specification; results (omitting the coefficients on the daily indicators) are in Panel B of Table 3.

The specifications on the left-hand side of Panel B find a consistent result; the interactions of the individual liquidity variables are significant in the same direction as the uninteracted coefficients in Panel A, and their importance increases. As aggregate liquidity worsens, these variables explain even more of the general relation between liquidity and pricing. For example, the coefficient on *BidAsk* in Panel A is -3.42, so on average, a 1 basis point wider spread associates with a 3.42 basis point lower price. In Panel B, the coefficient on the interaction of *BidAsk* with *AggLiq* is -0.34. Since *AggLiq* increases by about 6 basis points during the crisis peak (from 2 basis points pre-crisis to 8 at the peak of the crisis), this implies that the price discount associated with a 1 basis point wider spread increases by $(6)(0.34) = 2.04$ basis points at the peak of the crisis.

In the multiple regressions, issue size and share stripped stand out as contributors to the divergence of prices. We can get a back-of-the envelope sense of this contribution by combining the

point estimates from the regression with the summary statistics from Table 1. The average amount outstanding is 12.08 for bonds and 28.24 for 10-year notes, and $ln(12.08)-ln(28.24) = -0.85$. So the coefficient of 10.5 on the interaction of *AggLiq* with *ln(Out)* associates the 6 basis point increase in *AggLiq* with a $(10.5)(6)(-0.85) = 54$ basis point decrease in the price of bonds relative to 10-year notes. Also, 21 percent of the average bond is stripped, compared to just 2 percent of the average 10-year note, so the increase in *AggLiq* implies a $(-48.84)(6)(0.21-0.02) = 56$ basis point decrease in bond prices relative to 10-year notes.

These liquidity variables explain almost two thirds of the bond/note disparity. In the regression including only the bond/note pair indicator, the 32.36 loading on the indicator implies a divergence of 3.9% due to the decrease in aggregate liquidity (i.e. $32.36(6)(1-[-1]) = 3.9\%$). However, when we include the liquidity variables in the regression, the loading falls to 12.32, which implies just a 1.5% divergence not explained by these liquidity measures.

To summarize, we find that relative liquidity – as reflected in both the underlying drivers of liquidity such as issue size and the manifestations of liquidity such as bid-ask spreads – explains much of the general price difference in the cross section of Treasury securities, and in particular, explains most of the widening of prices to their extreme disparity during the crisis. In the next section we address the dynamics of this widening, in particular the feedback from the trading horizons of investors buying the different securities to the securities' subsequent liquidity.

4. Liquidity Feedback Generated by Investor Trades

With the insurer data, we can line up an investor's decision to buy or sell notes or bonds with, on the one hand, the securities' current relative prices, and on the other hand, the investor's general trading style as well as its current financial stress. We can therefore see how the demand for the more liquid security evolves as the price of liquidity changes and as the investor's liquidity needs change.

We gauge an insurer's trading style from its trading history and its financial stress from its balance sheet. The trading variables are designed to reflect liquidity needs arising from an insurer's style: *Churn* increases with an insurer's portfolio turnover and thus likely reflects its general demand for liquid securities, and *Holding Horizon* increases with an insurer's propensity to hold onto its investments, and thus likely decreases with its demand for liquid securities. The financial-stress variables include *Capital/Assets*, *RBC* and *Net Income*, which all increase with the health of the insurer and thus point to lower stress, and *Annuity*, which indicates that the insurer is potentially subject to the need for emergency liquidation of securities. As in Coval and Stafford (2007) and Merrill, Nadauld, Sherlund, and Stulz (2012), the financial stress captured in variables such as these likely increases the insurers' demand for liquidity in their investments.

The empirical question we address is how an insurer's relative demand for a security varies with the security's fitting error, which provides a useful estimate of a security's relative price. Accordingly, the dependent variable is the net purchase of security *i* by insurance company *j* in month *t*, expressed as a fraction of the total principal amount of the security, i.e. $NP_{i,j,t} = (Buys_{i,j,t} - Sells_{i,j,t})/Outstanding_{i,t}$. We aggregate across the days in each month due to the sparseness of off-the-run trades. Because we aggregate the trades within the month, we also average the fitting error within the month, so the regression takes the form

$$NP_{ijt} = \alpha + \beta FE_{it} + \varepsilon_{ijt}. \tag{5}$$

As before, *FE* is in price terms so a positive value indicates a relatively rich security. We estimate the regression first for the entire set of insurers, and then for each of the three insurer types separately. The regressions also include calendar month-of year fixed effects and insurer-type fixed effects.

Table 4 shows the parameter estimates for equation (5). The first column indicates that insurers as a group demand a Treasury security more when its price is relatively high. The next three

columns show that this result holds for each of the three classes of insurers, although the effect appears largest for the life insurers, showing an 8.27 coefficient. To give economic interpretation to this estimate, we aggregate over all 571 life insurers in our sample (Table 2), and find that a one percentage point increase in a security's relative price leads to a \$473,044 increase in net purchases of the security (571×827), for every \$1 billion outstanding of the issue, or roughly 0.05 percent of the total issue. This means that, on balance, insurers traded in the direction that would widen, rather than narrow, the price differentials among different Treasury securities.

To identify the role of the insurers' individual circumstances in their relative demands, we modify equation (5) to allow the intercept and slope coefficient to vary for each insurer, so that the equation becomes

$$NP_{ijt} = \alpha_j + \beta_j FE_{it} + \varepsilon_{ijt} \tag{6}$$

We moreover assume that the intercept and slope are linear functions of the trading-style and financial-stress circumstances of the insurer, collected in a vector, χ_j:

$$\alpha_j = a + b\chi_j, \quad \beta_j = c + d\chi_j \tag{7}$$

Substituting equation (7) into equation (6) yields

$$NP_{ijt} = a + b\chi_j + cFE_{it} + dFE_{it}\chi_j + \varepsilon_{ijt} \tag{8}$$

which we then estimate in a pooled regression. The main object of interest is the coefficient on the interaction term, *d*, as it tells us whether an insurer's particular circumstance makes it more or less likely to buy the relatively expensive securities.

The results from estimating equation (8) are shown in Table 5, with the insurer characteristics, χ_j, entering individually in the first six columns, and together in the last column. Among the individual results, the trading-style variables associate liquidity demand with buying the relatively expensive securities. Investors with shorter average holding periods (*Holding Horizon*) and more frequent portfolio turnover (*Churn*), both show an insurer's net purchases increase with relative price. The estimates on the insurer's financial variables also associate higher stress with demand for the more expensive securities; higher leverage and exposure to annuities both relate significantly to higher demand for liquid securities. In the multiple regression, these variables retain the sign and significance as in the simple regressions. Moreover, controlling for these characteristics drastically reduces the size of the coefficient on the fitting error, suggesting that the relationship between trading and prices is concentrated in certain types of insurers. To consider one comparison, the sensitivity of purchases to the fitting error is, at the point estimates, about three times greater for insurers with annuity exposure than for those without.[22]

These results show liquidity clienteles concentrating in the respective groups of securities as the price of liquidity grows. This gives new and direct empirical evidence of the clientele prediction from the Amihud and Mendelson (1986) model of the bid-ask spread; assets with higher bid-ask spreads are allocated in equilibrium to portfolios with the same or longer expected holding periods.[23] It is also in the spirit of Brunnermeier and Pedersen (2008), in that their model connects varying investor constraints to rapid increases in the price of liquidity.[24]

---

[22] This calculation is based on comparing the coefficient estimate on interaction of *Annuity* and the fitting error, and the coefficient on the fitting error alone.

