

FLOW AND STOCK EFFECTS OF LARGE-SCALE TREASURY PURCHASES

Stefania D'Amico
Thomas B. King

Division of Monetary Affairs
Federal Reserve Board
Washington, DC 20551

April 2011

ABSTRACT

Using a panel of daily CUSIP-level data, we study the effects of the Federal Reserve's program to purchase \$300 billion of U.S. Treasury coupon securities announced and implemented during 2009. This program represented an unprecedented intervention in the Treasury market and thus allows us to shed light on the price elasticities of Treasuries, theories of imperfect substitutability in the term structure, and the ability of large-scale asset purchases to reduce overall yields and improve market functioning. We find that the program resulted in a persistent downward shift in the yield curve (its "stock effect"), amounting to 20 to 30 basis points in intermediate-maturity zero-coupon yields. In addition, except at very long maturities, purchase operations caused an average decline in yields in the sector purchased of 3.5 basis points on the days when they occurred (the "flow effect"). The coefficient patterns generally support a view of segmentation or imperfect substitution within the Treasury market during this time.

Contact: stefania.d'amico@frb.gov; thomas.king@frb.gov. For helpful discussions and comments we thank other members of the Staff at the Federal Reserve Board and the Federal Reserve Bank of New York, as well as Greg Duffee, James Hamilton, Monika Piazzesi, Tsutomu Watanabe, Min Wei, Jonathan Wright, Dave Zervos, and seminar participants at the Bank of England, the Banco de España, Cal State Fullerton, the Federal Reserve System Real-Time Policy Conference, and the Federal Reserve Bank of San Francisco conference on Monetary Policy at the Zero Lower Bound. Shenje Hshieh, Alex Jutca, and Elizabeth Lewis provided excellent research assistance. Remaining errors are ours alone. The views expressed here do not reflect official positions of the Federal Reserve.

1. Introduction

During the crisis of 2008, policymakers took a number of extraordinary steps to improve the functioning of financial markets and stimulate the economy. Among the most important of these measures, in terms of both scale and prominence, were the Federal Reserve's purchases of large quantities of government-backed securities in the secondary market, conventionally known as the Large Scale Asset Purchase—or "LSAP"—programs. The LSAPs included debt obligations of the government-sponsored housing agencies, mortgage-backed securities (MBS) issued by those agencies, and coupon securities issued by the U.S. Treasury, and they collectively amounted to \$1.7 trillion over a period of about 15 months—the single largest government intervention in financial-market history. Given the unprecedented size and nature of these programs and the speed with which they were proposed and implemented, policymakers could have had, at best, only a very rough *ex ante* sense of their potential impact. The minutes of the December 2008 Federal Open Market Committee meeting summarized the prospects thus: "The available evidence indicated that [LSAP] purchases would reduce yields on those instruments, and lower yields on those securities would tend to reduce borrowing costs for a range of private borrowers, although participants were uncertain as to the likely size of such effects."

There was particular cause for skepticism regarding the program to purchase Treasury securities. The market for U.S. government debt is perhaps the largest and most liquid securities market in the world, and it was not obvious that even such a sizeable intervention—the \$300 billion purchased by the Fed constituted about 8 percent of the coupon market at the time—would have significant effects, given the array of other securities that serve as potential substitutes for Treasuries. Indeed, while the MBS and agency-debt LSAP programs were quickly judged successful, with conventional mortgage rates dropping about 1 percentage point after the announcement of these programs and remaining in the range of 5 percent for the duration of the crisis, the effects of the Treasury program were far from obvious. Treasury yields fell notably when the program was announced on March 18, 2009, but they retraced those declines in subsequent weeks. Some observers even speculated that the program could perversely serve to *increase* yields if the accompanying rise in reserve balances were seen as inflationary or if the Fed were viewed as accommodating fiscal expansion by "monetizing the debt."¹ On the other hand, there were those on the FOMC who argued for an even greater role for the Treasury program, relative to the MBS and agency-debt programs because the Treasury program was seen as having the ability to contribute to reductions in the cost

¹ For example, former Fed Governor Gramley told Reuters in June, 2009, "I don't think they can afford to go out and aggressively buy longer-term Treasuries or even step-up aggressively their purchases of mortgage debt... There is this fear going around in financial markets that the Fed is going to monetize the debt, and we're going to have big inflation. I don't believe that for a minute, but the perception is a reality that the Fed is going to have to deal with." (Bull, 2009).

of credit across a range of markets.² The issue received renewed attention in late 2010, when the FOMC announced further purchases of Treasury securities, first by reinvesting the principal payments on its MBS and agency debt portfolios into Treasury securities and then by expanding the size of its portfolio by an additional \$600 billion.

The FOMC's conviction that purchases of longer-term Treasury securities would reduce longer-term yields appears to have rested on a premise that the relative prices of financial assets are influenced in part by the relative quantities available to the public. Though perhaps intuitive, this view runs counter to the treatment of interest rates implied by the expectations hypothesis and canonical models of the term structure. Instead, such an idea would seem to require grounding in a theory of "imperfect asset substitutability," "portfolio rebalance," or "preferred habitat." Such theories have existed informally for decades (e.g., Culbertson, 1957; Modigliani and Sutch, 1966), but they have recently received greater attention as researchers have begun to supply them with rigorous foundations, as in the models of Andres *et al.* (2004) and Vayanos and Vila (2009). Understanding these mechanisms has policy importance beyond the immediate question of the LSAP programs. For example, under normal circumstances, the Fed conducts most of its open-market operations in Treasuries, the Treasury Department may worry about the interest-rate effects of introducing new supply, and foreign central banks conduct sizeable interventions in Treasuries to maintain their reserves.

Still, empirical work documenting these effects—that is, the extent to which changes in the relative supply of and demand for Treasury debt affect its pricing—is limited. The evidence that does exist typically relies on time-series or event-study methods examining aggregate measures of yields and outstanding debt (as in Bernanke *et al.*, 2004; Engen and Hubbard, 2005; Han *et al.*, 2007; Krishnamurthy and Vissing-Jorgensen, 2007; Greenwood and Vayanos, 2008; and Hamilton and Wu, 2010). While potentially informative, the nature of the data in such studies complicates identification and limits what can be learned about differences in impact across securities.

In this paper, we use a panel of daily CUSIP-level data on LSAP purchases and security prices to study the effects of changes in the supply of publicly available Treasury debt on yields.³ The differences in returns between the securities that were purchased in different amounts during the program identify the impacts of the purchases—that is, the Treasury price elasticities. The CUSIP-level data allow us to parse these reactions more finely than has been possible in previous studies. For example, we are able to examine differential effects of purchases across security characteristics such as maturity and liquidity. In addition, we estimate substitution effects across securities by constructing for each CUSIP buckets of Treasuries with similar remaining maturities and estimating the cross-elasticities of their prices. These cross-elasticities are

² See, for example, the minutes of the January 2009 FOMC meeting (Federal Reserve, 2009).

³ A CUSIP is a unique, security-specific identifying number.

crucial for determining how the aggregate level and term structure of interest rates were affected by the LSAP program and can themselves shed light on the relevance of theories, such as those featuring preferred habitat, that suggest imperfect substitution across securities.

An additional innovation of our approach is to distinguish between two types of impact that the LSAP program might have had—flow effects and stock effects. “Flow effects” are defined as the response of prices to the ongoing purchase operations and could reflect, on top of portfolio rebalancing activity due to the outcome of the purchases, impairments in liquidity and functioning of the Treasury market. Such market imperfections might allow even pre-announced withdraws of supply to have effects on prices when they occur. To estimate flow effects, we model the percentage change in each CUSIP’s price on each day that purchase operations occurred as a function of the amount of that CUSIP and the amounts of substitute securities purchased. Meanwhile, “stock effects” are defined as persistent changes in price that result from movements along the Treasury demand curve and include the market reaction to changes in expectations about future withdraws of supply. To estimate stock effects, we model the cumulative change in each CUSIP’s price between March 17, 2009 and October 30, 2009 (i.e., the cross section of Treasury returns) as a function of the total amount that the Fed purchased of that CUSIP and its close substitutes. Because, over the life of the program, purchased amounts could have responded endogenously to price changes, we instrument LSAP purchases with the purchased securities’ characteristics prior to the announcement of the program. By removing our estimated stock effects from the actual cross section of Treasury prices as of the end of the LSAP program, we are able to construct a counterfactual yield curve that represents what interest rates might have looked like if the program had never existed.

The results suggest that, on average across the yield curve, Treasury purchases reduced yields by about 30 basis points over the life of the program (the stock effect) and led to a further 3 to 4 basis point decline in purchased sectors on the days when purchases occurred (the flow effect). We find that these aggregate effects were driven largely by the responses of less liquid securities, such as those that were several issues off the run. The flow effects were concentrated in securities with remaining maturities of less than 15 years that were eligible for purchase on a given day. Within this set, coefficients across various types of security characteristics and subperiods are quite robust, although we find that the flow effects were more persistent for off-the-run bonds, which is consistent with the stock effect being mainly driven by this category of assets. The sample of securities that were ineligible for purchase exhibits some instabilities in its flow effects, but those results are consistent with the results for eligible securities over the second half of the sample, by which time Treasury market conditions had substantially improved.

We view these results as economically important. A decline in intermediate-maturity Treasury yields on the order of 30 basis points is economically significant by historical standards. Moreover, if this decline had indeed been passed through to private credit markets, it would have represented a substantial

reduction in the cost of borrowing for businesses and households. Although we do not test whether this pass-through actually occurred, the observation that most credit spreads declined during the life of the program suggests that at least some of it may have. It thus appears that the Treasury LSAP program was probably successful in its stated goal of broadly reducing interest rates, at least relative to what they would otherwise have been.

