Notes on Bonds: Liquidity at all Costs in the Great Recession

David Musto
Greg Nini
Krista Schwarz [*]

April 27, 2011

VERY PRELIMINARY AND INCOMPLETE

Abstract: We address the connection between market stress and asset pricing by analyzing a large and systematic discrepancy arising among off-the-run Treasuries. We first show that bonds traded for much less than notes with matching maturity and coupon, over five percent less in December 2008. We then ask how the small differences between these securities, in particular their liquidity, could project to so large a gap. We gauge the potential for long/short arbitrage with repo and fails data indicating the frictions arbitrageurs encountered, and then with daily transactions data we relate the demand for the expensive but liquid note to the cross section of insurers' liquidity needs.

[*]The Wharton School, University of Pennsylvania, 3620 Locust Walk, Philadelphia PA 19104. Department of Finance: David Musto, (215) 898-2323, musto@wharton.upenn.edu. Department of Risk & Insurance: Greg Nini, greg30@wharton.upenn.edu, (215) 898-0310. Department of Finance: Krista Schwarz, (215) 898-6087, kschwarz@wharton.upenn.edu. We are grateful to the Dean's Research Fund of the Wharton School data funding. We are particularly grateful to Michael Bulboff, who provided outstanding research assistance. All remaining errors are our own.

Notes on Bonds: Liquidity at all Costs in the Great Recession

Abstract: We address the connection between market stress and asset pricing by analyzing a large and systematic discrepancy arising among off-the-run Treasuries. We first show that bonds traded for much less than notes with matching maturity and coupon, over five percent less in December 2008. We then ask how the small differences between these securities, in particular their liquidity, could project to so large a gap. We gauge the potential for long/short arbitrage with repo and fails data indicating the frictions arbitrageurs encountered, and then with daily transactions data we relate the demand for the expensive but liquid note to the cross section of insurers' liquidity needs.

The market for U.S. Treasury debt is the largest, most liquid, and safest securities market in the world. The total amount of publicly-held, Treasury-issued debt currently stands at $9 trillion, and daily trading volume typically exceeds $500 million.[1] Not surprisingly, pricing anomalies in the Treasury market are infrequent, short-lived, and well-studied when they do occur. For instance, typically there is a pricing difference between off-the-run and on-the-run securities; the most recently issued coupon security of a particular maturity tends to be slightly more expensive than previously-issued securities of the same original maturity. However, Krishnamurthy (2002) shows that the trading profits from entering into a convergence trade that is short the on-the-run security and long the off-the-run security are largely offset by the cost of borrowing the on-the-run. He concludes that there are no consistent arbitrage profits to be made from these pricing differences.

In this paper, we document the occurrence of a large pricing anomaly in the Treasury market that created arbitrage profits that would not have been offset by borrowing costs. Specifically, we show that a large yield spread developed between securities originally issued as thirty-year bonds and securities originally issued as ten-year notes, even though the securities share the same maturity date. For a several month period toward the end of 2008, original issue bonds became substantially cheaper than original issue notes, even after adjusting for differences in coupons. We show that the pricing anomaly was large, reaching a level as high as 5 percent, and it remains even after correcting for the difference in funding costs as measured by repurchase agreement (repo) rates.

[1] The source for Treasury securities outstanding and trading volume is the Securities Industry and Financial Markets association, available at www.sifma.org.

The pricing anomaly happens during a period of significant market turmoil, when liquidity was in particularly high demand. The spread that we document is highly correlated with other measures of market liquidity, including spreads between on-the-run and off-the-run Treasury securities and average bid-ask spreads on all Treasury securities. We conjecture that the pricing anomaly we document is related to small differences in liquidity that became magnified during the financial crisis. We show that bonds generally have lower trading volume and wider bid-ask spreads, which widened further during the crisis, suggesting that old thirty-year bonds are less liquid than ten-year notes. There are three reasons for this difference in liquidity. First, smaller amounts of these bonds were originally issued. Second, the majority of bonds typically are stripped and held in stripped form for the remainder of their lives, whereas stripping of notes is much more limited (Jordan, Jorgensen and Kuipers (2000)). This further reduces the amount of the bond that is immediately available to trade. Third, we conjecture that the bond may be disproportionately held by longer-term investors.

We interpret the apparent mispricing of Treasury securities from the perspective of the "limits to arbitrage" literature. Although repo rates and bid-ask spreads did reduce the profits available from trading against the mispricing, we show that the magnitude of the discrepancy provided a clear arbitrage opportunity from the perspective of a hold-to-maturity investor. The "limits to arbitrage" literature presents an explanation for why an arbitrageur may still shy away from such a trade. The essence of the story - described in detail in papers such as Shleifer and Vishny (1998), Gromb and Vayanos (2002) and Vayanos and Vila (2009) - is that arbitrageurs are risk-averse, have a short-horizon, or are capital constrained. All of these frictions can prevent an arbitrageur from taking the perspective of a hold-to-maturity investor, making the mispricing less attractive. In particular, as noise traders move prices for reasons unrelated to an assets fundamentals, arbitrageurs are tempted to trade the security to take advantage of the mispricing. However, arbitrageurs have to assume the risk that the noise traders could make the mispricing worse in the short-run might, which subjects the arbitrageur to interim losses. Because the arbitrageur is risk-averse in the short-run or has limited capital, this risk can deter the arbitrageur from acting. This is true even if the mispricing would make the arbitrageur a sure profit, if he/she could hold out indefinitely.