[23] The clientele prediction differs from the positive and concave relationship between an asset's bid-ask spread and its return, which is consistent with our finding in section 3, equations (1) and (2), and has already been empirically tested in the literature.

[24] The Brunnermeier and Pedersen (2008) model relies on differences in margin requirements between securities. Krishnamurthy (2010) shows that repo haircuts for short-term Treasury securities averaged 2 basis points from the Spring of 2007 through the Spring of 2009, and haircuts for long-term Treasury securities ranged from 5 to 6 basis points over the same period.

To summarize, the insurance-company transactions show that investors select into securities according to their need for liquidity, and that this accelerates as the price of liquidity increases. Self-selection of insurers that trade less frequently into less liquid securities would tend to lower the securities' liquidity even further as they become even less traded. This is consistent with a feedback loop where an illiquidity discount begets a larger discount, due to the expectation that the security will trade less, and this expectation should correspondingly widen bid-ask spreads. For this mechanism to work, spreads should widen over time as trading slows down. Thus, the last empirical section asks how the changing rate of trading affects the change in bid-ask spreads over time.

5. Liquidity Provider Risk Pricing

A security's bid-ask spread is the price charged by liquidity providers to absorb temporary imbalances in security supply and demand. Panel B of Figure 4 shows a near quadrupling in the average bid-ask spread from before the crisis to its peak, and Panel A of Figure 5 shows a particularly dramatic spike in the average bid-ask spread for 30-year bonds as compared to 10-year notes. This variation across securities and over time raises two main questions. First, how much of the spike reflects a higher price charged to provide the same liquidity, and how much represents a larger quantity of liquidity being provided? And second, which elements of liquidity provision were responsible for the high spreads? Is it the price of carrying the position on the balance sheet until the next trade, or the price of bearing the security's interest rate risk exposure? We address these questions by examining the evolution of the relationship between bid-ask spreads and the two fundamental costs of providing liquidity: securities' interest rate risk and the expected time between trades.

A market maker builds the expected costs of its services, including financing a position and bearing its price risk, into the bid-ask spread charged to customers. We use two variables to capture these costs. The first is time-to-trade, *TTT*, as defined in Section 2.2. The realized time since the last

trade serves as a proxy for the expected time that the liquidity provider must wait until another trade arrives to take the position off his hands. We define *TTT* as the average of the time since the last sell and the time since the last buy to reduce the noise in this expectation, while still allowing it to pick up high-frequency changes in the trading environment. The second variable is the duration of the security, which represents the price risk per unit of time borne by the liquidity provider between trades, which we denote Dur_{it}. The intuition is that the cross section of the securities' interest rate risks is roughly proportional to the securities' durations.

Panel B of Figure 5 plots the monthly average levels of *TTT* and *Dur*. There is a strong upward trend in the time between Treasury transactions, more than doubling from mid-2008 to mid-2011. The drop in transaction frequency could result from either a general drying up of liquidity, or a drying up on this particular platform, TradeWeb. The database does not reveal TradeWeb's market share, but news reports and anecdotal evidence point strongly to a general trend of worsening market conditions, while the market share of electronic platforms such as TradeWeb has likely expanded.[25] The time between trades reaches a local peak around the end of 2008, as does the level of interest rate risk, which spikes again in mid-2009.

We relate bid-ask spreads to the two variables in the following regression model

$$Spread_{it} = \alpha_t + \beta_{1,t}TTT_{it} + \beta_{2,t}Dur_{it} + \varepsilon_{it} \quad (5)$$

where $Spread_{it}$ denotes the bid-ask spread for security *i* on day *t*. To allow the relation to change over time, we run a separate cross-sectional regression for each day *t* in the sample. The coefficient on

---

[25] For example, Bloomberg reported on March 29, 2015 (Kruger and McCormick, 2015), that Treasury trading by primary dealers shrank from 13 percent of the amount outstanding in 2007 to just 4.1 percent in 2015, and also that "…electronic trading of Treasuries has proliferated over the past decade and now accounts for 44 percent of all transactions." A July 13, 2015, Wall Street Journal article (Burne, 2015) cites data from the Securities Industry and Financial Markets Association (SIFMA) that show it would take 25 days for the entire universe of Treasury securities to turn over, compared to 8 days in 2005. So, there is abundant evidence that the increasing time between the electronic trades at TradeWeb reflects a general downward liquidity trend, rather than a trend away from electronic trading platforms such as TradeWeb.

TTT, $\beta_{1,t}$, provides an estimate of the price charged per unit of time, on date *t*, to finance a position until the expected arrival of next trade. The coefficient on Dur_{it}, $\beta_{2,t,}$ estimates the price charged over this same holding period to bear the position's interest rate risk. We present the results graphically in Figure 6, which has one bold line for each of the two coefficients along with lighter dashed lines showing confidence bounds. The graph shows monthly averages of the coefficient and confidence bounds from the daily regressions.

The time series of coefficients make two main points. First, the price charged per unit of financing time ($\beta_{1,t}$) surged when bid-ask spreads surged; both peak in September 2008, when Lehman failed, and both are higher in the first half of 2009 compared to the second half. Second, the price charged per unit of holding-period price risk ($\beta_{2,t}$) increased in late 2008 and again in early 2009, which contributed to the increase and slow decline of spreads in 2009, particularly for longer duration securities.

So to the question of why spreads increased so much, at the time when prices diverged so much, these results deliver an answer with two parts. First, the price charged by market makers for financing, i.e. for access to their balance sheets, spiked dramatically after the Lehman bankruptcy and again in early 2009. In combination with a rise in the quantity of financing needed, due to the slowdown in trading (Figure 5B), bid-ask spreads increased for particularly illiquid securities. Second, the price charged for interest rate risk rose, which increased spreads for more risky securities.

The spike in the price charged for financing can help explain the divergence between bond and note prices.[26] Not only did bonds trade less often than similar notes, the price charged for the lack of liquidity increased. To put the effect in perspective, in the first month of our sample (June 2008), the average estimated coefficients on *TTT* is 0.15, and at the peak of the bid-ask spread (March

---

[26] Interest rate risk cannot explain the difference in bid-ask spreads since the replicating portfolios have identical duration by construction.

2009), the coefficient becomes 1.58, a more than 10-fold increase. Based on the sample-average value of *TTT* (0.92 days), the increase in the coefficient implies a 1.3 basis point increase in the bid-ask spread for an average security. Combined with a widening of the gap in *TTT* between bonds and notes shown in Figure 5B, the estimates suggest that the difference in bid-ask spreads between bonds and notes reached about 4 basis points due to a slowdown in trading and an increase in the price charged by market makers to bear this risk. As shown in Figure 5A, the actual gap in bid-ask spreads reached nearly 15 basis points in early 2009, which suggests that additional factors beyond those in our model – such as alternative liquidity differences or scope for informed trading – also contributed to differences in bid-ask spreads across the bonds and notes.

