Journal of Financial Economics 66 (2002) 463–506
JOURNAL OF Financial ECONOMICS
www.elsevier.com/locate/econbase

The bond/old-bond spread ☆

Arvind Krishnamurthy

Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA

Abstract

I document the profits on a trade that is long the old 30-year Treasury bond and short the new 30-year Treasury bond, and is rolled over every auction cycle from June 1995 to November 1999. Despite the systematic convergence of the spread over the auction cycle, the average profits are close to zero. The difference in repo-market financing rates between the two bonds is a significant cost of carry in this trade. I show that variation in the bond/old-bond spread is driven by the Treasury supply of 30-year bonds as well as aggregate factors affecting investors' preference for liquid assets.

JEL classification: G14; G18; E43; E44

Keywords: Hedge funds; Short sales; CP-Bills spread; Aggregate liquidity; Liquidity crises

1. Introduction

Among the most actively traded bonds in the world is the 30-year Treasury bond. Almost since inception the 30-year Treasury bond has served as a benchmark for long term interest rates in the United States.[1] The Treasury has, until recently, auctioned 30-year bonds on roughly a six-month cycle. Upon issuance, a bond is dubbed the "new bond" and acquires benchmark status, replacing the bond that was

☆ I thank Phil Bond, Darrell Duffie, Andrea Eisfeldt, Ravi Jagannathan, Owen Lamont, Bob McDonald, Mitch Petersen, Todd Pulvino, Adriano Rampini, Paola Sapienza, Suresh Sundaresan, an anonymous referee, and seminar participants at the Bank of England, the Texas Finance Festival, Columbia University, University of Chicago, Northwestern University, University of Michigan, WFA 2001 conference, and the NBER 2001 Summer Institute for helpful comments. Amit Khandelwal and Maria Ponomareva provided excellent research assistance. All errors are my own.

E-mail address: a-krishnamurthy@northwestern.edu (A. Krishnamurthy).

[1] The Treasury's recent announcement of eliminating 30-year bonds has caused the market to search for other benchmarks. I discuss issues relating to this in the conclusion.

PII: S0304-405X(02)00207-6

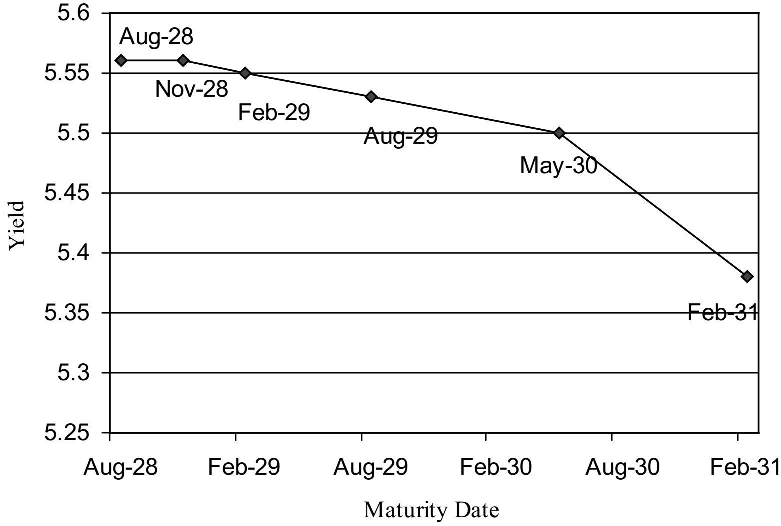

Fig. 1. The yield curve for the 30-year bond sector as of February 9, 2001.

issued six months prior (now the "old bond").[2] Fig. 1 plots the yield curve for the long bond sector as of February 9, 2001.

One of the striking characteristics of the long bond sector is the high premium attached to the new bond. In Fig. 1, the spread between the new and old bond is 12 bps, while the spread between the previous vintage of new and old bonds is only 3 bps (and only one bp going back further). This fact has motivated the following convergence trade:

- Short sell the new bond, upon issuance, at a spread of 12 bps to the old bond.
- Purchase the old bond in order to lock in a bet on the spread.
- Hold the spread position until the next auction date, and then unwind the position at the smaller spread of potentially 3 bps.

Since the spread between the new and old bond will converge toward zero as time passes, shorting the expensive new bond and purchasing the cheaper old bond has the potential to generate trading profits. Indeed, this particular convergence trade has entered the public eye, since it was written about extensively in the business press as one of long term capital management's (LTCM) significant positions.

"One of the fund's main strategies was to exploit tiny differences between the price of a newly issued ("on the run") 30-year American Treasury bond, and a similar one issued previously ("off the run"). There is little economic reason for these bonds to have different yields. Yet off-the-run Treasuries often trade slightly cheaper than on-the-run ones. LTCM bet that their yields would converge by buying off-the-run

[2] The new bond is also often called the "on-the-run bond," or the "benchmark bond," while the old bond is sometimes called the "off-the-run bond." I shall stick to the terminology of old and new bonds in this paper.

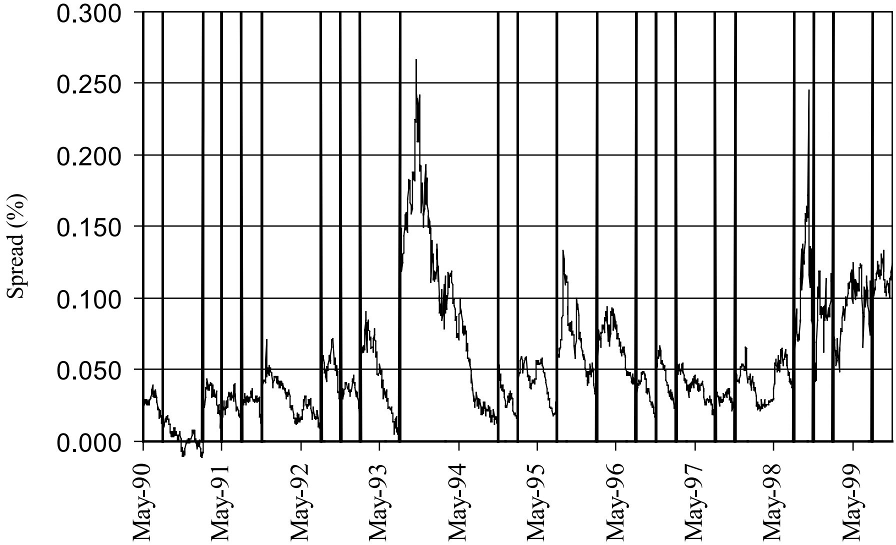

Fig. 2. Yield spread between bond and old-bond. The vertical lines mark auction dates.

Treasuries and selling their on-the-run counterparts short.'' (Economist, October, 17 1998.) Similar descriptions follow in the New York Times and the Financial Times.

Fig. 2 provides a historical record of the spread (old bond minus new bond). The vertical lines mark auction dates. The figure confirms that the spread does typically narrow between one auction date and the next. There are exceptions, notably the Fall of 1998, when we see that the spread moves out following the auction, before subsequently narrowing. This, of course, led to mark-to-market losses for funds such as LTCM holding the convergence trade.

The next two sections shows the profits on a convergence trade of purchasing the old bond and shorting the new bond at the auction date, holding this position until the next auction, then unwinding the position before re-establishing the position in the then current new bond and old bond. Central to the computation and indeed to the entire approach of this paper, is an explicit consideration of the costs of shorting the new bond. In order to establish a short position, an arbitrageur must borrow the new bond in the repurchase market. In doing so, the arbitrageur receives interest (at the *repo rate*) on the short sale proceeds. The repo rate is a market driven rate so that if the bond is in very high demand in the repo market, this interest rate falls below market interest rates. For an arbitrageur holding the convergence trade, a low repo rate (''specialness'') on the new bond translates into a nonnegligible carry cost. Using repo data on the old and new bonds from June 10, 1995 to November 15, 1999, I compute the daily profit/loss on holding the convergence trade over this period. Somewhat surprisingly, the average profit on this trade is close to zero, despite that as in Fig. 2, the spread does converge systematically over the auction cycle.

The culprit is the repo cost. If one ignores the repo market, the trade is profitable. Moreover, at times when the trade looks most profitable (high bond spread) the repo cost is also high. In the Fall of 1998, not only is the bond spread high, but also the repo cost is highest.

My conclusion from these profit results is that betting on convergence of the bond spread, based purely on the ability to commit capital for a long period of time, does not yield excess returns. In order to make profits on convergence, an arbitrageur must engage in some amount of "market timing." For example, betting on convergence only in the period prior to the Fall of 1998 is a winning strategy, but betting on convergence after this period loses money, even if held until the end of 1999. In this sense, convergence seems no different than any other investment strategy.

A high bond spread is consistent with no arbitrage profits, when arbitrage profits are altered to account for repo rates. Sections 5–7 turn to the question of what determines the spread between the old and new bonds. Existing theories ascribe the spread to liquidity differences. The new bond is widely viewed as more liquid than the old bond (e.g., see Fleming, 2001). Amihud and Mendelson (1986), Boudoukh and Whitelaw (1993), and Vayanos (1998) build models where transaction costs create illiquidity and show that an illiquid asset will trade at a discount to a liquid asset that depends on the liquidity preference of investors and the transaction costs. A key prediction of these models is that illiquid assets will offer a higher rate of return than liquid ones.

The existing literature largely verifies this prediction (see Amihud and Mendelson, 1991; Kamara, 1994; or Warga, 1992). Indeed for a larger set of on-the-run/off-the-run pairs of Treasury securities, Warga (1992) shows that investments in the off-the-run generates higher returns than the investment in the liquid on-the-run asset. However, I find that this result is not borne out once one explicitly accounts for the repo rates and the specialness of the new bond. Specialness means that an owner of the new bond can lend these new bonds in the market and earn an additional return equal to the difference between overnight interest rates and the repo rate. In my sample, when specialness is included, the investments in each of the old and new bonds have the same returns. Indeed this is just a restatement of the result that convergence profits are zero. This runs counter to existing theories.

Section 5 builds a model to reconcile the findings.[3] The model borrows from Duffie (1996) and embeds two assumptions. First, at the margin, investments in the two bonds are imperfect substitutes for some set of investors. Second, these investors do not participate fully in the repo market. Under these conditions, I show that the bond spread will be determined down by the marginal valuation of the new bond by these investors. I offer two pieces of evidence on the valuation of these investors.

In Section 6, I show that these investors have a liquidity preference to hold the new bond and that changes in liquidity demand affect the spread. I do this by instrumenting for changes in liquidity demand by using changes in the spread

[3] The theoretical literature on transaction costs and asset prices has had a difficult time justifying the magnitude of liquidity premia that empirical studies have identified (see for example Heaton and Lucas, 1996; Vayanos, 1998; Lo et al., 1999; Huang, 1999). I should note at the outset that this paper offers no new theoretical insights to bridge this gap.

between three-month Commercial Paper (CP) and three-month Treasury Bills (Bills).[4]

Commercial paper is a completely illiquid security. Essentially, it cannot be bought or sold prior to maturity. T-bills, on the other hand, are among the most liquid securities in the marketplace. If investor's prefer liquidity, they will require a premium on CP relative to T-bills. I find that the bond/old-bond spread is high at times when the CP-Bills spread is also high. Moreover this relation varies systematically across the auction cycle. Far from an auction date, the liquidity demand effect is strongest, while close to the auction it is small, but not zero. This is in line with the theory I develop, which suggests that far from an auction the new bond is more liquid than the old bond for a longer period of time, and therefore liquidity should be valued most strongly. The effects are also quantitatively important. Fixing time to be one year from an auction date, a 1% price change in CP relative to T-bills translates to 5% change in the price of the bond relative to the old bond.

Turning from demand to supply, I examine whether changes in the size of the new bond issue has explanatory power for the spread. The liquidity hypothesis rests on nonsatiation. Investors do not have enough new 30-year bonds to satisfy their desire to hold long term liquid assets. This implies that if the Treasury increases the amount of bonds, investors will have more liquidity vehicles and as such the premium on these bonds will fall. The results are supportive of this hypothesis and confirm that smaller issue sizes lead to wider bond/old-bond spreads.

2. Other related literature

While this paper is most directly related to the literature on liquidity in the Treasury bond and other asset markets that I have referenced earlier, it is distinguished from this literature in that it explicitly allows for shorting and utilizes data on repo rates.

Duffie (1996) presents a theoretical model of bond and repo markets, in which shorting is allowed, and both specialness and liquidity issues can co-exist.[5] My empirical analysis is motivated by this work. An important insight of Duffie (1996) is that repo market specialness will also be reflected as a price premium in the cash market. Jordan and Jordan (1997) and Buraschi and Menini (2002) confirm this relation in the data. To a large extent my results on profits—that is the spread in the

[4] While I am the first to use CP-Bill's to proxy for variation in liquidity demand, other papers use different instruments. In the context of the Treasury market, Kamara (1994) uses the volatility of interest rates to show that the spread between short-term Treasury notes and Treasury bills can in part be explained by changes in liquidity preference. His identification assumption is that higher volatility will induce greater liquidity demand. Although an objection one could raise to this identification is that volatility can induce arbitrageurs to cut back on their positions, causing spreads to be higher. In this case the cause is not liquidity demand changing.

[5] Specialness also has implications for term structure modeling. See Barone and Risa (1995) for an example in the context of valuation of floaters.

repo rates deducts a significant fraction of the profits from convergence of the bond/old-bond spread—also validate these conclusions. The case study of mispricing of Treasury long bonds by Cornell and Shapiro (1989) also reaches similar conclusions.

My analysis departs from these papers when I attempt to identify economic factors behind the price premium/specialness. Jordan and Jordan (1997), confirm a conjecture of Duffie (1996), that specialness is also correlated with a high bid-to-cover ratio in an auction (see also Sundaresan, 1994). Duffie's conjecture is that if dealers take short positions in the when-issued prior to the auction and if the bid-to-cover in the auction is unusually high, then many dealers will remain short after the auction. This will create demand for the collateral in the repo market driving repo rates down. This evidence says that when there are lots of shorts, the repo will be on special. This is a correlation between equilibrium prices and quantities in the repo market. The evidence is not helpful in my case because it does not identify the reason behind lots of shorts (high quantity). For example, if bidders expect that fund managers will have high liquidity demand for the new bond, they will setup large shorts in the when-issued and there will be high demand in the auction. Alternatively, if bidders expect that underwriters will be shorting the issue in the secondary market because of hedging demand, then there will be high demand in the auction as bidders will try to acquire the collateral to lend to the underwriters. Either fundamental leads to the same correlation.