In addition to the on-the-run/off-the-run spread that is studied extensively in Krishnamurthy (2002), other anomalies in the Treasury market have also been documented. Amihud and Mendelson (1991) and Kamara (1994) compare bills and notes with less than six months remaining to an identical maturity date, so that both are effectively zero-coupon securities. During their sample, the notes were consistently cheaper (traded at higher yields) than the bills. Amihud and Mendelson argue that the price differential represents a premium for the greater liquidity of bills, but Kamara (1994) suggests that the difference owes in part to the differential tax treatment that existed at the time. In the paper most closely related to ours, Strebulaev (2003) compares the yields of coupon securities (Treasury notes and bills) with different original-issue tenors but with identical maturity dates. Although Strebulaev confirms that bills tend to be more expensive than similar notes, he finds that standard liquidity proxies are not correlated with bill-note pricing differences. Moreover, he finds no evidence of systematic pricing differences within Treasury notes, leading him to conclude that liquidity differences are not the source of the note-bill anomaly.

Many pricing anomalies have been interpreted in the context of this type of argument, and it has potential to apply to the Treasury market as well. In fact, the Treasury market is an ideal setting to cleanly establish the existence of an arbitrage and test some empirical implications. Treasury securities all share identical credit risk and do not differ in priority, but some securities are more easily traded than others. Bonds are cheaper than notes, perhaps because some investors have a preference for the greater liquidity of notes, but arbitrage capital is normally available to keep the prices of notes and bonds aligned. The crisis was a time in which arbitrage capital was withdrawn, and this led the spread between note and bond yields to skyrocket. Treasury market anomalies (not confined to the bond-note spread) during the crisis are discussed in these terms by Hu, Pan and Wang (2010). By identifying specifically which Treasury securities are relatively cheap and which are relatively expensive, we can then investigate the types of participants that are either exploiting the arbitrage or acting as noise traders based on their trading activity in these particular Treasury securities at the time that the pricing divergence widened.

The plan for the remainder of this paper is as follows. In section 2, we describe the note-bond pricing anomaly that is the focus of this paper. Section 3 relates pricing anomalies in the crisis to characteristics of individual coupon securities, arguing that the cross-section of yields can be accounted for in terms of liquidity differences. Section 4 evaluates the characteristics of investors that were exploiting the arbitrage opportunity, and those that were making it worse. Section 5 concludes.

I. The Arbitrage

We begin by illustrating an example of the pricing anomaly that explore, shown in Figure 1. The figure shows the spread between the yields to maturity on two Treasury securities, both maturing on February 15, 2015. We take the difference between an original-issue 30 year bond and an original issue 10 year note. The bond was originally issued in 1985 with a coupon of 11.25 percent, and the note was originally issued in 2005 with a coupon of 4 percent. Throughout 2005, 2006, and much of 2007, the bond has a slightly higher yield-to-maturity than the note, by about few basis points, on average. In late 2007 and early 2008, however, the bond yield climbed substantially relative to the note yield. The difference spiked in the fall of 2008, when the spread of the bond yield over note yield climbed to 80 basis points, representing a price difference of about five dollars per 100 dollars face value. It is worth noting that both securities were well off-the-run during the time when the yield spread widened most notably.

This yield spread documented in Figure 1 does not necessarily represent an arbitrage opportunity, since the note has a lower coupon than the bond and thus a longer duration. But with an upward sloping yield curve, as was the case during 2008, the difference in coupons should result in a relatively *higher* yield to maturity for the note relative to the bond; Figure 1 shows that the note has a lower yield to maturity.

To conduct a more precise comparison between the pricing of the note and the bond, we create a synthetic portfolio of the bond and a Treasury STRIP to exactly match the cash flows of the note. Specifically, for a note with coupon rate C_n and a bond with coupon rate C_b (both maturing on the same day), we form a portfolio that puts weight C_n / C_b on the bond and weight $1 - C_n / C_b$ in a STRIP maturing on the maturity date of the bond. This portfolio will have identical cash flows to the note, which lets us compare the prices of two assets that generate identical cash flows. In our empirical analysis below, we show that the price of the note compared with the price of the bond and STRIP portfolio is usually close to zero but grew significantly during the period of the financial crisis before returning back close to zero in 2009.

A trading strategy that bet on convergence of the prices of two portfolios would have made positive profits. Moreover, these profits would be riskless, as long as the positions could be held until maturity, and as long as funding costs or other frictions in implementing the trade would not exceed the difference in returns on the portfolio. In particular, the cost of shorting the more expensive note could wipe out any profits created by convergence in the prices, although we show below that this is not the case. In the absence of such significant frictions, the strategy of betting on convergence could be scaled to produce enormous profits.

In this paper we consider nine bond-note pairs similar to the example displayed in Figure 1. In all cases, the original-issue note becomes expensive relative to the portfolio comprised of the original-issue bond and STRIP with identical maturity date, although the size of the pricing gap varies across the pairs. We explicitly incorporate the cost of forming the short position in the note using repo rates that account for any specialness in shorting a particular security. We interpret any remaining difference as potential arbitrage profits that would be available to a hold-to-maturity investment position.