There is also a notable change in the price to bear interest rate risk, as the coefficient on *Dur* roughly triples between the start of the sample and early 2009. This increase accounts for about three-quarters of the increase in the average security's spread, from trough to peak, shown in Figure 4B. The remainder is due to the increases in average *TTT* and the coefficient on *TTT*.[27]

In the June 2008 to March 2009 sub-sample, the regression results suggest that most of the changes in the distribution of bid-ask spreads are due to changes in prices rather than changes in quantities of financing or interest-rate risk. Panel B of Figure 6 shows the average predicted bid-ask spread for all 10-year notes in our sample as compared to the average predicted spread for all 30-year bonds. Bid-ask spreads for the less liquid securities are clearly higher, and this divergence increases in the crisis period.[28] The larger contribution of price vis-à-vis quantity is consistent with the empirical facts reported in the Federal Reserve Bank of New York's Treasury market liquidity overview (Adrian, 2013), which states "bid-ask spreads increased dramatically during the financial crisis, despite the sharp increase in volume. The increase in bid-ask spreads reflect uncertainty at that

---

[27] The duration of the average security remained roughly constant during the sample period.
[28] Since this analysis includes securities other than the maturity-matched pairs, the 30-year bonds have longer average duration, which contributes to the difference in predicted bid-ask spreads.

time, as well as the reduced balance sheet capacity of dealers." However, our results suggest that the increase in spreads had a more pronounced effect on illiquid securities, which resulted in spreads increasing more for 30-year bonds relative to 10-year notes.

The estimates shown in Figure 6 also shed some light on why average bid-ask spreads end up higher in late 2011 than they started before the Lehman bankruptcy, as shown in Figure 4. The estimated coefficients on *TTT* and *Dur* largely returned to their pre-crisis levels by the end of 2009, which we interpret as a reversion of prices charged for the *risk* of slower trading frequency and changes in interest rates to their pre-crisis levels. The rise in average bid-ask spreads is attributable, therefore, to the steady increase in *TTT* that happened over the period, particularly for 30-year bonds. The "new normal" environment appears to entail larger bid-ask spreads due to the higher costs of intermediation in the post crisis environment.

6. Conclusion

The importance of better understanding market liquidity pricing remains paramount amid episodes of extreme changes in liquidity; trading can concentrate in the most liquid securities and effectively evaporate from the least liquid securities. Normal arbitrage relationships have resumed since the crisis. However, unusual and sudden price movements continue to affect not only the Treasury market, but also corporate debt, equities, and derivatives. Market liquidity is often cited as an explanation, but the precise mechanism that relates a security's liquidity to its price, and the way in which the effects of this relationship can be exacerbated, is not well understood. Our setting is ideal for understanding this relationship since Treasury securities vary in few other dimensions, and we have straightforward methods for accounting for differences in the timing of cash flows.

Our analysis makes several contributions to the understanding of liquidity pricing. The first is simply the unprecedented magnitude of the price effect among perfectly-matched securities with

equal and nearly perfect credit quality, all within a single market. The low explicit cost of executing the arbitrage through to convergence points to a role for frictions and/or implicit costs affecting the security choices of investors. Second, with our data on investor transactions, we ask what drives the demand for the more-liquid securities and find that investors sort into securities based on their own expected need for market liquidity. As predicted in the model of Amihud and Mendelson (1986), investors with a shorter horizon amortize any transaction costs over only that timeframe, and hence have higher demand for more liquid securities. Combined with market makers who rationally price their services according to the risks that buying or selling a security imposes, the sorting of investors creates a liquidity feedback effect, as in the model of Dow (2004). Investors that expect to transact more frequently are attracted to securities with low transactions costs, and market makers are willing to keep bid-ask spreads low for securities in which investors transact more often. The heightened price of risk charged by market makers during the crisis exacerbates this phenomenon, as the difference in bid-ask spreads encourages additional segmentation. Indeed, insurance companies were net buyers of the relatively expensive, yet more liquid, securities during the height of the crisis.

In addition to providing direct evidence of the clientele effect modeled by Amihud and Mendelson (1986), our account suggests complementarities between bid-ask spreads and trading frequency that can lead to multiple equilibria, as suggested in Dow (2004). A security may be cheap, have a large bid-ask spread, and trade infrequently, or be rich, have a low bid-ask spread, and trade frequently.

Dependable price relationships break down in crises, and it is critical for investors and policymakers to understand why. Investors pursue strategies based on these relationships, central banks infer conditions in financial markets based on these relationships, and bank regulators look for signs of stress based on these relationships. We show that differences in liquidity, combined with the actions of investors and liquidity providers during a crisis, can contribute to sharp and long-lasting

deviations, even in the most typically-dependable pricing relationships. Our results support the idea that policies to enhance liquidity in specific securities, such as a re-opening or a securities lending program, could dampen the liquidity feedback effect that might otherwise spiral to an extreme. Strategies to bolster market resilience to liquidity shocks are especially relevant as investors' adjust to more conservative post-crisis market norms and liquidity providers navigate still evolving regulatory-related costs to market-making.

References

Adrian, Tobias, 2013. Treasury Market Liquidity: An Overview. Federal Reserve Bank of New York, *Presentation*.

Amihud, Yakov and Haim Mendelson, 1991. Liquidity, maturity, and the yields on U.S. Treasury securities. *Journal of Finance* 46, 1411-1425.

Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. Journal of Financial Economics, 17(2), 223-249.

Barclay, Michael J., Terrence Hendershott, and Kenneth Kotz,2006. Automation versus Intermediation: Evidence from Treasuries Going Off the Run. *Journal of Finance* 61(5), 2395-2413.

Burne, Katy,2015. Trading in Treasurys Stokes U.S. Unease. *The Wall Street Journal*, C1.

Coval, Joshua and Erik Stafford, 2007. Asset Fire Sales (and Purchases) in Equity Markets. *Journal of Financial Economics 86*, 479-512.

Donnelly, Barbara, 1988. Mutual Funds: New Ad Rules May Alter Focus of Bond Funds, *Wall Street Journal*, April 4, 1988.

Dow, James 2004. Is Liquidity Self-Fulfilling? *The Journal of Business* 77(4), 895-908.

Duffie, Darrell,1996. Special Repo Rates. *Journal of Finance 51*, 493-526.

Evans, Richard B., Christopher C. Geczy, David K. Musto and Adam V. Reed, 2009. Failure is an Option: Impediments to Short Selling and Options Prices, *Review of Financial Studies*, 22, 1955-1980.

Fleckenstein, Mathias, Francis Longstaff, and Hanno Lustig,2014. Why Does the Treasury Issue TIPS? The TIPS-Treasury Bond Puzzle. *Journal of Finance* 69 (5), 2151-2197.

Fleming, Michael, 2003. Measuring Treasury Market Liquidity. Federal Reserve Bank of New York, *Economic Policy Review*, September, 83-108.

Fleming, Michael J. and Eli M. Remolona, 1997. Price Formation and Liquidity in the U.S. Treasury Market: Evidence from Intraday Patterns Around Announcements, *Federal Reserve Bank of New York Staff Reports*, no. 27, 1-27.

Fontaine, Jean-Sébastien and René Garcia, 2012. Bond Liquidity Premia, *Review of Financial Studies*, 25 (4): 1207-1254.

Goldreich, David, Bernd Hanke, and Purnendu Nath, 2005. The Price of Future Liquidity: Time-Varying Liquidity in the U.S. Treasury Market. *Review of Finance* 9, 1-32.

