The Term-Premium Effect of the Federal Reserve's Asset Purchase Announcements[a]

Stefania D'Amico, Canlin Li, John S. Sears, and Michael E. Cahill[b]

First draft November 9, 2012; This draft May 5, 2014

ABSTRACT

The FOMC's asset purchase announcements and the releases of the operational details provide a sequence of events shedding light on the relative importance of the duration-risk and local-supply channels for the transmission of the Fed's balance sheet policy shocks to Treasury yields. Using intraday security-level price data, we first document local-supply and duration-risk effects. Then, we measure local-supply and duration-risk *surprises* to quantify these channels' price impact. The impact averages about -8 basis points per $100-billion surprise, is almost evenly split between the two channels, did not decline across programs, and is similar for purchases and sales, a result potentially relevant for the unwinding of these programs. These findings differ from previous studies that did not isolate the unexpected component of the Fed's asset purchase announcements, and have therefore increasingly underestimated their impact, as the Fed signaled its intentions to a greater extent ahead of each consecutive program.

Keywords: Yield curve, quantitative easing, LSAP, preferred habitat, limits of arbitrage

JEL Codes: E43, E52, E58, E65, G12, G14

[a] Correspondence to: Canlin Li, Division of Monetary Affairs, Federal Reserve Board, 20th and C Streets NW, Washington, DC 20551 USA. canlin.li@frb.gov; 202-452-2227. 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, and in particular Jim Clouse, Thomas King, Matthew Raskin and Min Wei, as well as Greg Duffee and Jonathan Wright.

[b] Federal Reserve Bank of Chicago, Federal Reserve Board, Stanford University, and Goldman Sachs, respectively.

1. Introduction

In the most recent literature on the different channels through which central banks' Treasury purchase programs affect interest rates (e.g., Krishnamurthy and Vissing-Jorgensen 2011, D'Amico et al. 2012, Benerjee et al. 2012, and Li and Wei 2013), two channels emerge as most likely: duration-risk and local-supply/scarcity effects.[1] The duration-risk channel is associated with the general notion of interest rate risk and predicts that these programs affect Treasury yields across the entire maturity spectrum, with larger effects in longer-duration securities. The local-supply channel derives its rationale from the preferred-habitat approach and predicts that the impact is larger for securities where the shortage of supply is bigger, independently of their durations. In theoretical term-structure models where the supply of Treasury bonds matters, as in Vayanos and Vila (2009), Greenwood and Vayanos (2014), and King (2013), both channels impact the term premium component of Treasury yields by altering the market prices of risk.

In a recent speech, former Federal Reserve Chairman Bernanke emphasized that large-scale asset purchase (LSAP) programs and forward rate guidance "affect longer-term rates through somewhat different channels".[2] In particular, he noted that while forward rate guidance affects longer-term rates primarily by changing investors' expectations of future short-term interest rates, LSAPs "most directly affect term premiums." However, he also stressed that LSAPs are a less familiar tool, as the FOMC has "much less experience with policies designed to operate on term premiums." It is the objective of this paper to shed some light on how LSAPs' shocks affect term premiums, to better understand the transmission mechanism of this policy tool. Therefore, our contribution is twofold. First, we provide a measure of LSAPs' shocks, and second, we illustrate through which channels these shocks propagate to the term-structure of interest rates.

[1] Bauer and Rudebusch (forthcoming) stress the importance of a third channel in their study of the first LSAP: the signaling channel; but it is not found to be a relevant driver of the announcements' impact after the first LSAP, and we will illustrate this point in detail in the robustness section.

[2] See Speech November 19, 2013. Other FOMC members (for example, Kohn, 2009, Willimas, 2011, and Yellen, 2011) made similar points in previous speeches.

It is important to identify the channels through which central banks' balance sheet policy accommodation works for several reasons. First, it is crucial for the calibration of these policies and their eventual unwinding. A better understanding of the channels can allow a central bank to maximize or minimize the impact of purchases and/or sales on Treasury yields, depending on the stance of monetary policy. Second, determining the relative importance of these supply channels across multiple purchase programs illustrates how their efficacy has evolved over time, which is instrumental to comprehend whether these channels are always operating or are exceptional mechanisms prompted by the disruption of normal market functioning. If these channels' impact did not decline, then they may be key factors in the determination of Treasury securities prices and should be included in the widely-used asset pricing models. Finally, if the term premium component of interest rates and the channels impacting this component are important for the transmission mechanism of unconventional monetary policy, then they should be incorporated in traditional macro/monetary economic models, where usually monetary policy operates only through the expectation component of interest rates.