II. Apparent Pricing Anomalies

Our analysis begins by recognizing that the period between August 2007 and May 2009 represented a period of significant market turmoil. For example, Hu, Pan, and Wang (2010) document that deviations in Treasury yields from a smooth yield curve hit a record high in the weeks following the Chapter 11 filing of Lehman Brothers in September 2008. Hu, Pan, and Wang construct a measure of illiquidity based on the average deviation of Treasury prices from those based on a smooth yield curve and show that this measure provides a useful proxy for illiquidity and is a priced risk factor. We adopt their measure of illiquidity and show that deviations from the smooth curve were systematic: original-issue 30-year bonds became cheap relative to original-issue 10 year notes. The systematic nature of the pricing deviations leads us to create the bond-note pairs that we explore further in the next section.

II.1 Cheap and Expensive Securities

This subsection addresses the question of which securities became relatively cheap during the crisis and which were relatively expensive. To address this question, we compare actual Treasury prices with those implied from a parametric zero-coupon yield curve fitted to the set of all coupon securities. We use the parameter estimates provided by the Federal Reserve Board, who every day fit the six-parameter model of instantaneous forward rates of Svensson (1994) to observed prices on coupon treasury securities.[2] With the parameter estimates, we can compute the fitted price of each security on every calendar day and compute the difference between observed prices and the fitted price. We denote the difference as the pricing error, which by construction has mean close to zero across all securities.[3] We use the CRSP daily Treasury database for our treasury security prices.

Figure 2 shows the average pricing error for all securities which were originally issued as thirty-year bonds, ten-year notes and five-year notes. Prior to the summer of 2007, average pricing errors were close to zero, and there is very little difference between thirty-year bonds and ten-year notes. Beginning in the fall of 2007 and extending through early summer of 2009, a notable pattern emerges. The thirty-year bonds became cheap relative to the smooth curve, and the ten-year notes became expensive. Notably, the pricing errors on the five-year notes do not

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

[3] The mean is not exactly zero because prices are a non-linear function of forward rates.

show systematic time series deviations. As we will show next, part of the pattern can be explained by the amount of the bonds outstanding, but there will still remain a significant pattern that thirty-year bonds became cheap relative to ten-year notes.

To further explore the determinants of the pricing error, we estimate a panel regression of the pricing errors on individual securities onto a variety of security characteristics. The regression is of the form

$$e_{it} = \alpha + \beta' X_i + \varepsilon_{it},$$

where e_{it} denotes the pricing error for the ith security on day t, and X_i is a vector of bond-specific characteristics. We use two sets of independent variables. First, as two proxies for liquidity, we include the size of the issue and the quoted bid-ask spread. We measure size as the log of the original amount of the bond issued, and use the log of the dollar value of the bid-ask spread.[4] Second, we include dummy variables indicating whether the security was originally issued as a thirty-year bond, a ten-year note, or a five-year note, with the excluded category including seven-year, three-year, and two-year notes. The regression includes observations during 2005 through 2010 and includes all coupon securities with remaining time to maturity of at least one year and no more than ten years; the smooth yield curve fits best within this range. We also run the regression on a sub-sample of observations during the crisis period, which we define as lasting from the fourth quarter of 2007 through the second quarter of 2009. Since the same security appears many times in the sample, standard errors account for potential serial correlation in residuals.

[4] We will eventually account for buybacks and re-openings by allowing the amount outstanding to vary by day.

The results are shown in Table 1. The coefficient on the amount issued is positive, suggesting that larger issues tend to be relatively expensive. Similarly, securities with larger bid-ask spreads tend to trade at higher prices, although the effect is fairly small and only significantly different from zero during the crisis period. We view these results as suggesting that differences in liquidity, as proxied by issue size and bid-ask spreads, lead to systematic differences pricing, with more liquid issues trading cheaper than less liquid issues. Interestingly, the effect of issue size and bid-ask spreads is much stronger during the crisis, suggesting that liquidity differences were exacerbated during the crisis. Even after controlling for these liquidity proxies, the dummy variables for the original-issue term of the security confirm that ten-year notes became expensive relative to thirty-year bonds. The difference in estimated coefficients on the thirty-year bonds and the ten-year notes is large and statistically significantly different from zero in both samples. During the crisis, the difference in coefficients exceeds 1, meaning that, on average over the seven quarters, bonds were more than 1 percent cheaper than notes.

We do not have a compelling reason why the notes became rich relative to the bonds, but we conjecture that unobserved differences in liquidity are the underlying source, which was exacerbated during the crisis. Although an interesting area for future research, the underlying reason is unimportant for our subsequent analysis. At maturity of the proposed trading strategy, both securities are equally liquid, so from the perspective of a hold-to-maturity investor, any liquidity differences do not matter. What matters for us is that bonds systematically cheapened relative to notes, which creates the potential arbitrage that we explore.

III. The Arbitrage Strategy

In this section, we describe how we construct two portfolios with identical cash flows and show that, in normal times, the two portfolios have very similar prices. We then document the pricing anomaly that emerges and show that funding costs did not reach levels that would overwhelm the arbitrage profits.

III.1 The Bond-Note Pairs

We construct nine pairs of securities with the same maturity dates that were originally issued as ten- and thirty-year Treasuries. We consider only nominal, non-callable Treasury securities, of which the February 2015 securities (described above) are an example. Prior to 1985, the U.S. Treasury Department exclusively issued callable thirty-year securities, so we use only bonds issued after 1984. We also restrict our sample to notes that were issued prior to the summer of 2008, so that all of the bond/note pairs exist during the peak of the financial crisis. With these restrictions, we are left with nine bond/note Treasury pairs with identical maturity dates ranging from February 2015 to May 2018.