Gorton, Gary and Andrew Metrick, 2012. Securitized Banking and the Run on Repo. *Journal of Financial Economics* 104, 425-451.

Greenwich Associates, 2014. Global Trends and Competitive Analysis of Multi-Dealer Platforms. *Market Trends Report.*

Gürkaynak, Refet, Brian Sack, and Jonathan Wright, 2006. The U.S. Treasury Yield Curve: 1961 to the Present. *Finance and Economics Discussion Series Paper*, 2006-28.

Huang, R. D., and H. R. Stoll, 1997, The Components of the Bid-Ask Spread: A General Approach, *Review of Financial Studies*, 10, 995-1034

Hu, Xing, Jun Pan and Jiang Wang, 2013. Noise as Information for Illiquidity. *Journal of Finance* 68 (6), 2341-2382.

Jordan, Bradford D. and Susan Jordan, 1997. Special Repo Rates: An Empirical Analysis, *Journal of Finance 52*, 2051-2072.

Kamara, Avraham, 1994. Liquidity, Taxes, and Short-Term Yields. *Journal of Financial and Quantitative Analysis 29*, 403-417.

Keane, Frank, 1996. Repo Rate Patterns for new Treasury Notes. *Current Issues in Economics and Finance.*

Krishnamurthy, Arvind, 2010. How Debt Markets Have Malfunctioned in the Crisis, *Journal of Economic Perspectives*, 24 (1), 3-28.

Krishnamurthy, Arvind, 2002. The Bond/Old-Bond Spread. *Journal of Financial Economics 66*, 463-506.

Kruger, Daniel and Liz McCormick, 2015. How DIY Bond Traders Displaces Wall Street's Hot Shots. *Bloomberg News*, Article.

Longstaff, Francis, 2004. The Flight-to-Liquidity Premium in U.S. Treasury Bond Prices, *Journal of Business* 77(3), 511- 526.

McMenamin, Robert, Zain Mohey-Deen, Anna Paulson, and Richard Rosen, 2012. How Liquid are U.S. Life Insurance Liabilities? *Chicago Fed Letter*, Number 302.

Merrill, Craig B., Taylor D. Nadauld, Rene M. Stulz, and Shane Sherlund, 2012. Did Capital Requirements and Fair Value Accounting Spark Fire Sales in Distressed Mortgage-Backed Securities? National Bureau of Economic Research Working Paper, 18270.

Mitchell, Mark, and Todd Pulvino, 2012. "Arbitrage crashes and the speed of capital." *Journal of Financial Economics* 104(3): 469-490.

Pasquariello, Paolo and Clara Vega, 2009. The on-the-run Liquidity Phenomenon. *Journal of Financial Economics* 92, 1-24.

Schwarz, Krista, 2016. Mind the Gap: Disentangling Credit and Liquidity in Risk Spreads. *The Wharton School*, Working Paper.

Svensson, Lars, 1994. Estimating and Interpreting Forward Rates: Sweden 1992-4*, National Bureau of Economic Research,* Working Paper, 4871.

Figure 1. Price Difference between Maturity-Matched Treasury Securities

Panel A: Individual Maturity-Matched Pairs
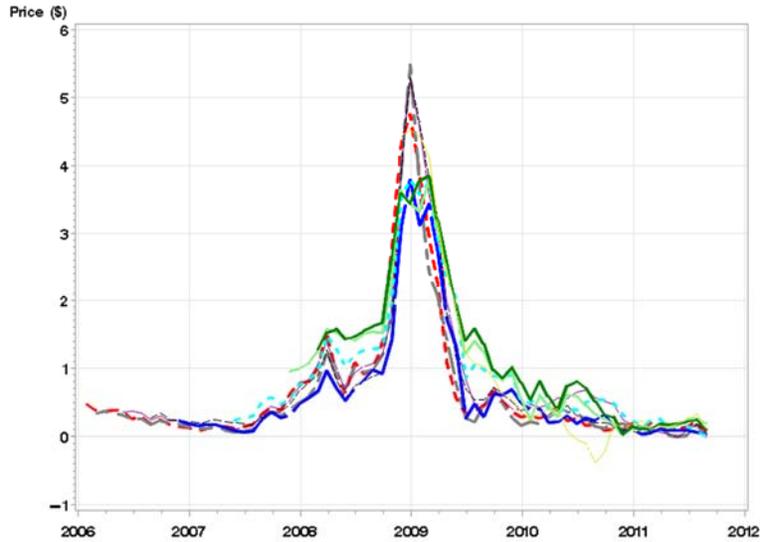
Panel B: Average of Maturity-Matched Pairs
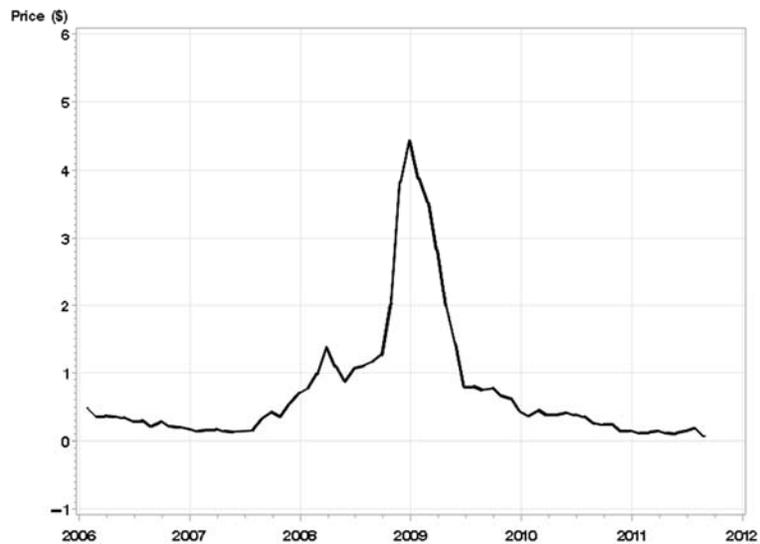
Figure 1. Panel A shows the monthly-average in price difference of maturity-matched Treasury security portfolios with identical cash flows; each pair is represented by a separate line. Each price difference subtracts the price of a 10-year original issue note from that of a replicating portfolio that consists of a 30-year original-issue bond plus a STRIP that matures on the same date. Each series starts with the date that the pair's 10-year original issue note becomes off-the-run. The maturity dates of the nine pairs range from February 2, 2015 to May 15, 2015. Panel B shows the cross-sectional average over the spreads show in Panel A.