Using a new dataset and a new identification procedure, we first attempt, for each purchase program conducted by the Federal Reserve (Fed) from 2009 to 2012, to disentangle the local-supply effect from the duration-risk effect and show that the location (in maturity space) of the supply shocks matters on top of the total duration risk removed from the market. Next, we construct measures of duration-risk and local-supply surprises and quantify the relative importance of these two channels for the transmission mechanism of this monetary policy tool. Finally, we show how the local-supply and duration-risk channels have continued to drive the impact of the Fed asset purchase announcements on Treasury yields across consecutive programs, during periods characterized by market conditions and risk sentiment very different from those prevailing at the time of the first LSAP.

In particular, the new dataset consists of intraday price quotes on all outstanding U.S. nominal Treasury securities from 2009 to 2012. On average, we have high-frequency information for about 200 CUSIPs at each point in time over a sample of almost 4 years. The identification procedure exploits not only the prices' reactions to the FOMC announcements regarding the total size of each asset purchase program, but also the reaction to the New York Fed Open Market Trading Desk (the Desk) releases of the programs' operational details, which provided the intended distribution of purchases and sales across maturity sectors. For each of these programs, we try to carefully control for the pre-announcement market expectations in order to estimate both the total stock surprise, that is, the unexpected component of the total size of the announced program, and the maturity distribution surprise, that is, the unexpected component of the purchases' allocation to each maturity sector. The availability of these two surprises allows us to measure the supply 'shock' local to each maturity sector and consequently the 'shock' to the aggregate duration risk measured by the unexpected change in ten-year equivalents.[3] Clearly, the quality of these shocks' measurement depends on the ability to control for market participants' expectations about the size and maturity distribution of each asset purchase program. To this purpose, we use the Desk Primary Dealer Survey (PDS) results compiled by the New York Fed before each FOMC announcement, supplemented by market commentaries from the same primary dealers.

Our empirical results suggest that the duration-risk and local-supply 'shocks' together can explain most of the variation in the reaction of Treasury yields to the Fed purchase program announcements and each separately has about 25 to 50 percent explanatory power. In terms of impact on the 10-year nominal Treasury yield, we find that the average effect of each channel across programs is about -4 basis points per $100 billion of unexpected purchases. This implies that the duration-risk and local-supply channels have similar importance in the transmission mechanism of the Fed asset purchases to the term-structure of Treasury yields. Finally, we find that, once the pre-announcement market expectations are carefully controlled for, there does not appear to be evidence that the effects of these two channels have declined across consecutive programs, suggesting that they may be key factors in the determination of Treasury securities prices rather than exceptional mechanisms triggered by market disruption or extremely high risk aversion. These findings differ from previous studies that did not isolate the unexpected component

[3] The ten-year equivalents are the amount of ten-year Treasury notes that an investor would have to buy to replicate the same duration of a specific portfolio under consideration.

of the Fed's asset purchase announcements, and have therefore increasingly underestimated their impact, as the Fed signaled its intentions to a greater extent ahead of each program that came after the first LSAP.

A number of papers estimate the impact of central banks' quantitative easing (QE) programs on the term-structure of interest rates, but none focus on constructing the unexpected component of this policy in order to better isolate its impact on bond yields. This requires estimating how changes in securities' available supply translate into changes in yields, controlling for expectations about future monetary policy and macroeconomic conditions, which poses some empirical challenges. Event studies have been so far the most common approach (e.g., Gagnon et al. 2011, Krishnamurthy and Vissing-Jorgensen 2011, Joyce et al. 2011), where the policy impact is computed as the sum of yield changes around few specific QE-related events. This approach becomes increasingly more problematic after the introduction of the first QE program, as the central bank signaled its intentions well before the actual policy announcements and strengthened the conditionality of the QE program to macroeconomic variables. As a consequence, the identification of the relevant events for the event-study becomes extremely hard as, for example, any macroeconomic data releases or any policy makers' speech and interview can alter market participants' expectations. Controlling for pre-announcement market expectations using the PDS helps avoiding these limitations.[4]

Further, while a number of studies have previously examined the effects of LSAP announcements on a few constant-maturity U.S. Treasury yields (Gagnon et al. 2011, Krishnamurthy and Vissing-Jorgensen 2011, Neely 2010, Rosa 2012, and Swanson 2011), none has employed data at the individual security level and none has exploited the Desk releases of the operational details.[5] Observing how the price reactions to these announcements differ across duration/maturity and liquidity characteristics of the

[4] Alternative strategies consist in either estimating supply effects in the pre-QE sample while controlling for macro-economic factors (e.g., Hamilton and Wu 2012, Li and Wei 2013, and D'Amico et al. 2012) and then infer the estimated effects for the QE period assuming stability of the parameter estimates, or exploiting the cross-sectional dimension of Treasury security data and control for market-wide factors (e.g., D'Amico and King 2013, Joyce and Tong 2012).