For each pair, we construct a portfolio that is short the note, long a fraction of the bond to match the coupons of the note, and long a Treasury STRIP to match the principal payment at maturity (as discussed above). This portfolio is constructed to have zero cash flows after origination, so it should not have any cost or benefit at origination. We view any money received at origination as an arbitrage opportunity.

III.2 Accounting for the Cost of Funding

In the classic "convergence trade" that we describe above, an arbitrageur would take a long position in the cheap security (the bond) and a short position in the expensive one (the note). In reality, it is often expensive to short some securities, a friction discussed by Duffie (1996) and Krishnamurthy (2002), among others. The only way to take a short position in a Treasury issue is to enter into a repo contract where one lends out cash and takes the security as collateral. The lender can then sell the collateral immediately, betting that the price will fall, intending to buy it back at the close of the repo contract, hopefully at a cheaper price. An investor wishing to bet on an anomaly in the Treasury market must short the expensive security in this way. At the same time, the investor can buy the cheap security, typically using the security as collateral to borrow money to finance the purchase. In most repo transactions, any Treasury security is considered to be acceptable collateral, and the corresponding interest rate on the loan is known as the general collateral (GC) interest rate.

In some cases, repo cash borrowers may deliver any Treasury security as collateral, leading particularly expensive issues to not be delivered in GC agreements. However, some repo agreements specify the precise issue that must be used as collateral and must be returned at the end of the repo contract. When one security is unusually expensive, demand from investors wishing to short it can drive down the repo rate on that security to a level below the GC rate, and the security is referred to as "special." Securities that are expensive in the cash market are typically "special" in the repo market, meaning that the cost of shorting them is particularly high. When Treasury securities become special, the repo rate on the particular security is known as the security's special repo rate, which will be lower than the GC repo rate. When this happens, an investor betting on an anomaly in the Treasury market will receive a lower interest rate on his/her loan of cash (collateralized by the expensive security) than he/she must pay to borrow to buy the cheap security, which will be at the GC repo rate. The spread between GC repo and specials rate could in principle wipe out the profitability of the convergence trade.

Indeed, Krishnamurthy (2002) shows that the profits on the convergence trade between on-the-run and off-the-run bonds are roughly wiped out by the gap between the corresponding repo rates. Although the spread between these two bonds systematically converges over time, the average profits of this trade are close to zero due to the cost of shorting the newly issued bond. Krishnamurthy argues therefore that there is no genuine arbitrage opportunity.

It is particularly important for us to account for funding costs, since anecdotal evidence suggests that funding became quite difficult during the crisis. Strains in the repo market likely made it hard to short comparatively expensive Treasury securities. Additionally, an institutional feature of the Treasury repo market made investors relatively reluctant to lend out their securities when GC rates because very low, as they did following Lehman's bankruptcy. Specifically, the lack of a penalty for failing to deliver on a repo transaction created a bound of zero on the specials rate, which could have prevented the market from clearing without excess demand or supply. [5]

[5] A market participant will lend funds against a security that is priced "special" only to meet an obligation to deliver that security. Until May 2009, the penalty for a failure to deliver a security into a transaction was that the security was to be delivered the next day at the same price. This is equivalent to giving the buyer of the security an interest free loan. This would be preferable to borrowing at a negative specials rate. So specials rates cannot normally go below zero. Due to massive fails in the repo market, and the resulting drop in securities lend via repos, the Treasury Market Practices Group (TMPG), a self-governing industry group, proposed a penalty fails rate, which was backed by the Federal Reserve. The explicit penalty in failing to deliver a security was introduced in May 2009 as Max(3-FFT,0), where FFT is the base of the Federal Reserve's target rate. In a zero policy rate environment, this rule levies a 3 percent penalty rate on a fail.

Using data on repo transactions from a large interdealer broker, we show that the profits available from the bond-note convergence trade we propose would have been much larger than the costs of funding the trade. Figure 3 plots the monthly return on the convergence trade (ignoring funding costs) along with the level of the special-GC spread (the funding cost) for the bond-note pair maturing on February 15, 2015. Although funding costs do rise during the crisis, the picture shows that the pricing differences were substantially larger than the funding costs. Even at the peak divergence in prices of the underlying securities, the repo funding costs remain below 15 basis points per month. Monthly returns, however, are much larger, in some cases exceeding 2 percent at the peak. Per $1000 principal, funding costs reach a maximum of $1 per month. Raw returns peak at $14.1 per $1000 principal in December 2008, following Lehman's bankruptcy filing in September 2008.

III.3 Time Series Pattern of Arbitrage

We next explore the time periods when the arbitrage grew to its widest levels, focusing on aggregate liquidity and limits to arbitrage. We conjecture that the risk aversion of potential arbitrageurs increased and arbitrage capital was withdrawn from the market. If so, the pricing error should be correlated with other systemic liquidity indicators.

To investigate this further, we run a daily time-series regression of the average pricing error across our nine bond-note pairs on several measures of aggregate liquidity. We use the LIBOR-OIS spread, the repo bid-ask spread, and the GC repo rate. The results are shown in Table 2. To account for the significant serial correlation in the pricing errors, we use Newey-West standard errors with a lag-length of 30.