Figure 2. Notes versus Bonds, November 20, 2008
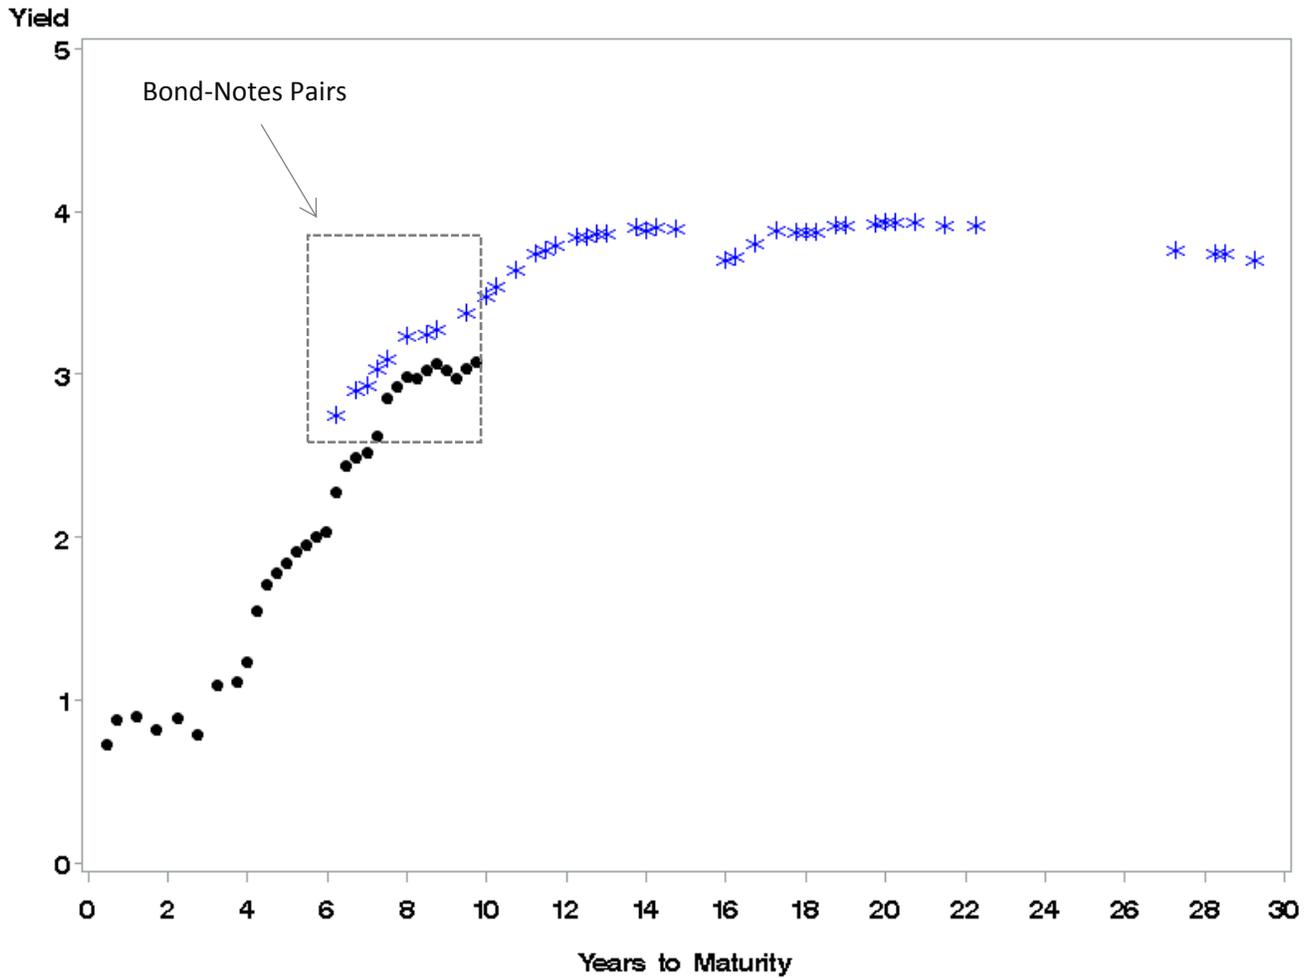
Figure 2. This figure shows the yield curve for 10-year and 30-year original-issue Treasury notes and bonds on November 20, 2008. Each point represents the actual yield-to-maturity of a different security on this date. The blue stars represent Treasury bonds. The black dots represent Treasury notes. The yield is given in percentage points on the vertical axis and years-to-maturity are in years on the horizontal axis.

Figure 3. Overnight Treasury Funding Market

Panel A: GC and Specials Rates

Panel C: Treasury Specials Transactions and Treasury Fails

Panel B: Cumulative Net Funding Costs
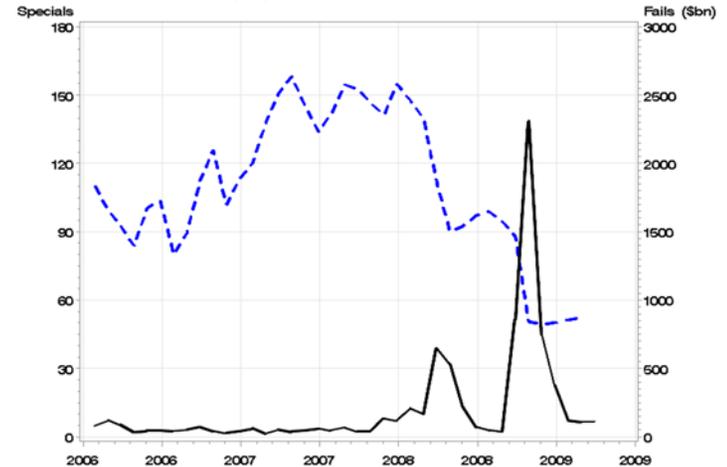
Figure 3. All panels in this figure represent monthly averages. Panel A shows overnight repo rates in percentage points, annualized. The solid black line is the overnight special repo rate averaged across the off-the-run Treasury notes comprising the nine matched portfolios in our sample, using transaction rates from a large inter-dealer broker. The dotted blue line is the overnight Treasury GC repo rate, obtained from Bloomberg. Panel B shows the running sum of the GC minus specials rate from Panel A, the cumulative net funding cost to the bond/note trade. In Panel C, the dashed blue line shows the number of overni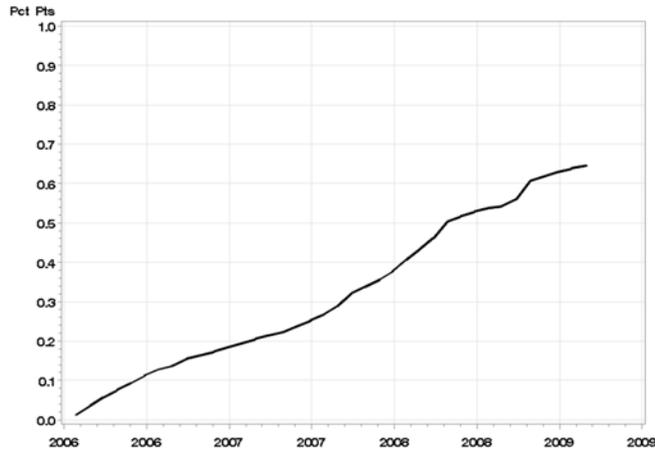ght Treasury special repo transactions in our sample, referencing the scale on the left vertical axis. The solid black line shows the volume of Treasury fails in $ billions, referencing the scale on the right vertical axis, as reported by primary dealers to the Federal Reserve Bank of New York.