Buraschi and Menini (2002) look at the term structure of repo rates, and finding gross violations of the expectations hypothesis, conclude that there must be a liquidity risk premium. One issue they raise is that if technical repo factors that are idiosyncratic to the repo market are responsible for specialness, then the risk should not be priced. Since I show that aggregate liquidity factors play an important role in the bond/old-bond spread, one would expect the risk to be priced.

D'Avalio (2002), Lamont and Thaler (2002), and Geczy et al. (2002) investigate stock price anomalies in the context of short selling costs. D'Avalio (2002) and Geczy et al. (2002) provide direct evidence that the short rebate (the analogue of the repo cost for stocks) can help explain these anomalies. Lamont and Thaler (2002) impute short rebates from option prices and show that the estimated short rebate can help explain market mispricing in carve-outs (i.e. Palm/3Com). The evidence in these papers on stock market shorting costs is complementary to the bond market evidence I provide.

My paper also relates to recent theories on convergence trading and hedge funds. These papers argue that volatility in spreads (say, in the Fall of 1998) is due to the fact that arbitrage investors (i.e. hedge funds) become capital constrained (see for example, Shleifer and Vishny, 1997; Xiong, 2001; Yuan, 1999). The time variation in my profit series and correlation with the stock market does coincide with these theories. However, while these theories ascribe crises to shortages of arbitrage capital ("liquidity supply"), I show that there is substantial variation in the bond spread that owes purely to changes in liquidity demand. For example, in the Fall of 1998, both bond/old-bond spread, and the CP-Bills spread widened substantially. Moreover the spread in repo rates between bond and old bond also widened substantially. This is a sign that arbitrageurs were actively involved in taking the convergence position. Thus, it remains a question as to how much of this rise in spreads was driven by arbitrage capital and how much by factors such as liquidity demand.

My approach of identifying liquidity demand is related to Chordia et al. (2000) (see also Huberman and Halka, 2001; Hasbrouck and Seppi, 2001; Fleming, 2001). Their paper looks for comovement in micro-structure motivated liquidity measures, and argues that since there is much comovement during periods such as the Fall of 1998, there must be some aggregate liquidity effects. Since I use the CP-Bills spread to measure liquidity demand, my analysis moves one step closer to establishing what sort of macroeconomic effects are behind this comovement. I connect the liquidity literature to the macroeconomics literature on the CP-Bills spread (see Stock and Watson, 1989; Friedman and Kuttner, 1992, 1993; Kashyap et al., 1993). However, unlike this literature, I show that variation in CP-Bills is due to variation in investor liquidity demand.

3. Trade mechanics

The convergence trade involves shorting the new bond and purchasing the old bond on an auction date, financing these positions in the repo market, and unwinding the position on the next auction date. Before moving into details, it will be helpful if I briefly describe how an arbitrageur establishes short positions by using the repo market. Much of what I discuss regarding the repo market is explained in far greater detail in Duffie (1996) (see also Stigum, 1989). I cover the minimum required to motivate the empirical analysis.

3.1. The repo market

To create a short position in a bond, an arbitrageur must execute both a sale as well as a reverse-repo. In the Treasury market the sale occurs today, with settlement tomorrow. Tomorrow, when it is time to deliver on the sale, the arbitrageur borrows the bond from an investor who owns it, agreeing to return the bond to the investor the following day. The borrowing transaction is known as a reverse-repo. A typical overnight reverse will have the arbitrageur depositing cash equal to the value of the bonds ($P(t)$) with the party with which the reverse is conducted, and receiving in return the bonds. Settlement on this transaction is the same day.[6] If the short is reversed tomorrow, the arbitrageur purchases back the bonds for settlement the following day and delivers the bonds against the maturing overnight reverse, receiving back the cash that was deposited plus interest.

The interest rate paid on the cash is denoted the overnight repo or financing rate ($f(t)$). In Treasury bonds, the repo market is an active market with both direct as well as brokered transactions. According to Stigum (1989), the brokered market

[6] Most repo's also require a haircut to be left with the repo dealer as a credit margin. For the analysis that follows I assume haircuts are 0%.

transacts upwards of \$10 billion per night, with brokers taking an average spread of 5 bps, while the direct market transacts several multiples of this amount.

An important aspect of the repo market is market-wide limitations on the number of new shorts established in a particular bond any given day. Because of the fact that every short has to be matched by a borrowing of the bond in the repo market, the shorts are tied to the quantity of bonds available for borrowing. For example, if we take a case where the repo market clears only once a day, then the theoretical maximum number of bonds that can be borrowed is equal to the issued amount of bonds. Thus if the float of outstanding bonds today is \$10 billion, then it is not possible for investors to increase short positions by more than \$10 billion today. In practice, there are two considerations that affect this constraint. First, it is actually tighter, as only a fraction of this \$10 billion of bonds are available for lending in the repo market. Second, the repo market clears throughout the day so that this smaller float can service a larger number of shorts.

In a normal market, when the supply constraint does not bind, the overnight financing rate will be equal to the overnight riskless interest rate. An investor who has a bond that is available for lending is indifferent between lending this bond at the riskless rate, and keeping the bond until tomorrow. This is because he could always take the funds received on lending the bond, invest them at the riskless rate, and be in precisely the same position as if he had not lent the bonds. When the supply constraint binds, the financing rate falls below the riskless rate in order to ration the scarce supply of bonds available for borrowing. An investor owning these bonds earns a premium on lending the bonds and is able to borrow funds at the low rate of $f(t)$ and invest this at the riskless interest rate of $r(t)$. This situation is known as a ‘‘special,’’ and the extent of specialness can be measured as

$$s(t) = r(t) - f(t). \tag{1}$$

Fig. 3 represents these two alternatives. S is the supply curve for bonds available for lending in the repo market. D represents demand in a normal market. In this case the supply constraint is not binding and $r(t) = f(t)$. D' represents demand when the bond is on special. In this case the supply constraint limits short positions and $f(t) < r(t)$ to ration the scarce supply of bonds.

The repo market is really an overnight rental market for bonds. The rental price is $s(t) = r(t) - f(t)$. There is a finite supply of the bonds available for renting. When demand exceeds this supply, the rental price rises. The owner of the bond then becomes entitled to an additional benefit of ownership. In addition to any capital gains or coupons, the owner is able to rent his bonds out in the repo market and capture the additional benefit of $s(t)$.

It is important to recognize that the specialness is security specific. Only bonds that are in high demand in the repo market will have repo rates below overnight riskless rates; bonds that are not having repo rates close to the riskless rate. The repo rate on a generic non-special bond is know as the general collateral rate.

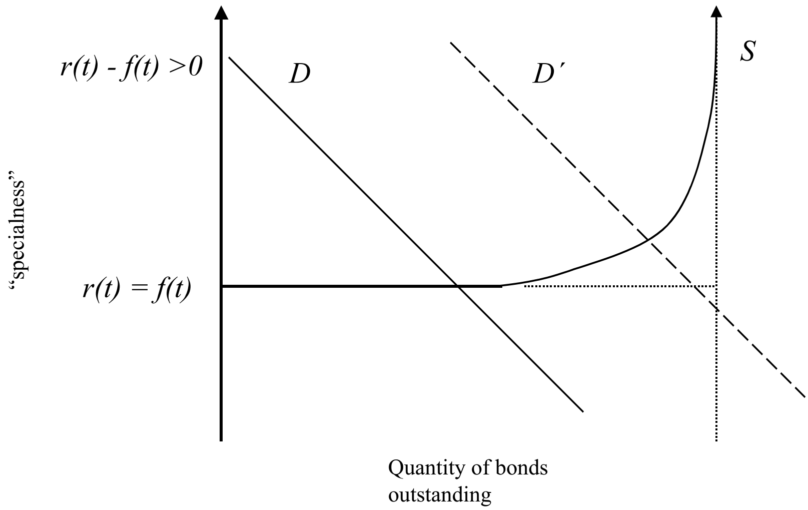

Fig. 3. Diagram of equilibrium in the repo market.

3.2. The convergence trade

Define the following quantities for one auction cycle:

- Let $t_n, n = 1, \ldots, N$, correspond to the sequence of business days on which the two bonds are traded. Thus, if n is a Friday and $n+1$ a Monday, then $t_{n+1} - t_n$ is equal to 3. t_1 corresponds to the first trading day following the auction, and t_N is the first trading day following the next auction.
- Let $\theta(t_n)$ be the number of units of old bonds held from date t_n until date t_{n+1}. Likewise, let $\hat{\theta}(t_n)$ be the holdings of new bonds.
- The price of the old bond at t_n is $P(t_n)$. This price includes accrued interest, and is for standard settlement, which in the US Treasury market means the next business day (i.e. t_{n+1}).
- The repo rate for either borrowing or lending the old bond from t_{n+1} until t_{n+2} is $f(t_n)$.
- Let $y(t_n)$ be the yield to maturity of this bond, and let, $DP(t_n) = (\partial P(y, t_n)/\partial y)|_{y=y(t_n)}$. That is DP is the derivative of P with respect to a change in the yield to maturity.
- Finally, define all of these same variable for the new bond as well, distinguished with a "hat."

The profit from purchasing $\theta(t_n)$ units of the old bond on t_n, financing this purchase at the repo rate of $f(t_n)$ and unwinding the position at t_{n+1} is:

$$\theta(t_n)\left(P(t_{n+1}) - P(t_n) - P(t_n)f(t_{n+1})\frac{t_{n+2} - t_{n+1}}{360}\right). \tag{2}$$

The first term, $P(t_{n+1}) - P(t_n)$, is the capital gain from holding this bond. The second term is the interest that must be paid to finance the purchase of the bond. In writing Eq. (2), I have assumed no coupon payments occur on t_{n+1}. Note that, because of settlement details, this cash flow is realized on date t_{n+2}.

Likewise the profit on a short position of $\hat{\theta}(t_n)$ units of the new bond on t_n, with a reverse repo executed at $\hat{f}(t_n)$ is

$$-\hat{\theta}(t_n)\left(\hat{P}(t_{n+1}) - \hat{P}(t_n) - \hat{P}(t_n)\hat{f}(t_{n+1})\frac{t_{n+2} - t_{n+1}}{360}\right). \tag{3}$$

The total profit is just the sum of the above two expressions. Let $\pi(t_n)$ be this total profit.

It is common in establishing these trades to choose relative positions sizes so that profits are invariant to an equal level change in the yield of each bond, and are only responsive to a change in the yield spread. Thus,

$$\hat{\theta}(t_n)\widehat{DP}(t_n) = \theta(t_n)DP(t_n). \tag{4}$$

This still leaves one free variable, $\theta(t_n)$. Rather than choosing this to be constant across time, I have chosen to set the product of $\theta(t_n)$ and $DP(t_n)$ to be constant over time. This implies that the sensitivity of profits to a change in the yield spread will be constant. No doubt a more complex strategy could be devised in which $\theta(t_n)$ was dependent on some variables that forecast risk and/or return. Precisely, I have chosen $\theta(t_n)DP(t_n)$ to be equal to one thousand. The scaling means that a ten basis point change in the spread will translate to a profit of approximately $1{,}000 \times 0.001 = \$1$. In the sample, this implies that θ averages about \$75 face value of bonds.

4. Profits

The data comprises the period June 12, 1995 to November 15, 1999 and includes 1,157 trading days. Bond prices for both new and old 30-year bonds are from Bloomberg (GovPX) and Bridge. The price data is taken as the midpoint of bid and offer quotes of a sample of dealers. The repo market overnight financing rates for these bonds is from a dealer. Summary statistics are reported in Table 1.

Table 1
Data summary
Summary statistics for 30-year bond and old-bond data. Data covers the period June 12, 1995 to November 15, 1999 and includes 1,157 trading days. Bond prices are from Bloomberg (GovPX) and Bridge. The Price data is taken as the midpoint of bid and offer quotes of a sample of dealers. The repo market overnight financing rates for these bonds is from a dealer.

Number of auction cycles	11
Average size of issue	\$11.40 bn
Average days between auctions	82
Average yield spread (old bond–bond)	6.25 bps
Average new bond repo rate	4.31%
Average old bond repo rate	4.55%
Average general collateral repo rate	5.22%

There are 11 auction cycles in the sample. The smallest auctioned amount is the 6.625 of February 15, 2027 bond at \$10.45 bn. There is one reopening. The 6.125 of November 15, 2027 was reopened to create a total issue amount of \$22.51 bn. Both old and new bonds are on special as they have financing rates significantly below the general collateral rate. More importantly for my analysis, there is a significant spread between bond and old-bond repo rates.

4.1. Average profitability

Profits are computed as per Eqs. (2) and (3). In order to minimize the effects of bid–ask bounce, stale quotes, and other transaction data specific issues, the daily profits are aggregated to a weekly figure. This is done by rolling over daily profits at the riskless rate (GC repo rate), and arriving at a profit figure for each week (ending Friday).

The average profits are 0.0016. Let me put this number in perspective. This is a weekly profit number on a position which is long an average of \$75 of old bonds and short \$75 of new bonds. Thus the 0.0016 number is the sum of the gain/loss for the week on being long \$75 of old bonds plus the weeks gain/loss on being short \$75 of new bonds. One way of interpreting this number is to divide by 75 to arrive at weekly return on the position size, and then multiply by 52 to annualize this figure. Doing this gives a figure of 11.1 bps per annum per \$1 of bond.

The profit estimate is surprisingly low. In fact, with five years of data, one cannot reject the hypothesis that true profits are equal to zero (the standard error of the estimate of this mean is 34 bps, and the autocorrelation in weekly profits is -0.033). This is despite the fact that I have assumed all transactions are done at the midpoint, rather than explicitly factoring in bids and asks. Indeed, the only real-world cost I have introduced into the computation is that of financing these bonds.

As another point of comparison, if the financing costs on each of the bonds was equal, and set to the overnight GC repo rate, this same number would be 33.3 bps per annum per \$1 of bond.

4.2. Time variation in profits

The solid line in Fig. 4 is the eight-week moving average profits/losses from the strategy. Note the drop in profits in the latter part of the period, coinciding with the financial crises of the Summer and Fall of 1998. Average profits for the period ending May 1998 was 0.0125 per week (87 bps per annum), while in the period after, average profits fell to -0.0184 per week (-128 bps per annum). One can see this quite clearly from Fig. 2. In the period after Summer 1998, the bond spread was higher but did not have the same systematic convergence of earlier periods, so that the convergence strategy did not perform as well.

