THE "CHERRY-PICKING" OPTION IN THE U.S. TREASURY BUYBACK AUCTIONS

Bing Han
Francis A. Longstaff
Craig Merrill

Current version: September 2005.

Bing Han is with the Fisher College of Business at The Ohio State University. Francis Longstaff is with the UCLA Anderson School and the National Bureau of Economic Research. Craig Merrill is with the Marriott School of Management at Brigham Young University. We are grateful for valuable comments and assistance from Tony Bernardo, Sushil Bikhchandani, David Hirshleifer, Andrew Karolyi, Val Lambson, Brett Myers, Kjell Nyborg, and Ross Valkanov. All errors are our responsibility.

ABSTRACT

We study an important recent series of buyback auctions conducted by the U.S. Treasury in retiring $67.5 billion of its debt. We find that the Treasury was successful in buying back large amounts of illiquid debt while suffering only a small volatility-related market-impact cost. Although the Treasury had the option to "cherry pick" from among the bonds offered, we find that the Treasury was actually penalized for being "spread too thin" by including multiple bonds in a buyback auction. We find evidence that the Treasury may have attempted to minimize its interest expense rather than its buyback costs in these auctions. There is no evidence, however, that the Treasury used its "timing" option to exploit auction participants.

1. INTRODUCTION

One of the most dramatic recent events in the Treasury bond market was the surprise announcement in January 2000 of the Treasury's first buyback program for its long-term debt in 70 years.[1] Through this program, the Treasury retired $67.5 billion of its debt in 45 separate buyback operations. The introduction of this program was in response to the budget surpluses of the late 1990s as well as to the Treasury's goal of having greater control over its debt structure. Market participants supported the buyback program enthusiastically and individual buybacks were always oversubscribed by wide margins.

The buyback auctions differed from the standard Treasury auctions used to issue bills and bonds in several important ways. Foremost among these was that the bonds involved were older and less-liquid issues, contrasting sharply with the usual type of Treasury issuance auction for highly-liquid on-the-run bills and bonds. Thus, the Treasury faced the risk of suffering huge market-impact costs in buying back such massive amounts of its illiquid debt. To address this problem, the Treasury designed a unique structure for the buyback auction that gave the Treasury the option to "cherry pick" among the various bonds offered by auction participants. Specifically, the Treasury buyback announcement specified the *total* amount of debt to be bought back from a list of as many as 26 bonds, but did not specify the amount for *individual* bonds. Thus, the Treasury could pick and choose among the different bonds offered since they were perfect substitutes for buyback purposes. This option provided at least partial protection to the Treasury against higher buyback costs resulting from the effects of bond-specific illiquidity.

How well did the "cherry-picking" option actually work in reducing the Treasury's buyback costs? On one hand, the Treasury clearly used the option since it did not simply follow a $1/N$ policy among bond issues in buying back its debt. Thus, the Treasury's exercise of the option could have reduced its buyback costs. On the other hand, the option also had the potential to increase the winner's curse problem faced by buyback auction participants and, in turn, affect equilibrium bidding strategies and prices in a way that might actually increase the Treasury's buyback costs. Whether the Treasury was well served by the "cherry-picking" option in the buyback auctions is an important open issue not just for the U.S., but also for a growing number of countries that are considering, or have already implemented, buyback programs of their own

[1]On the day after the announcement, long-term Treasury markets had their largest single-day rally since Black Monday, October 19, 1987. The 30-year Treasury bond rallied by 5 points ($5 per $100 notional amount) and dropped 32 basis points in yield.

such as Australia, Canada, Finland, Germany, Iceland, Italy, the Netherlands, Norway, Poland, South Africa, South Korea, and Sweden.

This paper studies the buyback auctions and the "cherry-picking" option from the perspective of the Treasury's buyback costs. A key aspect of this study is the use of high-frequency intraday data from the Treasury bond market. This has the important advantage of allowing us to measure precisely the Treasury's buyback cost by comparing winning offers with the nearly-simultaneous secondary market price of the bonds.

We first examine the average buyback costs incurred by the Treasury. The average winning offer is only 4.38 cents per $100 notional amount higher than the secondary market ask price of the bond. This average buyback cost is remarkably small given the huge notional amount of bonds involved in the buyback auctions. In fact, this average cost is only slightly larger than the cost for auctioning on-the-run bonds in the standard type of Treasury auctions, and is smaller than the usual bid-ask spread observed in the markets for much smaller quantities. These results suggest that the buyback auction structure was generally successful in minimizing the market-impact cost of buying back less-liquid debt issues. The results also indicate, however, that the Treasury's costs are directly related to the volatility of bond prices, suggesting that at least some of these costs represent compensation to auction participants for bearing the risk of the winner's curse.

We next turn to the key issue of how the "cherry-picking" option affected the results of the buyback auctions. Surprisingly, we find that the Treasury's buyback cost is an *increasing* function of the number of bonds included in each auction. Thus, rather than benefiting from its option, the Treasury actually appears to have been penalized for being "spread too thin" in these buyback auctions. Intuitively, this could occur if auction participants built the full value of the "cherry-picking" option into their offering prices, but the Treasury then failed to fully exploit its option in accepting offers.