Figure 4. Security Liquidity Characteristics

Panel A: Average Price Deviations by Security Size Quartile
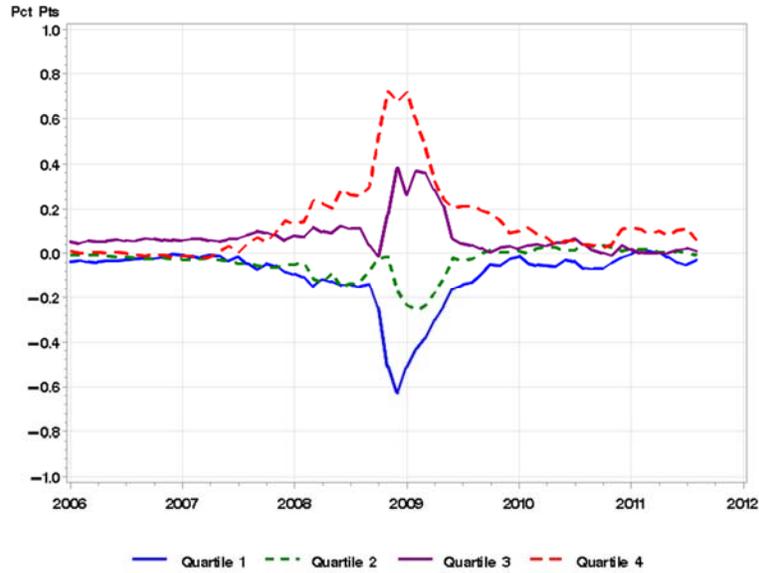
Panel B: Average Bid-Ask Spread
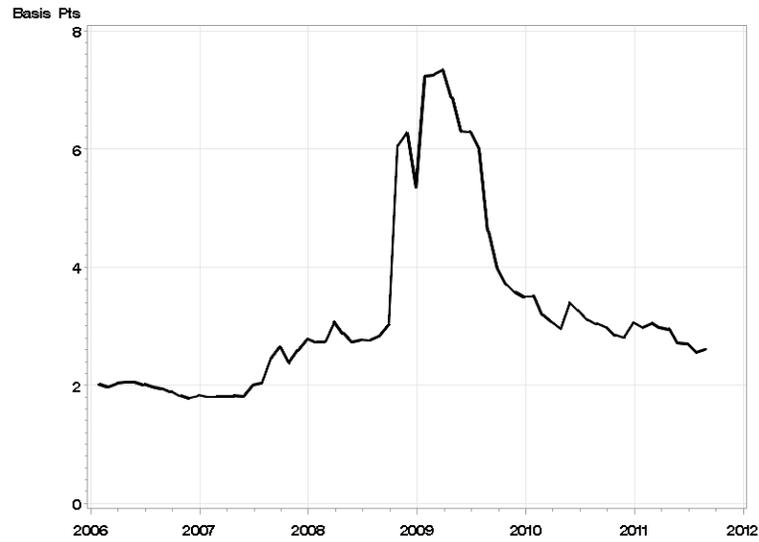
Figure 4. Panel A shows the monthly average fitting error for each security size quartile. The sold blue line represents the quartile with the smallest securities, and the dashed red line represents the quartile with the largest securities. The vertical axis is measured in percentage points. Panel B shows the monthly average bid-ask spread (in price terms) for all off-the-run Treasury security observations in our sample. The data are from the TradeWeb platform. The vertical axis is measured in basis points.

Figure 5. Liquidity Fundamentals

Panel A: Average Bid-Ask Spread
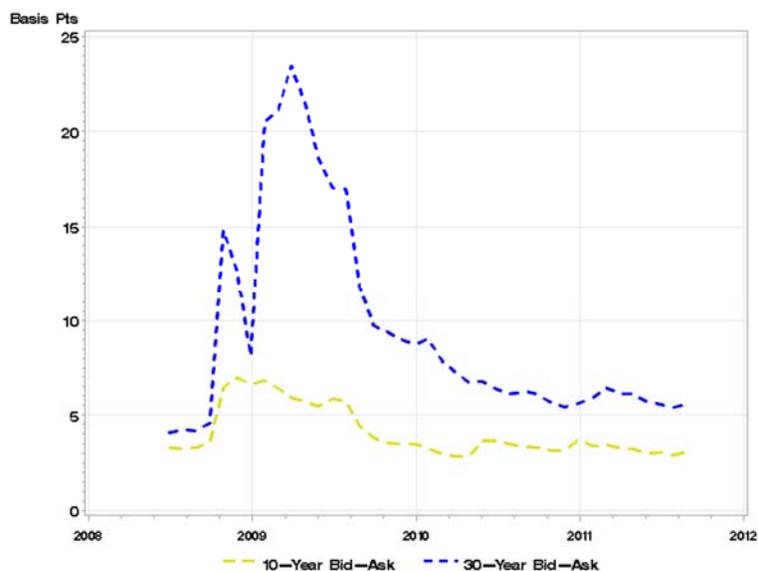
Panel B: Average TTT and Duration
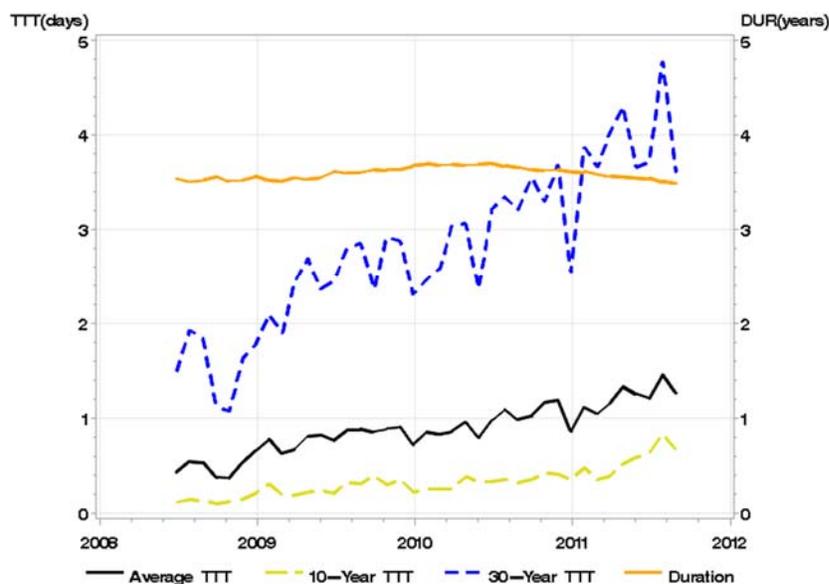
Figure 5. This figure shows the monthly average of the *Bid-Ask Spread*, *TTT* and *Duration* for various coupon-paying Treasury securities in our sample. Panel A shows the *Bid-Ask Spread*. Panel B shows *TTT* and *Duration*. In both Panels A and B, the dashed yellow and blue lines represent average values over 10- and 30-year Treasury securities, respectively. In Panel B, the solid black line shows average *TTT* , which is measured in days, referencing the scale on the left vertical axis, over all bonds and notes in our sample. This panel also shows the average level of *TTT* for 10-year notes, and the average level of *TTT* for 30-year bonds in our sample. The solid orange line shows average *Duration* referencing the scale on the right vertical axis. *Duration* is measured in years.