Both the stock- and flow-effect results provide support for portfolio-rebalancing theories, as they demonstrate that Treasury rates are sensitive to the amount of Treasury debt available to the public. This is consistent with the widely held view that Treasury securities play a special role in the global economy and thus are not perfect substitutes for other types of debt.⁴ Our results further indicate that, on the days when a security was eligible to be bought, purchases of securities with similar maturities had almost as large effects on its yield as did purchases of the security itself—that is, the cross- and own elasticities for flow effects were nearly identical—while purchases of maturities further away had smaller effects. This supports the view that Treasuries of similar maturities are close substitutes but that substitutability diminishes as maturities get farther apart, consistent with imperfect substitutability across the term structure. In addition, we find that certain types of Treasury securities exhibit greater evidence of segmentation, and our sub-sample analysis seems to suggest that segmentation is determined by the intrinsic characteristics of the assets rather than market conditions, which is supportive of “preferred habitat” theories. For example, we generally reject equality of the own- and cross-elasticities in far-off-the-run securities, suggesting that limits to arbitrage may play an even greater role among those securities.⁵ In other words, in segmented portions of the market preferred-habitat investors’ demand seems to dominate arbitrageurs’ demand. The following section of the paper provides an overview of the Treasury LSAP program. Section 3 reviews some of the theory that could motivate the LSAP programs and some previous evidence on the link between Treasury supply and interest rates. Section 4 presents our empirical work, with sub-section 4.1 discussing some general issues, sub-section 4.2 considering stock effects, and sub-section 4.3 considering flow effects. Section 5 offers concluding thoughts.

2. Details of the Treasury LSAP Program

The program to purchase up to \$300 billion of Treasury coupon securities was announced in the FOMC statement on March 18, 2009. Although policymakers had floated the possibility of Treasury purchases prior to this announcement, market participants did not appear to place high odds on the decision

⁴ Some possible reasons for this specialness include the exceptional liquidity and safety that are associated with the Treasury market, Treasuries’ heavy use as collateral in repurchase agreements, and their preferential treatment as assets in assigning regulatory and credit ratings for financial institutions.

⁵ Limits to arbitrage can originate from arbitrageurs’ limited risk-bearing capacity and/or end-investors’ slow moving capital (e.g. Mitchell, Pedersen, and Pulvino, 2007; Greenwood and Vayanos, 2008; Gurkyanak and Wright, 2010; and Lou, Yan and Zhang, 2010).

being taken, and the announcement generally came as a surprise.⁶ The first operation under the Treasury LSAP program was conducted on March 25. Purchases continued at a pace of about \$10 billion per week over the subsequent five months. On August 12, 2009 the Committee announced that it would purchase the full \$300 billion (eliminating the ambiguity of the “up to” language) and that it would wind down the program in October. After this announcement, the pace of purchases gradually slowed to minimize any potential disruption that might have resulted from a sudden closing of the program.

The logistics of the purchase operations were as follows. Every-other Wednesday, the Desk announced the broad maturity sectors in which it would be buying over the subsequent two weeks and the days on which it would be conducting these operations. These maturity sectors included securities spanning ranges of between one year (at the short end of the yield curve) and 13 years (at the long end), with an average range of about four years. Auctions took place from Monday of the first week through Friday of the second week and typically settled on the following day.⁷ At 10:15 on the morning of each auction, the Desk published a list of CUSIPs that were eligible for purchase, which generally included nearly all securities in the targeted sector,⁸ and began accepting propositions from primary dealers. Propositions included the amount of each CUSIP that the dealer was willing to sell to the Desk and the price at which it was willing to sell. At 11:00 AM, the auction closed. The Desk then determined which securities to buy from among the submitted bids based on a confidential algorithm and published the auction results within a few minutes. Market participants were not aware in advance of the total amount to be purchased or of the distribution of purchases across CUSIPs. Notably, settlement of the winning bids did not occur until the following day, so that dealers could, in principle, have submitted propositions for securities they did not own and, if they won, purchased these securities to settle the next day with the Desk.

Overall, purchases of nominal securities under the Treasury LSAP program included 160 unique CUSIPs, spanning remaining maturities of about two to thirty years. \$300 billion represented about 3 percent of the total stock of outstanding Treasury debt and about 8 percent of the outstanding coupon securities as of the time of the announcement. Table 1 and Figure 1 provide some additional statistics summarizing the characteristics of purchased securities. Most purchases were concentrated in the 2- to 7-year sectors, although, as a percentage of total outstanding Treasuries within each sector, purchases across maturities were less concentrated. Coupon rates and vintages of securities purchased were roughly similar to

⁶ For instance, the morning of the FOMC announcement, Bloomberg reported that Goldman Sachs and several other banks believed that policymakers would not introduce such a program and that the March FOMC statement would be largely similar to the January statement:

<http://www.bloomberg.com/apps/news?pid=20601087&sid=alt7yEi9XAhc&refer=home>

⁷ In practice, the Desk avoided conducting Treasury operations on Fridays, preferring to reserve these days for agency purchases. They also avoided conducting purchases in a given sector on days when Treasury auctions were occurring in that sector.

⁸ The securities that were excluded were the cheapest-to-deliver in futures markets, those with high scarcity value in the repo market, and those for which 35 percent of the issue has already been purchased under the LSAP program.

the averages of all outstanding Treasuries. The maturity of securities bought was a bit longer than average, but the yields on purchased securities were notably higher than average—seemingly by too great a margin to be accounted for solely by their slightly longer maturities, especially given that a relatively high fraction (approximately 30 percent) of purchases were on-the-run issues, which generally have lower yields. This suggests that, consistent with contemporary commentary, the Desk deliberately purchased securities that were underpriced, a claim that we will illustrate more formally below.

Figure 2 shows the behavior of Treasury yields over the period of the program. After the initial announcement, medium- to long-dated yields fell by as much as 50 basis points, with yields in the 5- to 10-year sector declining the most. However, the decline was short-lived—by early May most yields had returned to their pre-announcement levels, and they shot up further in June as the economic outlook improved, policy uncertainty increased, and investors shed duration following a rise in mortgage rates. Although some of these increases reversed by October, most yields were still 20 to 40 basis points higher at the end of the program than they were before it started, and indeed they increased by the greatest amount in precisely the 5- to 10-year portion of the term structure where purchases were concentrated. These increases led some observers (e.g., Thornton, 2009) to conclude that the LSAPs had been ineffective in reducing interest rates. Of course, as we demonstrate below, such reasoning ignores other factors that may have been influencing yields over this period, including the possibility that the distribution of purchases itself was responding endogenously to relative changes in Treasury prices.

Another notable pattern over this period—and one that may itself have been due in part to the Treasury LSAP program—is the improvement in liquidity in the Treasury market. Traders had pointed to reduced liquidity as an important factor putting upward pressure on some yields in the weeks leading up to the introduction of the program. As Table 2 illustrates, almost every measure of liquidity improved between the first and second halves of the program’s life. Average trading volumes increased by 20 percent, the yield premium paid for on-the-run (i.e., the most recently issued) 10-year note over off-the-run securities with comparable remaining maturities fell by a quarter, and failures to deliver securities into repurchase agreements on Treasuries declined by 80 percent. The final column of the table shows the average residuals that result from fitting a smooth curve, using the functional form proposed by Svensson (1994), to the cross-section of yields on each day. These yield curve “fitting errors,” which can be interpreted as a measure of unexploited price discrepancies, declined by about half between the two sub-periods.

3. Theory and Evidence on the Effects of Treasury Supply

3.1 Theoretical Motivation

In this paper, we ask whether changes in the stock of Treasuries during the LSAP program affected the yields on Treasuries. Perhaps surprisingly, the arbitrage-free models of the term structure of interest rates that have become standard in the finance literature do not generally allow for this possibility. To see why, consider the simple example of a frictionless market with risk-neutral traders, where the short-term (instantaneous) risk-free rate evolves according to some exogenous stochastic process. Because of risk neutrality, a strict version the expectations hypothesis, in which there are no term premiums, holds, and the rate of maturity n at time t is given by

$$r_t^n = \frac{1}{n} \int_0^n E_t r_{t+s}^0 ds \quad (1)$$

Note that this condition involves no role for fluctuations in the outstanding supply of longer-term debt, provided that such fluctuations do not influence the instantaneous risk-free rate. Indeed, if we further suppose that the short-term rate follows the mean-reverting process

$$dr_t^0 = \phi(k - r_t^0) dt + \sigma dB_t \quad (2)$$

where B_t is Brownian motion, then the integral in (1) reduces to an affine function of the time- t short-term rate:

$$r_t^n = a(n) + b(n)r_t^0 \quad (3)$$

where $b(n) = (1 - \exp(-\phi n))/\phi n$, and $a(n) = k(1 - b(n))$. Under the expectations hypothesis, asset purchases should have no effect on yields because they do not appear anywhere in this equation.

Relaxing the assumption of risk neutrality allows for departures from the expectations hypothesis but is not, by itself, sufficient for the supply of debt to matter. Since the prices of securities with different durations exhibit different sensitivities to changes in expectations of the path of short rates, they involve different levels of risk and, when the marginal investor is risk-averse, will command different risk premia. Thus, under risk aversion, securities of different maturities are not perfect substitutes. Indeed, under risk aversion, allowing either time variation in volatility or in the risk-aversion parameter itself generally changes the functional form of $a(n)$ and $b(n)$ and introduces additional linear terms in equation (3), giving rise to the class of so-called affine term-structure models. Still, these extensions do not generally create an explicit role for quantities to matter. While *estimations* of such models could of course capture the effects of changing supply through the use of latent variables, that possibility provides little scope for testing whether supply matters or for shedding light on exactly how it might matter, given that those factors rarely have clean interpretations. Indeed, the usual way that unobserved factors enter in affine models—by affecting short-

term rates, risk, or the price of risk—does not appear to capture fully the types of mechanisms that economists seem to have in mind when they consider supply effects.