[5] D'Amico and King (2013) is the first study to use CUSIP-level data to identify the local supply effects, but the empirical analysis is not based on intraday changes around FOMC announcements and is focused only on the first Treasury asset purchase program.

Treasury securities is essential to the identification of the channels. Even more crucial to their identification is the use of new information not only about the total size of the program, but also about the distribution of these purchases across maturity sectors. The only paper that uses disaggregated five-minute interval data on individual securities to analyze QE announcements is Joyce and Tong (2012), which examines the Bank of England's asset purchases. However, that paper does not focus on the reaction to the operational details and therefore cannot separately identify the unexpected component of the total size and maturity distribution of each QE program. With the exception of the case study in D'Amico et al. (2012), none of the previous works has focused on the yield reaction to the maturity distribution of purchases. However, while these authors use only a graphical analysis of a single event and a few securities, we quantify the price reactions across all Treasury securities outstanding to multiple events over a period of about four years, which is essential to extrapolate how the price reaction has evolved over time and allows a more precise identification. The closest relation to our study is the concurrent paper by Benerjee et al. (2012), which studies how the announced operational changes to the Bank of England QE program affected gilt yields.

The following section discusses in detail each announcement employed in the event studies and describes the intraday security-level yield reaction to each announcement. Section 3 provides the basic theoretical framework that motivates the construction of our surprise measures and the chosen empirical specification. Section 4 details the computation of the local-supply and duration-risk shocks. Section 5 presents the estimation of the duration-risk and local-supply effects. Section 6 analyzes the robustness of our results. Section 7 concludes.

2. Description of the announcements and corresponding security-level event studies

In this section we outline the sequence of events that we employ in our empirical analysis. For each event, we provide a description of the timing and content of both the FOMC announcement and accompanying Desk technical note detailing the program operations. Further, using high-frequency data on all Treasury securities outstanding, we compute and plot the reactions to these announcements.

Our dataset consists of intra-day CUSIP-level prices obtained from the Thomson Reuters Tick History database at minutely frequency. The availability of detailed high-frequency information allows us to examine the price reaction to each FOMC and Desk announcement separately within the same day, which is important for the identification of the two channels. In addition, the direction and magnitude of the high-frequency reactions across all the outstanding CUSIPs makes it possible to analyze how the price response varies across the entire duration/maturity spectrum, which is also crucial to the identification of the channels.

The first LSAP (LSAP1) announcement

On March 18, 2009 at 2:15 p.m. the FOMC announced its decision to bring its maximum purchases of agency MBS to $1.25 trillion as well as of agency debt to $200 billion; and to purchase up to $300 billion of longer-term Treasury securities over the subsequent six months. According to market commentaries, those policy actions were more aggressive than expected both in size and scope. In particular, the Treasury purchases were largely unexpected.[6] Without prior notice, the Desk announced at 2:44 p.m. that the Treasury purchases would be concentrated in the 2- to 10-year sector of the nominal Treasury curve, and that it would purchase both nominal and inflation-indexed Treasury securities. Changes over this half-hour interval (between 2:15 and 2:44 p.m.) in market expectations about the maturity distribution of purchases could have affected the Treasury yields' behavior in a way that can be revealing about the respective roles of the local-supply and duration channels. This is because, based on the 2:15 p.m. announcement, all the securities with longer-term maturities (that is, usually with maturities beyond two years) may have been perceived as equally likely candidates for purchases by the Federal Reserve. However, following the 2:44 p.m. Desk statement, investors should have assigned much smaller probability to the purchase of securities with remaining maturities above 10 years. This change in the perceived maturity distribution of purchases would affect both the expected average duration of future

[6] The March 2009 Primary Dealer Survey conducted by the Desk one week before the FOMC meeting indicated that the average probability associated to the Treasury purchase program being announced was 49 percent and not necessarily at the upcoming March meeting.

purchases and the future supply available to private investors in each maturity bucket. That is, because of their concentration in the 2- to 10-year maturity sector, the assets purchased would have a smaller average duration than purchases more heavily weighted towards the 10- to 30-year sector, and the securities in the excluded sector relative to those in the 2- to 10-year sector (as a percentage of the outstanding stock of Treasuries) would also become relatively less scarce. Therefore, following the Desk announcement, in the 10- to 30-year sector, the smaller-than-expected scarcity effect moves prices in the same direction as the smaller-then-expected duration effect (both should drive prices down and yields up); while, in the 2- to 10-year sector, the larger-than-expected scarcity effect should move prices up, and the smaller-than-expected duration effect should drive them down. This implies that if both channels are operating and have similar importance, the second announcement by the Desk should have no material impact on the prices of Treasury securities included in the purchase sector (2- to 10-year maturities) but should exert a potentially sizable negative impact on the prices of securities outside the purchase sector (with maturities beyond ten years). This is clearly illustrated in Figure 1.