To explore this possibility, we study whether Treasury buyback decisions may have been influenced by factors not directly related to minimizing buyback costs. Interestingly, we find that the Treasury had a strong propensity to buy back bonds with higher coupons and longer maturities. This suggests that the Treasury may have been trying to minimize its interest expense rather than just trying to minimize its buyback costs. This would be analogous to a firm following the IRR rule rather than the NPV rule in making capital budgeting decisions.

We also examine whether the Treasury used the "timing" option provided to it by the buyback auction structure. Since the auctions closed at 11 a.m., but auction results were not announced until up to two hours later, the Treasury had the option to condition its decisions on price movements subsequent to the close of the auction.

We find no evidence that the Treasury exploited its "timing" option in this way.

Other recent work on Treasury auctions includes Bihkchandani and Huang (1989, 1993), Cammack (1991), Spindt and Stolz (1992), Umlauf (1993), Simon (1994), Nyborg and Sundaresan (1996), Gordy (1999), Bihkchandani, Edsparr, and Huang (2000), Nyborg, Rydqvist, and Sundaresan (2002), and Keloharju, Nyborg, and Rydqvist (2005). This paper extends and complements this literature in a unique way by focusing on buyback auctions rather than the usual type of Treasury issuance auctions. Because these buyback auctions differ in a fundamental way from the usual issuance auctions, our results offer a new perspective on the Treasury auction process.

The remainder of the paper is as follows. Section 2 describes how the Treasury buyback auctions were structured. Section 3 discusses the data and the empirical approach. Section 4 presents results on the average buyback cost. Section 5 conducts an analysis of the buyback costs. Section 6 examines the Treasury's buyback strategy. Section 7 makes concluding remarks.

2. THE TREASURY BUYBACK AUCTIONS

From March 2000 to March 2002, the U.S. Treasury conducted a series of 45 buyback or procurement auctions for its outstanding debt. Like standard Treasury auctions used to issue debt, these buybacks were sealed-offer share auctions. Unlike standard Treasury auctions, however, each buyback involved between 6 and 26 different bonds. This feature is important since these bonds were essentially perfect substitutes from the perspective of the Treasury's stated objective of buying back a specific total notional amount of its debt (across all bonds). Thus, these buyback auctions represent a unique type of auction with both multi-item and multi-unit auction features. These auctions are also interesting from an informational and risk perspective. Because of the active secondary market for Treasury bonds in which bond prices are almost continuously observable, these buybacks can be viewed as almost pure common-value auctions.

At the program's inception, the Treasury anticipated that four primary benefits would result from the buybacks. First, buybacks would enhance market liquidity by allowing the Treasury to issue benchmark bonds at regular intervals in greater volume. Second, buybacks would allow greater control over the maturity structure of Treasury debt. Third, the buybacks would provide an additional cash management tool. Fourth, the buybacks would allow the Treasury to reduce its interest expense by purchasing "off-the-run" debt and replacing it with lower-yielding "on-the-run" debt.

To initiate a buyback, the Treasury issued a buyback auction announcement one to two days prior to the auction. The announcement described which bonds were to

be included in the buyback as well as the maximum total notional amount (across all bonds) to be bought back. The Treasury reserved the right to buy back less than the maximum amount, but never exercised this option. Intuitively, it is clear why the Treasury would be extremely unlikely to exercise this option. To buy back less debt than was planned would likely be viewed by the market as an auction failure in the same way that an undersubscribed issuance auction would be viewed as an extremely negative signal about the Treasury's ability to sell debt. In a repeated game setting, the cost of an auction failure to the Treasury could easily far outweigh the market-impact costs of an individual buyback auction. In buying back the total desired amount of debt, however, the Treasury usually accepted offers for only a subset of the bonds that it had announced were eligible for buyback.

Any institution approved to conduct open market operation transactions with the Federal Reserve Bank of New York was eligible to submit offers (price and quantity) in the buyback operation. Others who desired to participate could submit offers through the approved institutions. Thus, the number of eligible direct participants in the auction was fixed and finite (although relatively large). In contrast to the debt sales auctions conducted by the Treasury, there was no provision for non-competitive participation.[2]

Offers had to be submitted by the closing time indicated in the buyback auction announcement which was always 11:00 a.m. Offers were binding and the Treasury announced results within two hours after the closing of the auction at 11:00 a.m. Offers were either accepted or rejected at the prices submitted in the offers. For accepted offers, bidders sold their bonds back to the Treasury at their offered price in the quantity offered (discriminatory auction). The public announcement of the results included both the notional amount of offers accepted for each bond as well as the weighted-average price of accepted offers.