Figure 6.   Monthly Regressions of Bid-Ask Spread on TTT and Duration

Monthly Regression Coefficient Estimates (Dependent Variable: Bid-Ask Spread$_{it}$)
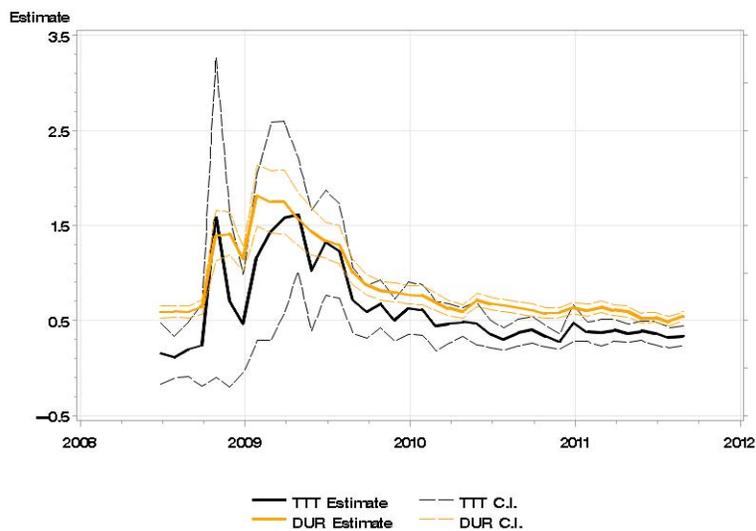
Figure 6. The figure shows the monthly-averages of coefficient estimates and 95 percent confidence intervals for a series of daily, cross-sectional multivariate regressions of the $Bid - Ask\ Spread_{it}$ onto TTT_{it} and Dur_{it} for each security i on each day t, where the $Bid - Ask\ Spread_{it}$ is measured in basis points. The solid, black line shows the point estimates for TTT. The solid, orange line shows the point estimates for Dur. The corresponding color, thin dashed lines show the confidence intervals for the respective coefficient estimates.

Table 1. Treasury Security Characteristics

Panel A: Means and Standard Deviations

	Bid-Ask Spread (basis points)		Amount Out ($bn)		Share Stripped (share of security)		Age (years)		TTM (years)		Volume/Day ($mn)		Coupon (percent)		Duration (years)		Time-to-Trade (Days)	
	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev
All Coupons	3.24	3.11	23.88	11.03	0.04	0.09	5.08	6.58	4.29	3.53	44.80	75.95	4.31	2.19	3.62	2.53	0.92	1.78
2-Year	1.61	0.65	32.12	7.78	0.00	0.00	0.94	0.50	1.06	0.50	41.33	132.90	2.82	1.62	1.05	0.49	0.35	0.61
3-Year	1.90	0.64	31.38	7.68	0.01	0.01	1.26	0.76	1.74	0.76	19.33	57.68	2.40	1.46	1.70	0.74	0.46	0.82
5-Year	2.58	1.18	20.35	8.42	0.01	0.02	2.31	1.33	2.69	1.33	82.24	206.64	3.65	0.93	2.55	1.22	0.79	1.27
10-Year	3.44	1.67	28.24	11.92	0.02	0.03	4.53	2.69	5.48	2.69	85.43	191.01	4.50	0.76	4.80	2.21	0.33	0.76
30-Year	6.89	6.14	12.08	6.14	0.21	0.15	19.69	2.84	10.32	2.85	16.94	94.01	8.55	1.06	7.40	1.63	2.86	3.25
STRIPS	11.65	8.28	0.85	0.65	N/A	N/A	N/A	N/A	4.84	2.84	4.13	19.51	N/A	N/A	4.84	2.84	6.18	5.36

Panel B: Correlations

	Bid-Ask Spread	Amount Out	Share Stripped	Age	TTM	Volume/Day	Coupon	Duration	Time-to-Trade
Bid-Ask Spread	1.00								
Amount Out	-0.33	1.00							
Share Stripped	0.35	-0.50	1.00						
Age	0.43	-0.61	0.60	1.00					
TTM	0.38	-0.24	0.42	0.24	1.00				
Volume	-0.10	-0.02	0.03	0.00	0.01	1.00			
Coupon	0.45	-0.75	0.68	0.76	0.36	0.03	1.00		
Duration	0.39	-0.24	0.52	0.32	0.94	0.01	0.52	1.00	
Time-to-Trade	0.20	-0.41	0.31	0.36	0.09	0.02	0.36	0.09	1.00

Table 1. Panels A and B show summary statistics for our liquidity proxies, which differ for each individual Treasury security. Panel A shows the means and standard deviations of our variables in aggregate (All Coupons), and also by original issue maturity bucket. The share of a security that is stripped and the dollar amount outstanding are from the Treasury's Statement of Monthly Debt. Bid-ask spreads are from the TradeWeb trading platform. Trading volume is that of U.S. insurance companies.

Table 2. Insurer Characteristics

Panel A: Means and Standard Deviations

# of	Assets ($bn)		Capital/Assets (%)		Annuity (%)		RBC (%)		Net Income (%)		Churn (ratio)		Holding Horizon (years)	
Insurers	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev	Mean	St Dev
All Insurers 2327	3.21	17.11	45.65	25.83	4.41	20.54	50.04	50.00	76.29	42.53	0.08	0.14	2.53	1.49
Health 330	0.35	0.91	59.09	20.41	0.00	0.00	31.70	46.53	75.82	42.82	0.09	0.20	2.10	1.11
Life 571	10.85	32.97	30.79	27.28	18.38	38.74	53.36	49.89	71.02	45.37	0.10	0.18	2.50	1.64
P&C 1426	0.90	4.64	48.39	23.66	0.00	0.00	52.90	49.92	78.43	41.13	0.07	0.09	2.63	1.49

Panel B: Correlations

	Assets	Capital/Assets	Annuity	RBC > Median	Net Income > 0	Churn	Hold Horiz
Assets	1.00						
Capital/Assets	-0.22	1.00					
Annuity Focus	0.38	-0.24	1.00				
RBC > Median RBC	-0.03	0.43	-0.01	1.00			
Net Income > 0	-0.01	0.12	-0.05	0.11	1.00		
Churn	0.24	-0.12	0.13	-0.08	0.00	1.00	
Hold Horizon	-0.10	0.12	-0.05	0.14	0.01	-0.34	1.00

Table 2. This table shows summary statistics for insurer-level characteristics. *Annuity* is an indicator variable equal to 1 if an insurer has a business focus on annuities, and 0 otherwise, defined based on the insurer writing more than 50% of its business in annuities. *RBC* is an indicator equal to 1 if an insurer has an RBC ratio above the median insurer ratio, and 0 otherwise. *Net Income* is an indicator variable equal to 1 if net income is positive, and 0 otherwise. *Churn* is the ratio of an insurer's trading volume relative to holdings. *Hold Horizon* is the volume-weighted average number of years that an insurer holds a Treasury security. The sample period is January 1, 2006 through September 30, 2011.