In Figure 1, the red dots show the yield changes of all outstanding nominal Treasury securities from 2:00 to 2:30 p.m. as a function of their modified duration on March 18, 2009. Since this time-window includes only the reaction to the FOMC announcement about the total size of the program, we would in general expect those yield changes to be monotonically decreasing in duration, as bonds with longer duration generally bear larger exposure to interest rate risk and as such can experience a larger change in their risk premiums. The yield change reaction is indeed convex in duration but with a few important exceptions. The discontinuity around the 5-year duration sector is caused by the fact that newly issued 5- and 7-year notes reacted less strongly than the deeply off-the-run 30-year bonds. The importance of the local supply effects becomes clear following the second announcement by the Desk. The blue dots in Figure 1 show the yield changes from 2:00 to 3 p.m., about 15 minutes after the Desk announcement, and again at 4 p.m., shown in green. All securities excluded from the purchase range (i.e. outside of the 2-10-year maturity bucket, which in our sample corresponds to about 1.9-7.8 years of duration) reversed a large part of their initial yield declines after the Desk announcement, while those included in the purchase range largely maintained (or amplified) their earlier responses, even though, in the aggregate, a smaller amount of dollar duration was going to be taken out of the market. This would suggest that in the 2-10-year maturity sector, the local supply effect (elicited by a larger-than expected reduction in the supply of these securities relative to those in the 10-30-year sector) counterbalanced the smaller-than-expected total duration effect, leaving yields unchanged or lower following the second announcement. Further, we find it extremely striking that the turning point is around the 7.8 years duration, which, being the average duration of 10-year Treasury securities, represents the threshold between the maturity ranges included and excluded from the Fed purchase program.

The Reinvestment policy announcement

On August 10, 2010, at 2:15 p.m., the FOMC announced that it would keep the face value of its System Open Market Account (SOMA) holdings constant by reinvesting principal payments from agency debt and agency MBS in longer-term Treasury securities. This time, the announcement contained a footnote indicating that the Desk would issue shortly thereafter a technical note containing operational details on the announced transactions. At 2:45 p.m., the Desk indicated that it would again concentrate purchases in the 2- to 10-year sector of the nominal Treasury curve, and would refrain from purchasing securities for which there was a heightened demand or for which the SOMA already held large concentrations. According to market commentaries, only some market participants had anticipated an announcement regarding the reinvestment of principal payments, and among those only a few expected this reinvestment to be in Treasuries rather than in other securities.[7] Therefore, similar to the March 2009 episode, the announcement of the reinvestment program came largely as a surprise, and the potential for a price reversal elicited by the Desk's announcement could again throw considerable light on the impact of Treasury operations and their channels. Further, the price reversal in some securities could have been exacerbated by the new part of the statement indicating that the Desk would refrain from purchasing

[7] Similarly to market commentaries, the August Primary Dealer Survey indicated that the average probability of agency MBS reinvestments to be in Treasury securities was 19 percent.

securities that were already held in the SOMA portfolio in high concentration, indicating that those securities potentially had a lower probability of being bought.

Figure 2 shows the yield changes during these events. Following the FOMC announcement of the reinvestment program, the initial yield reaction was quite similar in shape to the one prevailing after Desk's announcement on March 18, 2009, reflecting learning by market participants about the maturity distribution from the first round of purchases. In other words, it seems reasonable that based on the previous experience with LSAP purchases, market participants expected the maturity distribution of the Fed purchases under the reinvestment program to reflect the actual distribution of the first LSAP's purchases. However, following the second announcement by the Desk, securities with maturities beyond the purchase range more than reversed their initial yield declines, reflecting in part the relatively large concentration of 10- to 30-year Treasury securities in the SOMA holdings.[8] In contrast, yields of securities included in the purchase range continued to decrease later in the day. This pattern suggests once more the likely existence of local supply effects, as the surprise in the maturity distribution of purchases had significant effects on yields that were not monotonic in their durations.