A second, and less obvious source, of this change in profits was the spread in repo rates. From Eqs. (1) and (2), one can see that the difference between these repo rates is a carry cost in holding the convergence trade. The dashed line in Fig. 4 is the

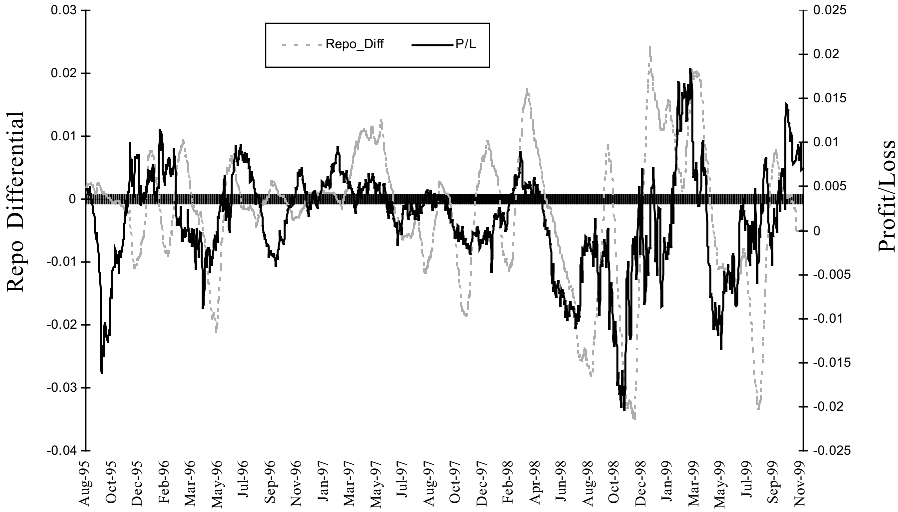

Fig. 4. Eight-week moving average P/L and repo differential on convergence strategy.

Table 2
Profits by auction cycle
The second column reports profits from the convergence trade of being short the 30-year bond and long the old-bond, broken down by auction cycle. The last two columns report the average spread in yields and repo rates between the bond and old-bond, for each auction cycle.

Cycle	Profits (bps per annum)	Average bond spread	Average repo spread
1	187	3.3	6
2	98	7.0	16
3	−8	6.8	20
4	115	3.7	9
5	128	4.1	−5
6	145	4.0	−23
7	−48	3.0	70
8	−41	4.1	45
9	−286	11.1	158
10	−22	8.9	−130
11	−39	9.0	59

moving average differential (new – old) in the repo rates between the two bonds. The correlation coefficient between the two series plotted in Fig. 4 is 0.40.

Table 2 gives a breakdown of the profits by auction cycle. In the period prior to May 1998, the spread in repo rates between new and old bonds averaged 5 bps (old > new). In the period after, the average spread jumped to 55 bps. In the period from May 1998 to November 1998 (cycle 9), the spread averaged 158 bps. The variation in repo rates is of the same magnitude as the variation in profits, underlining its importance in the profitability of the convergence trade.

Table 3
Profits from timing strategy
Profit from timing strategy where the convergence trade is entered for one week if the bond spread at the beginning of the week exceeds the cutoff reported in column 1. Column 2 reports the average profits over the weeks that the trade is entered into. The last column is the number of data points in the average.

Cutoff	Profits (bps per annum)	N (weeks)
0	11.1	232
3	16.4	196
5	26.7	127
7	38.9	68
9	214.2	52
11	336.0	28
13	139.1	4

The profits reported above arise from following a brute-force convergence strategy. Enter the trade for each auction cycle and roll it over into the next auction cycle. In practice, an arbitrageur engaging in this trade will adjust the size of her trade depending on state variables. Table 3 presents the result from a simple market timing strategy: Enter the trade if the bond spread exceeds a cutoff value and hold it as long as this true. The table reports profits for a number of different cutoffs. As one can see, profits are higher when a higher cutoff is chosen, suggesting a "smarter" arbitrageur can achieve higher profits. There are two caveats here. First, the high profit samples (cutoffs over 9 bps) are fairly small. For 50 weeks the standard error of the mean estimate is about 75 bps. Second, eyeballing the bond spread data, this strategy seems to be picking off a few observations where the bond spread jumps up for a few weeks before settling back down. Since I picked these cutoffs after observing the data, there is a selection bias.

4.3. Adjusting profits for systematic risk

The fall in profits in Fall of 1998 suggests that profits are related to aggregate events such as the liquidity crises. In this case, average profit is not the appropriate benchmark of profitability, and these profits need to be adjusted down for bearing systematic risk.

Tables 4 and 5 report the results of regressions of the weekly profits on the CRSP weekly returns, using both a linear and a nonlinear specification. In model 1, profits are regressed on excess stock returns, yielding a small, but positive β coefficient. Model 2 (cutoff $= 0$) clarifies this further. The coefficients on the independent variable are allowed to vary depending on whether or not the return in any week is positive or negative. The coefficient on the positive moves is close to zero, while the coefficient on negative moves is higher.[7] Table 5 clarifies this further, by regressing profits on the payoff of one week put-like securities. In model 3, profits are regressed

[7] Mitchell and Pulvino (2001) find a similar relation between profits and aggregate returns in their study of the profitability of risk (merger) arbitrage.

Table 4
Profits and aggregate returns
This table reports regressions where the dependent variable is the weekly profit from the convergence strategy, and the independent variable is the weekly excess return on the CRSP universe of stocks over the federal funds rate. Model 1 is a linear regression where the independent variable is the CRSP excess return. Model 2 allows the coefficient on positive CRSP excess returns to differ from that on negative CRSP excess returns.

Model 1:

$$\pi_t = \alpha + \beta\, CRSP_t + \varepsilon_t,$$

π_t = weekly profit from the convergence strategy,
$CRSP_t$ = weekly excess return on CRSP.

α	β	R^2
0.0002 (0.03)	0.462 (1.12)	0.0075

Model 2:

$$\pi_t = \alpha + \beta^H CRSP_t^H + \beta^L CRSP_t^L + \varepsilon_t,$$

$CRSP_t^H = CRSP_t$ if $CRSP_t \geqslant 0$ and 0 otherwise,
$CRSP_t^L = CRSP_t$ if $CRSP_t < 0$ and 0 otherwise.

α	β^H	β^L	R^2
0.006 (0.52)	0.114 (0.17)	0.86 (0.90)	0.009

Number of observations = 232. t-statistics in parentheses are based on robust standard errors (White, 1980).

Table 5
Estimating the size of the embedded put
This table reports regressions where the dependent variable is the weekly profit from the convergence strategy, and the independent variable is a function of the weekly excess return on the CRSP universe of stocks over the federal funds rate. Model 3 is a linear regression where the independent variable is the payoff on one-week put option on the CRSP return struck one standard deviation out-of-the-money the CRSP. Model 4 instead uses the payoff on a digital put as independent variable.

Model 3:

$$\pi_t = \alpha + \beta^L CRSP_t^L + \varepsilon_t,$$

$CRSP_t^L = \min[-\sigma(CRSP) - CRSP_t, 0]$,

$\sigma(CRSP)$ is the standard deviation of weekly CRSP returns.

α	β^L	R^2
0.0034 (0.48)	−1.59 (−0.73)	0.019

Model 4:

$$\pi_t = \alpha + \beta^L CRSP_t^L + \varepsilon_t,$$

$CRSP_t^L$ = dummy variable for $CRSP_t < -\sigma(CRSP)$.

α	β^L	R^2
0.0071 (0.99)	−0.0458 (−1.85)*	0.019

Number of observations = 232. *Significant at 10% level. t-statistics in parentheses are based on robust standard errors (White, 1980).

on the payoff of a put option struck one standard deviation out of the money. The results are weak, but encouraging. The best results are in model 4, where profits are regressed against the payoff on a digital option that pays \$1 if the weekly CRSP return is one standard deviation below zero. The results are stronger, but still only significant at the 10% level. I think the improved results owe mainly to the fact that losses on the convergence strategy are naturally bounded (i.e. the spread between bond and old bond does not get too far apart). On the other hand, the returns on the stock market can be quite negative.

There are two points to take away from Table 5. First, there is a correlation with the market implying that there is some systematic risk involved in this convergence trade and that profits should be corrected down for bearing this risk. Second, the correlation with the market resembles an out of the money put option. I expand on this in the discussion that follows.

If the systematic risk embedded in this trade can be expressed as a linear function of the market return, then the correct profit measure must deduct β times the market risk premium. In my case, the systematic risk is a nonlinear function of the market return, so the correction is a little more involved.

A simple way to correct for the nonlinear systematic risk is to value the embedded option and price it out using an asset pricing model. From model 4 in Table 5, we see that the convergence strategy profits on average 0.0071 per week, but has an embedded short position of 0.0458 digital put options. Using a standard option pricing model with sample average interest rate and stock market volatility, the option is valued at 0.156.[8] Thus,

$$\text{Cost of option} = 0.156 \times 0.0458 = 0.0072. \tag{5}$$

Subtracting this option cost from 0.0071 gives an excess profit of 0.001 (5 bps per annum). Notice that this is approximately 5 bps below the sample average profits. The reason is that the fraction of the profits that can be directly attributed to taking systematic risk has been deducted.

4.4. Discussion

I draw a few conclusions from these results. In contrast to implications of the rhetoric surrounding hedge-fund convergence strategies, convergence is not the sole source of profits for the bond/old-bond trade over my sample. While profits are positive over the early part of the sample, they are significantly negative over the latter part. Thus more important than taking sure bets on convergence is timing the market. In other words, taking positions in order to benefit from rates of convergence faster than what the market has priced in via the repo rates. In this sense, convergence strategies seem to be no different than any other market timing strategy. Finally, the results do lend credence to a commonly held view that hedge funds who put on these convergence trades were simply selling insurance against low probability events. Such events have come to transpire since the Fall of 1998.

[8] See Appendix A of the earlier version of this paper for the details.

This still leaves many questions unanswered, however. What are these low probability events and why did convergence breakdown in the latter half of the period? Many observers have described this as a liquidity crisis, but it is unclear what exactly is meant by this term. There are a number of recent papers that ascribe these events to a shortage of informed capital. During this time, arbitrageurs saw their capital diminish and were forced to cut back on arbitrage activities causing a widening of spreads. Is there any merit in this explanation, as applied to the bond/old-bond spread? Moreover, while there seems to be a correlation with the aggregate stock market, what is behind this? The rest of this paper delves a little more deeply into the long-term bond market, seeking to answer these questions by identifying economic factors that drive the variation in the bond/old-bond spread.

5. A model

Let us consider a simple representation of the convergence trade in order to arrive at expressions relating convergence to repo rates. The two assets under consideration have prices at time t given by $P(X,t)$ and $\hat{P}(X,t)$. $X(t)$ is a vector summarizing the state of the world at time t. At date T these two assets will converge in price so that,

$$P(X,t) = \hat{P}(X,t), \forall t \geqslant T. \tag{6}$$

However, at $t < T$ these assets can have different prices. Let us also suppose that the financing costs for these bonds in the repurchase market are given by $f(X,t)$ and $\hat{f}(X,t)$, respectively. For simplicity, assume that the assets pay no coupons for any $t \leqslant T$.

Consider a version of the convergence trade in which $\theta(t) = \hat{\theta}(t) = 1$. Since the assets are almost identical, their durations will also be close to identical and hence similar to the strategy of the previous section. At date 0, a trader enters into the convergence trade by shorting the expensive asset and purchasing the cheaper asset. Both sides of this trade are financed in the repurchase market, so that the trade is completely self-financing. Without loss of generality, suppose that the expensive asset is earmarked by the hat. At date t, she shorts one unit of the expensive asset at price $\hat{P}(X,t)$, investing the short proceeds at the financing rate of $\hat{f}(X,t)$, while simultaneously purchasing one unit of the cheaper asset at $P(X,t)$ and financing this at $f(X,t)$. The profit from this trade at date $t+1$ is given by

$$\pi(X,t+1) = \Delta P(X,t) - \Delta\hat{P}(X,t) - P(X,t)f(X,t) + \hat{P}(X,t)\hat{f}(X,t), \tag{7}$$

where, $\Delta P(X,t) = P(X,t+1) - P(X,t)$, and likewise for $\Delta\hat{P}(X,t)$. To save notation I denote the financing cost in terms of the time increment. For example, if it is 5% and each time period is one week, then $f = 0.05 \times 7/360$.

A1 (Risk-neutral arbitrageurs). *There is a large class of risk-neutral, wealth unconstrained, arbitrageurs in this economy who trade at prices* $P(X,t)$ *and* $\hat{P}(X,t)$ *and borrow and lend against the bonds at* $f(X,t)$ *and* $\hat{f}(X,t)$. *As a result,*

$$\mathrm{E}_t[\pi(X,t+1)] = 0. \tag{8}$$

The condition requires that there are no excess profits at any date t by investing in the convergence trade. If there are, the arbitrageurs would enter the market and drive these profits to zero. Eq. (8) can be rewritten as

$$\begin{aligned} \mathrm{E}_t[\Delta\hat{P}(X,t) - \Delta P(X,t)] &\equiv \mathrm{E}_t[\Delta SPREAD(X,t)] \\ &= (\hat{P}(X,t)\hat{f}(X,t) - P(X,t)f(X,t)). \end{aligned} \tag{9}$$

This is essentially what we verified in the previous sections.[9] This can be written so that the actual change in the spread is,

$$\Delta SPREAD(X,t) = (\hat{P}(X,t)\hat{f}(X,t) - P(X,t)f(X,t)) + \varepsilon(X,t), \tag{10}$$

where $\varepsilon(X,t)$ is a zero mean innovation that is uncorrelated with the other right-hand-side variables. Summing this expression from 0 to T, taking expectations, and using the terminal condition that $SPREAD(X,T) = 0$, yields,

$$\hat{P}_0 - P_0 = \sum_{t=0}^{T-1} \mathrm{E}_0[P(X,t)f(X,t) - \hat{P}(X,t)\hat{f}(X,t)]. \tag{11}$$

This expression is fairly intuitive. It tells us that if $f(X,t) > \hat{f}(X,t)$, then $\hat{P}_0 - P_0 > 0$. That is, the bond that has a lower financing cost in the repo market is also the one which is more expensive in the cash market. The relation is quite similar to uncovered interest parity in the foreign exchange market. The two bonds are like two currencies, with interest rates equal to the negative of the financing rates. Then $\hat{P}_0 - P_0$ is equal to the expected depreciation of the expensive bond from 0 until T.