The results suggest that the pricing error is significantly correlated with the measures of aggregate liquidity. The coefficient on the LIBOR-OIS spread is large and positive, indicating that broader funding strains were correlated with the anomaly in the Treasury market. The coefficient on the repo bid-ask spread is also positive, suggesting that the strains in the repo market also happened coincidentally with the pricing anomalies. Finally, the GC rate is significantly negative, which corroborates the notion that a lower GC spread makes lending expensive securities in the repo market less attractive, which in turn prevents arbitrageurs from bringing prices back into line.

IV. Investor Response?

The pricing of Treasury securities in the crisis represented an arbitrage opportunity. Based on the "limits to arbitrage" paradigm, we suggest that this reflects a lack of arbitrage capital willing to take short-run risk to wait for the long-run gain. In this section, we explore the trading behavior of insurance companies, who are potential long-term investors that could profit from the arbitrage.

IV.1 Trading and Holdings Data

We have a dataset consisting of transactions-level data showing all buys and sells of Treasury securities for all U.S. registered insurance companies, who report such transactions in Schedule D within their statutory regulatory filings. For each trade in the dataset, we know the insurance company conducting the trade, along with the date, size, and direction of the transaction. We also have several characteristics of the insurers, including several measures of their capital, including their financial strength ratings and their size. We also construct three additional variables based on the trading history of each insurer. In sum, we consider six cross-sectional characteristics of each insurer:

(i) *Buy-and-hold indicator*. A dummy that is one if that insurer is a "buy and hold" insurer, i.e. never conducts a sell transaction of Treasury securities.
(ii) *Horizon*. This is the average number of days that an insurer holds a given Treasury security.
(iii) *Churn*. This is the ratio of transactions volume relative to holdings over all Treasury securities for each investor. A lower value corresponds to a less active trader.
(iv) *Size*. This is the amount of assets held by the insurer.
(v) *Investment Grade*. A dummy that is one if the insurer's best rating is classified as investment grade.
(vi) *Premium-to-Asset* ratio. This is a leverage measure for each insurer.

For each insurer *i* in month *t*, we construct the net purchases of notes less the net purchases of bonds, which we denote as $NP_{i,t}$. This is a measure of the propensity to engage in the arbitrage trade, and we relate the measure to the size of the pricing error and the cross-sectional characteristics of each insurer. In particular, we conduct a regression of the form:

$$NP_{i,t} = \alpha_i + \beta_i PE_t + \varepsilon_{i,t} \tag{1}$$

We further assume that the intercept and slope coefficients in this regression are linear functions of the characteristics of the insurer, collected in a vector X_i:

$$\alpha_i = a + b' X_i, \quad \beta_i = c + d' X_i \tag{2}$$

Substituting (2) into (1) gives:

$$NP_{i,t} = a + b' X_i + cPE_t + d' X_i PE_t + \varepsilon_{i,t} \tag{3}$$

which we then estimate as a pooled regression. The main object of interest is the interaction coefficient, d. This tells us whether a particular characteristic of an insurer makes the insurer more or less likely to buy the note when it becomes particularly expensive.

Table 3 reports the results from estimating equation (3) with each of the six different insurer characteristics separately. As can be seen, longer-horizon investors are on net sellers of the expensive note, which suggests that they are behaving as arbitrageurs. This is consistent with the finding of Coval and Stafford (2007) in equity markets. Non-investment grade and more leveraged insurers are also sellers of the expensive note, which means that these are the insurers who were in effect exploiting the arbitrage opportunity. Perhaps they had to sell Treasuries quickly in order to raise cash, and chose to do so by selling the relatively expensive and liquid notes.

V. Conclusions

In normal times, the pricing of different Treasury securities is internally consistent. Two different Treasury coupon securities with different coupon rates but the same maturity date will have almost identical yields. Indeed, one can form a portfolio combining either one of these coupon securities with a set of STRIPS such that the portfolio has exactly the same payoffs as the other security. The portfolio and the security should—and normally do—have almost exactly the same price; otherwise one could create riskless profits that should not exist in a well-functioning market. However, starting with the onset of the financial crisis in August 2007, and then accelerating after the collapse of Lehman in the fall of 2008, these arbitrage relationships broke down dramatically. Bonds that were originally issued as thirty-year bonds that had 6-9 years to maturity became much cheaper than bonds originally issued with a ten-year maturity, even though both had nearly the same maturity date.

In the canonical theoretical model of persistent arbitrage opportunities, Shleifer and Vishny (1997) show that risk-aversion and bounded capital can explain why arbitrageurs are limited in their ability to prevent the emergence of pricing anomalies. In their model, "noise traders" have a liquidity-based motivation for trading that may cause prices to deviate from their fundamental value. Arbitrageurs trade against the noise traders to offset the deviations, but risk-aversion and limited capital can prevent the arbitrageurs from completely offsetting the divergence. The model explains why pricing discrepancies, and apparent arbitrage opportunities, can persist for some time. This paper aims to give some empirical content to the Shleifer and Vishny (1997) model by characterizing the nature of the noise traders and arbitrageurs and offering clues as to their motivation.

Studying the unusual pricing of Treasury securities at times of market stress gives us useful insights into the behavior of fixed income markets at times when there are distressed asset sellers. The Treasury market environment allows for particularly clean analytical results and interpretation of these issues, but the lessons learned should have applicability to other fixed income securities and perhaps even to different asset classes.

References

Amihud, Yakov and Haim Mendelson (1986). Asset Pricing and the Bid-Ask Spread. *Journal of Financial Economics 17*, 223-249.

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

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

Gurkaynak, 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.

Gromb, Denis and Dimitri Vayanos (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66, 361-407.