The second LSAP (LSAP2) announcement

On November 3, 2010 at 2:15 p.m., the FOMC announced its decision to expand its holding of longer-term Treasury securities by $600 billion by the end of the second quarter of 2011, and, differently from the previous two announcements, *simultaneously* the Desk released its operating policy detailing the intended distribution of purchases for the nominal securities across seven maturity sectors. From the release of the operational details, market participants learned that the 10- to 30-year maturity sector would receive only 6 percent of purchases, compared with an allocation of about 15 percent in the two previous programs. It is plausible to assume that based on past experience, they would have expected a larger weight on the 10- to 30-year sector, and therefore the announced distribution of purchases was more

[8] The Federal Reserve's SOMA portfolio held at the time of the announcement 18.5 percent of the outstanding amount in the 10-30-year maturity sector, 13 percent of the 7-10-year maturity sector and 14 percent of the 5.5-7-year maturity sector.

heavily weighted towards short- and medium-term securities than some market participants had expected. This in turn would have implied a smaller-than-expected reduction in the average duration, driven both by a smaller-than-expected reduction in the available supply of the 10- to 30-year sector and a larger-than-expected reduction in the available supply of the 2- to 10-year sector. Since the FOMC and the Desk's announcements were done simultaneously, differently from the previous two episodes, we will not be able to observe the price behavior in between the two announcements; but, if the local supply channel is operating, we should observe prices moving in opposite directions below and above the 10-year maturity threshold to reflect the negative purchase surprise in the 10- to 30-year sector and the positive purchase surprise in the 2- to 10-year sector. This type of response is evident in Figure 3.

Differently from the previous two pictures, Figure 3 shows the yield responses only in one time-window (that is, from 2:00 to 4:00 p.m.) around the FOMC announcement, because in this case, as already explained, the Desk's announcement was contemporaneous to that of the FOMC. It can be seen that the yields of short- and intermediate-term securities decreased, while the yields of securities with duration beyond 8 years increased after the announcement, most likely reflecting a smaller-than-expected reduction of the available supply in the 10- to 30-year sector and a larger-than-expected reduction of the available supply in the 2- to 10-year sector. In other words, these two wide humps elicited by yields moving in opposite directions below and above the 7-year duration threshold (at that time the duration of the 10-year benchmark was 8.5 years) cannot be explained by a supply effect mainly driven by the duration risk channel, suggesting once again that the local supply channel may be at work.

The Maturity Extension Program (MEP) announcement

On September 21, 2011, at 2:23 p.m. the FOMC announced its intention to extend the average maturity of its holdings of securities by purchasing $400 billion of Treasury securities with remaining maturity between 6 and 30 years and selling an equal amount of Treasury securities with remaining maturity of 3 years or less, which is known as the MEP. Once again, the Desk simultaneously released the intended distribution of purchases, but this time across five rather than seven maturity sectors, to reflect the fact that purchases were going to occur only in the medium- and long-term sections of the yield curve. Most likely, based on previous purchase operations, in which the 10- to 30-year sector of the yield curve received at most 15 percent of the purchases, this time market participants could have been surprised by the large fraction of 29 percent of purchases allocated only to the 20- to 30-year maturity range. Therefore, differently from past experience, in this program the announced distribution of purchases was more heavily weighted towards longer-term securities than expected. This would imply a larger-than-expected removal of duration risk from the market, driven by a larger-than-expected reduction of the available supply to private investors in the 20- to 30-year sector (i.e. a positive surprise). This combined with sales at the front end of the curve, which translate into an increased availability of securities with maturity of 3 years or less (i.e. a negative surprise), would suggest that if there is a local supply effect we should observe yields increasing at the front-end of the curve in response to the negative surprise and yields decreasing at the long-end of the curve in response to the positive surprise.

Again, this pattern can be observed in Figure 4, which illustrates the yield response in the same time window used for LSAP2 (that is, from 2:00 to 4:00 p.m.) to the FOMC and Desk announcement about the MEP. In this case, the yields in the short- and medium-term sectors of the yield curve increased a little, while yields beyond the 4-year duration decreased, with the largest decline being observed at the 20- to 30-year maturity range. This would be consistent with larger-than-expected sales of securities at shorter maturities (that is, a negative surprise) and larger-than-expected purchases at longer maturities (that is, a positive surprise), suggesting again a combination of local supply and duration risk effects.

The MEP extension (MEP2) announcement

On June 20, 2012, at 12:30 p.m. the FOMC announced its decision to extend the MEP through the end of 2012, resulting in the purchase and sale of $267 billion of additional Treasury securities. Contemporaneously, the Desk released the operational details indicating that the maturity buckets and the weights associated to each bucket were the same as in the initial MEP. However, it also stated the suspension, for the duration of this program, of its practice of rolling over maturing Treasury securities held in the SOMA portfolio into new issues at auction. It is worth noting that, in terms of effect on the SOMA portfolio, redeeming maturing Treasury securities is nearly identical to selling securities that are approaching maturities. This new component of the announcement could have potentially surprised the market participants, as the decision to redeem securities maturing in the second half of 2012 allowed the Fed to increase the total size of future purchases beyond the $205 billion of securities with maturity of 3 years or less already held in SOMA (which during the MEP were sold to finance the purchases). Therefore, it is conceivable that market participants had used this amount to guide their expectations about future purchases, underestimating the size of the MEP extension.[9] This implies that the announced size and duration of purchases could have exceeded expectations, because through the securities redemptions the Fed was going to purchase a larger-than-expected amount in the 6- to 30-year sector and implicitly sell or redeem a larger-than-expected amount in the 3 years or less maturity sector, which should lead Treasury yields to increase in this sector and to decrease in the in the 6- to 30-year sector.