One way of getting changes in bond supply to affect pricing in a theoretical model is to suppose that investors prefer holding certain types of securities. Greenwood and Vayanos (2008) and Vayanos and Vila (2009) develop models formalizing this idea. In their theory, preferred-habitat investors have preferences for certain maturities, independent of their risk and expected returns.⁹ While arbitrageurs can profit by buying securities that are in low demand and selling securities that are in high demand, risk aversion prevents them from engaging in this process until expected returns are exhausted. Thus, exogenous fluctuations in demand for different types of securities can have effects on prices. A simple version of the Vayanos-Vila model generates a solution of the form

$$r_t^n = a(n) + b(n)r_t^0 + c(n)\delta_t \quad (4)$$

where δ_t is a “demand factor” that, for any given configuration of the yield curve, maps into specific quantities purchased by preferred-habitat investors. Higher levels of risk aversion on the part of arbitrageurs result in larger values of $c(n)$. Since LSAP purchases affected the demand for securities in certain sectors, such purchases could be modeled as shifts in δ_t . We do not provide explicit tests of Vayanos and Vila or any other asset-pricing model in this paper, but this intuition captures the type of mechanism we have in mind.¹⁰ If a channel like this one exists, withdrawing Treasury supply through programs like the LSAP would be able to affect the level and slope of the yield curve (implying a change in term premiums).

Importantly, the significance of the demand factor in this model depends on the degree to which arbitrage is able to occur, as determined by the prevalence of the preferred-habitat investors and the risk aversion of the arbitrageurs. (More broadly, risk aversion in this framework could be thought of as a stand-in for capital constraints or other limits to arbitrage.) This suggests that demand shocks may have greater effects on the prices of less-liquid assets. Though not emphasized in previous work, another consequence of a market-segmentation assumption is that changes in supply could have effects when they occur, even if those changes are expected in advance. Just as risk-averse (or otherwise constrained) arbitrageurs cannot perfectly eliminate the effects of demand factors across the yield curve at any point in time, they also cannot perfectly eliminate the effects of demand factors at a point on the yield curve over time. Under imperfect

⁹ The classic example of such a preference is an insurance company or pension fund that wants to match a particular set of long-term liabilities with a portfolio of assets that have similar cash flows or a bond mutual fund that is restricted to hold longer-term maturities. More broadly, it is also possible to imagine investors with preferences for short- or medium-term securities, securities with especially high liquidity (such as on-the-run issues), or securities with particular coupons. Andres *et al.* (2005) provide an alternative model that also allows for preferred habitat. However, it is difficult to square their framework with empirical work on Treasury prices because it abstracts from many features of the Treasury market, including the possibility of secondary-market trading, which is obviously crucial when studying daily movements in the yield curve.

¹⁰ While we believe that this channel likely captures the first-order effect of LSAPs, it does not rule out the possibility that LSAP purchases could also change expectations of short-term interest rates through a signaling channel.

arbitrage, if some investors (such as the preferred-habitat investors in the Vayanos-Vila model) respond only to spot prices and not to expectations of possible returns, spot prices will change when supply is withdrawn, even if that withdraw is perfectly anticipated. This could be one mechanism generating “flow effects” of LSAP purchases.

3.2 Previous Evidence

A number of previous studies have examined empirically the extent to which changes in outstanding Treasury supply affect the nominal yield curve by examining various circumstances in which supply was added or withdrawn by the government. Tarhan (1995) examined market prices around Federal Reserve purchases of longer-term securities for the System Open Market Account (SOMA) portfolio in the usual course of its open-market operations. Engen and Hubbard (2005) and Krishnamurthy and Vissing-Jorgenson (2008) both used time series analysis to examine the extent to which various yields track levels and forecasts of government debt/GDP ratios.¹¹ Kuttner (2006) used “excess returns” regressions in the style of Cochrane and Piazzesi (2005) to test whether the Fed’s SOMA purchases had significant effects on term premia. Finally, Greenwood and Vayanos (2008) and Hamilton and Wu (2010) used the aggregate term structure of outstanding Treasury debt to explain the time-series patterns in longer-term yields. Results from these studies generally suggest negative effects of supply on yields, but there is substantial variance across the magnitudes and statistical significance of the estimates. All of them use a fairly limited amount of the data that are potentially relevant (e.g., by looking only at a handful of yields and some aggregated measure of outstanding supply characteristics), which limits what they can say about variation across different classes of securities. It is also often unclear how the results in these studies might be affected by expectations of supply versus the amount of supply that is actually outstanding—that is, possible flow and stock effects are often conflated. By using CUSIP-level data, we hope to be able to address these issues. In addition, outside of some recent event-study analyses,¹² our paper is the first to focus explicitly on the consequences of fluctuations in available Treasury supply during the financial crisis period, a time when theory suggests higher levels of risk aversion and illiquidity may have led supply effects to be particularly pronounced.

Another type of evidence comes from the case studies of particular episodes that have involved relatively large or rapid changes in Treasury supply. One such episode that is well known is the Federal Reserve’s attempt to decrease long-term yields relative to short-term yields in the early 1960s—termed “Operation Twist”—which involved the sale of short-term Treasury debt and the purchase of long-

¹¹ Engen and Hubbard (2005) also review related literature in macroeconomics and public finance that studies whether government deficits are related to the level of interest rates (i.e., tests of Ricardian equivalence). They argue that this literature is generally inconclusive.

¹² These event studies include Gagnon et al. (2010) and Neely (2010). These studies generally do not distinguish the effects of the Treasury program from those of the other MBS and agency-debt programs, and their methodology does not allow them to consider stock and flow effects as we define them.

term debt. This program was analyzed in several contemporaneous studies¹³ and more recently by Hakim and Rashidian (2000) and Swanson (2010). These studies have generally not found large effects of the program.

Bernanke *et al.* (2004) studied the responses of the yield curve during several more-recent cases of government intervention in Treasury market and concluded that such interventions could have significant effects on yields, and they cited this evidence as providing possibilities for monetary policy when short-term interest rates are constrained by the zero lower bound. One of these cases in particular is worth noting, because it came close to the LSAP program in mechanics, if not in scale. Namely, under Treasury Department's repurchase of long-term debt during the early 2000's, the government bought back \$67.5 billion of bonds (about 1.2 percent of outstanding Treasury debt, but 10.4 percent of the marketable debt with maturity greater than ten years), entirely in off-the-run issues with original maturities of 30 years. In many details, the operations were similar to LSAP purchases—for example, the broad sector for each operation was announced in advance, but Treasury could choose which securities to purchase from among submitted bids within that sector.¹⁴ Bernanke *et al.* (2004) argued that the buyback program had significant effects (although they did not provide precise estimates of the magnitudes). Longstaff (2004) found also that the Treasury buyback program had significant effects on spreads between Treasuries and other government-guaranteed debt (RefCo bonds), and Greenwood and Vayanos (2010) argue that the purchases served to decrease longer-term Treasury yields relative to short-term. On the other hand, Han et al. (2005) found small and statistically insignificant differences in yield changes between bonds that were purchased and those that were not.¹⁵

Overall, the existing evidence seems consistent with the hypothesis that changes in Treasury supply affect yields, perhaps through preferred-habitat type channels. However, it is difficult to get a precise idea of the magnitude of these effects or of how they vary with security characteristics because most previous analysis is done at an aggregated level. No previous study that we know of exploits variation in quantities and prices for individual securities to account for timing issues (flow versus stock effects), substitutability across the term structure, and specific security characteristics, as we do below.

¹³ Modigliani and Sutch (1966), Ross (1966), and Holland (1969).

¹⁴ Importantly however, in contrast to the LSAP program, the Treasury buyback program was designed to increase new issue size and frequency, improving liquidity of the on-the-run and preventing average maturity from rising, and the Treasury specifically attempted to *minimize* the effects of operations on market prices, since increases in prices would have driven up the cost of the purchases.

¹⁵ Bernanke et al. (2004) also examine two other episodes that may pertain to the effects of changing Treasury supply on yields—the initiation of large purchases of Treasuries by Asian central banks in 2002 and the market perception that the Fed might undertake an LSAP-type program in 2003. Kuttner (2006) also provides a narrative overview of several instances of large interventions in Treasury markets and suggests that most of these interventions have had economically significant effects. Finally, Greenwood and Vayanos (2010) argue that U.K. pension reforms in 2004 had significant effects on the term structure of gilt yields.

4. Empirical Work

To organize the discussion of our results, we note four qualitative hypotheses that are generated by the preferred-habitat view, on which we will offer evidence:

- H1. Government purchases of Treasury securities have significant effects on Treasury yields.
- H2. Those effects are largest for the particular securities that are purchased, somewhat smaller for securities that have similar maturities to the purchased securities, and minimal for securities that have much different maturities than the purchased securities.
- H3. These differences in responses are more pronounced when purchases take place in portions of the market that are more likely to be segmented—that is, among securities that are not widely traded or for which many good substitutes do not exist.
- H4. Among such securities, even anticipated purchases might not be fully priced until they actually occur, resulting in persistent price changes on the day of operations.

As noted, we will analyze Treasury price changes both using panel data on the days when purchases occurred (flow effects) and cross-sectionally over the period of the program as a whole (stock effects). Both sets of tests share the same data sources and variable definitions, described below.

4.1 Data and Specification Issues

Our data consist of daily observations on the universe of outstanding nominal Treasury coupon securities from March 17 through October 30, 2009.¹⁶ To simplify the analysis, we exclude TIPS and securities with remaining maturities of less than 90 days, leaving us with an unbalanced panel of 204 CUSIPs (including 44 securities that were never purchased under the LSAP program). Our dependent variables are percentage price changes in each of these securities (measured at end-of-day). Our independent variables are constructed from the security-level amounts purchased and total outstanding amounts.¹⁷

A central element in our approach is the possibility that the price of a given security may also move in response to purchases of other securities for which it is a substitute. Toward this end, we define buckets of substitutes for each security. Although in principle we could choose the size and number of the substitute

¹⁶ We examine only the initial round of Treasury purchases announced and conducted in 2009, not the second round (the so-called QE2) that commenced in August, 2010. This program was still on-going as of the writing of this paper and we are therefore unable to apply all of our analysis to it. In addition, some of the details differed between the first and second round of purchases—for example, in the second round the New York Fed announced the full maturity distribution of purchases in advance—making it harder to compare across the two episodes.