The common-value element of these auctions is inherent in the timing of the offers and the acceptance decision. All bidders were required to submit offers by 11:00 a.m. on the day of the auction. Furthermore, all participants could observe the secondary market value of the bonds at the time of offer submission. What was unknown to the bidders, however, is what the value of the bond would be at the time that the results of the auction were announced. The uncertainty about the price was substantial. In particular, the standard deviation of price changes in the secondary market between offer submission and announcement of results is 26.09 cents per $100 notional amount. As we show later, this number is very large relative to the expected benefit to auction participants from an accepted offer. When the auction results were announced, all

[2]For details about the rules governing Treasury buyback operations see *Federal Register*, Vol. 65, No. 12, January 19, 2000 (available at ftp://ftp.publicdebt.treas.gov/gsrfr1192000.pdf).

participants could directly observe the secondary market price of the bonds which was clearly the common value of the bonds to the bidders. Thus, this auction comes close to representing a pure common-value auction. Note that the downside to a bidder of not having an offer accepted was relatively minor. In this situation, the bidder could have simply sold the bonds in the secondary market at the current secondary market price.

There may also be a small private-information component to these auctions. Auction participants were dealers in Treasury securities. Thus, they might be expected to have private knowledge of clients who had a need or desire to sell securities that were included in the buyback operations. This information could influence a dealer's offer. While the secondary market established the common value of the bonds accepted in the auction, the submitted offers varied according to participants' forecasts of future bond values and whatever private information they may have possessed. The observed average difference between accepted offer prices and secondary market prices is a measure of the cost of the auction to the Treasury.

3. THE DATA

Our study uses a dataset for intraday pricing in the U.S. Treasury market provided by GovPX Inc. This widely-used Treasury bond data source consolidates quote and trade information from the interdealer broker market for U.S. Treasury securities. The dataset contains records of the best indicative bids and offers (in both price and yield) along with information about individual trades (trade size, price, yield, and an indicator for which side initiated the trade). Each quote is time stamped to the second. The information is available either directly from GovPX or through financial service distributors including Bloomberg, Reuters, Bridge, and Telerate. The GovPX dataset is used by industry analysts, dealers, traders, brokers, as well as investors.

Since the buybacks were procurement auctions, our focus is on the expected cost to the Treasury of the auction, where the cost is defined as the difference between the average accepted offer price and the corresponding secondary market ask price for the bonds in each buyback.[3] With this measure of the Treasury's cost, we can then explore how the "cherry-picking" option affected the outcome of the buybacks.[4]

[3]We use the ask price quote for all of the results presented in the paper. The results remain virtually the same when we use the average of the bid and ask quotes.

[4]The average accepted offer prices are reported by the Treasury in the announcement of auction results. Buyback and buyback result announcements can be found at http://www.publicdebt.treas.gov/of/ofbuybakpr.htm.

As the measure of the secondary market price for each bond, we use the last market quotation in the GovPX dataset available at the time that the auction results were announced (the auction result announcement is time stamped to the nearest minute).[5] The mean (median) time between the last market quotation and the time of the buyback announcement is 57.96 (39.00) seconds. A total of 95.26 percent of the last market quotes occur within two minutes of the release of the buyback announcement. All of the last market quotations are within 20 minutes of the buyback announcement. Thus, we can compare auction results to virtually simultaneous prices of the bonds in the secondary market.[6]

Table 1 presents descriptive statistics for the Treasury buybacks. Altogether, there were 45 Treasury buyback operations between March 2000 and April 2002. Of these, 20 occurred during 2000, 22 during 2001, and 3 in 2002. All of the bonds included in these buybacks were originally issued as 30-year bonds between 1980 and 1997. In fact, every 30-year bond issued between 1980 and 1997 was included in at least one buyback operation. In total, the 45 buyback operations involved exactly 45 unique bonds.

There were always multiple eligible bonds in each operation. The number of bonds in a redemption operation ranges from 6 to 26, with a median of 11 and mean of 11.4. The maximum total par amount of bonds that the Treasury proposed to buy back in any single operation (as stated in the Treasury redemption operation announcement) is between 0.75 to 3 billion dollars, with a mean of 1.5 billion dollars. The Treasury always redeemed the maximum amount, but did not always spread repurchases evenly among all the eligible bonds. In fact, Table 1 shows that in virtually every auction, there were a number of bonds for which the Treasury accepted no offers at all. (Note that in every auction, market participants submitted offers for every bond listed on the buyback announcement). On average, only about 70 percent of the eligible bonds had at least some offers accepted.[7] Thus, the Treasury clearly did not simply buy bonds back using a $1/N$ rule based on the amount outstanding, amount offered, or

[5]We also use the next quotation after the release of the auction results as the measure of the secondary market value of the bonds. This approach gives results almost identical to those we report.

[6]Three inflation protected bonds (TIPS) were eligible for buyback in the April 18, 2000 auction. This is the only time that TIPS appeared in any buyback. Since GovPX quotes are not available for these bonds, we exclude them from the analysis. Also, bond quotations for part of the day of the 37th buyback are missing from the dataset. Thus, we exclude the 37th buyback from the analysis.