Hu, Xing, Jun Pan and Jiang Wang (2010). Noise as Information for Illiquidity. NBER Working Paper, 16468.

Jordan, Bradford D., Randy D. Jorgensen and David R. Kuipers (2000). The Relative Pricing of U.S. Treasury STRIPS: empirical evidence. *Journal of Financial Economics 56*, 89-123.

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

Krishnamurthy, Arvind (2002). The bond/old-bond spread. *Journal of Financial Economics 66*, 463-506.

Shleifer, Andrei, and Robert W. Vishny (1997). The Limits of Arbitrage, *Journal of Finance 52*, 35-55.

Strebulaev, Ilya (2002). Liquidity and Asset Pricing: Evidence from the U.S. Treasury Securities Market, Working Paper, Stanford University.

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

Vayanos, Dimitri and Jean-Luc Vila (2009). A preferred-habitat theory of the term structure of interest rates, Working Paper, London School of Economics.

Figure 1 – The Arbitrage

This figure presents the time series of the difference between the yields to maturity on two Treasury securities: an original-issue 30 year bond and an original-issue 10 year note. Both securities mature on February 15, 2015. The bond was, originally issued in 1985 with a coupon of 11.25 percent; the note was originally issued in 2005 with a coupon of 4 percent.

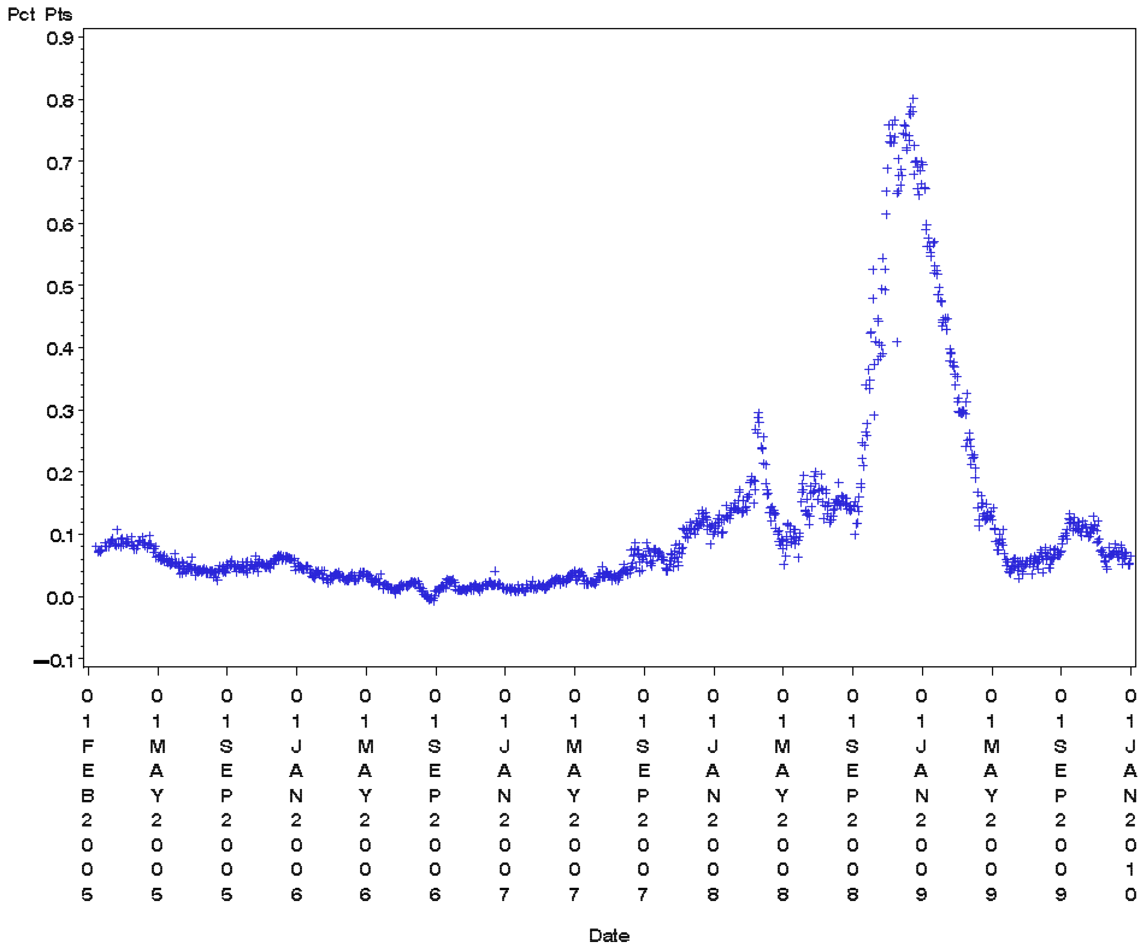

Figure 2 – Pricing Errors by Original-Issue Maturity

This figure presents the one-month rolling averages of the pricing errors across three original-issue maturity buckets: thirty-year bonds, ten- year notes, and five-year notes. The pricing error is defined as the difference between the actual price of the security and the fitted price based on a smooth forward rate yield curve. The vertical axis is measured in percentage points.