¹⁷ Purchased amounts by CUSIP are publicly available on the Federal Reserve Bank of New York's website, and amounts outstanding are based on information from the Treasury Department. Daily pricing data come from NPQS.

buckets in a variety of ways, a division into three buckets based on remaining-maturity ranges seemed to provide a good combination of parsimony and flexibility. In particular, for each security i , we define our most narrow bucket to include all securities having remaining maturities within two years of security i 's maturity. We refer to these securities as “near substitutes” for security i . The second bucket, which we call “mid-substitutes” for i , includes all securities having remaining maturities that are between two and six years different from security i 's. The third bucket (“far substitutes”) includes all securities having remaining maturities between six and fourteen years different from security i 's. We denote the dollar amount of each bucket purchased by the Desk on day t by Q_{ijt}^{sub} , where j indexes the degree of substitutability. In addition, we let Q_{it}^{own} denote the amount of security i purchased. We normalize Q_{ijt}^{sub} and Q_{it}^{own} by the total amount of securities outstanding that have remaining maturities within two years of security i (that is, the sum of a security's own amount outstanding and the amount of its near substitutes), which we denote by S_{it} . This normalization generates coefficients that all take the same units, allowing us to compare the effects of a given dollar amount of purchases across different sectors.¹⁸ We denote the normalized values by q_{ijt}^{sub} and q_{it}^{own} . Finally, for completeness, let O_{it} denote the amount outstanding of security i .

Because coupon rates and maturities vary considerably across the universe of Treasury securities that we work with, we conduct our regressions in price space, rather than in yield space.¹⁹ Converting equation (4) to prices and applying Ito's lemma gives an equation for the bond return dp_t^n/p_t^n that is linear in the levels of and shocks to the short-term interest rate and demand factor. This motivates the following general form for both our flow- and stock-effect regressions:²⁰

$$R_{it} = \beta q_{it}^{own} + \sum_j \gamma_j q_{ijt}^{subs} + \boldsymbol{\phi} \mathbf{x}_{it} + \varepsilon_{it} \quad (5)$$

where R_{it} is the gross return (change in principal plus accrued coupon payments) on Treasury security i during period t and \mathbf{x} is a vector of control variables. In our flow-effect regressions we use daily data, while the stock-effect regressions are cross-sectional, so that t represents the entire sample period in that case. The constituents of \mathbf{x} also differ between the flow- and stock-effect models.

The parameter β reflects the *own*-price elasticity of Treasury securities, while the parameters γ_j primarily reflect the *cross*-elasticities of Treasury security prices with respect to other Treasury securities.

¹⁸ In economic terms, the rationale for this denominator is that it proxies for the relevant measure of supply within a given sector. We tried various alternative normalization schemes and generally found results consistent with those reported below.

¹⁹ However, the central results also hold if we use yields as the dependent variable.

²⁰ Of course, this motivation for the functional form is only heuristic—the securities in our dataset are not zero-coupon, and we do not impose the restrictions on the coefficients that would be required by a formal affine term-structure model.

These latter elasticities depend on the degree of substitutability between different Treasury issues, which in turn depends upon the liquidity of the Treasury market and the ability of participants to arbitrage away price inefficiencies in this market. The own-price response β is of some interest, as its magnitude is indicative of the purchases' effects on the amounts by which an individual security's yield could deviate from those of similar securities (i.e., yield-curve fitting errors). The cross-responses γ_j , however, are likely to be much more important in terms of the aggregate level and term structure of interest rates. This is because the purchase of a particular security affects that security's yield alone through the β term, but it affects *every* security's yield through the applicable γ_j terms.

4.2 Stock Effects

Specification

By “stock effects” we mean the impact that the LSAP program had on prices by permanently reducing the total amount of Treasury securities available for purchase by the public. Of course, expectations of such effects should have been impounded into Treasury prices as soon as the market became aware of the program, before any purchases took place—presumably, this mechanism accounted for much of the 25 to 50 basis point drop in Treasury yields on the day the program was announced. Thus, it is crucial to account for expectations when measuring stock effects. However, we note that these effects matter prior to the conclusion of the program. In other words, while there may be temporary price fluctuations reflecting changing expectations of future purchases, these expectations become irrelevant once the total actual amounts and distribution of purchases is revealed. Thus, all else equal, the difference in price changes across two securities between the time the program was announced and the time it was concluded should depend only on the relative amount of each security that was *actually* purchased over the life of the program.

With this in mind, our regressions for the stock effects use the cross section of total price changes for all nominal Treasury coupon securities between March 17 and October 30, 2009. Some previous studies of LSAPs, such as Gagnon *et al.* (2010), have tried to identify their effects by looking at the reaction of prices within specific event windows around important announcements. The difficulty with this approach (apart from the difficulty of specifying the appropriate windows) is that it relies solely on changes in expectations of purchases that occur within the windows—if market participants had some expectation of purchases prior to the windows, or if they changed their expectations any time outside the windows, or if they waited until purchases actually occurred to fully impound their effects, the event study will not capture the true effects of the program. Instead, our approach relies solely on cross-sectional variation for identification and is therefore less susceptible to this sort of timing critique.

However, there is an obvious danger of endogeneity in our exercise—if the Fed was deliberately targeting securities that were underpriced, purchases may have been higher among issues whose yields rose

the most during the life of the program. To control for this possibility, we use two-stage least squares. In the first stage, we instrument the LSAP purchase amounts of each security using information available before the program was announced. Specifically, our instruments are: the residuals (“fitting errors”) from a yield curve estimated on March 17 using the Svensson functional form; the percentage of the security held by the SOMA portfolio as of March 17; a dummy variable for whether the security was on-the-run as of March 17; and a dummy variable indicating whether the security had less than two years remaining until maturity on March 17. We also include remaining maturity and remaining maturity squared as exogenous variables in both the first- and second-stage regressions to account for possible secular changes in the slope and curvature of the yield curve during our period, such as could have resulted from macroeconomic conditions and new Treasury issuance.

Table 3 reports the results of a regression of actual LSAP purchases (as a percentage of the par value of each security outstanding) on these instruments. All of the coefficients are statistically significant at the 1 percent level. The coefficients on the maturity variables suggest that purchases depended strongly on remaining maturity and, controlling for other factors, peaked around the twelve-year sector (in percentage-of-issue terms). The yield-curve fitting errors have a positive sign, confirming the conjecture that the Desk tended to purchase securities that were underpriced (i.e., had higher yields) than other securities with similar remaining maturities. The Fed was less likely to purchase securities that it already owned, presumably reflecting its self-imposed 35 percent limit, and, as was suggested in Table 1, was more likely to purchase on-the-run than off-the-run issues. Finally, the Fed purchased virtually nothing with maturity of less than two years.²¹

Two further complications arise with our instruments. First, if we believe that the specification presented in Table 3 is the right one—the Desk likely determined how much to buy as a fraction of the amount of that CUSIP that was outstanding. However, in our second-stage regressions, we want to normalize by the amount outstanding in the *sector*— Q_i^{own}/S_i , not Q_i^{own}/O_i . We thus use q_i^{own} as the dependent variable in the first stage but, to maintain consistency, we weight each of our four security-level instruments by O_i/S_i . Second, as noted above, we also want to account for the possible effects of substitute purchases. Purchases of near substitutes are subject to the same endogeneity concerns as own purchases, so we also instrument for this variable.²² As instruments, we simply average each of the four instrumental variables listed above over the bucket of near substitutes for each security, weighting by amounts

²¹ One interpretation of these results is that savvy market participants could have predicted with a fair degree of accuracy which securities would be purchased the most. Indeed, it appears that this occurred to some extent. For example, on the day the program was announced, yields on securities that were ultimately purchased fell by an average of 39 basis points, while yields on securities that did not end up being purchased fell by only 27 basis points, even though the announcement gave no specific indication of which securities would be purchased.

²² We do not use the mid- or far-substitute categories in the cross section because of the high degree of collinearity, especially given our inclusion of the remaining maturity variable.

outstanding. We include both the security-specific and sector-average instruments in both of our first-stage equations.

In the second-stage regression, we use instrumented purchases from the first stage as independent variables and the cumulative changes in Treasury prices as the dependent variable. In summary, our baseline two-stage system takes the form

$$\begin{pmatrix} q_i^{own} \\ q_i^{sub} \end{pmatrix} = \boldsymbol{\theta} \mathbf{x}_i + \mathbf{u}_i \quad (6)$$

$$R_i = \alpha + \beta \hat{q}_i^{own} + \gamma \hat{q}_i^{sub} + \phi_1 m_i + \phi_2 m_i^2 + \varepsilon_i \quad (7)$$

where R_i is security i 's gross return, \mathbf{x}_i is the vector of instruments, hats indicate instrumented values from the first stage, and m_i is the remaining maturity as of March 17. Because we are using a cross section, we exclude securities that matured or were issued while the program was in progress, leaving us with 148 observations.

Finally, in order to examine how our results vary with liquidity and other security characteristics, we want to allow the second-stage coefficients to differ across securities. In particular, we will divide the sample by security type, maturity, and vintage. The small number of observations makes running separate regressions on these sub-samples problematic, and, moreover, there is no particular reason to think that the first-stage equation or the second-stage remaining-maturity coefficients should differ across them. Therefore, in the second stage, we run a single regression but use interactive dummies to allow the coefficients on own and substitute purchases to differ across the subsamples:

$$R_i = \alpha + \beta_k \hat{q}_i^{own} + \gamma_k \hat{q}_i^{sub} + \phi_1 m_i + \phi_2 m_i^2 + \varepsilon_i \quad (8)$$

where k indexes the security groups.

Results

The results of the second-stage regressions, with gross returns as the dependent variable, are presented in the first column of Table 4. Both own purchases and near-substitute purchases have positive and statistically significant effects on returns, although the effects of own purchases appear to be considerably larger. However, a likely source of misspecification in these results is that, if individual yield curve fitting errors are not persistent, the prices of securities with positive fitting errors would tend to fall relative to other securities and those with negative fitting errors would tend to rise, even in the absence of LSAP purchases—in other words, fitting errors might be correlated with the second-stage error term.²³ In addition, there may be other information embedded in the initial prices of securities that reflects expectations of future returns. To control for these possibilities, we run an alternative specification in which we include the initial percentage

²³ However, in our sample fitting errors are quite persistent.

price fitting error and the initial log price as exogenous regressors. (We also include both variables in the first-stage regression.) These results are shown in the second column. Both of the new regressors are highly significant, arguing for their inclusion. In this specification, the estimated values of both the own and near-substitute coefficients fall notably, but they retain their basic qualitative pattern and statistical significance.