Figure 5 illustrates how the shape of the reaction to the MEP extension is very similar to that one of the initial MEP, except that this time its overall magnitude was smaller because the size of the surprise was smaller. Most likely, also in this case, larger-than-expected sales (that is, a negative surprise) pushed yields up at shorter maturities and larger-than-expected purchases (that is, a positive surprise) drove yields down at longer maturities, implying the same transmission mechanism of the impact of this program to the term structure of interest rates as during the first MEP.

Overall, our first look at the shape and magnitude of the reactions to the various Federal Reserve asset purchase programs across all nominal outstanding CUSIPs points to the existences of both a local supply effect as well as a duration risk effect. However, to formally test these conjectures we first make our best attempt at measuring the individual local-supply and duration-risk shocks and second estimate their effects on nominal Treasury yields.

---

[9] The amount of $205 billion is obtained by simply computing the amount of SOMA holdings that at that time were expected to fall below 3 years over a 6-month period (which was the announced length of the program), and was known by most market participants as SOMA holdings are public information.

3. Basic theoretical framework

We are going to outline a basic theoretical framework to understand the main drivers of changes in Treasury yields around the Fed's asset purchase announcements. And specifically, we are going to illustrate how this type of policy, which is designed to operate on term premiums, affects Treasury yields through channels that are quite different from the expectation hypothesis.

According to simple multi-factor affine term-structure models, the yield to maturity, $r_{\tau,t}$, of a bond with remaining maturity τ at time *t*, can be decomposed in an expectation component, which reflects the expected path of short-term interest rates, r_t, over the life of the bond, and a term premium component, $tp_{\tau,t}$, which is the extra return that investors require to hold a longer-term security:

$$r_{\tau,t} \approx \frac{1}{\tau}\int_{t}^{t+\tau} E_t(r_u)du + tp_{\tau,t} \qquad (1)$$

$$tp_{\tau,t} = \boldsymbol{A_\tau} \sum \boldsymbol{\lambda_t}. \qquad (2)$$

As shown in equation 2, the term premium is determined by the sensitivity of the maturity-τ bond to the product of the quantity and market-price of risk, and in particular, in our multiple-factor setup: $\boldsymbol{A_\tau}$ is a matrix of the sensitivities to the various risk factors, $\sum$ is a variance-covariance matrix of the factors' innovations (assumed to be constant), and $\boldsymbol{\lambda_t}$ is a time-varying vector of market prices of risk.

During the Fed's asset purchase announcements new information is released about the size and duration of these programs as well as their implementation, and this information may induce investors to change their expectations about the available Treasury supply's size and maturity composition. Then a natural question is how these two key variables enter in the above equation, and thus through which channels they affect Treasury yields. To answer these questions and better organize our empirical investigation we build on the preferred-habitat term-structure model of Vayanos and Vila 2009 (V&V).

Suppose that there are *N* distinct Treasury securities outstanding and the Fed determines not only the short rate r_t (assumed to follow an exogenous process) but also the stock and maturity composition of Treasury debt available to private investors, Q_t-Q_t^{Fed}, by choosing the total amount of purchases, Q_t^{Fed}, and a vector of purchase allocations (weights) to each security within a maturity sector *k*, $\mathbf{w}_{t,k} = (w_{t,k}(1) \ldots$

$w_{t,k}(N)$), such that Q_t^{Fed} is given by a linear combination of "quantity risk factors," $\mathbf{q}_t = (q_{1,t} \ldots q_{K,t})$, and weights, $\mathbf{w}_{t,k}$, determining how the k^{th} factor maps into the net supply of each security.[10] Similarly to the simplified version of V&V developed in Greenwood and Vayanos (2014), we treat the securities' demand coming out of the Fed as exogenous and impose that in equilibrium private investors must hold the remaining available supply of each security.[11] In this case, as shown in V&V the yield to maturity takes the affine form:

$$r_{\tau,t} = A_r(\tau)r_t + \boldsymbol{A_q(\tau)q_t} + C(\tau), \qquad (3)$$

where $A_r(\tau)$ is the sensitivity of the maturity-τ yield to short-rate fluctuations, and $\mathbf{A_q}(\tau) = (A_{q(1)}(\tau), \ldots, A_{q(K)}(\tau))$ is the sensitivity of the maturity-τ yield to each of the supply factors. In this basic framework, where the term-structure is a linear function of multiple factors (i.e., the short-term rate and the supply of Treasury securities in each maturity-k sector), each risk factor has its own market price, which in the V&V's formulation has the following representation:[12]

$$\lambda_{r,t} = \alpha\sigma_r \int_0^N x_{\tau,t} A_r(\tau) d\tau \qquad (4)$$

$$\lambda_{q_k,t} = \alpha\sigma_{q_k} \int_0^{N_k} x_{\tau,t} A_{q_k}(\tau) d\tau \qquad (5)$$

where α is the investors' coefficient of risk aversion and $x_{\tau,t}$ is the dollar value of the security's share in the private-sector portfolio.