[7]There are also extreme cases. For example, the amount accepted was nonzero for all eligible bonds in two operations, but there were also three redemption operations where only one-fifth of the eligible bonds were bought back. The dispersion seems to be primarily driven by whether the bonds involved in the redemption were callable

the number of bonds on the list. Furthermore, even when the Treasury accepted some offers for a particular bond, the amount accepted was nearly always only a fraction of the total amount offered. The amount offered was on average 4.4 times that of the amount accepted.

The maturity dates of the buyback bonds ranged from February 15, 2010 to April 15, 2032. At the time they were considered for redemption, the time to maturity for these bonds was between 10 to 30 years. The average time to maturity was 18.19 years. The coupon rates for the bonds eligible for buyback ranged from 3.375 to 14.000 percent. Callable bonds and bonds issued in the early 1980's had the highest coupon. The Treasury always chose a block of bonds (i.e., next to each other in issuance date) for buyback in the same operation. Thus, bonds eligible for buyback in a given operation were all very similar in coupon and maturity. Furthermore, the Treasury followed the policy of cycling through the available debt issues in a way that made it relatively easy to forecast which issues might be included in subsequent buyback announcements. Finally, virtually all of the buyback bonds were all selling at premium, and the average bid-ask spread was 6.20 cents per $100 notional amount (or about 1/16th).

4. THE AVERAGE COST

In this section, we examine the average cost to the Treasury of buying back its debt. To calculate the cost to the Treasury from an individual auction, we first calculate the buyback cost for the individual bonds in that auction. As described, the cost of buying back a bond is the difference between the weighted average accepted price and the corresponding price for the bond in the secondary market at the time the auction results are announced. As the measure of the cost for each auction, we take the average (weighted by the notional amount bought back) of the costs for all bonds bought back in that auction. These averages are reported in the last column of Table 1.

The average cost over all buybacks is 4.38 cents (all costs are measured per $100 notional amount of the bonds). The average cost for the individual buybacks ranges from a minimum of -39.3 cents to a maximum of 120.8 cents. The average cost is positive for 59 percent of the auctions. The standard deviation of the mean of the average costs is 3.83 cents.

---

or not. On average, only 30 percent of the bonds in the nine operations that involve callable securities were bought back, while the ratio was 80 percent for the other operations.

To put the average cost to the Treasury into perspective, observe from Table 1 that the average bid-ask spread for auctioned bonds is 6.20 cents per $100 notional amount. Thus, the 4.38 cent point estimate of the average cost to the Treasury is only about 70 percent of the average bid-ask spread. This result is remarkable as it implies that the Treasury suffered virtually no market-impact costs in buying back its debt. For example, recent empirical evidence by Babbel, Merrill, Meyer, and Villiers (2004) estimates that large financial institutions face market-impact costs of more than 10 cents for transactions larger than the typical daily trading volume of about $100 million for individual off-the-run Treasury bonds. In contrast, the average notional amount of each bond bought back by the Treasury in an auction was $184.4 million.

The average buyback cost also compares well with estimates of the market-impact costs in the usual types of Treasury auctions of on-the-run bills and bonds. For example, Goldreich (2005) estimates that the average underpricing of notes and bonds in Treasury auctions during the 1991-2000 period is on the order of 3.50 cents per $100 notional amount. Cammack (1991), Nyborg and Sundaresan (1996), and others provide similar estimates.

5. AN ANALYSIS OF THE BUYBACK COSTS

As the above results show, the buyback auctions allowed the Treasury to buy back its debt at a very low average cost. In this section, we examine the extent to which the unique features of the buyback auction design such as the "cherry-picking" option contributed to this outcome.

In the idealized situation in which market participants provided atomistic offers across a continuum of prices for all bonds, the optimal way for the Treasury to exploit the "cherry-picking" option would have simply been to equate the marginal buyback costs across bonds in each auction. In reality, of course, offers were likely to be "lumpy" and not uniformly distributed across bonds. Thus, it may not have been possible for the Treasury to equate its marginal buyback costs across bonds. With this caveat in mind, and as a preliminary to the analysis of buyback costs, we first examine the extent to which marginal buyback costs were equalized.

The marginal buyback costs can be calculated by simply taking the difference between the maximum price paid by the Treasury for a bond and the market ask price for the bond immediately before the announcement of auction results. As shown in Table 1, there is considerable variation in the marginal buyback costs across bonds in the individual buyback auctions. The standard deviation of the marginal buyback costs ranges from 0.3 cents to 16.3 cents across the auctions. The average value of these standard deviations is 6.4 cents. Thus, not surprisingly, these results indicate

that the marginal buyback costs are not equated across bonds in the buyback auctions. Again, however, this does not necessarily imply that the Treasury failed to maximize the value of the "cherry-picking" option.

Ideally, we would like to have the complete listing of offers presented at each buyback auction to determine how well the Treasury used the "cherry-picking" option to minimize its buyback costs. Unfortunately, this data is not available to us. In its absence, however, there are several ways in which we can identify the effects of the Treasury's option on the buyback costs.