The coefficient of 0.07 on substitute purchases suggests that buying 1 percent of a security's near substitutes (about \$10 billion for the average security in our sample) increased the price of that security by 0.07 percent. For a typical ten-year Treasury, with a modified duration of 7 years, this translates into a yield change of -1 basis point. The coefficient on own purchases implies that if the same dollar amount had been used entirely to purchase a single security, the price of that security would have risen by 0.61 percent; again taking the ten-year Treasury as representative, its yield would have fallen by about 9 basis points. The strong statistical and economic significance of these results supports hypothesis H1—purchases have effects on yields. In addition, the smaller magnitude of the substitute coefficient than the own coefficient suggests imperfect substitution across securities within the same sector, supporting hypothesis H2.

We next break our data into subsamples, using the interactive-dummy technique described above. In particular, we consider possible differences between notes and bonds, securities with relatively long maturities versus those with shorter maturities, and recently issued securities versus older issues. The distinction between notes and bonds is potentially interesting largely as a proxy for other security features. The bonds all had original maturities of 30 years and most of them are old issues that tend to have much smaller trading volumes and lower liquidity than notes, as well as higher coupons.²⁴ To distinguish longer and shorter maturities, we split the sample at the middle of the yield curve, 15 years. (About 90% of the securities in our sample were in the less-than-15-year sector.) To distinguish securities of recent versus older vintage, we split the sample in to securities that are more than five issues off the run and those that are less than six. In all of these regressions, we include controls for initial prices and fitting errors, given the better fit of that model in the pooled regression above, although the qualitative results are robust to these controls.

Table 5 presents the results for these various subsamples. The magnitude and significance of substitute purchases vary across the subsamples, but the differences are not economically large and we cannot reject that any of these coefficients are equal to the value of 0.07 presented in the previous table. The effects of own purchases are only statistically significant among the subsamples of bonds, far-off-the-run securities, and shorter-maturity securities. However, we cannot reject that the own coefficients are the same for notes and bonds and for short and long maturities. The one instance in which we do find significant differences in the own coefficients is when we split according to security vintage—purchases of far off-the-

²⁴ That the market perceived important differences between the two types of securities was obvious immediately prior to the announcement of the LSAP program, when the yields on 5- to 10-year maturity bonds were as much as 50 basis points higher than those on comparable-maturity notes. See Gurkaynak and Wright (2010).

run securities have large and highly significant effects, while purchases of near on-the-run securities have none at all. Consistent with the significance of this difference, this specification also fits the data better in terms of adjusted R^2 than any of the others we tried. Proximity to on-the-run status is a proxy for liquidity, suggesting that a great deal of the stock effect of the LSAP program was due to the liquidity of the securities purchased. Overall, these results generally support hypotheses H1 – H3.

Counterfactual yield curve

To summarize the stock effects of the LSAP program, we construct a counterfactual yield curve using the best-fitting results presented in Table 5. In particular, by using the actual value of purchases of each security and its near substitutes, together with the coefficients for the appropriate sub-sample, we compute the estimated amount by which the price of that security changed as a result of LSAP purchases. Subtracting this value from the actual price at the end of the program gives the counterfactual price of each security that would have obtained if the LSAPs had not occurred. The corresponding yields are shown as the red squares in Figure 3, with red + and – signs indicating the 95 and 5 percent confidence bounds for each security. (These are calculated by finding the confidence interval around the fitted value of each security’s price and then transforming that value into a yield.) The blue dots in the figure show the actual yields on October 30, 2009. The difference between the red and blue marks represents the stock effects of the LSAP program on yields to maturity. For almost all securities, the counterfactual yields lie significantly above the actual yields. The largest effect on average appears to be at the very short end of the curve, but these responses are estimated with some noise and are driven by the very short duration of securities in that region. (The estimated effect on *prices* at the short end is very small.) Beyond maturities of about a year, the largest effects are around horizons of about 6 to 8 years and 11 to 14 years, consistent with the relatively high proportion of securities that were purchased in these sectors and the relatively high coefficients on off-the-run securities, which dominate those sectors. Averaging over all securities, we estimate that LSAP purchases shifted the level of yields by about 30 basis points, for an average elasticity of about 1 basis point for every \$10 billion bought.

To see these effects in another way, we obtain smooth zero-coupon yields by fitting prices to a Svensson yield curve, with securities weighted by their inverse duration as in Gurkaynak et al. (2007). The 5 percent and 95 percent confidence bands around the counterfactual prices are treated in the same way. These results are shown, together with the actual Svensson curve on October 30, 2009, in Figure 4. The difference between the solid red and solid blue lines represents the stock effects of the LSAP program on zero-coupon yields. From this picture, we can see that the effects were negative and statistically significant over all maturities of less than about 25 years. Indeed, over horizons of about 5 to 15 years, the shift in the yield curve is roughly uniform at 20 to 30 basis points. (The positive response of zero-coupon yields at the very

long end of the yield curve is consistent with the observation, from Figure 3, that yields to maturity at that horizon were essentially unaffected by the program.)

4.3 Flow Effects

In this section, employing a panel, we test whether Treasury LSAP purchases had effects on Treasury prices around the times when purchases occurred. Again, we define such responses as the “flow effects” of the program. Because the sectors of purchase operations were announced in advance and both the list of CUSIPs and sizes of these operations were fairly predictable, one might expect that examining yield changes as function of contemporaneous purchases would reveal no statistically significant responses. While this may well have been the case at the aggregate level, however, it need not be the case at the CUSIP level. Because the particular CUSIPs that were purchased and the amounts of these purchases were not known in advance to the market, yield differentials should have emerged on the days of purchases between those that are purchased and those that are not (assuming the demand for Treasury securities is sufficiently elastic). In addition, market illiquidity and other technical factors could cause yields to move in response to purchases, even if those purchases were perfectly predictable.²⁵ Our measures of flow effects will include both of these phenomena, as well as the possible dynamic described in hypothesis H4.

Specification

Our flow-effect regressions are the following special case of equation (5):

$$R_{it} = \alpha_i + \delta_t + \beta q_{it}^{own} + \sum_{j=1}^3 \gamma_j q_{ijt}^{sub} + \varepsilon_{it} \quad (9)$$

where α_i is a security-specific fixed effect, δ_t is a time dummy, and ε_{it} is an error term. The panel structure of our data, by allowing for both CUSIP-level fixed effects and daily time dummies, enables us to control for unobservable variables that might otherwise cloud the analysis, such as the level of and changes in short-term interest rates and other factors, like Treasury auctions, that could shift relative demand and supply in a small portion of the nominal coupon market. Most purchase operations were concentrated in maturity windows that were only 2 to 5 years wide, so that these “level” effects will be local to a relatively limited portion of the curve. (In other words, one factor should be enough to explain movements in a small sector of the yield curve, while if we allow for additional factors the risk of overfitting is highly increased).

Because the maturity sectors within which securities were purchased on any given day were announced in advance, we may expect that securities within those sectors might have reacted differently to the purchase operations than securities that were outside the purchased sectors. To examine this possibility, we split the sample into (1) observations of securities on days when those securities were within the

²⁵ Lou, Yan and Zhang (2010) demonstrate a similar effect around Treasury auctions.

announced purchase sectors (defined as “eligible” securities) and (2) observations of securities on days when purchases took place in different sectors (defined as “ineligible”). These subsets are mutually exclusive and exhaustive within the set of days on which purchases took place, though the same CUSIP can appear in both groups on different days.

Finally, it is possible that, because of settlement lags or other microstructure issues, the effects of purchases are not fully realized until the day after they occur. Moreover, we are interested in the persistence of the flow effects and would like to test for possible reversion in prices in the days following operations. To check for these possibilities, we also look at returns on the days after purchase operations by running regressions of the form

$$R_{it+1} = \alpha_i + \delta_t + \beta q_{it}^{own} + \sum_{j=1}^3 \gamma_j q_{ijt}^{sub} + \varepsilon_{it+1} \quad (10)$$

where R_{it+1} is the return in the day after the operation.

Results for Eligible Securities

We begin by analyzing the results for the eligible securities. In Table 6 we report the baseline results. Initial tests suggested that the coefficients were not stable for securities with very long remaining maturities, so we report a sample split at the midpoint maturity of 15 years. Focusing on the first column of the table, which pertains to eligible securities with remaining maturities of less than 15 years, the coefficient of 0.276 on own purchases implies that, on average, purchasing \$1 billion of Treasuries increased the price on the securities purchased by about 0.02 percent; for a representative 10-year security with a duration of 7 years, this translates into a yield decrease of about 0.3 basis points per billion dollars purchased. On the days when a security was eligible to be bought, purchases of its “near substitutes” had almost as large effects on its yield as did purchases of the security itself, pointing to a very high degree of substitutability among these securities.²⁶ However, the coefficients are somewhat smaller for mid-substitutes, consistent with hypothesis H2. Applying the aggregate coefficients to averages of the dependent variables, and multiplying by the average inverse modified duration, we find that the typical effect of each operation was on the order of -3.5 basis points for the sector being purchased. (Operations averaged about \$5 billion in size.) The second column of Table 6 shows that these results did not generally hold for issues with maturities greater than 15 years.

In the remainder of this section, we focus only on securities with less than 15 years to maturity, given that that is where the great majority purchases occurred and where most of the statistical significance seems to be. Within this sub-sample, Table 7 splits the data into various subsamples to examine the stability

²⁶ Because it was rare to purchase securities with maturities more than six years apart in the same operation, the far-substitute variable is omitted from this regression.

of the coefficients. First, we split the sample into purchases that occurred during the first half of the LSAP program (March 25 – July 6) and those that occurred during the second half (July 7 – October 29). As noted earlier, liquidity in the Treasury market was substantially better during the second half of the sample. Thus, if the price responses to LSAP purchases were due to impediments to market clearing and price discovery resulting from poor market functioning, we would expect the results to be substantially weaker in the second sub-sample. The first two columns of the table show that there is no evidence of this among securities that were eligible for purchase—the coefficients are nearly identical for the two sub-periods and are very close to the pooled results reported in the first column of Table 6.