If we substitute equations (4) and (5) in equation (1), the yield to maturity can be rewritten as linear combination of the main sources of risk and market prices of risk that are affected by the Fed's

---

[10] A maturity-k sector can include more than one maturity. In the empirical specification we split the yield curve in 7 maturity sectors that coincide with the operational sectors employed by the New York Fed.

[11] It is straightforward to see that this model can be easily mapped into V&V, where the preferred-habitat sector consists of the Fed and the private investors are the arbitrageurs. This interpretation is similar to Hamilton and Wu (2011), Kaminska, Vayanos, and Zinna (2011), Greenwood and Vayanos (2014), and Kaminska and Zinna (2014). However, here the net supply shocks are function of shocks either to the total amount of purchases, Q_t^{Fed}, or to the purchase allocations $\mathbf{w}_k$, or both, since $q_{k,t}= w_{t,k} * Q_t^{Fed}$.

[12] See equations 27 and 28 in V&V or equivalently equation 11 in Greenwood and Vayanos (2014).

asset purchase programs, as they are most directly impacted by changes in each security's share $x_{\tau,t}$ that in equilibrium private investors must hold:

$$r_{\tau,t} \approx \frac{1}{\tau}\int_t^{t+\tau} E_t(r_u)du + A_r(\tau)\sigma_r\lambda_{r,t} + \sum_{k=1}^{K} A_{q_k}(\tau)\sigma_{q_k}\lambda_{q_k,t} \quad (6)$$

$$= EH_\tau + A_r(\tau)\alpha\sigma_r^2 \int_0^N x_{\tau,t}A_r(\tau)d\tau + \alpha\sum_{k=1}^{K} A_{q_k}(\tau)\sigma_{q_k}^2 \int_0^{N_k} x_{\tau,t}A_{q_k}(\tau)d\tau, \quad (7)$$

where in the last equation, for simplicity, we set $EH_\tau = \frac{1}{\tau}\int_t^{t+\tau} E_t(r_u)du$. To understand how a change in $x_{\tau,t}$ translates into a change in the yield to maturity around the Fed's asset purchase announcements, we have to compute the derivative of $r_{\tau,t}$ with respect to the security's share:

$$\frac{\partial r_{\tau,t}}{\partial x_{\tau,t}} = A_r(\tau)\sigma_r\frac{\partial\lambda_{r,t}}{\partial x_{\tau,t}} + \sum_{k=1}^{K} A_{q_k}(\tau)\sigma_{q_k}\frac{\partial\lambda_{q_k,t}}{\partial x_{\tau,t}}. \quad (8)$$

Equation (8) illustrates in a very simple way that the yield change is determined by the sum of two distinct effects: one related to the change in the market price of interest-rate risk and another one related to changes in market prices of supply risk specific to each sector *k*. Further, it is also easy to note that the expectation component, *EH*, being independent of $x_{\tau,t}$, does not play any direct role in explaining yield changes during these announcements within this framework.[13] However, it is reasonable to think that during these announcements, the quantity of risk may also change, although for simplicity we assumed it to be constant. For example, interest-rate uncertainty, σ_r, could decline if asset purchase announcements are interpreted as signaling a further commitment to keep short-term rates low for a longer period of time. This type of signaling channel would affect the term premium component by driving down an element of $\sum$ in equation (2). In our reduced-form setup one cannot separate this effect from changes in the market price of risk; therefore, this second order effect might contaminate the impact of the two channels under consideration.