5.1. Why is there a spread?

One of the lessons of the model is that absence of arbitrageur profits is not a strong enough condition to pin down the bond spread. As long as there is a corresponding spread in repo rates, a high bond spread is perfectly consistent with A1. *But this just restates the puzzle*. Why is the normal state of affairs a high bond spread and a high spread in repo rates, as opposed to no spread in either bonds or repo rates? After all, from a cash flow standpoint, the new and old bonds are virtually identical.

Since Eq. (8) does not pin down the bond spread, we need a condition determined by a different set of investors to pin this down. This is not to say Eq. (8) does not hold. The empirical results of Section 4 suggest that arbitrageur profits are zero so that Eq. (8) does hold. It is just that there must be an additional condition determined by a different set of investors that involves only bond spreads or only repo spreads that also holds.

[9] Note that A1 suppresses all issues of aggregate risk premia in asset prices. As we saw in the previous sections, returns on the bond/old-bond convergence trade are weakly correlated with aggregate returns. The simplification is made primarily for expositional purposes.

5.2. Which models pin down a spread?

Market participants often point out that liquidity differences can account for differences in the Treasury bond prices. Kamara (1994) offers evidence that the spread between short term Treasury notes and Treasury bills can in part be explained by liquidity. Indeed, the idea that liquidity differences between assets can lead to different prices is an old one. Amihud and Mendelson (1986), Boudoukh and Whitelaw (1993), and Vayanos (1998) build models where transaction costs create illiquidity and show that an illiquid asset will trade at a discount to a liquid asset that depends on the liquidity preference of investors and the transaction costs. A key prediction is that illiquid assets will offer a higher rate of return than liquid ones.

There are two difficulties in directly applying this logic to the bond/old-bond setting. First, the results of Section 4 suggest that Eq. (8) holds. The returns to purchasing the new bond and lending it out to capture the specialness premium is equal to the returns to purchasing the old bond and lending it out. Trading at prices in my data and lending at repo rates in my data does not give any credence to the idea that the illiquid asset provides a higher return. Note that this is in contradiction to earlier studies, such as Warga (1992), in which the illiquid asset yields a higher return than the liquid asset.[10] The reason is that these studies neglect the returns to lending the bonds out in the repo market. Since the new bond is more special than the old bond, the profits on lending the new bond makes up for the fact that the new bond provides a lower capital gain.

Second, since the papers I reference above assume a short sales constraint and thereby do away with repo markets, the question arises as to whether introducing a repo market will negate liquidity premia. In the example below, I show that even if we assume that only arbitrageurs can actually trade at the prices in my data sample, so that to other investors the old bond actually does look less liquid, this illiquidity assumption has no pricing implications.

Take the following example, suppose that there is a group of investors who wish to hold long-term liquid Treasury bonds in their portfolios. A standard way of modeling liquidity considerations is to introduce transaction costs. As such suppose that only arbitrageurs can trade at prices P and $\hat{P}$ (so that Eq. (8) holds) but that other investors only trade new bonds at P but pay some transactions costs when trading old bonds so that their effective prices differ from $\hat{P}$. The expected return to holding a unit of the new bond and simultaneously lending these bonds in the repo market is given by

$$\mathrm{E}_t[\Delta P(X,t)] - P(X,t)f(X,t). \tag{12}$$

[10] Amihud and Mendelson (1991) consider the convergence strategy of trading same maturity Treasury notes against Treasury bills. As in Kamara (1994), there is a spread between these two assets that can be partly ascribed to liquidity. Amihud and Mendelson (1991) include practical costs such as bid–ask spreads and shorting costs in determining arbitrage profits. The authors do not have data on repo rates, but recognize that this will affect the profits. They show that a 50 bps shorting cost would be consistent with absence of arbitrage profits.

Eq. (8) can be written as:

$$\mathrm{E}_t[\Delta P(X,t)] - P(X,t)f(X,t) = \mathrm{E}_t[\Delta \hat{P}(X,t)] - \hat{P}(X,t)\hat{f}(X,t). \tag{13}$$

If the liquidity investors value their return as, $\mathrm{E}_t[\Delta P(X,t)] - P(X,t)f(X,t)$, then via the arbitrageurs indifference condition, their investment is exactly *as if* they could trade at the arbitrageurs prices of $\hat{P}$ in the illiquid old bond. In other words, with just the assumptions outlined above, illiquidity of the old bond has absolutely no pricing implications, and the bond spread is still indeterminate.

The reason for this is clear. The valuation of the investment for the liquidity driven investor is $\mathrm{E}_t[\Delta P(X,t)] - P(X,t)f(X,t)$, which is the same term that enters the arbitrageurs objective.

5.3. A model in which investors value liquidity and do not fully participate in the repo market

Given these difficulties in appealing to the results of the models referenced above, let me turn to a model in which repo markets and liquidity considerations do co-exist. The model borrows from Duffie (1996). I outline a partial equilibrium model. For the interested reader, I have worked out a fully specified general equilibrium model in Appendix A.

As Duffie (1996) notes that there are a number of institutions that take long positions in bonds but are restricted from participating in the repo market and lending these bonds. This clearly violates A1, however by itself it has no bite. All things being equal, these investors should simply purchase the cheaper old bonds at $\hat{P}(X,t)$, thereby forgoing the premium on lending the new bonds. This action would push the bond spread toward zero as opposed to increasing it. On the other hand if these same investors preferred for liquidity reasons to hold the new bond, then this preference could drive a wedge between these bond prices. A common example of these investors are bond-market mutual funds. These funds typically try to track or beat bond indices that are constructed to include on-the-run liquid bonds. Thus flows in and out of these funds constitute changes in demand for the on-the-run securities. Second, these funds typically have mandates that limit the amount of leverage they can achieve. A leverage limitation is mechanically equivalent to limiting the mutual fund to repo out only a portion of owned securities.

Consider this last situation. Suppose that investors can only lend out a fraction $\lambda < 1$ of the bonds in their portfolios. Defining specialness (as before) as the difference between $r(X,t)$, the overnight general collateral rate, and $f(X,t)$, the specific collateral financing rate, their excess return (over $r(X,t)$) can be rewritten as

$$\Delta P(X,t) + \lambda P(X,t)s(X,t) - P(X,t)r(X,t). \tag{14}$$

Thus if the bond is on special, the investors only capture λ of the specialness premium in their excess return.

Now the question is whether $\lambda < 1$ allows for an equilibrium in which liquidity preferences of investors will affect spreads. As such, imagine an equilibrium in which the marginal liquidity investor assigns a positive utility value, $v(X,t)$, to holding an additional unit of exposure to the liquid new bond rather than the illiquid old bond. Then this indifference condition gives

$$\begin{aligned} &\mathrm{E}_t[\Delta P(X,t)] + \lambda P(X,t)s(X,t) - P(X,t)r(X,t) + v(X,t) \\ &\quad = \mathrm{E}_t[\Delta \hat{P}(X,t)] + \lambda \hat{P}(X,t)\hat{s}(X,t) - \hat{P}(X,t)\hat{r}(X,t). \end{aligned} \tag{15}$$

($v(X,t)$ can also be thought of as a convenience yield to holding the new bond.)

Let me show that these assumptions are consistent with a determinate bond spread. Indifference for the arbitrageur gives

$$\begin{aligned} &\mathrm{E}_t[\Delta P(X,t)] + P(X,t)s(X,t) - P(X,t)r(X,t) \\ &\quad = \mathrm{E}_t[\Delta \hat{P}(X,t)] + \hat{P}(X,t)\hat{s}(X,t) - \hat{P}(X,t)\hat{f}(X,t). \end{aligned} \tag{16}$$

Combining these last two conditions, allows us to determine specialness,

$$P(X,t)s(X,t) - \hat{P}(X,t)\hat{s}(X,t) = \frac{v(X,t)}{1-\lambda}. \tag{17}$$

Eq. (11) can be written as

$$P_0 - \hat{P}_0 = \sum_{t=0}^{T-1} \mathrm{E}_0[\hat{P}(X,t)(r(X,t) - \hat{s}(X,t)) - P(X,t)(r(X,t) - s(X,t))].$$

For $P(X,t)$ close to $\hat{P}(X,t)$, we can rewrite this purely in terms of the difference in specialness of the two bonds,

$$P_0 - \hat{P}_0 \approx \sum_{t=0}^{T-1} \mathrm{E}_0[P(X,t)s(X,t) - \hat{P}(X,t)\hat{s}(X,t)] = \sum_{t=0}^{T-1} \mathrm{E}_0\left[\frac{v(X,t)}{1-\lambda}\right]. \tag{18}$$

5.4. Testing the model

Eq. (18) is a model that describes how $v(X,t)$ and λ affect the bond spread. It embodies two assumptions. First, it requires that $\lambda < 1$ and ties variation in the bond spread to variation in λ. A direct test of this limited participation assumption would require me to identify independent variation in λ. I have no direct test of this assumption. However, it should be clear that liquidity considerations will not matter unless $\lambda < 1$. Thus, I can only appeal to theory – if $\lambda = 1$, for the reasons outlined above, the bond spread is indeterminate and none of the hypotheses that follow would make any sense.

The second assumption that is made is that the indifference condition for the marginal liquidity investor is what determines the bond spread. $v(X,t)$ represents the answer to the question: What is the value for one day of holding an additional unit of the liquid new bond for the marginal investor?

The valuation of $v(X,t)$ is proportional to the financing spread between new and old bonds. A useful way to think about what affects $v(X,t)$ is to return to Fig. 3, as it is revisited in Fig. 5 below. To illustrate, let us suppress old bonds and imagine that old bonds are plentiful so $\hat{s}(X,t) = 0$, and $v(X,t)$ only corresponds to the $s(X,t)$ for the new bond.

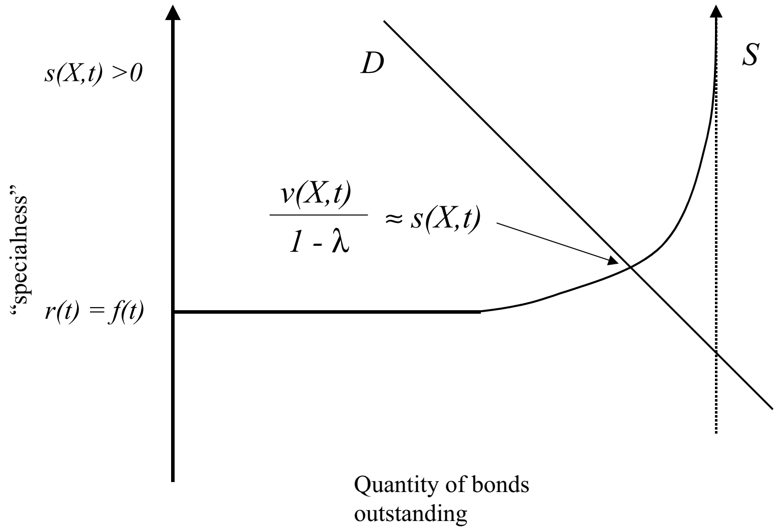

Fig. 5. Diagram of the equilibrium valuation of liquidity.

There is a supply of new bonds available for renting overnight in the repo market. Likewise there are investors who receive a benefit by renting the new bond. Ordering these investors from highest to lowest benefit gives us the demand curve. The intersection of supply and demand gives equilibrium, with the equilibrium price representing exactly what we are after. That is, the valuation of an additional unit of the liquid new bond for the marginal investor.

Two hypotheses present themselves from the figure. First,

H1 (*Differential liquidity hypothesis*). *Increases in investor preference for liquidity are correlated with increases in the spread between old and new bonds.*

If investors prefer to hold the new bond over the old bond because it is more liquid, as their preference for liquidity rises, their demand for the new bond will rise as well. D shifts rightwards causing $v(X,t)$ to rise and the bond spread to increase.

H2 (*Investor nonsatiation hypothesis*). *Increases in the auctioned amount of new Treasury bonds will be correlated with decreases in the spread between old and new bonds.*

An increase in the supply will cause S to shift right and, as long as demand crosses supply where supply is somewhat inelastic, the bond spread to rise as well. Essentially, if investors are willing to accept a negative relative return of $v(X,t)$ to hold the new bond it must be that the supply of this bond does not satiate their demand for it. For example, if in aggregate investors needed only \$10 bn of new bonds for liquidity purposes and the supply to them is \$20 bn, then we would expect that $v(X, t)$ be zero in equilibrium. Thus implicit in H1 is that supply is somewhat inelastic and that market equilibrium features nonsatiation.

Both the model of Boudoukh and Whitelaw (1993) and that of Holmstrom and Tirole (2001) have this sort of feature. In both papers investors have a demand for a liquid asset to meet liquidity needs. However there is a shortage of this liquid asset causing the price of it to rise above other assets. In Boudoukh and Whitelaw this occurs because the government issues the liquid asset and acts as a monopolist in extracting consumer surplus from the liquidity investors. In Holmstrom and Tirole the shortage occurs because the corporate sector is unable to create sufficient liquid assets for aggregate liquidity needs.

Finally, it is theoretically possible that the increase in S shifts D as well (and more), and the increase in supply leads to the opposite correlation. For example, if liquidity investors see that the Treasury issues more new bonds and this increase deepens the market for new bonds so that it is perceived to be even more liquid, then demand for the bond as a liquidity vehicle could increase, swamping the effects of the increased supply. Thus H2 should rightly be viewed as increases in supply dominate any induced changes in demand, so that spreads fall upon increased issuance.