Table 3. Fitting Errors and Treasury Security Liquidity

Dependent Variable: Fitting Error$_{i,t}$

Panel A: Characteristics that Differ by Security

Bid-Ask$_{i,t}$	-3.42***									-3.48***	-3.40***
	(0.31)									(0.26)	(0.26)
ln(Out)$_{i,t}$		20.77***							17.17***	16.87***	15.92***
		(3.07)							(4.86)	(4.13)	(4.25)
ShrStrip$_{i,t}$			-97.20***							-94.52***	-94.36***
			(13.62)							(18.69)	(17.64)
ln(Age)$_{i,t}$				-5.23***					0.47	4.04***	3.26***
				(0.86)					(1.04)	(1.04)	(1.04)
Volume$_{i,t}$					38.23***					15.50***	14.99***
					(7.78)					(3.90)	(3.90)
Coupon$_{i,t}$						-3.59***			-1.98**	0.68	1.62**
						(0.38)			(0.86)	(0.80)	(0.74)
ln(TTM)$_{i,t}$							-0.02		3.99***	8.27***	7.30***
							(1.69)		(1.48)	(1.52)	(1.51)
Pair$_{i,t}$								49.41***			24.95***
								(12.12)			(7.78)
R-Squared	0.084	0.087	0.079	0.043	0.002	0.055	0.000	0.048	0.099	0.191	0.201
Observations	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873

Dependent Variable: Fitting Error$_{i,t}$

Panel B: Security Characteristics Interacted with Aggregate Liquidity

Bid-Ask$_{i,t}$	-1.30*									0.17	0.23
	(0.70)									(0.97)	(0.90)
ln(Out)$_{i,t}$		-33.63***							-21.31***	-20.94***	-18.81***
		(5.57)							(7.28)	(6.58)	(7.00)
ShrStrip$_{i,t}$			98.04***							53.37**	57.89***
			(18.99)							(22.98)	(22.10)
ln(Age)$_{i,t}$				9.82***					0.53	-0.06	1.76
				(1.77)					(2.62)	(2.45)	(2.52)
Volume$_{i,t}$					-15.49					15.77	15.08
					(17.61)					(11.45)	(11.33)
Coupon$_{i,t}$						6.32***			4.47***	0.79	-0.36
						(0.71)			(1.65)	(1.87)	(1.82)
ln(TTM)$_{i,t}$							0.34		-5.20***	-8.15***	-7.26***
							(2.40)		(1.99)	(2.07)	(2.23)
Pair$_{i,t}$								-55.86***	11.78***	-0.31***	-12.94
								(18.39)	(3.42)	(0.11)	(13.30)
AggLiq$_t$*Bid-Ask$_{i,t}$	-0.34***										-0.26**
	(0.09)										(0.11)
AggLiq$_t$*ln(Out$_{i,t}$)		16.64***							0.31	11.53***	10.50***
		(2.42)							(1.14)	(2.97)	(3.13)
AggLiq$_t$*ShrStrip$_{i,t}$			-61.58***							-47.33***	-48.84***
			(8.40)							(10.93)	(10.22)
AggLiq$_t$*ln(Age)$_{i,t}$				-4.59***						1.07	0.07
				(0.74)						(1.05)	(1.11)
AggLiq$_t$*Volume$_{i,t}$					16.80**				-2.02**	0.56	0.81
					(6.98)				(0.79)	(4.24)	(4.21)
AggLiq$_t$*Coupon$_{i,t}$						-3.02***				-0.19	0.54
						(0.31)				(0.86)	(0.83)
AggLiq$_t$*ln(TTM)$_{i,t}$							-0.11		2.51***	4.24***	3.59***
							(1.19)		(0.87)	(0.93)	(0.99)
AggLiq$_t$*Pair$_{i,t}$								32.36***			12.32**
								(8.84)			(6.29)
R-Squared	0.088	0.196	0.166	0.115	0.003	0.133	0.000	0.096	0.216	0.295	0.311
Observations	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873	207,873

Table 3. This table presents a panel regression of fitting errors on security characteristics. The *Fitting Error*, the dependent variable, is defined as the difference between the actual price of the security and the fitted price based on a smoothed yield curve, measured in basis points. *Volume* is trading volume of U.S. insurance companies. *Volume* for each security on each day is the ratio of the trading volume for that security relative to the sum of volume for all securities on that day. *Coupon* is measured in percentage points per year. *Pair* is a dummy equal to 1 for a note and -1 for a bond in cases where both the note and bond have the same maturity date, and 0 for all other securities. Panel A shows univariate and multivariate results for characteristics that differ by security. Panel B shows results for specifications that include interaction effects between security characteristics and *AggLiq$_t$*, which is a time series proxy for market-wide liquidity. The variable *AggLiq$_t$* is the average bid-ask spread over all securities in our sample on day *t*. Calendar day fixed effects are included in each regression. Standard errors (in parentheses) account for clustering within security *i* and arbitrary heteroskedasticity; *** indicates significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Table 4. Insurer Transaction Patterns

	Dependent Variable: Net Purchase of Treasury$_{i,j,t}$			
Insurer Type:	All	Health	Life	P&C
Fitting Error$_{ijt}$	3.21***	1.72***	8.27***	1.57**
	(0.76)	(0.27)	(2.67)	(0.65)
R-Squared	0.0039	0.0119	0.0045	0.0024
Observations	31,895,768	4,298,496	7,683,239	19,914,033

Table 4. This table presents a pooled regression of insurance companies' monthly net purchases of Treasury security i for insurer j in month t on the fitting error for each security i at time t. Net purchases are measured per \$1 billion of the security's original issue size. The fitting error is defined as the difference between the actual price of the security and the fitted price based on a smoothed yield curve, measured in basis points. Month-of-year fixed effects and insurer type fixed effects are included in all regressions. The sample period is January 2006 through August 2011. Standard errors are shown in brackets beneath the coefficient estimates; *** indicates significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.

Table 5. Who Engages in the Arbitrage?

	Dependent Variable: Net Purchases of Treasury$_{i,j,t}$						
Capital/Assets$_{j,t}$ * FE$_t$	-0.12***						-0.01***
	(0.03)						(0.00)
Capital/Assets$_{j,t}$	0.87						0.12
	(1.10)						(0.14)
Annuity$_j$ * FE$_t$		16.27***					1.25***
		(3.67)					(0.40)
Annuity$_j$		-344.19**					-32.82**
		(137.96)					(14.61)
RBC$_{j,t}$ * FE$_t$			-2.04				0.11
			(1.57)				(0.18)
RBC$_{j,t}$			-28.12				-6.62
			(55.24)				(6.55)
Net Income$_{j,t}$* FE$_t$				1.44			0.34*
				(1.64)			(0.18)
Net Income$_{j,t}$				-9.72			-0.78
				(63.25)			(6.74)
Churn$_j$ * FE$_t$					29.04***		1.84***
					(5.62)		(0.67)
Churn$_j$					-44.42		10.38
					(191.94)		(23.07)
Hold Horizon$_j$ * FE$_t$						-6.11***	-0.36**
						(1.30)	(0.15)
Hold Horizon$_j$						45.90	5.11
						(45.41)	(5.33)
FE$_t$ (Fitting Error)	8.64***	2.47***	4.30***	2.25*	0.92	7.91***	0.55**
	(1.53)	(0.78)	(1.12)	(1.35)	(0.88)	(1.26)	(0.24)
R-Squared	0.0039	0.0040	0.0040	0.0039	0.0040	0.0040	0.0044
Observations	31,895,768	31,895,768	31,895,768	31,895,768	31,895,768	31,895,768	31,895,768

Table 5. This table presents results from a pooled regression of an insurance company's monthly net purchases of Treasury security *i* for insurer *j* in month *t* on the fitting error for each security *i* at time *t* interacted with various characteristics of the insurance company. The fitting error is defined as the difference between the actual price of the security and the fitted price based on a smoothed yield curve, measured in basis points. Net purchases are measured per \$1 billion of the security's issue size. The sample period is January 1, 2006 through September 30, 2011. Month fixed effects and insurer type fixed effects are included in all regressions. Standard errors are in parentheses; *** indicates significance at the 1 percent level, ** at the 5 percent level, and * at the 10 percent level.