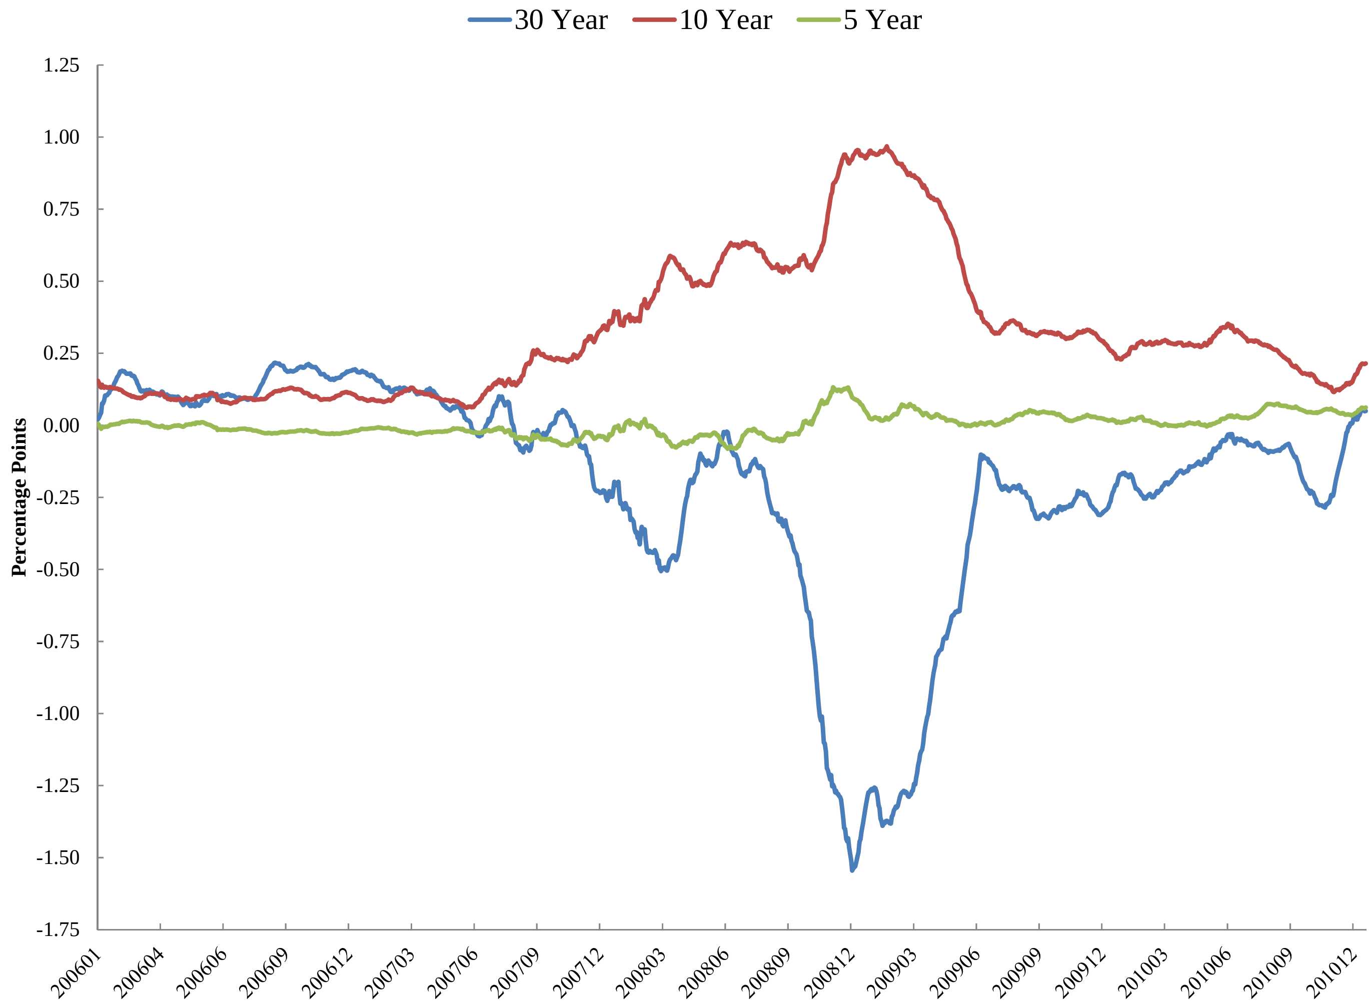

Figure 3 – Arbitrage Profits vs Funding Costs

This figure presents the monthly return on the convergence trade (ignoring the special-GC spread) and the level of the special-GC spread for the bond-note pair maturing on February 15, 2015.