Intuitively, it is clear that the ex post value of the Treasury's "cherry-picking" option should be an increasing function of the number of bonds over which it can choose. Thus, if the Treasury were to use its option optimally, one might expect that the realized buyback cost would be a decreasing function of the number of bonds on the buyback list. To test this hypothesis, we regress the buyback costs from the auctions for which we have data on the number of bonds in each auction. Recall that the number of bonds in a buyback auction varies from 6 to 26. The inclusion of a large number of bonds in each buyback auction is a key distinction between these auctions and the usual type of Treasury auction in which only one particular issue in a maturity range is auctioned. The extensive time-series variation in the number of bonds in each buyback operation can be used to identify the relation between buyback costs and the number of bonds from which the Treasury can "cherry pick."

To shed additional light on the buyback costs, we also include the average volatility of the prices of the bonds involved in a buyback as an explanatory variable in the regression. Theory suggests that participants in common-value auctions should adjust their Bayesian-Nash equilibrium strategies in response to the degree of uncertainty about the value of the auctioned item. In these buyback auctions, where the uncertainty relates to the future market value of the bonds being auctioned, rational auction participants should increase their offer prices as bond price volatility increases. As a measure of the price volatility of the bonds, we compute the standard deviation of 20-minute price changes during the five days previous to the buyback announcement date (and convert it to a one-day volatility by multiplying it by the square root of 21). As the volatility measure for each auction, we use the average of the volatility measures for the individual bonds bought back in the auction.

To explore the extent to which the scale of the auction affected the Treasury's buyback costs, we also include as an explanatory variable the total par or notional amount of bonds to be bought back (in $ millions) as per the Treasury's buyback announcement, normalized by the number of bonds eligible for buyback in that auction. Thus, this measure captures the average amount per bond that the Treasury intends to auction.

To test for the possibility that learning occurred in a way that affected the Trea-

sury's costs over time, we also include the number of the buyback as an additional explanatory variable. The motivation for this variable stems from the literature on sequential auctions (Milgrom and Weber (1982), Weber (1983), Ashenfelter (1989), and many others). This literature suggests that auction participants may learn over time and resolve underlying informational asymmetries. In this context, the Treasury observes all of the prices and quantities offered, but only releases summary information. Thus, as the buyback program progressed, auction participants may have been able to learn about other participants' information sets or supply functions. If so, there could be a trend in the expected cost faced by the Treasury over time.

Finally, in conducting this analysis, it is also important to control for the potential effects of extraneous factors. For example, the Treasury states that one of the motivations for buying back its debt is to replace higher-cost off-the-run bonds with on-the-run bonds. Thus, it is possible that accounting issues could also influence the Treasury's decisions and, therefore, the average cost of the auctions. To control for this possibility, we include the average coupon rate, average maturity, and a dummy variable for callable bonds as additional explanatory variables in the regression. We also control for the possibility of persistence in costs by including the lagged cost as an additional explanatory variable. Finally, we include the ratio of the total notional amount of bonds offered to the total notional amount of bonds accepted for each auction. This variable controls for the degree to which an auction is oversubscribed.

Table 2 reports the results from the regression. The coefficient for the number of bonds included in each auction is positive and significant. Specifically, the regression coefficient implies that the Treasury's expected cost increased by about 3.10 cents for every additional bond included in a buyback auction. This result is very counterintuitive since it indicates that the Treasury actually did *worse* as its opportunities to "cherry pick" increased.

In contrast, the coefficient for the average notional amount of bonds being auctioned is not significant. Thus, the size or scope of the individual buyback auctions does not appear to affect the Treasury's cost directly. This is consistent with recent evidence by Keloharju, Nyborg, and Rydqvist (2005) on the price-quantity relation in Treasury auctions.

The coefficient for price volatility is both positive and significant. This supports the hypothesis that auction participants adjust their offer prices in response to forecasted volatility in the value of the bonds being auctioned. These results parallel those reported by Cammack (1991) and Nyborg, Rdyqvist, and Sundaresan (2002) who find a similar relation between expected auction revenue and volatility in Treasury debt issuance auctions.

The regression coefficient for the time trend variable is positive but not significant. Thus, there is no evidence of learning in this series of sequential auctions. This result is

interesting in itself given the evidence that there are trends in other types of sequential auctions.[8] With the exception of the lagged cost variable, none of the control variables are significant.

6. TREASURY BUYBACK STRATEGY

The previous results are perplexing since they suggest that the Treasury did worse as the number of bonds included in an auction increased. One potential explanation for these results might be that market participants built the ex ante value of the Treasury's option into offering prices, but that the ex post option exercise strategy followed by the Treasury did not maximize the value of the "cherry-picking" option. To explore this possibility, we examine whether there is any evidence of the Treasury pursuing alternative option exercise strategies in their buyback decisions.