The middle columns of Table 7 split the sample by security type, that is, notes versus bonds. Again, the results for the subsamples are generally similar to each other and to the results presented in Table 6 in terms of sign, magnitude, and significance. Similarly, the last two columns, which split the sample into securities more than five and less than six issues off the run, show no major differences. The modest exception is that the samples of bonds and far-off-the-run securities show somewhat smaller effects of substitute purchases. Since these are the least liquid segments of the Treasury market, this finding is consistent with hypothesis H3.

Results for Ineligible Securities

Table 8 displays results for securities that were ineligible for purchase, comparably to Table 6. In the aggregate, these responses display little economic or statistical significance. However, a further split of this sample reveals a more interesting pattern. Namely, as shown in Table 9, the coefficients on all of the substitute purchases are negative in the first half of the sample and positive in the second half. (As in the previous subsection, we focus here on the less-than-15-years sector of the yield curve, where most purchases took place.) During the second half of the sample, the coefficients on near- and mid-substitute purchases are close to those for the eligible sample, as we would expect given that there was generally little qualitative difference between eligible and ineligible securities. Thus, the first half of the sample for the ineligible securities is the puzzling piece of the data. It seems possible that the differences between eligible and ineligible securities during this time may have had to do with the liquidity impairments in the market noted above, although the exact mechanism that would result in this pattern is unclear. Another possible explanation is that dealers initially anticipated being able to sell more to the Fed than they actually were able to sell and thus sold securities (including securities that had not been eligible) after LSAP operations in order to maintain a portfolio target. Such an effect would likely have dissipated by the second half of the sample, as participants learned the pattern of the Desk's operations.

Table 10 shows that the basic patterns described above for the eligible and ineligible securities in the first and second halves of the sample do not generally depend on the liquidity characteristics of the securities

considered, as proxied by the split into notes and bonds. In particular, the coefficients on purchases on eligible securities are almost always positive and significant, with fairly consistent magnitudes,²⁷ while the negative coefficients for ineligible securities in the first half of the sample appear irrespective of security type.

Results for the Day After Purchases

Table 11 turns to the question of what happened on the days after LSAP operations took place. For comparison, the sample breakdown and independent variables are the same as those used in Table 10, but now the dependent variable is the security return on day $t+1$. Consider first the sample of eligible securities, presented in the left-hand sets of columns. For eligible note securities, prices almost uniformly reversed the increases they experienced on the days of purchases—the coefficients are of roughly similar magnitudes to those reported in the top panel of Table 10, but they are all negative (although they are not individually significant in the second half of the sample). This suggests that flow effects among notes were short-lived. For eligible bonds, on the other hand, prices actually increased further on the days after purchases. Indeed, tests using price changes on subsequent days (not shown) suggest that the effects of LSAP purchases on these bonds may have never fully reversed, perhaps contributing to the large stock effect observed in this category.²⁸

These results are broadly consistent with hypothesis H4—that even anticipated purchases could have significant effects on prices. A brief spike and retreat in prices, such as occurred among notes, can be explained by settlement, clearing, and rebalancing frictions that do not necessarily constitute material departures from market efficiency. But a *persistent* increase in prices following a purchase that was announced in advance would seem to call for a more substantial explanation, such as preferred habitat. That this pattern is evident among the less-traded securities further supports the idea that such a mechanism may be at work.

Finally, turning to the day-after results for ineligible securities, the coefficients for notes are similar to those in the eligible sample, again suggesting good substitutability (as we would expect) across these groups. The sign and significance of the coefficients on ineligible bonds do not show a clear pattern, but the coefficient magnitudes are small compared to most of the other samples we have reported. In general (also taking into account Table 10), it does not appear that the prices of ineligible bonds increase with their

²⁷ The sample of eligible purchases of bonds in the second half of the period contains only 33 observations and consequently the coefficients are not statistically significant. However, we also cannot reject that they are equal to their full-sample counterparts.

²⁸ Our finding that purchases had significant effects on prices the days after they occurred raises the possibility that our baseline regressions are misspecified (since they do not control for these effects). However, when we reestimated those equations controlling for lagged purchases and lagged price changes, the coefficient estimates were essentially unchanged.

eligible counterparts following purchases. This is somewhat puzzling but could again be consistent with the relatively weak liquidity for these securities.

Robustness to Error Correlation

Tables 12 and 13 present results for the baseline samples and key subsamples using clustered standard errors. Because it seems plausible that the regression errors are correlated across maturity, we allow for clustering within one-year maturity buckets for each security. This adjustment does not alter any of our results, either for the baseline breakdown or for the subsamples—indeed, in some cases the clustered standard errors are smaller than in the baseline case (suggesting negative correlation within clusters). We also clustered by security type (not shown) and did not observe any notable differences with the results reported above.

5. Conclusion

In this paper, we have used CUSIP-level data to estimate the flow and stock effects of the Federal Reserve’s 2009 program to purchase nearly \$300 billion of nominal Treasury coupon securities. We find that both types of effect were statistically and economically significant. Specifically, we estimate that the average purchase operation temporarily reduced yields by about 3.5 basis points and that the program as a whole shifted the yield curve down by up to 30 basis points, with both effects concentrated at short to medium maturities. It thus seems likely that the Treasury LSAP program met the Federal Reserve’s objectives of improving Treasury market liquidity and contributing to a reduction in the cost of credit.


More broadly, our results provide support for preferred-habitat and portfolio-balance theories of the term structure. Consistent with such theories, we find that (1) withdrawals of Treasury supply decrease yields by an economically meaningful amount; (2) these decreases are generally biggest for the specific securities being bought and for securities of similar maturities but smaller for securities with much different maturities; (3) particularly for stock effects, the discrepancies between own purchases and substitute purchases are larger among less-liquid securities (off-the-run bonds); and (4) also among off-the-run bonds, the flow effects are persistent, suggesting that they are not just due to short-run rebalancing or microstructure-related distortions. Our study is the first to specifically treat the effects of substitute cross-elasticities, but the overall magnitudes of our stock-effect estimates (combining own and cross effects) are roughly comparable to what would be implied by Treasury price elasticities found in some other research, such as Kuttner (2006) and D’Amico et al. (2011). As far as we are aware, no other study has estimated flow effects as we have defined them. It is perhaps surprising that these effects should be so large in most subsamples, given that most details of the purchases were announced in advance. There is certainly room for additional work to

understand whether similar effects hold in other markets and in other periods and, if so, exactly what mechanisms are behind them.


Finally, we caution that the environment that produced our data involved exceptional conditions in the Treasury market that could make extrapolation of these results to other situations problematic. We provided some evidence that our main flow-effect results did not depend greatly on the market conditions over the course of 2009. But we are not able to conduct a similar test for stock effects, and we still cannot rule out that other special circumstances during that period might have helped to drive the results—for example, Gurkaynak and Wright (2010) conjecture that increased risk aversion during the financial crisis could have reduced arbitrage activity and led to greater market segmentation. Testing how stock and flow effects differ across risk and liquidity regimes is an interesting area for future research and could further illuminate the factors that influence pricing and functioning in the Treasury market.

Figure 1. Maturity Distribution of Nominal Treasury LSAP Purchases

Billions of Dollars




Percentage of Sector



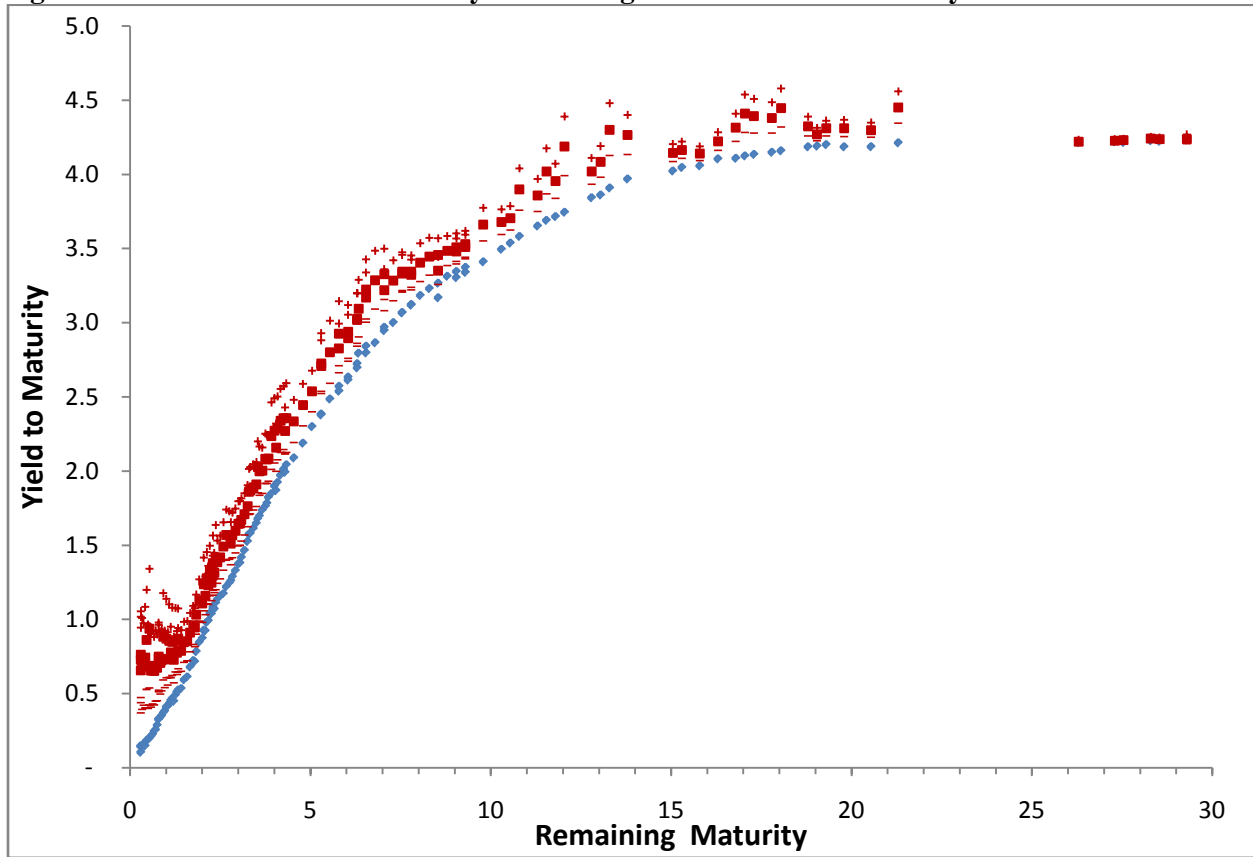
Note: Denominator is par value of outstanding securities in each sector as of July 13, 2009. No securities were outstanding with remaining maturities between 22 and 26 years.