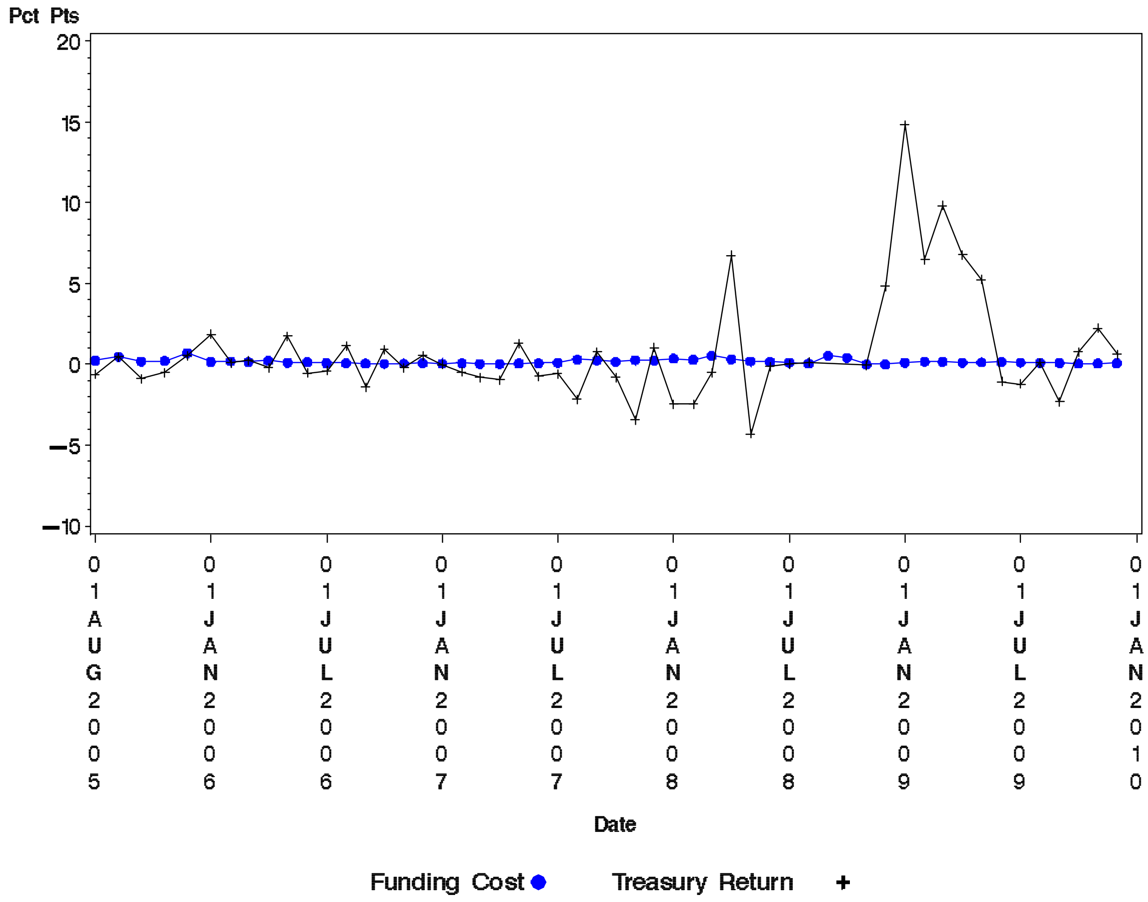

Table 1 – Cross-Sectional Characteristics of Pricing Errors

This table presents a regression of pricing errors on several bond characteristics: ln(outstanding) is the log of the dollar amount of the bond outstanding, ln(bid-ask) is the log of the dollar difference in quoted bid and ask prices, and the other three variables are dummy variables indicating the original issue maturity of the bond. The pricing error is defined as the difference between the actual price of the security and the fitted price based on a smooth forward rate yield curve. The sample period is January 1, 2005 through December 31, 2010. The crisis period is from September 1, 2007 to June 30, 2009. Standard errors (in parentheses) account for clustering within bond cusip and arbitrary heteroskedasticity; ** (*) denotes estimates that are statistically significantly different from zero at the 1(5)-percent level.

	Dependent Variable: Pricing Error	
	Full Sample	Crisis Period
Intercept	-2.232**	-4.505**
	(0.747)	(1.792)
ln(outstanding)	0.136**	0.281**
	(0.044)	(0.107)
ln(bid-ask)	0.034	0.092**
	(0.022)	(0.036)
Original issue 30-year	-0.144	-0.378**
	(0.075)	(0.121)
Original issue 10-year	0.342**	0.671**
	(0.041)	(0.079)
Original issue 5-year	0.070**	0.147**
	(0.023)	(0.054)
R-Square	.231	.407
Observations	149,228	46,192

Table 2 – Time Series Characteristics of Pricing Errors

This table presents a daily time regression of average portfolio pricing errors for nine bond-note pairs on several macro measures of liquidity: the average bid-ask spread on repo transactions (Repo B/A Spread), the repo rate for general collateral Treasuries (GC Repo Rate), and the spread between Libor and the overnight indexed swap rate (Libor-OIS Spread). For each bond-note pair, the pricing error is the price difference between the note and a bond plus a STRIP that gives the identical cash flows to the note. The sample period is January 1, 2005 through December 31, 2010. The crisis period is from September 1, 2007 to June 30, 2009. Newey-West Standard errors (with 30 lags) are in parentheses; ** (*) denotes estimates that are statistically significantly different from zero at the 1(5)-percent level.

	Dependent Variable: Portfolio Pricing Error			
Repo B/A Spread	15.80**			-2.58
	(7.67)			(2.61)
GC Repo Rate		-4.45**		-3.89**
		(1.06)		(1.01)
Libor-OIS Spread			14.03**	9.19**
			(2.90)	(2.11)
R-Square	.231	.406		
Observations	149,231	46,194		

Table 3 – Who Engages in the Arbitrage

This table presents a pooled regression of net purchases of notes less bonds on the pricing error interacted with various characteristics of the insurance companies. Standard errors (in parentheses) account for clustering within insurer and arbitrary heteroskedasticity; ** (*) denotes estimates that are statistically significantly different from zero at the 1(5)-percent level.

	Dependent Variable: Net Purchases
Buy and hold	-12.8**
	3.51
Horizon	7.2**
	1.80
Churn	-3.9**
	1.07
Assets	-3.5**
	1.22
I-grade	-19.2**
	5.00
Premium/Assets	9.2**
	2.12