In doing this, we focus on the acceptance ratio which is defined as the amount of each bond accepted by the Treasury to the total amount of bonds offered in the corresponding buyback auction. Intuitively, the acceptance ratio provides direct information about the Treasury's buyback strategy. For example, if the Treasury were to follow a strategy of buying back debt equally across bonds, then the acceptance ratio would be equal across bonds within an auction. Alternatively, if the Treasury bought back bonds proportionally to the amount in which they were offered, the acceptance ratio would be perfectly correlated with the ratio of the amount of each bond offered to the total amount of bonds offered.

Our empirical approach consists of regressing the acceptance ratio on variables proxying for potential option exercise strategies. The first variable is simply the offered ratio which is defined as the ratio of the amount offered for each bond to the total amount of bonds offered in the individual buyback auctions. As described above, if the distribution of offers was relatively homogeneous across bonds, then the buyback cost minimizing strategy would tend to result in bonds being accepted in roughly the same proportion in which they were offered.

To examine alternative possibilities, however, we also include measures such as the coupon rate and maturity of each bond, as well as a dummy variable for whether

[8]For example, Milgrom and Weber (1982) show that in a model with affiliation, the expected revenue from a sequential auction can increase over time. In contrast, Ashenfelter (1989) finds that there is a downward trend over time in sequential auctions of identical lots of wine or art. This is the well-known "declining-price" anomaly (or "afternoon effect"). Other empirical studies that document a similar declining trend include Ashenfelter and Genesove (1992), McAfee and Vincent (1993), and Beggs and Graddy (1997).

the bond is callable. As described in the previous section, if the Treasury were to attempt to minimize its accounting interest expense rather than its buyback costs, these variables might be correlated with the acceptance ratio.

As an additional explanatory variable, we include the change in the price of the bond between the 11 a.m. close of the buyback auction and the time at which the Treasury announced the auction results. As shown in Table 1, the time between the two events ranged from 17 to 115 minutes and averaged about 55 minutes. Because of the lag between the 11 a.m. cutoff for submitting offers and the time when the Treasury chose to announce its decisions, the Treasury had a potentially valuable "timing" option. For example, the Treasury could have tilted the acceptance ratio towards bonds with prices that had increased relatively more than those of other bonds during the delay period. In this sense, the Treasury could have potentially exploited market participants who were in effect giving a call option to the Treasury with a time horizon of up to two hours.

The results from the regression are reported in Table 3. As shown, the offered ratio is by far the most significant explanatory variable. The slope coefficient for this variable is 0.296 with a t-statistic of 14.05. This result is broadly consistent with the hypothesis that offers are relatively homogeneous and that the Treasury tended to "cherry pick" in a way that minimized its buyback costs.

On the other hand, Table 3 shows that other factors affected the Treasury's buyback decisions. In particular, both the slope coefficients for the coupon rate and the maturity are positive and highly significant. The signs of these coefficients are consistent with the view that the Treasury had some propensity to also buy back bonds that lowered their accounting interest expenses. As discussed in the introduction, this is analogous to the Treasury putting some weight on an IRR rule rather than using the NPV rule exclusively in buying back its debt.[9]

Finally, Table 3 shows that the Treasury had no propensity to accept more of a bond that had increased in value relative to the others during the auction delay period. This strongly suggests that the Treasury did not attempt to exploit its "timing" option at the expense of market participants. Intuitively, this is consistent with the notion that Treasury auctions are repeated games and that predatory behavior in one auction might have long-lasting negative reputational consequences for the Treasury.

[9]Note that these results are not inconsistent with those in Table 2 in which the average coupon rate and maturity of the bonds included in a buyback auction have no significant explanatory power for the average buyback costs across auctions. The intuition for this is that the earlier results focus on averages across auctions. In contrast, this regression addresses the allocation of demand for individual bonds within an auction and uses coupon rates and maturities of individual bonds rather than averages across all bonds.

7. CONCLUSION

We study how costly it was to the U.S. Treasury to retire $67.5 billion of its debt through the buyback program of 2000 to 2002. Using high-frequency intraday data, we find that the Treasury paid an average of only 4.38 cents per $100 notional amount more than the prevailing market ask price to buy back its debt. This cost is about two thirds the size of the usual bid-ask spread for the bonds that were bought back. On average, the Treasury buyback program appears to have been very effective in retiring debt with a minimum of "market-impact" costs.

The results also indicate, however, that there are features of the unique buyback auction structure used by the Treasury that might not have been optimal for the Treasury. In particular, we find that the average cost to the Treasury was an increasing function of the number of bonds included in each buyback. Since increasing the number of bonds in a buyback allowed the Treasury a greater option to "pick and choose," this evidence suggests that the cost of the option outweighed the benefits the Treasury was able to extract from it. We find evidence that the Treasury may have attempted to minimize its accounting interest costs at the expense of maximizing its "cherry-picking" option.

REFERENCES

Ashenfelter, Orley. "How Auctions Work for Wine and Art." *Journal of Economic Perspectives*, 1989, 3, pp. 23-36.

Ashenfelter, Orley and Genesove, David. "Testing for Price Anomalies in Real-Estate Auctions." *The American Economic Review*, 1992, 82(2), pp. 501-505.