Figure 2. Selected Treasury Yields during the Treasury LSAP Program




Notes: Yields are constant-maturity. Highlighted region shows the period (March 18 – October 30, 2009) during which the Treasury LSAP program was in effect.

Figure 3. Stock Effect of the Treasury LSAP Program on Yields to Maturity



Notes: The blue dots are nominal yields on each of the 148 securities outstanding in our cross section as of October 30, 2009, based on the 148 securities in our cross section. The red squares are the counterfactual yield point estimates on that day, using the same set of Treasury prices but with the estimated effects of LSAP purchases removed according to the coefficients in the last column of Table 5. The red + and - signs are the 95 and 5 percent confidence bands around the counterfactual yields.

Figure 4. Stock Effect of the Treasury LSAP Program on Zero-Coupon Yields



Notes: The blue line is the nominal yield curve as of October 30, 2009, based on the 148 securities in our cross section and constructed using the Svensson (1994) function. The solid red line is the counterfactual yield curve on that day, using the same set of Treasury prices but with the estimated effects of LSAP purchases removed according to the coefficients in the last column of Table 5. Dotted lines show Svensson curves based on the 95 and 5 percent confidence bands around the counterfactual prices.

Table 1. Characteristics of Nominal Treasury LSAP Purchases

	Average of LSAP purchases	Average of all outstanding coupon securities
Remaining maturity	6.5 years	5.7 years
Coupon	3.7%	3.8%
Yield	2.4%	1.9%
Time since issued	4.0 years	3.9 years
% On-the-run	29.0%	4.9%
% Notes	79.5%	82.8%

Note: All figures are dollar-weighted.

Table 2. Indicators of Liquidity in the Nominal Treasury Market

	Daily market volume	10-year on-the-run premium	Fails to deliver	Average absolute fitting errors of Svensson curve
March 25-July 6 2009	\$100 bil	39 bp	\$73 bil	6.4 bp
July 7-October 29 2009	\$120 bil	29 bp	\$15 bil	3.3 bp

Notes: The table reports averages of daily values. Volume data come from Bloomberg, and fails-to-deliver data come from the FR 2004 reports. The on-the-run premium is the difference between the yield on the on-the-run 10-year note and 10-year value from a Svensson curve fit to off-the-run securities.

Table 3. LSAP Purchases as a Function of pre-LSAP Information

Intercept	0.051*** (0.015)
Remaining maturity	0.015*** (0.003)
Remaining maturity squared	-0.0006*** (0.00009)
Svensson fitting error	0.208*** (0.038)
% of issue held by Fed	-0.271*** (0.083)
On-the-run dummy	0.091*** (0.024)
< 2-years dummy	-0.044*** (0.014)
# Obs	148
Adjusted R ²	0.431

*Notes: The dependent variable is the total amount of each security purchased under the LSAP program as a percentage of the total amount of that security outstanding. All independent variables are as of March 17, 2009. Standard errors in parentheses. *** indicates statistical significance at the 1 percent level.*

Table 4. Stock Effects (IV)—Pooled

	Baseline	Controlling for initial prices
Own Purchases (IV)	2.17*** (0.43)	0.61*** (0.21)
Purchases of near substitutes (IV) (maturity w/in 2 yrs of own)	0.13** (0.07)	0.07** (0.03)
Remaining maturity	-0.003** (0.001)	-0.0005 (0.0007)
Remaining maturity squared	0.00003 (0.00004)	-0.00004 (0.00002)
Log initial price	---	0.019*** (0.007)
Initial % fitting error	---	-0.58*** (0.07)
Intercept	0.011*** (0.002)	-0.08** (0.03)
# Obs	148	148
Adj. R ²	0.695	0.857

Notes: The dependent variable is the cumulative percentage holding return from March 17 to October 30, 2009. Regressions are 2SLS with March 17-dated variables used as instruments. All purchase variables are normalized by the total quantity of near substitutes outstanding. Asymptotic standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (), 5 percent (**), and 1 percent (***) levels.*

Table 5. Stock Effects (IV)—Subsamples

	Notes	Bonds	> 15 years	< 15 years	Near on-the-run	Far off-the-run
Own Purchases (IV)	0.35 (0.32)	0.68*** (0.24)	0.38 (0.31)	0.66*** (0.21)	-0.03 (0.34)	1.23*** (0.34)
Purchases of near substitutes (IV) (maturity w/in 2 yrs of own)	0.08** (0.03)	0.09* (0.05)	0.04 (0.07)	0.04 (0.03)	0.17*** (0.06)	0.11*** (0.04)
Remaining maturity	-0.0006 (0.0007)		-0.00007 (0.0008)		-0.0022** (0.0011)	
Remaining maturity squared	-0.00003 (0.00002)		-0.00005* (0.00003)		0.00002 (0.00004)	
Log initial price	0.015 (0.0098)		0.018*** (0.007)		0.025** (0.0099)	
Initial % fitting error	-0.56*** (0.07)		-0.57*** (0.06)		-0.52*** (0.08)	
Intercept	-0.06 (0.04)		-0.07** (0.03)		-0.11** (0.05)	
# Obs	148		148		148	
Adj. R ²	0.903		0.904		0.914	

Notes: The dependent variable is the cumulative percentage holding return from March 17 to October 30, 2009. Regressions are 2SLS with March 17-dated variables used as instruments. All purchase variables are normalized by the total quantity of near substitutes outstanding. Asymptotic standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) levels.

**Table 6. Flow Effects on Day of Purchase
(eligible securities)**

	< 15y to maturity	>15y to maturity
Own Purchases	0.276*** (0.053)	-0.106 (0.098)
Purchases of:		
Near substitutes (maturity w/in 2 yrs of own)	0.240*** (0.048)	-0.124*** (0.044)
Mid-substitutes (maturity 2 to 6 years away)	0.170*** (0.045)	-0.050* (0.026)
# Obs.	923	145
# CUSIPS	146	23
Adj. R ²	0.976	0.985

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only securities that were eligible to be purchased in a given operation are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (), 5 percent (**), and 1 percent (***) levels.*

**Table 7. Flow Effects on Day of Purchase, by Subsamples
(eligible securities with remaining maturity < 15 years)**

	Mar 25 – Jul 6	Jul 7 – Oct. 29	Notes	Bonds	Near on-the- run	Far off-the- run
Own Purchases	0.3442*** (0.094)	0.2975*** (0.089)	0.2669*** (0.068)	0.2498*** (0.090)	0.2318** (0.107)	0.2488*** (0.065)
Purchases of:						
Near substitutes (maturity w/in 2 yrs of own)	0.2863*** (0.086)	0.3038*** (0.083)	0.2503*** (0.062)	0.1694** (0.083)	0.2435** (0.105)	0.1584*** (0.057)
Mid-substitutes (maturity 2 to 6 years away)	0.1989*** (0.082)	0.2037** (0.073)	0.2088** (0.055)	0.0929 (0.080)	0.2501*** (0.092)	0.0744 (0.055)
# Obs.	563	360	769	154	249	674
# CUSIPS	131	121	123	23	53	114
Adj. R ²	0.974	0.975	0.976	0.986	0.986	0.977

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (), 5 percent (**), and 1 percent (***) levels.*

**Table 8. Flow Effects on Day of Purchase
(Ineligible securities)**

	<15y to maturity	>15y to maturity
Purchases of:		
Near substitutes (maturity w/in 2 yrs of own)	0.0665*** (0.018)	-0.0268 (0.053)
Mid-substitutes (maturity 2 to 6 years away)	0.0047 (0.0099)	-0.007 (0.021)
Far substitutes (maturity 6 to 14 years away)	-0.0238** (0.008)	0.0021 (0.003)
# Obs.	8008	1104
# CUSIPS	181	23
Adj. R ²	0.52	0.96

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only securities that were not eligible to be purchased on days when operations occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) levels.

**Table 9. Flow Effects on Day of Purchase, by Sub-Period
(Ineligible securities, remaining maturity < 15 years)**

	Mar 25 – Jul 6	Jul 7 – Oct. 29
Purchases of:		
Near substitutes (maturity w/in 2 yrs of own)	-0.127*** (0.025)	0.384*** (0.031)
Mid-substitutes (maturity 2 to 6 years away)	-0.143*** (0.015)	0.202*** (0.017)
Far substitutes (maturity 6 to 14 years away)	-0.153*** (0.014)	0.093*** (0.011)
# Obs.	4529	3479
# CUSIPS	167	172
Adj. R ²	0.51	0.57

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) levels.