6. Data and the linear specification

Let us begin with a linear model for the demand and supply curves of Fig. 5. The supply schedule is given as

$$s(Q) = A + X^{\mathrm{S}} b^{\mathrm{S}} + c^{\mathrm{S}} Q, \tag{19}$$

and the demand is

$$s(Q) = X^{\mathrm{D}} b^{\mathrm{D}} - c^{\mathrm{D}} Q. \tag{20}$$

$s(Q)$ is the specialness difference between new and old bonds, and Q is the quantity of new bonds transacted in the repo market. X^{D} and X^{S} are instruments for shifts in demand and supply. The reduced form of this model is

$$s = \frac{A}{1 + c^{\mathrm{S}}/c^{\mathrm{D}}} + X^{\mathrm{D}} \frac{b^{\mathrm{D}}}{1 + c^{\mathrm{D}}/c^{\mathrm{S}}} + X^{\mathrm{S}} \frac{b^{\mathrm{S}}}{1 + c^{\mathrm{S}}/c^{\mathrm{D}}}. \tag{21}$$

Substituting this back into Eq. (18) leads to

$$\begin{aligned}\hat{P}_t - P_t &= \sum_{j=t}^{T} \mathrm{E}_t\left[\frac{V(X,j)}{1-\lambda}\right] \\ &= \sum_{j=t}^{T} \mathrm{E}_t\left[\frac{A}{1 + c^{\mathrm{S}}/c^{\mathrm{D}}} + X^{\mathrm{D}} \frac{b^{\mathrm{D}}}{1 + c^{\mathrm{D}}/c^{\mathrm{S}}} + X^{\mathrm{S}} \frac{b^{\mathrm{S}}}{1 + c^{\mathrm{S}}/c^{\mathrm{D}}}\right].\end{aligned} \tag{22}$$

Let us use the approximation $P - P_0 \approx -DP(y_0)(y - y_0)$, where the y's are bond yields, and we are taking the first order Taylor expansion around $P_0(y_0)$. Then we can write

$$\hat{P}_t - P_t = DP(y_t - y_0) - \widehat{DP}(X,t)(\hat{y}_t - \hat{y}_0). \tag{23}$$

I make the simplifying assumption that $DP(X,t) = \widehat{DP}(X,t) = K$, which is assumed constant. Since P and $\hat{P}$ are virtually identical, there is not too much lost in making this assumption. Additionally, in the period from 0 until T, the duration of the securities will not change much if they are 30-year bonds, and the auction cycles are only six months.[11] The price premium can be rewritten in terms of the yield spread between old and new bond:

$$\hat{P}_t - P_t \approx Constant + DP(y_t - \hat{y}_t) = Constant + DP \times spread(t),$$

where $spread(t)$ is the yield spread between the two bonds. Substituting this into Eq. (22) gives us the following specification:

$$spread(X,t) = f(T-t) + X^{\mathrm{D}}\beta^{\mathrm{D}}(T-t) + X^{\mathrm{S}}\beta^{\mathrm{S}}(T-t) + \varepsilon(t). \tag{24}$$

First note that $f(\cdot)$, $\beta^{\mathrm{D}}(\cdot)$ and $\beta^{\mathrm{S}}(\cdot)$ are functions of $T-t$, the time to the next auction date. This follows from observing that in Eq. (22) we are summing expectations from t to T. Second, the observations will be overlapping (i.e. $\varepsilon(t)$ will be correlated across time). I will return to this issue shortly.

6.1. Bond data

The Treasury used to auction 30-year bonds on a quarterly cycle. In the Fall of 1993, they announced that this cycle was likely to be extended to six months. At the time, the spread between the new bond and old bond increased dramatically (to around 25 bps, see Fig. 2), and since then the spread has at times been high and volatile. I extend my sample back in order to include this period and provide better estimates. All bonds issued in the 1990s, beginning at May 16, 1990, are included.

The data consists of weekly observations of the spread for 22 auction cycles over this ten-year period. The way to think about this data is that each observation of the spread is earmarked by an auction cycle and the number of days from the end of that particular auction cycle. For example the observation on May 6, 1994 is 8 bps, which is for the tenth auction cycle and occurs 193 days before the end of that auction cycle. There are 497 observations, which gives an average of 23 data points per cycle, although there is a lot of variation in the length of each cycle. All of my specifications are motivated using the auction cycle as the unit of analysis. They describe how the bond spread should behave over the auction cycle.

[11] Recall that in setting up this trade, it is common to choose $\theta(t)$ and $\hat{\theta}(t)$ such that $\theta(t)DP(t) = \hat{\theta}(t)DP(t)$. In my presentation of Section 4, I suppressed θ's, keeping them constant and equal to one. If I was to reintroduce θ's, and kept the trade neutral with respect level changes, and kept θDP constant, the relations I would get would not require me to assert this simplifications.

6.2. CP-Bills and liquidity demand

H1 implies that changes in the preference of investors to hold liquid assets should affect the bond spread. To test these predictions, we need an instrument that is correlated with changes in liquidity demand.

Commercial Paper (CP) is a short-term unsecured obligation of corporations. For a variety of institutional reasons, the CP market is very illiquid.[12] An investor in CP who liquidates before maturity will have to pay a cost of around 1% of face. Considering that the tenor of CP is typically 30–180 days, the liquidation cost makes CP essentially a buy and hold security. Treasury Bills (T-bills) on the other hand are among the most actively traded and liquid securities in the world.

I use the spread between three month CP (nonfinancial) and three month T-Bills as a proxy for variation in liquidity demand. If the marginal investor in CP is one who may face a liquidity need over the next three months, then the spread between CP and T-bills will reflect changes in liquidity demand.

There are some issues regarding the validity of this instrument. First, if the marginal investor in CP does not care about liquidity (that is, if the investor will in no state of the world have to liquidate his CP investment in order to finance a sudden cash need) it is obvious that CP will not carry a liquidity premium.

The largest investors in CP are money market mutual funds (30–40%) and commercial banks (15–25%) (see Stigum, 1989). These institutions hold large amounts of both CP as well as T-bills and cash (see, e.g, Kashyap and Stein, 2000). Liquidity management for these financial institutions is a fact of life and is a prominent subject of writings of both academics as well as bankers. If these investors are the ones who set prices, it seems safe to say that CP will carry a liquidity premium.

However, if the price setters are the issuers of CP and these agents are active enough, their actions can eliminate liquidity premia. By far the largest active issuers of CP are non-bank financial corporations (e.g., GECC, GMAC. See Stigum, 1989). Now suppose that the liquidity preference of financial institutions rises and they demand a higher premium to hold CP, then it is possible that the issuing corporations will sharply curtail their sale of CP, and perhaps even switch sides to become buyers. If this action is strong enough, then CP will not reflect the increased liquidity preference of investors. As a practical matter, short term inflexibilities, such as switching to issuing medium term notes as a financing source, will limit the reactiveness of issuers. In addition aggregate CP issuance does not reflect the volatility that this reactiveness would suggest. The most striking component of the time series is just growth. Although not conclusive, Friedman and Kuttner (1993) show that the ratio of CP issuance to T-bills (to normalize) is positively correlated with the CP-Bills spread, suggesting that issuers increase rather than decrease CP issuance when CP-Bills spreads are high.

[12] The primary reason is that the CP market is very heterogenous in terms of issuers and maturities. Since each issue is relatively small, making a market in a specific issue is not cost effective for a dealer. See Stigum (1989).

A second concern is that since commercial paper is a corporate obligation, the CP-Bills spread reflects expected future default risk (Friedman and Kuttner, 1993). While undoubtedly true, this fact does not invalidate the use of CP-Bills as an instrument; it means that the instrument is noisy. Controlling for credit spreads would sharpen my results. In the regressions based on monthly data, I use the spread between commercial paper rated P2 (the second highest category) and P1 (the highest) by Moody's Investor Service as a control.[13]

The CP and T-bills data is weekly for the Friday ending each week. I use three month nonfinancial CP obtained from the Federal Reserve (BOG) web site. It is based an a sampling of dealers. I use three month T-bills also from the Fed web site. It is based on secondary market prices sampled from dealers.

6.3. Bond issuance

Over my ten-year sample, there are 22 auction cycles. The smallest issue is \$10 bn, while the largest is \$32 bn. On the other hand, the bulk of issues are around \$12 bn, with only five issues around \$20 bn or above. I use this variation as the measure of shifts in the supply curve.

The inclusion of the supply of new bonds also controls for correlation between CP-Bills and the bond spread driven by common Treasury issuance of securities. If investors have downward sloping demand for T-bills versus CP, an increase in T-bill supply will cause the spread between T-bills and CP to contract. If at the same time, the Treasury also increases supply of new 30-year bonds, then we would also see that the spread between bond and old bond would contract. Including the issue size of new bonds controls for this effect.

6.4. Results

Regressions using the specification Eq. (24) are reported in Table 6. The first column uses just the CP-Bills spread and bond issuance. Both dependent variables come in significant and with the predicted sign. Increases in the CP-Bills spread lead to higher bond spreads, while increased in bond issuance lead to smaller bond spreads. In addition the coefficient on $T - t$ is negative, confirming the pattern of the bond spread falling over the auction cycle.

The next two columns are regressions where the independent variables are interacted with $(T - t)/90$, the number of quarters remaining until the next auction date. Under the hypothesis that both CP-Bills and bond/old-bond spreads are driven by liquidity preference, one would expect that the size of the liquidity premium would vary with the number of days until the next auction. For example, one can expect changes in liquidity preference to have a greater impact on the spread 180 days from the auction than 30 days from the auction (because the new bond is liquid for that much longer.) Without more assumptions, theory does not pin down the exact time structure of this liquidity preference.

[13] The monthly P2–P1 data was provided by Ken Kuttner of the New York Fed.

Table 6
Linear specification
This table presents regression based on the specification equation (24),

$$spread(X,t) = f(T-t) + \beta^{\mathrm{D}}(T-t) \times CP\text{-}Bills(t) + \beta^{\mathrm{S}}(T-t) \times BondIss(t) + \varepsilon(t).$$

The data is weekly. The dependent variable is the level of the new bond–old bond spread. Independent variables are time to next auction, the CP-Bills spread and the size of the new bond issue. T-statistics are based on standard errors corrected for overlapping observations based on Newey and West (1987).

Constant		0.006	
		(0.66)	
$(\frac{T-t}{90})$	0.032***	0.064***	0.083***
	(7.83)	(8.36)	(8.27)
$(\frac{T-t}{90})^2$			−0.005**
			(−2.30)
CP-Bills	0.083***	0.131***	0.153***
	(4.8)	(4.36)	(8.69)
(*CP-Bills*) $\frac{T-t}{90}$		−0.189***	−0.255***
		(−6.15)	(−7.09)
(*CP-Bills*) $(\frac{T-t}{90})^2$		0.049***	0.070***
		(5.35)	(5.23)
BondIss	−0.001**		
	(2.36)		
(*BondIss*) $\frac{T-t}{90}$		−0.0015***	−0.0015***
		(−6.73)	(−6.87)
N	497	497	497
R^2	0.30	0.54	0.55

*,**,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses.

To get at this time structure, I include terms that are linear and quadratic in time. The fit improves dramatically (R^2 rises from 0.30 to 0.54), and the t-statistics rise as well. In the last column I also include $((T-t)/90)^2$ to make sure that the correlation with time is not spurious.

Fig. 6 plots the estimated $\beta^{\mathrm{D}}(T-t)$ from the regression reported in the middle column. The middle line is the estimate from the regression. The outer lines are the one standard error bounds on the estimates.

The regularities which are most striking is the fact that $\beta^{\mathrm{D}}(0) > 0$, and that $\beta^{\mathrm{D}}(T-t)$ is quite high for large $T-t$. $\beta^{\mathrm{D}}(0) > 0$ means that even at the auction date changes in liquidity preference can have an effect on the bond spread. This suggests that the old and new bond are not perfect substitutes even after the new bond rolls to become off-the-run. Thus, identifying T (the convergence date) as the next auction date is not correct, and in fact the convergence date is some time after. One would certainly

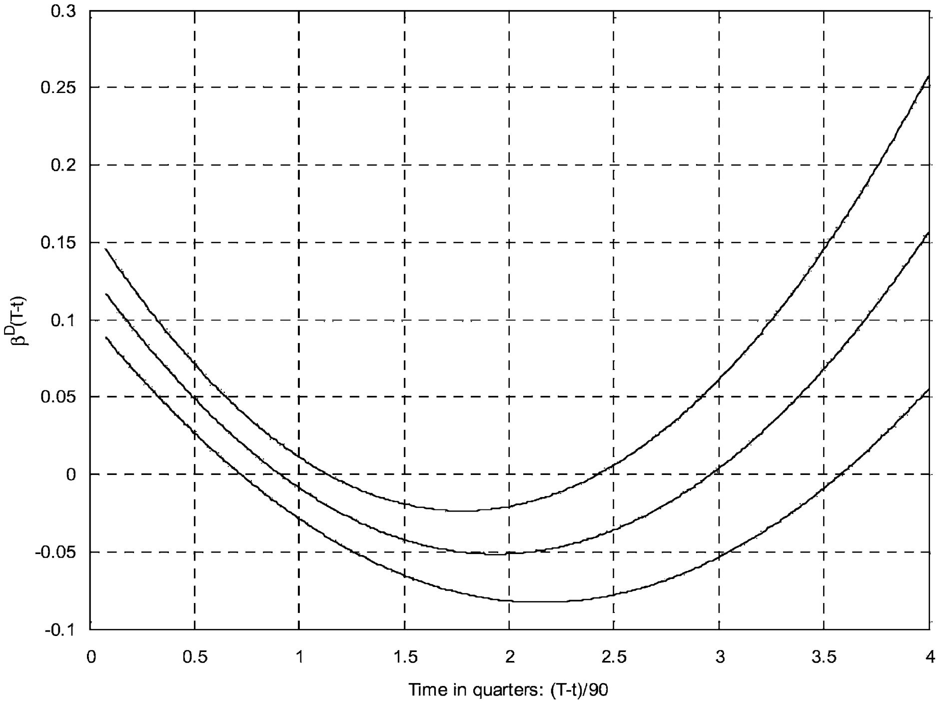

Fig. 6. Effect of one bp change in the CP-Bill spread on the bond/old-bond spread, as a function of time until the next auction date. The middle line is the estimate, with the other two lines forming a one standard deviation envelope.

expect a larger coefficient for β^{D} when we are far from the auction, but the magnitudes in the figure seem surprising. One way to look at this is that long times between auctions are times when the market is particularly sensitive to liquidity issues. For example, a long time between auctions may stir expectations that the bond is going to be phased out.

It seems quite surprising that β^{D} dips around two quarters. At face value, the latter means that six months from the auction date, liquidity is negatively valued in the new bond. It turns out that the negative valuation only occurs in some of the regressions and is perhaps due to the quadratic specification. In the next section I move to a nonlinear specification, and, while the dip remains in some of the regressions it is barely discernible in others.