Babbel, David F.; Merrill, Craig B.; Meyer, Mark F. and de Villiers, Meiring. "The Effect of Transaction Size on Off-the-Run Treasury Prices." *Journal of Financial and Quantitative Analysis*, 2004, 39(3), 595-611.

Beggs, Alan and Graddy, Kathryn. "Declining Values and the Afternoon Effect: Evidence from Art Auctions." *The RAND Journal of Economics*, 1997, 28(3), pp. 544-565.

Bikhchandani, Sushil and Huang, Chi-fu. "Auctions with Resale Markets: A Model of Treasury Bill Auctions." *Review of Financial Studies*, 1989, 2, pp. 311-339.

Bikhchandani, Sushil and Huang, Chi-fu. "The Economics of Treasury Securities Markets." *Journal of Economic Perspectives*, 1993, 7, pp. 117-134.

Bikhchandani, Sushil; Edsparr, Patrick and Huang, Chi-fu. "The Treasury Bill Auction and the When-Issued Market: Some Evidence." Working paper, MIT, 2000.

Cammack, Elizabeth. "Evidence on Bidding Strategies and the Information in Treasury Bill Auctions." *Journal of Political Economy*, 1991, 99, pp. 100-130.

Goldreich, David. "Underpricing in Discriminatory and Uniform-Price Treasury Auctions." Working paper, London Business School, March 2005.

Gordy, Michael B. "Hedging Winner's Curse With Multiple Bids: Evidence From The Portuguese Treasury Bill Auction." *Journal of the European Economic Association*, 1999, 81(3), pp. 448-465.

Keloharju, Matti; Nyborg, Kjell G. and Rydqvist, Kristian. "Strategic Behavior and Underpricing in Uniform Price Auctions: Evidence from Finnish Treasury Auctions." *Journal of Finance*, 2005, forthcoming.

McAfee, Preston R. and Vincent, Daniel. "The Declining Price Anomaly." *Journal of Economic Theory*, 1993, 60(1), pp. 191-212.

Milgrom, Paul and Weber, Richard. "A Theory of Auctions and Competitive Bidding." *Econometrica*, 1982, 50, pp. 1089-1122.

Nyborg, Kjell and Sundaresan, Suresh. "Discriminatory Versus Uniform Treasury Auc-

tions: Evidence from When-Issued Transactions." *Journal of Financial Economics*, 1996, 42, pp. 63-104.

Nyborg, Kjell; Rydqvist, Kristian and Sundaresan, Suresh. "Bidder Behavior in Multi Unit Auctions - Evidence from Swedish Treasury Auctions." *Journal of Political Economy*, 2002, 110, pp. 394-424.

Simon, David. "Markups, Quantity Risk, and Bidding Strategies at Treasury Coupon Auctions." *Journal of Financial Economics*, 1994, 35, pp. 43-62.

Spindt Paul and Stolz, Richard. "Are US Treasury Bills Underpriced in the Primary Market?" *Journal of Banking and Finance*, 1992, 16, pp. 891-908.

Umlauf, Steven. "An Empirical Study of the Mexican Treasury Bill Auction." *Journal of Financial Economics*, 1993, 33, pp. 313-340.

Weber, Richard. "Multi-Object Auctions," in Richard Engelbrecht-Wiggans, Martin Shubik, and Robert M. Stark eds., *Auctions Bidding, and Contracting*. New York University Press, New York, 1983.

Table 1

Descriptive Statistics for the Treasury Buybacks. Bonds denotes the number of bonds eligible for buyback in each auction. Not bought denotes the number of bonds for which no offers were accepted. Total offered and total accepted are the aggregate amounts across all bonds in an auction and are measured in $ millions. Delay is the number of minutes between the close of the auction and the Treasury's announcement of auction results. Variation in marginal is the standard deviation taken over all bonds bought back of the marginal buyback cost for those bonds. Average cost is the average buyback cost for bonds bought back weighted by amount bought back.

Buyback Number	Date	Bonds	Not Bought	Total Offered	Total Accepted	Delay	Variation In Marginal	Average Cost
1	Mar 9 2000	13	4	8627	1001	111	0.029	0.073
2	Mar 16 2000	11	0	6444	1001	104	0.050	0.188
3	Apr 20 2000	14	2	8524	2000	40	0.045	0.300
4	Apr 27 2000	26	2	10829	3000	50	0.086	0.390
5	May 17 2000	12	2	9116	2001	72	0.052	−0.060
6	May 25 2000	13	1	8114	2000	63	0.064	−0.393
7	Jun 22 2000	12	1	7340	2000	54	0.072	0.181
8	Jun 29 2000	13	2	7023	2000	36	0.035	−0.006
9	Jul 20 2000	8	1	4440	1500	32	0.128	−0.014
10	Jul 27 2000	6	1	3641	1000	30	0.025	−0.036
11	Aug 17 2000	11	3	6879	1500	40	0.071	0.047
12	Aug 24 2000	10	7	4950	750	35	0.058	−0.051
13	Sep 21 2000	11	1	5888	1500	29	0.056	−0.159
14	Sep 28 2000	10	8	5660	1000	31	0.003	0.056
15	Oct 19 2000	10	1	4785	1501	25	0.105	−0.165
16	Oct 26 2000	11	2	5308	1500	28	0.064	0.081
17	Nov 9 2000	11	1	4796	1250	95	0.061	0.037
18	Nov 16 2000	10	7	4903	1000	88	0.018	0.004
19	Dec 7 2000	11	5	4984	1250	115	0.067	−0.135
20	Dec 14 2000	9	2	4194	1251	97	0.068	0.071
21	Jan 18 2001	12	2	5242	1750	93	0.093	0.143
22	Jan 25 2001	10	6	5158	1000	90	0.112	−0.343
23	Feb 21 2001	12	3	5536	1750	87	0.058	−0.010