**Table 10. Flow Effects on Day of Purchase, by Security Type and Sub-Period
(remaining maturity < 15 years)**

Notes

	Eligible Securities		Ineligible Securities	
	Mar 25 – Jul 6	Jul 7 – Oct. 29	Mar 25 – Jul 6	Jul 7 – Oct. 29
Own Purchases	0.445*** (0.147)	0.245** (0.122)	---	---
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	0.396*** (0.148)	0.314*** (0.107)	-0.213*** (0.041)	0.486*** (0.049)
Mid-substitutes (maturity 2 to 6 years away)	0.346*** (0.144)	0.260*** (0.086)	-0.244*** (0.037)	0.266*** (0.038)
Far substitutes (maturity 6 to 14 years away)	---	---	-0.061*** (0.040)	0.063*** (0.039)
# Obs.	442	327	3891	2960
# CUSIPS	108	111	144	149
Adj. R ²	0.979	0.972	0.477	0.548

Bonds

	Eligible Securities		Ineligible Securities	
	Mar 25 – Jul 6	Jul 7 – Oct. 29	Mar 25 – Jul 6	Jul 7 – Oct. 29
Own Purchases	0.381*** (0.159)	0.108 (0.129)	---	---
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	0.270*** (0.109)	0.011 (0.134)	-0.208*** (0.049)	0.084* (0.046)
Mid-substitutes (maturity 2 to 6 years away)	0.178* (0.104)	-0.161 (0.134)	-0.077*** (0.019)	0.074*** (0.020)
Far substitutes (maturity 6 to 14 years away)	---	---	-0.077*** (0.014)	0.041*** (0.012)
# Obs.	121	33	638	519
# CUSIPS	23	10	23	23
Adj. R ²	0.978	0.994	0.946	0.939

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (), 5 percent (**), and 1 percent (***) levels.*

**Table 11. Flow Effects on Day after Purchase, by Security Type and Sub-Period
(remaining maturity < 15 years)**

Notes

	Eligible Securities		Ineligible Securities	
	Mar 25 – Jul 6	Jul 7 – Oct. 29	Mar 25 – Jul 6	Jul 7 – Oct. 29
Own Purchases	-0.379*** (0.143)	-0.145 (0.116)	---	---
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	-0.478*** (0.145)	-0.152 (0.108)	-0.464*** (0.039)	-0.135*** (0.049)
Mid-substitutes (maturity 2 to 6 years away)	-0.620*** (0.139)	-0.106 (0.087)	-0.436*** (0.035)	-0.009 (0.039)
Far substitutes (maturity 6 to 14 years away)	---	---	-0.308*** (0.039)	0.134*** (0.038)
# Obs.	442	327	3886	2957
# CUSIPS	108	111	144	149
Adj. R ²	0.986	0.974	0.573	0.507

Bonds

	Eligible Securities		Ineligible Securities	
	Mar 25 – Jul 6	Jul 7 – Oct. 29	Mar 25 – Jul 6	Jul 7 – Oct. 29
Own Purchases	0.377*** (0.145)	0.58*** (0.117)	---	---
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	0.157 (0.098)	0.556*** (0.119)	-0.047 (0.046)	0.086* (0.047)
Mid-substitutes (maturity 2 to 6 years away)	0.111 (0.082)	0.472*** (0.125)	-0.057*** (0.018)	-0.046*** (0.020)
Far substitutes (maturity 6 to 14 years away)	---	---	-0.049*** (0.014)	0.015*** (0.013)
# Obs.	121	33	638	519
# CUSIPS	23	10	23	23
Adj. R ²	0.99	0.99	0.94	0.93

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (), 5 percent (**), and 1 percent (***) levels.*

Table 12. Flow Effects on Day of Purchase, All Securities
Clustered Standard Errors

	Eligible Securities		Ineligible Securities	
	<15y to maturity	>15y to maturity	<15y to maturity	>15y to maturity
Own Purchases	0.2763*** (0.044)	-0.1063 (0.089)	---	---
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	0.2403*** (0.038)	-0.1238*** (0.029)	0.0665*** (0.029)	-0.0268 (0.017)
Mid-substitutes (maturity 2 to 6 years away)	0.1700*** (0.053)	-0.0501** (0.014)	0.0047 (0.016)	-0.007 (0.037)
Far substitutes (maturity 6 to 14 years away)	---	---	-0.0238** (0.009)	0.0021 (0.003)
# Obs.	923	145	8008	1104
# CUSIPS	146	23	181	23
Adj. R ²	0.976	0.985	0.519	0.968

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) levels.

Table 13. Flow Effects on Day of Purchase, by Security Type and Vintage
(Eligible securities with remaining maturity < 15 years)

Clustered Standard Errors

	Bonds	Notes	Near on-the-run	Far off-the-run
Own Purchases	0.2498** (0.06)	0.2669*** (0.058)	0.2318** (0.098)	0.2488*** (0.051)
Purchases of:				
Near substitutes (maturity w/in 2 yrs of own)	0.1694** (0.03)	0.2503*** (0.051)	0.2435** (0.073)	0.1584*** (0.029)
Mid-substitutes (maturity 2 to 6 years away)	0.0929 (0.035)	0.2088** (0.059)	0.2501*** (0.059)	0.0744 (0.026)
# Obs.	154	769	249	674
# CUSIPS	23	123	53	114
Adj. R ²	0.986	0.976	0.986	0.977

Notes: The dependent variable is the daily percentage price change in each outstanding CUSIP. Only days when LSAP purchases occurred are included. Fixed effects and daily time dummies not shown. Standard errors in parentheses. Asterisks indicate statistical significance at the 10 percent (*), 5 percent (**), and 1 percent (***) levels.

References

- Andres, J.; Lopez-Salido, J. D.; and Nelson, E., 2004. "Tobin's Imperfect Asset Substitution in Optimizing General Equilibrium." *Journal of Money, Credit, and Banking* 36(4): 665-90.
- Bernanke, B. S., 2009. "The Federal Reserve's Balance Sheet." Speech delivered at the Federal Reserve Bank of Richmond 2009 Credit Markets Symposium, Charlotte, North Carolina, April 3.
- Bernanke, B. S.; Reinhart, V. R.; and Sack, B. P., 2004. "Monetary Policy Alternatives at the Zero Bound: An Empirical Assessment." *Brookings Papers on Economic Activity* (2): 1 – 78.
- Bull, A., 2009. "FED Focus—Fed seen extending, not increasing Treasury buys." Thomson Reuters, June 16.
- Cochrane, J. and Piazzesi, M., 2005. "Bond Risk Premia." *American Economic Review* 95(1): 138-60.
- Culbertson, J. M., 1957. "The Term Structure of Interest Rates." *Quarterly Journal of Economics* 71(4): 485-517.
- Doh, T., 2010. "The Efficacy of Large-Scale Asset Purchases at the Zero Bound." FRB Kansas City *Economic Review* (Q2): 5-34.
- Engen, E. and Hubbard, R. G., 2005. "Federal Government Debt and Interest Rates." In Gertler and Rogoff (eds.), *NBER Macroeconomics Annual* 2004. Cambridge, MA: The MIT Press.
- Federal Reserve, 2008. Minutes of the Federal Open Market Committee, December 15-16, 2008. Washington, DC. <http://www.federalreserve.gov/monetarypolicy/files/fomcminutes20081216.pdf>
- Federal Reserve, 2009. Minutes of the Federal Open Market Committee, January 27-28, 2009. Washington, DC. <http://www.federalreserve.gov/monetarypolicy/files/fomcminutes20090128.pdf>
- Gagnon, J.; Raskin, M.; Remache, J.; and Sack, B., 2010. "Large-Scale Asset Purchases by the Federal Reserve: Did They Work?" FRB New York Staff Report 441.
- Greenwood, R. and Vayanos, D., 2008. "Bond Supply and Excess Bond Returns." NBER Working Paper 13806.
- Greenwood, R. and Vayanos, D., 2010. "Price Pressure in the Government Bond Market." *American Economic Review* 100(2): 585-90.
- Gurkyanak, R.; Sack, B.; and Wright, J., 2007. "The U.S. Treasury Yield Curve: 1961 to the Present." *Journal of Monetary Economics* 54(8): 2291 - 304.
- Gurkyanak, R. and Wright, J., 2010. "Macroeconomics and the Term Structure." Mimeo, April.
- Hakim, S. R. and Rashidian, M., 2000. "Testing for Segmentation in the Term Structure: Operation Twist Revisited." *Quarterly Journal of Business and Economics* 39(1): 3-21.
- Hamilton, J. and Wu, J., 2010. "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment." Mimeo, August.
- Han, B.; Longstaff, F. A.; and Merrill, C., 2007. "The U.S. Treasury Buyback Auctions: The Cost of Retiring Illiquid Bonds." *Journal of Finance* 62(6): 2673-93.
- Holland, T. E., 1969. "'Operation Twist' and the Movement of Interest Rates and Related Economic Time Series." *International Economic Review* 10(3): 260-5.
- Krishnamurthy, A. and Vissing-Jorgensen, A., 2007. "The Demand for Treasury Debt." NBER Working Paper 12881.
- Kuttner, K., 2006. "Can Central Banks Target Bond Prices?" NBER Working Paper 12454.

- Longstaff, F. A., 2004. "The Flight-to-Liquidity Premium in U.S. Treasury Bond Prices." *Journal of Business* 77(3): 511-26.
- Lou, D.; Hongjun Y.; and Zhang, J., 2010. "Anticipated and Repeated Shocks in Liquid Markets." Working paper, mimeo.
- Modigliani, F. and Sutch, R., 1966. "Innovations in Interest Rate Policy." *American Economic Review* 56(1): 178-197.
- Neely, C. J., 2010. "The Large Scale Asset Purchases Had Large International Effects." FRB St. Louis Working Paper 2010-018A.
- Ross, M., 1966. "'Operation Twist': A Mistaken Policy?" *Journal of Political Economy* 74(2): 195-99.
- Stroebel, J. C. and J. B. Taylor, 2009. "Estimated Impact of the Fed's Mortgage-Backed Securities Purchase Program." NBER Working Paper 15626.
- Svensson, L. E. O., 1994. "Estimating and Interpreting Forward Rates: Sweden 1992-4," NBER Working Paper 4871.
- Swanson, E. T., 2011. "Let's Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2." Federal Reserve Bank of San Francisco working paper.
- Tarhan, V., 1995. "Does the Federal Reserve Affect Asset Prices?" *Journal of Economic Dynamics and Control* 19: 1199-222.
- Thornton D. L., 2009. "The Effect of the Fed's Purchase of Long-Term Treasuries on the Yield Curve," FRB St. Louis *Economic Synopses* (25).
- Vayanos, D. and Vila, J-L., 2009. "A Preferred-Habitat Model of the Term Structure of Interest Rates." NBER Working Paper 15487.
- Yellen, J., 2009. Panel discussion for the Federal Reserve Board/*Journal of Money, Credit, and Banking* conference on "Financial Markets and Monetary Policy." Washington, DC, June 5.