Fig. 7 plots the raw data in a simple fashion. On the vertical axis is the level of the bond spread divided by the CP-Bills spread, and on the horizontal axis is the number of quarters until the auction date. The shape resembles the estimate without ever falling below zero. However it is still surprising that the liquidity valuation is not monotonically increasing over time, as one would expect from my static model. This is perhaps due to dynamics in the supply and demand for repo over the auction cycle

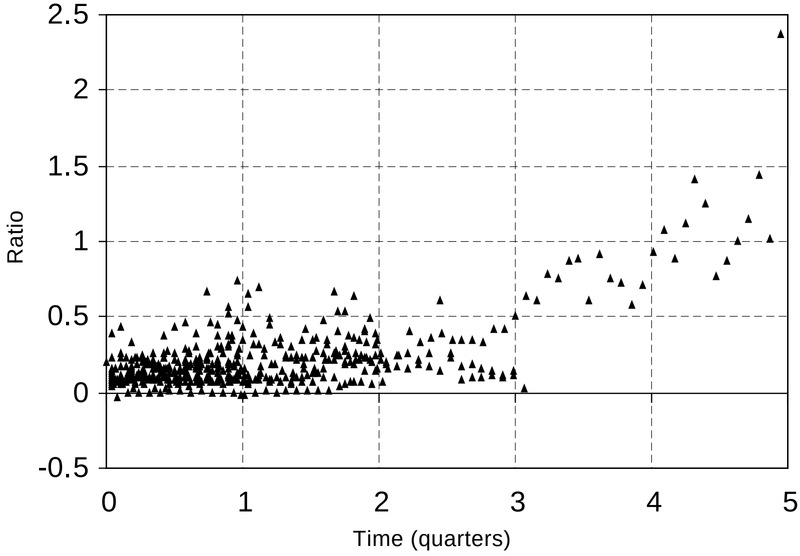

Fig. 7. Graph of the raw data. The ratio of the bond/old-bond spread to the CP-Bill spread versus quarters until the next auction.

(in particular around the when-issued period) that I am not capturing in the static model.

In terms of magnitudes, four quarters from the auction, the value of β^{D} is around 0.15. This implies that a ten basis-point change in the CP-Bills spread is associated with a 1.5 basis-point change in the bond spread. Seven days from the auction, a ten basis-point change in CP-Bills is associated with a one basis-point change in the bond spread. While small in yield terms, in price terms the numbers are much more comparable. Indeed the natural measure of a liquidity premium should be in price terms. The ratio of durations of a 30-year bond and 90-day security is roughly 50 to one. Taking an average yield sensitivity of one basis-point for every ten basis-points, this implies that a 1% change in the price of CP-Bills is met by a 5.0% change in the relative price of bonds to old bonds. Since the new bond is more liquid for about six months (an auction cycle), while T-bills are more liquid than CP for three months, this is in the ballpark of what one would expect.

6.5. Results based on differencing the financial data

The estimates in Table 6 are based on regressing levels of the bond spread on the CP-Bills spread. The resulting standard errors are autocorrelated, with correlation coefficients around 0.60. I do report autocorrelation and heteroskedasticity-consistent standard errors. On the other hand, by differencing the financial data, much of the autocorrelation can be eliminated.

Table 7
Changes in bond spread and CP-Bill spread
The dependent variable in all regressions is the weekly change in the spread between the old and new bond. Independent variables are all some combination of

$\Delta(\textit{CP-Bills})$ = change in spread between three month CP and T-Bills,

$Lag\Delta(\textit{CP-Bills})$ = last weeks change in the spread,

$\frac{T-t}{90}$ = number of days to next auction divided by 90.

Constant	−0.0009**	−0.001**	−0.0009*	−0.0008*	−0.0008**
	(−2.002)	(−2.054)	(−1.933)	(−1.781)	(−2.002)
$\Delta(\textit{CP-Bills})\ \frac{T-t}{90}$	0.0072	0.0083*			
	(1.585)	(1.825)			
$Lag\Delta\ (\textit{CP-Bills})\ \frac{T-t}{90}$		0.0055			
		(1.384)			
$Lag\Delta(\textit{CP-Bills}) + \Delta(\textit{CP-Bills})$			0.0067**	0.0232***	0.0231***
			(2.00)	(2.985)	(2.8366)
$(Lag\Delta(\textit{CP-Bills}) + \Delta(\textit{CP-Bills}))\ \frac{T-t}{90}$				−0.0458***	−0.0472***
				(−2.946)	(−2.735)
$(Lag\Delta(\textit{CP-Bills}) + \Delta(\textit{CP-Bills}))\ (\frac{T-t}{90})^2$				0.0168***	0.0175***
				(2.969)	(2.715)
$\frac{T-t}{90}$					2e − 5
					(1.23)
$(\frac{T-t}{90})^2$					−6e − 8
					(−0.96)
N	472	472	472	472	472
R^2	0.0068	0.0102	0.0085	0.0437	0.0479

*,**,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses. All t-statistics are based on robust standard errors (White, 1980).

Table 7 reports the results of regressions based on the following specification:

$$\Delta spread(X,t) = \alpha + \beta^{\mathrm{D}}(T-t) \times \Delta(\textit{CP-Bills})(t) + \varepsilon(t). \quad (25)$$

The resulting errors have a negative autocorrelation of −0.09.[14] The first column reports the regression of the change in the bond spread on the change in the CP-Bills spread, scaled by days until next auction. The results are not significant, but the coefficient is of the right sign. The second column includes the lagged value of the

[14] I have also estimated the regression using the Cochrane-Orcutt method to adjust for the autocorrelation. The results are unaffected.

change in the CP-Bills spread. The coefficient on the contemporaneous change is now significant at the 10% level. Surprisingly, the coefficient on the lagged term is of the same magnitude as the contemporaneous change. While this seems to suggest predictability in future spread changes, the more likely culprit is data problems. The CP-Bills spread is sizeable and fairly volatile. The bond/old-bond spread is much more sticky in its changes. In fact, a number of the spread changes are zero. If the reported bond/old-bond spread is adjusted less frequently than the true bond/old-bond spread, we can expect to see lagged values of CP-Bills have some predictive content for current changes in the bond/old-bond spread. The rest of the regressions use the sum of the contemporaneous change in CP-Bills and last week's change in CP-Bills as the liquidity variable.[15]

The third through last columns report regressions where the dependent variable is this sum, scaled in the last two columns by $(T-t)/90$. The results reinforce the conclusion based on the levels regressions. Increases in CP-Bills increase the bond spread. The effect varies across the auction cycle in a quadratic fashion, with a dip around six months.

7. Nonlinear specification

The assumption of a linear supply curve in Eq. (19) is not an accurate representation of supply in Fig. 5. From the earlier discussion of the repo market, the supply constraint should only matter after demand exceeds a certain point. I now alter Eq. (19) to the following supply curve,

$$s(Q) = A + c_1^{\mathrm{S}} Q + c_2^{\mathrm{S}} \max[Q - Q^*, 0]. \tag{26}$$

This curve is linearly increasing upto the point Q^*, and then increases with a higher slope past this point. In accordance with Fig. 5, I will associate Q^* with the issued amount of new bonds.

Taking the same demand schedule from Eq. (20), the reduced form of this model has two regimes.

$$s = \begin{cases} \frac{A}{1+c_1^{\mathrm{S}}/c^{\mathrm{D}}} + X^{\mathrm{D}} \frac{b^{\mathrm{D}}}{1+c^{\mathrm{D}}/c_1^{\mathrm{S}}} & \text{if } s < s^*, \\ \frac{A}{1+(c_1^{\mathrm{S}}+c_2^{\mathrm{S}})/c^{\mathrm{D}}} + X^{\mathrm{D}} \frac{b^{\mathrm{D}}}{1+c^{\mathrm{D}}/(c_1^{\mathrm{S}}+c+2^{\mathrm{S}})} - Q^* \frac{c_2^{\mathrm{S}}}{1+(c_1^{\mathrm{S}}+c_2^{\mathrm{S}})/c^{\mathrm{D}}} & \text{if } s \geqslant s^*, \end{cases} \tag{27}$$

where, s^* is a constant.

Repeating the earlier steps of moving from s to $spread(X,t)$ leads to the following specification:

[15] There are three ways to look at this. First, if my guess of sticky prices is right, then using the lagged variable is appropriate to estimate the liquidity effect. Second, if it is wrong and actually the lagged value has no predictive content, the estimate will be biased down, as I am adding noise to the independent variable. Third, if the true model is one in which the lagged values have predictive power for future demand for liquidity, then using the lagged values and interpreting the coefficients as a liquidity effect is the correct thing to do.

$$\mathit{spread}(X,t) = \begin{cases} f^H(T-t) + X^{\mathrm{D}}\beta^{\mathrm{D},H} + \varepsilon(t) & \text{if } \mathit{spread}(X,t) < \mathit{spread}^*, \\ f^V(T-t) + X^{\mathrm{D}}\beta^{\mathrm{D},V} - Q^*\beta^{\mathrm{S},V} + \varepsilon(t) & \text{if } \mathit{spread}(X,t) \geqslant \mathit{spread}^*. \end{cases} \tag{28}$$

I use the mnemonics "H" and "V" to correspond to parameters in the horizontal regime (*spread* < *spread**), and vertical regime (*spread* ⩾ *spread**). As before, the β's may be functions of $T-t$. X^{D} are demand shifters. Q^* is the supply shifter and is the issued amount of new Treasury bonds.

7.1. Results

I start by repeating the exercise of the previous section, but with the nonlinear specification. I use GMM to jointly estimate the parameters in each region along with the switching point, *spread**. Lags of CP-Bills are included as instruments. Table 8 reports the results.

Comparing Table 8 to Table 6, we see that the fit is uniformly better when we move to the nonlinear specification. The best fit is in the last column with an R^2 of 0.84. Fig. 8 graphs the time dependence of the CP-Bills coefficients based on the fit of the last column. The results are similar to Fig. 6. The quadratic shape is preserved, while the dip around six months is no longer as pronounced. The encouraging change is that the figure matches the scatter plot of Fig. 7 much better with the nonlinear specification.

The only substantive change is that the coefficients on bond issuance are no longer significant. The reason for this is that the variation in the size of bond issues is large, and moves equilibrium from the V to the H regime, rather than moving it within one of the regimes. As a result, the two regime regression does not pick up the effects of issue size on the spread. Indeed, I will verify this in Table 13.

7.2. CP-Repo and Repo-Bills

The three-month general collateral repo rate is the rate on borrowing/lending against general treasury collateral for a term of three months. An investment in three-month repos is less liquid than T-bills, because the investor typically has to reverse out of the position and pay the larger bid–ask in the repo market. However it is far more liquid than CP because these bid–ask spreads are much smaller than the 1% paid on liquidation of CP investments.

In this section I decompose the CP-Bills spread into CP-Repo and Repo-Bills. Table 9 presents these estimates based on the two regime specification of Eq. (28). The general collateral repo data begins 5/24/91, thereby reducing the number of observations in the sample.

In the first column I report regressions including only the levels of both CP-Bills and CP-Repo. Both come out significant and the coefficients have the expected sign. In each of the next two columns I use the quadratic specification. Again the results show that both spreads have explanatory power for the bond spread. The R^2 is only

Table 8
Nonlinear specification
This table presents GMM estimates for the two-regime specification of Eq. (28).

$$spread(X,t) = \begin{cases} f^H(T-t) + X^{\mathrm{D}}\beta^{\mathrm{D},H} + \varepsilon(t) & \text{if } spread(X,t) < spread^*, \\ f^V(T-t) + X^{\mathrm{D}}\beta^{\mathrm{D},V} - Q^*\beta^{\mathrm{S},V} + \varepsilon(t) & \text{if } spread(X,t) \geqslant spread^*. \end{cases}$$

The data is weekly. The dependent variable is the level of the bond/old-bond spread. Independent variables are time to next auction, the CP-Bills spread and the size of the new bond issue. t-statistics are from standard errors corrected for overlapping observations.

Constant^V	0.027	0.079***
	(1.23)	(2.54)
$(\frac{T-t}{90})^V$	0.023***	0.003
	(3.36)	(0.23)
CP-Bills	0.092***	0.082***
	(4.52)	(3.03)
(*CP-Bills*) $\frac{T-t}{90}$		−0.089***
		(−3.94)
(*CP-Bills*) $(\frac{T-t}{90})^2$		0.04***
		(4.35)
BondIss	−0.0001	
	(−0.08)	
(*BondIss*) $\frac{T-t}{90}$		−0.001
		(−0.61)
Constant^H	0.029***	0.030***
	(12.09)	(12.04)
$(\frac{T-t}{90})^H$	0.006***	0.005***
	(3.12)	(2.52)
*spread**	0.062	0.0615
N	495	495
R^2	0.79	0.84

*,**,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses.

slightly smaller than the R^2 from using only CP-Bills (Table 8). The last column includes both CP-Repo and Repo-Bills using the quadratic specification. There seems to be an independent role for both spreads, although the time dependence coefficients are close to zero for Repo-Bills.

Both Duffie (1996) and Jordan and Jordan (1997) note that repo specialness is correlated across issues. The fact that Repo-Bills comes in significant reconfirms

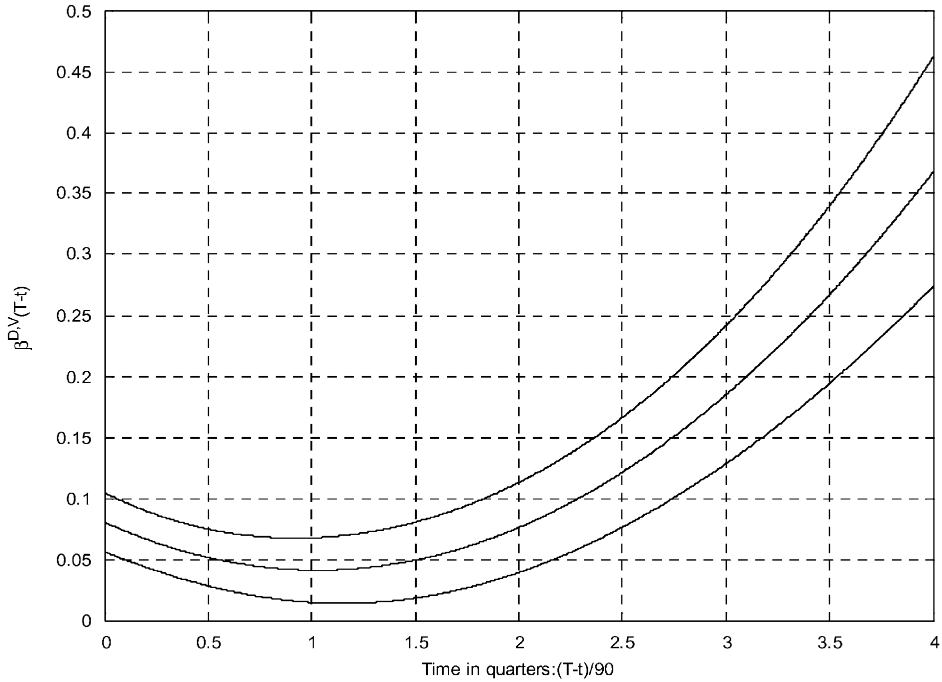

Fig. 8. Results from nonlinear specification. Effect of a one bp change in the CP-Bill spread on the bond/old-bond spread, as a function of time until the next auction date. The middle line is the estimate, with the other two lines forming a one standard deviation envelope.

their observation. Many authors have described a "flight-to-quality" as an event in which investors shift their preferences towards holding only Treasury securities. The results also confirm these views. It is also interesting that CP-Repo contains information above and beyond this flight-to-quality or correlated repo specialness effects. However, without more theory it is hard to economically interpret the decompositions. They are both liquidity proxies, but it is unclear what independent information they each contain.