Table 1 Continued

Descriptive Statistics for the Treasury Buybacks.

Buyback Number	Date	Bonds	Not Bought	Total Offered	Total Accepted	Delay	Variation In Marginal	Average Cost
24	Mar 1 2001	11	0	5490	1750	85	0.062	0.201
25	Mar 22 2001	11	1	3750	1750	95	0.085	0.112
26	Mar 29 2001	10	6	6197	1000	93	0.092	0.217
27	Apr 19 2001	12	2	8040	2000	28	0.064	0.187
28	Apr 26 2001	12	1	7138	2001	26	0.128	−0.254
29	May 17 2001	12	3	5775	1750	22	0.046	−0.066
30	May 24 2001	10	6	6733	750	40	0.043	0.130
31	Jun 21 2001	12	5	8415	1750	24	0.083	0.125
32	Jun 28 2001	12	1	7643	1750	20	0.054	0.119
33	Jul 19 2001	11	5	6079	1500	33	0.028	0.099
34	Jul 26 2001	10	8	6078	1000	23	0.019	−0.005
35	Aug 16 2001	12	5	10452	1751	23	0.039	0.138
36	Aug 23 2001	12	2	6765	1750	18	0.049	−0.077
37	Oct 18 2001	11	6	5632	1500	66	–	–
38	Oct 25 2001	10	8	5044	1000	24	0.163	0.254
39	Nov 15 2001	12	2	7617	1750	92	0.108	−0.083
40	Nov 29 2001	10	1	4654	1500	89	0.068	−0.548
41	Dec 13 2001	12	3	5393	1500	84	0.062	1.208
42	Dec 20 2001	10	7	4906	1250	78	0.056	−0.154
43	Apr 18 2002	13	5	6213	1000	25	0.078	0.006
44	Apr 23 2002	10	2	5088	1500	20	0.043	0.048
45	Apr 25 2002	12	4	5186	1500	17	0.026	0.067
Average		11.40	3.31	6237	1500	54.89	0.064	0.0438

Table 2

Buyback Cost Regression Results. This table reports the results from the regression of the weighted average cost paid by the Treasury in each auction on the indicated explanatory variables. Number of bonds denotes the number of bonds eligible to be bought back listed in the Treasury buyback announcement. Average volatility is the average daily volatility of the bonds bought back during the five days previous to the buyback announcement. Average par amount is the ratio of the total par amount to be bought back divided by the number of bonds for each buyback, measured in $ millions. The number of the buyback ranges from 1 to 45. The coupon and maturity variables are the average values over bonds bought back in each auction. The callable variable is a dummy variable that takes value one if the bonds bought back in an auction are callable. The offer-accept ratio is the ratio of the total dollar amount of bonds offered to the total dollar amount bought back in each auction.

Variable	Coefficient	t-Statistic
Intercept	−1.9520	−0.89
Number of Bonds	0.031	1.99
Average Volatility	0.499	2.08
Average Par Amount	0.001	0.55
Buyback Number	0.002	0.46
Lagged Cost	−0.314	−1.94
Average Coupon	0.022	0.15
Average Maturity	0.032	0.73
Callable Dummy	0.264	0.84
Offered/Accepted	0.074	1.47
N		43
R^2		0.365

Table 3

Acceptance Ratio Regression Results. This table reports the results from the regression of the acceptance ratio for each bond on the indicated explanatory variables. The acceptance ratio is the ratio of the total amount of a bond accepted by the Treasury in an auction divided by the total amount of bonds offered in that auction. The offering ratio is the ratio of the total amount offered for the bond in an auction divided by the total amount of bonds offered in that auction. The coupon and maturity variables are the respective values for the individual bonds, where maturity is measured in years. The callable variable is a dummy variable that takes value one if the bond is callable. Price change is the change in the value of the bond between the end of the auction and the last price immediately preceeding the announcement by the Treasury of the auction results.

Variable	Coefficient	t-Statistic
Intercept	−0.085	−5.22
Offered Ratio	0.296	14.05
Coupon	0.006	5.21
Maturity	0.002	4.08
Callable Dummy	−0.018	−4.34
Price Change	0.001	0.32
N		513
R^2		0.321