7.3. Adjusting for credit risk

If liquidity demand is orthogonal to credit risk, then controlling for default risk should not change my results, and could provide sharper estimates in my regressions. The next regression uses the P2–P1 short term credit spread as a control for changes in default risk. I move to a monthly frequency because the P2–P1 data is monthly.

I use a two-step procedure to orthogonalize the credit component. The first-step regression is

$$\textit{CP-Bills}(t) = \alpha + \beta \times (P2 - P1)(t) + \varepsilon(t).$$

Table 9
CP-Repo and Repo-Bills
This table presents GMM estimates for the vertical and horizontal regime specification equation (28). The dependent variable is the level of the bond/old-bond spread. Independent variables are time to next auction, the CP-Repo spread, Repo-Bills and the size of the new bond issue. T-statistics are from standard errors corrected for overlapping observations.

$Constant^V$	0.022	0.099***	0.048***	0.067***
	(1.11)	(14.31)	(2.84)	(5.30)
$(\frac{T-t}{90})^V$	0.023***	−0.004	0.01	0.007
	(3.44)	(−0.43)	(0.79)	(0.47)
CP-Repo	0.053**	0.057***		0.067***
	(2.06)	(4.15)		(3.26)
$(CP\text{-}Repo)\ \frac{T-t}{90}$		−0.185***		−0.160***
		(−8.81)		(−5.21)
$(CP\text{-}Repo)\ (\frac{T-t}{90})^2$		0.064***		0.054***
		(9.72)		(7.49)
Repo-Bills	0.112***		0.148***	0.073***
	(6.95)		(3.49)	(2.86)
$(Repo\text{-}Bills)\ \frac{T-t}{90}$			−0.09*	0.01
			(−1.90)	(0.18)
$(Repo\text{-}Bills)\ (\frac{T-t}{90})^2$			−0.002	−0.01
			(−0.1)	(−0.9)
BondIss	0.0003			
	(0.27)			
$(BondIss)\ \frac{T-t}{90}$		0.0002	0.0002	0.0002
		(0.52)	(0.72)	(0.37)
$Constant^H$	0.03***	0.029***	0.031***	0.031***
	(12.95)	(13.37)	(12.60)	(12.03)
$(\frac{T-t}{90})^H$	0.006	0.006	0.007	0.007
	(3.57)	(3.77)	(3.77)	(4.00)
*spread**	0.062	0.0615	0.0647	0.0649
N	442	442	442	442
R^2	0.81	0.85	0.84	0.87

*,** ,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses.

Table 10
CP-Bills and P2–P1 spreads
This table reports results from the following regression:

$$CP\text{-}Bills(t) = \alpha + \beta \times (P2 - P1)(t) + \varepsilon(t).$$

α	β^L	R^2	N
0.293***(6.86)	0.627 (3.67)***	0.157	115

*** Significant at 1% level. t-statistics in parentheses are autocorrelation and heteroskedasticity consistent (Newey and West, 1987).

Table 11
Adjusting for credit risk
This table presents GMM estimates for the vertical and horizontal regime specification equation (28). The data is monthly. The dependent variable is the level of the bond/old-bond spread. Independent variables are time to next auction, the size of the new bond issue and the liquidity variable described in the text. T-statistics are from standard errors corrected for overlapping observations.

$Constant^V$	$(\frac{T-t}{90})^V$	Liq	$BondIss$	$Constant^H$	$(\frac{T-t}{90})^H$
0.036*	0.021***	0.123***	0.001	0.029***	0.004**
(1.82)	(2.58)	(4.21)	(0.59)	(10.39)	(2.09)
$spread^* = 0.062$, $R^2 = 0.78$, $N = 115$.					

*,**,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses.

Table 10 reports these results. As expected, the spreads are significantly positively correlated. The second step is to use the residuals from this regression as the indicator of liquidity demand in the specification equation (28). Table 11 reports the estimates. The liquidity variable still comes in significant and with about the same magnitude, confirming that CP-Bills has a liquidity component independent of credit risk considerations. Using the quadratic specification gives similar results, although the time dependence is not as strong as in previous estimates.

7.4. Repo market hedging demand

The primary reason to borrow bond collateral in the repo market is to short sell bonds. The convergence strategy by arbitrageurs of shorting the new bond against a long position on the old bond is to create a speculative position. However, another important reason agents short sell bonds is to create a hedging position. Dealers who underwrite long-term bond issues typically hedge their market risk by taking short positions in the new bond. The reason they carry their shorts in the new as opposed to the old bond is because of the superior liquidity of the new bond.

To test for this source of liquidity demand, I construct two monthly series. The first corresponds to the aggregate amount of domestic (US Dollar) and Eurodollar bond issues by corporations over the entire sample. Only issues with maturities

Table 12
Corporate and agency issuance
This table presents GMM estimates for the vertical and horizontal regime specification equation (28). The data is monthly. The dependent variable is the level of the bond/old-bond spread. Independent variables are time to next auction, the size of the new bond issue and the liquidity variable described in the text. T-statistics are from standard errors corrected for overlapping observations.

$Constant^V$	0.037**	0.037***	0.041**
	(2.32)	(2.59)	(2.46)
$(\frac{T-t}{90})^V$	0.006	0.009	0.035***
	(0.34)	(0.54)	(2.77)
CP-Bills	0.059*	0.056*	0.102***
	(1.74)	(1.75)	(3.87)
$(CP\text{-}Bills)\ \frac{T-t}{90}$			−0.13***
			(−4.69)
$(CP\text{-}Bills)\ (\frac{T-t}{90})^2$			0.04**
			(2.21)
$(BondIss)\ \frac{T-t}{90}$	0.0005	0.0006	−0.0008
	(0.58)	(0.81)	(−1.25)
$\sum ActualAgencyIss$		0.0029	0.0031
		(1.33)	(1.37)
$\sum ActualCorpIss$		−0.0082***	−0.0079***
		(−2.96)	(−4.40)
$Constant^H$	0.025***	0.025***	0.025***
	(10.11)	(11.70)	(11.63)
$(\frac{T-t}{90})^H$	0.003	0.003	0.003
	(1.36)	(1.33)	(1.19)
*spread**	0.044	0.044	0.044
N	115	115	115
R^2	0.65	0.69	0.79

*,**,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with t-statistics in parentheses.

longer than 15 years are included. To compare these for hedging purposes to the 30-year bonds, each bond is duration weighted before aggregating across bonds. This aggregate is then divided by the duration of the new 30-year bond. All data is from Bloomberg. While bond issuance is a daily activity, it tends to be sporadic. A week of inactivity is followed by a week of high issuance. Because I cannot be sure during which interval a dealer will be hedging the issue, the issuance variable is aggregated to a monthly frequency. This series is fairly volatile. For example, there were no long term bonds issued in July 1990, while there was in excess of \$5 bn of bonds issued the previous month. The second series is the same except includes mortgage and other agency issues.

Although, in principle, we would like the sum of expected issuance until the next auction date as the explanatory variable in the regression, I have used the sum of actual issuance instead.

The results of including this variable in the regression are reported in Table 12. The corporate issuance variable comes in negative, which is the opposite of what we would have expected. The agency issuance variable is positive, but not significant. One way to read this evidence is that long corporate issuance is not an independent repo market shifter, and that actually both bond spread and issuance are driven by some common macro factor (good times), while agency issuance, which is less cycle driven, is being picked up in the regression. It is also possible that issuance and hedging demand are less of a factor in the 30-year bond, and these effects may be easier to see in the ten-year bond.

7.5. Old and new bond issue sizes

One caveat to the finding of effects on new bond issue size is that since the issue size is constant across an auction cycle, the sample size of 115 is misleading and actually these results arise from variation only across auction cycles. Over the ten-year sample, there are 22 auction cycles. The smallest issue is \$10 bn, while the largest is \$32 bn. On the other hand, the bulk of issues are around \$12 bn, with only five issues around \$20 bn or above. So it is fair that really these results arise from five sets of observations.

Table 13 reports similar regressions to that of Table 6, but with the sample reduced to 22 auction cycles. The dependent variable is the average value of the bond spread for each auction cycle. Similarly independent variables are computed as averages over the auction cycle. Changes in liquidity demand are controlled for using just the average CP-Bill spread, as opposed to the quadratic specification of before. There are too few observations to provide meaningful results with the more involved specification. The issue size of the old bond is also included in the regressions. Since the bond spread is the spread between new and old bonds, for the same reason that the new bond issue size is significant, one can expect the size of the old bond to affect the spread.

The results largely confirm the previous findings. The coefficient on the new bond is of the same magnitude as before and significant. The coefficient on CP-Bills is close to the 0.1 number that we see in previous regressions. The only new finding in this table is the coefficient on old bond issue size. The results are weak, but suggestive. Larger old bond issues will translate into a larger spread, since effectively the old bond is less tight in the repo market. On the other hand the coefficient is smaller than that of the new bond, which should not be too surprising since the old bond and older bonds are closer substitutes.

Table 13
Spreads and bond supply
This table presents regressions measuring the effects of a change in the supply of new bonds on the bond spread. The dependent variable is the series of average bond spreads during each of 22 auction cycles. Independent variables are

$$\begin{aligned} \textit{CP-Bills} &= \text{average CP-Bills spread in an auction cycle,}\\ \textit{New bond} &= \text{new bond issue size in billions,}\\ \textit{Old bond} &= \text{old bond issue size in billions,}\\ \textit{Issuance} &= \text{duration weighted corporate issuance in billions,}\\ \frac{Days}{90} &= \text{number of days in auction cycle divided by 90.} \end{aligned}$$

All regressions account for autocorrelation in errors using the Cochrane-Orcutt procedure, iterated twice. I retain the first observation. The *t*-statistics are based on White (1980) robust standard errors.

$\frac{Days}{90}$	0.011**	0.019***	0.021***
	(2.46)	(4.03)	(4.73)
CP-Bills	0.099**	0.083*	0.092*
	(2.68)	(1.81)	(1.92)
Issuance	−0.003		−0.002
	(−0.6)		(−0.52)
New bond		−0.002**	−0.003**
		(−2.66)	(−2.65)
Old bond		0.001*	0.001
		(1.9)	(1.61)
N	22	22	22
R^2	0.154	0.355	0.361

*,** ,*** Significance at the 10%, 5%, and 1% levels respectively. Estimates are reported, with *t*-statistics in parentheses.

8. Discussion and conclusions

Textbook finance tells us that prices are set by the marginal investor, who we usually take to be a smart, wealthy and sophisticated trader, whose actions ensures no opportunities for arbitrage. Demands such as the preference for liquidity by one group, or the demand for collateral by another, rarely has an effect on prices. The bond/old-bond spread presents an environment in which we *can* observe the demand of the subgroup, despite the fact that the arbitrageur is an active market participant. The reason is that the arbitrageur preserves price relations across a combination of two markets. While, individually, each of the bond and the repo market is free to reflect other demands.

To understand variation in the bond spread, we need to turn to identifying who this subgroup is and why their demand varies as it does. The hypothesis that I have forwarded and provided evidence for is that investors have preference for a liquid asset that gives them exposure to long term risk, and that at current quantities, this liquidity demand is not satiated. This hypothesis can account simultaneously for the fact that, (1) changes in preference for liquid assets is correlated with the bond spread, and (2) changes in the supply of the new bond has price effects.

If investors desire to hold a liquid long-term asset, and this demand is not satiated with the new bond, a further possibility under this hypothesis is that some investors will substitute into using the old bond as a liquidity vehicle as well. Although weaker, there is evidence for this implication. There is usually a positive spread between old and older bonds. The evidence in Table 13 implies that decreases in the supply of the old bond are correlated with decreases in the bond/old-bond spread. Each of these points is consistent with the hypotheses I have forwarded.

My results are interesting for a few reasons. First, referring back to Fig. 2, convergence seems to have broken down since the Summer of 1998. Hedge funds putting on the convergence trade lost money over this period. Why? The hypotheses I have forwarded gives two reasons. First, liquidity was more valued in the latter period. The CP-Bills spread averaged 18 more bps from May 1998 to November 1999 than during the rest of the period. This compares to an increase in the average bond spread of 5 bps over this period. Second, because of anticipated budget surpluses, the Treasury announced a reduction in debt issuance. This and increases in expectations surrounding long-bond phase-out increased the bond spread over this period.

Second, my results suggest that eliminating the 30-year Treasury bond will have negative welfare consequences. If the long bond is uniquely valued by investors as a liquidity vehicle, then removing this bond should reduce welfare. Computing these costs requires an estimate of the convenience yield in holding the bond. The estimation is the subject of current work.

While the liquidity hypothesis I have forwarded seems to fit the data best, it leads to deeper questions. Who exactly are the investors who demand long-term liquid assets? What are their portfolios composed of, and tangibly why do they alter their preference for both the long-term (T-bonds) and short-term (T-bills) liquid assets simultaneously? Answering these questions seems important not only to provide further support for the hypothesis, but also to help us gauge the magnitude of preference parameters required to justify observed liquidity premia.

Finally, these results are relevant for the macroeconomics literature on the forecasting power of the CP-Bills spread for economic conditions. The existing explanations for the forecast power of CP-Bills rely on changes in the *supply* of CP, which take as given imperfect substitutability of CP for bills in investors' portfolios. All of my results stem from the *demand* side, as the liquidity hypothesis identifies changes in the CP-Bills spread as arising from changes in investors liquidity demand. On the one hand, since the results validate the common assumption of imperfect substitutability, they provide greater support to these explanations. On the other hand, I have shown that changes in CP-Bills can be driven by changes in investor's liquidity preference. Holmstrom and Tirole (2001) and Krishnamurthy (2002) develop models in which agents shift their preferences towards liquid assets, and thereby drive up liquidity premia, ahead of downturns. Kashyap and Stein (2000) provide evidence on buffer-stocking behavior of banks ahead of recessions that is consistent with these theories. Thus, a competing explanation of the forecast power of the CP-Bills spread could be that the equilibrium behavior of spreads is the result of forward looking agents increasing liquidity demand ahead of downturns. After saying this, I am hesitant to push the argument much further at this stage. More works needs to be done before judging this explanation. I will need to extend m sample back at least another 20 years (to coincide with the rest of the literature), construct a liquidity demand variable, and see how well it forecasts economic activity. This remains as future work.

Appendix A. Bond and repo market equilibrium

Let us consider the following model, which borrows from Duffie (1996). Time is indexed as $t = 0, 1, 2$. There is a single consumption good.

There are two traded assets. Both assets are bonds. As in the Treasury market, I shall assume that assets are next day settle. That is purchase at t means payment and delivery at $t+1$. A liquid bond costs P_0 at $t = 0$, and pays off P_1 units of consumption at $t = 2$. An illiquid bond costs 1 at $t = 1$ and pays 1. The second bond is illiquid because liquidating it involves payment of a transaction cost of α_i, where i indexes the type of agent (described below).

There is also a repo market in which agents can borrow and short bonds. The rental cost in this shorting market is given as s. The repo market is same day settle. For example, bonds can be borrowed at $t = 1$ for a rental price of s, and must be returned at $t = 2$.

There are two types of agents in this economy. Investors (I) have an endowment of $E \gg 0$ units of the consumption good at $t = 1$ which they wish to transfer to $t = 2$. Their preferences are

$$U^{\mathrm{I}} = v(c_2) + c_1, \tag{A.1}$$

where $v(c_2)$ is concave and increasing with $\lim_{c_2 \to \infty} v'(c_2) = 1$. As mentioned, the illiquid bond involves a transaction cost. I shall assume that $\alpha_{\mathrm{I}} = 1$, so that the illiquid bond is never an investment option for I-agents. Thus, this agent can only smooth consumption by purchase of the liquid bond. I-agents can also lend out any bonds they purchase in the repo market. However doing so involves a private transaction cost per-unit of bond of $\lambda(\phi^{\mathrm{I}})$, where ϕ^{I} is the amount of bonds lent.

An investor chooses a pair $(\theta^{\mathrm{I}}, \phi^{\mathrm{I}})$, where θ^{I} is the amount of bonds purchased. This results in

$$c_1 = E - P_0\theta^{\mathrm{I}} + (s - \lambda(\phi^{\mathrm{I}}))\phi^{\mathrm{I}} \tag{A.2}$$

$$c_2 = \theta^{\mathrm{I}} P_1. \tag{A.3}$$

Thus the investor solves the program (PI),

$$\max_{\theta^{\mathrm{I}},\phi^{\mathrm{I}}} \quad v(c_2) + c_1 \tag{A.4}$$

$$\text{s.t.} \quad \theta^{\mathrm{I}} \geqslant 0 \tag{A.5}$$

$$\phi^{\mathrm{I}} \leqslant \theta^{\mathrm{I}}. \tag{A.6}$$

The second agent type are arbitrageurs (A). These agents have no cost of trading the illiquid bond ($\alpha_{\mathrm{A}} = 0$), and trade this against the liquid bond. The arbitrageur has preferences,

$$U^{\mathrm{I}} = c_2. \tag{A.7}$$

An arbitrageur shorts ϕ^{A} of the liquid bond at $t = 0$. A borrows ϕ^{A} in the repo market to deliver against his short at $t = 1$. Thus, given next day settlement on bonds, the total inflow at $t = 1$ is $\phi^{\mathrm{A}}(P_0 - s)$. Simultaneously at $t = 0$, A purchases $\phi^{\mathrm{A}}(P_0 - s)$ of the illiquid bond. Thus, the payment against the illiquid bond matches against the inflow from the short position.

At $t = 1$, A unwinds both positions, selling the illiquid bonds and purchasing back the liquid bond. This results in a date 2 flow of

$$c_2 = -\phi^{\mathrm{A}} s + \phi^{\mathrm{A}}(P_0 - P_1). \tag{A.8}$$

The arbitrageur solves (PA)

$$\max_{\phi^{\mathrm{A}} \in \mathscr{R}} \quad c_2. \tag{A.9}$$

Finally let us close the model by introducing supplies of the two assets. I assume that a government has a large endowment of consumption goods at $t = 2$. It issues two assets against this endowment. $\bar{\theta}$ of the liquid bond at price P_0 at $t = 0$, each of which is a claim on P_1 units of consumption goods at $t = 2$. Additionally, it offers illiquid bonds elastically at a price of one at $t = 0$. Each of these pay one as well at $t = 2$.

Market clearing conditions are as follows. First, the bond market must clear at $t = 0$:

$$\theta^{\mathrm{I}} - \phi^{\mathrm{A}} - \bar{\theta} = 0. \tag{A.10}$$

That is I's purchase must be matched by short sales by A, and issuance by the government. Additionally, the repo market clearing is the bonds lent must be equal to bonds borrowed,

$$\phi^{\mathrm{I}} = \phi^{\mathrm{A}}.$$

Given these conditions, by Walras' Law, the market for the consumption good clears.

An equilibrium is a collection of decisions, $(\theta^{\mathrm{I}}, \phi^{\mathrm{I}}, \phi^{\mathrm{A}})$, and prices, (P_0, s). The decisions are solutions to PI and PA given prices, and at these decisions, the bond and repo market clear.

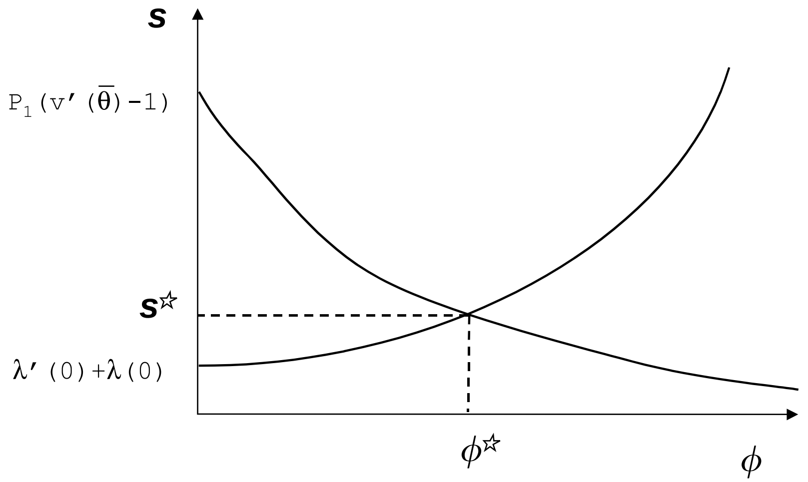

Fig. 9. Diagram of repo market clearing condition.

Let us analyze equilibrium now. The first-order-condition for the arbitrageur gives

$$P_0 = P_1 + s. \tag{A.11}$$

This is just the two-period model version of what I wrote in the text. Likewise the F.O.C. for the investor gives

$$P_1 v'(\theta^{\mathrm{I}}) - P_0 = 0, \tag{A.12}$$

and,

$$s - \lambda'(\phi^{\mathrm{I}}) - \lambda(\phi^{\mathrm{I}}) = \mu, \tag{A.13}$$

where μ is the multiplier on the constraint that $\phi^{\mathrm{I}} \leqslant \theta^{\mathrm{I}}$. There are two cases to consider. First, if λ is small so that $\phi^{\mathrm{I}} = \theta^{\mathrm{I}}$ at all prices and $\mu > 0$, we are in a situation of indeterminacy and there is no equilibrium. The other case is in which $\mu = 0$ so that the F.O.C's define the solution.

In this case, since $P_0 = P_1 + s$, we can write the two conditions as,

$$s = P_1(v'(\bar{\theta}) + \phi^{\mathrm{I}}) - 1) \quad \text{DEMAND}, \tag{A.14}$$

$$s = \lambda'(\phi^{\mathrm{I}}) + \lambda(\phi^{\mathrm{I}}) \quad \text{SUPPLY}. \tag{A.15}$$

With $v(\cdot)$ concave, and $\lambda(\cdot)$ convex, these two conditions define a unique equilibrium. Existence is guaranteed under mild assumptions.

Fig. 9 illustrates equilibrium in the repo market. On the quantity axis, ϕ is the total amount borrowed/lent in the repo market. On the vertical axis is the rental cost in the repo market of s. This corresponds to the specialness premium we discussed in the text.

Note the following comparative statics. As $\bar{\theta}$ rises, since $v'(\cdot)$ is decreasing, the demand shifts to the left and s^* and ϕ^* fall. This corresponds to the experiment of changing the auctioned amount of Treasury securities. Second, if $v'(\cdot)$ rises for every argument, this means that investors' liquidity preference rises, and demand in the repo market rises and s^* and ϕ^* rise. This is the second case I look at in the regressions. Note that since the bond repo market is over-the-counter, it is close to impossible to measure short interest of ϕ^*. A simultaneous equations approach may be possible in the stock market where this is more observable. Finally, note that as λ rises, the cost of lending rises, and s^* rises. This is the implication that I am unable to test because I have no measured variation in λ.

A last point worth making is that the model I have just laid out pins down equilibrium using λ as a private cost of participating in repo markets. Another important restriction in practice is that the repo market does not clear infinitely often at date 1, as has been implicitly assumed above. One could introduce this into the model, and allow for one or a few rounds of clearing (as in models in monetary economics) in which case there will in equilibrium be some failures on short delivery. This would involve a private fail cost borne by the arbitrageur and limit his short positions. This is an alternative way to arrive at equilibrium. The new implication of this equilibrium would be to link specialness to the number of rounds of clearing and to variation in fail costs. Practically, I have no way to measure variation in either of these parameters, so this approach yields no more testable implications than the approach I have taken.

References

Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid–ask spread. Journal of Financial Economics 17, 223–249.

Amihud, Y., Mendelson, H., 1991. Liquidity, maturity, and the yields on U.S. Treasury securities. Journal of Finance 46, 31–53.

Barone, E., Risa, S., 1995. Valuation of floaters and options on floaters under special repo rates. Unpublished working paper, Istituto Mobiliare Italiano.

Boudoukh, J., Whitelaw, R., 1993. Liquidity as a choice variable: a lesson from the Japanese government bond market. Review of Financial Studies 6, 265–292.

Buraschi, A., Menini, D., 2002. Liquidity risk and special repos: how well do forward repo spreads price future specialness. Journal of Financial Economics 64, 243–284.

Chordia, T., Roll, R., Subrahmanyam, A., 2000. Commonality in liquidity. Journal of Financial Economics 56, 3–28.

Cornell, B., Shapiro, A., 1989. The mispricing of U.S. Treasury bonds: a case study. Review of Financial Studies 2, 297–310.

D'Avalio, G., 2002. The market for borrowing stock. Journal of Financial Economics 66, this issue.

Duffie, D., 1996. Special repo rates. Journal of Finance 51, 493–526.

Fleming, M., 2001. Measuring treasury market liquidity. Unpublished working paper, Federal Reserve Bank of New York.

Friedman, B., Kuttner, K., 1992. Money, income, prices, and interest rates. American Economic Review 82, 472–492.

Friedman, B., Kuttner, K., 1993. Why does the paper-bill spread predict real economy activity. In: Stock, J., Watson, M. (Eds.), New Research in Business Cycles: Indicators and Forecasting. University of Chicago Press, Chicago.

Geczy, C., Musto, D., Reed, A., 2002. Stocks are special too: an analysis of the equity lending market. Journal of Financial Economics 66, this issue.

Hasbrouck, J., Seppi, D., 2001. Common factors in prices, order flows and liquidity. Journal of Financial Economics 59, 383–411.
Heaton, J., Lucas, D., 1996. Evaluating the effect of incomplete markets on risk sharing and asset prices. Journal of Political Economy 104, 443–487.
Holmstrom, B., Tirole, J., 2001. LAPM: a liquidity based asset pricing model. Journal of Finance 56, 1837–1867.
Huang, M., 1999. Liquidity shocks and equilibrium liquidity premia. Unpublished working paper, Stanford University.
Huberman, G., Halka, D., 2001. Systematic liquidity. Journal of Financial Research 24, 161–178.
Jordan, B., Jordan, S., 1997. Special repo rates: an empirical analysis. Journal of Finance 52, 2051–2072.
Kamara, A., 1994. Liquidity, taxes, and short-term treasury yields. Journal of Financial and Quantitative Analysis 29, 403–417.
Kashyap, A., Stein, J., 2000. What do a million observations on banks say about the transmission of monetary policy. American Economic Review 90, 407–428.
Kashyap, A., Stein, J., Wilcox, D., 1993. Monetary policy and credit conditions: evidence from the composition of external finance. American Economic Review 83, 78–98.
Krishnamurthy, A., 2002. Collateral constraints and the amplification mechanism. Journal of Economic Theory, forthcoming.
Lamont, O., Thaler, R., 2002. Can the market add and subtract: mispricing in tech-stock carve-outs. Unpublished working paper, University of Chicago.
Lo, A., Mamaysky, H., Wang, J., 1999. Asset prices and trading volume under fixed transaction costs. Unpublished working paper, Massachussetts Institute of Technology.
Mitchell, M., Pulvino, T., 2001. Characteristics of risk and return in risk arbitrage. Journal of Finance 56, 2135–2175.
Newey, W., West, K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. Econometrica 55, 703–708.
Shleifer, A., Vishny, R., 1997. The limits of arbitrage. Journal of Finance 52, 35–55.
Stigum, M., 1989. The repo and reverse markets. Richard Irwin, Burr Ridge, IL.
Stock, J., Watson, M., 1989. New indexes of coincident and leading economic indicators. NBER Macroeconomics Annual, pp. 351–394.
Sundaresan, S., 1994. An empirical analysis of U.S. treasury auctions: implications for auction and term structure theories. Journal of Fixed Income 4, 35–50.
Vayanos, D., 1998. Transaction costs and asset prices: a dynamic equilibrium model. Review of Financial Studies 11, 1–58.
Warga, A., 1992. Bond returns, liquidity, and missing data. Journal of Financial and Quantitative Analysis 27, 605–617.
White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.
Xiong, W., 2001. Convergence trading with wealth effects: an amplification mechanism in financial markets. Journal of Financial Economics 62, 247–292.
Yuan, K., 1999. Asymmetric price movements and borrowing constraints: a REE model of crisis, contagion and confusion. Unpublished working paper, University of Michigan.