The main determinant of the first term in (8) is

[13] Similarly, it can be observed in equation (A.10) of Greenwood and Vayanos (2014) that the expression for C(τ) does not depend on the security share $x_{\tau,t}$.

$$\frac{\partial \lambda_{r,t}}{\partial x_{\tau,t}} = \alpha \sigma_r A_r(\tau),$$

therefore, for each program at time *t*, the total change in the market price of interest-rate risk is given by

$$\int_0^N \frac{\partial \lambda_{r,t}}{\partial x_{\tau,t}} d\tau = \alpha \sigma_r \int_0^N A_r(\tau) d\tau, \qquad (9)$$

which is nothing more than the overall sensitivity (across all *N* securities outstanding) to interest-rate risk, that is, the private sector's *aggregate* exposure to this risk. In this simple framework, we label the portion of the change in the yield term premium due to changes in the aggregate duration risk, the *duration-risk effect*. Similarly, for each maturity-*k* sector, the main component of the second term in (8) is

$$\frac{\partial \lambda_{q_k,t}}{\partial x_{\tau,t}} = \alpha \sigma_{q_k} A_{q_k}.$$

therefore, for each program at time *t*, the total change in the market price of supply risk is given by

$$\int_0^{N_k} \frac{\partial \lambda_{q_k,t}}{\partial x_{\tau,t}} d\tau = \alpha \sigma_{q_k} \int_0^{N_k} A_{q_k}(\tau) d\tau. \qquad (10)$$

Since A_{q_k} measures the sensitivity of the maturity-τ yield to the k^{th} supply factor, this expression measures the overall exposure of the N_k securities to supply risk in a particular sector. As a consequence, we label the portion of the change in the yield term premium due to fluctuations in a specific supply factor q_k, the *local-supply effect*.

Equations (7) to (10) imply that if we want to explain the *cross-sectional* variation in the actual yield changes of all Treasury securities outstanding around the Fed's asset purchase announcements, we can ignore the terms that are constant across securities (i.e., α, σ_r, σ_{q_k}), but we have to measure the *unexpected* changes in $\int_0^N x_{\tau_i,t} A_r(\tau_i) d\tau_i$ and $\int_0^{N_k} x_{\tau_i,t} A_{q_k}(\tau_i) d\tau_i$, which is the first main contribution of our empirical methodology. In turn, the availability of empirical proxies for the exogenous changes in these key variables within a time-window ω around the Fed's announcements, allows us to quantify the separate impact of the duration-risk and local-supply channels on Treasury yields, the second main contribution of this study. This is done by estimating the following cross-sectional specification (which is the empirical equivalent of equation 7):

$$\Delta_t^{t+\omega} r_{\tau_i} = C + A_r(\tau_i)\Delta_t^{t+\omega}\left(\sum_{i=1}^{N} x_{\tau_i,t}\, A_r(\tau_i)\right) + \sum_{k=1}^{K} A_{q_k}(\tau_i)\Delta_t^{t+\omega}\left(\sum_{i=1}^{N_k} x_{\tau_i,t} A_{q_k}(\tau_i)\right) + \varepsilon_i, \quad (11)$$

where $\Delta_t^{t+\omega}$ indicates changes from *t* to *t+ω* and for simplicity we have omitted the constants α, σ_r^2, and $\sigma_{q_k}^2$. By comparing equations 7 and 11, it is easy to see that under our specification, any announcement's impact on the expectation component (*EH*) of the yield to maturity will be absorbed in the constant *C*. And because of the empirical evidence on the signaling effect at the zero-lower bound, provided in Section 6.1, we impose that this term does not vary across maturities.

It should be stressed that in equation (11) we are already allowing for quantity fluctuations at the individual-security level rather than at maturity level, as indicated by τ_i (i.e., the remaining maturity of security *i* at time *t*), to reflect the higher level of disaggregation of our data and facilitate comparisons to the variables' specification in the next section. In particular, we will describe in detail our empirical proxies for $\Delta_t^{t+\omega}\left(\sum_{i=1}^{N} x_{\tau_i,t}\, A_r(\tau_i)\right)$ and $\Delta_t^{t+\omega}\left(\sum_{i=1}^{N_k} x_{\tau_i,t} A_{q_k}(\tau_i)\right)$, as well as the restrictions that these proxies entail on the functional form of $A_r(\tau_i)$ and $A_{q_k}(\tau_i)$.

4. Computations of the individual security's local-supply and duration-risk surprises

Financial markets are inherently forward looking and react only to the new information contained in any announcements. Therefore, a rigorous event-study analysis requires a careful specification of both the expected and unexpected components of the announcement. Given this study's focus on local-supply and duration-risk shocks, this calls for the estimation of market participants' *expectations* about three main variables: the probability of the announcement to occur (P), the total size of the program (Q), and the weight (w_k) associated to each of the *K* maturity buckets across which the purchases are distributed. We obtain a measure of each of these variables from the Primary Dealer Survey (PDS) conducted by the New York Fed one week before each FOMC announcement, supplemented with the available information from market commentaries published before each FOMC.

4.1 Measuring the local-supply shocks

Our procedure consists of three steps, as we start from computing the stock surprises for the maturity sectors employed in the Desk announcements, then we move from maturity-sector surprises to security-level surprises, and finally we group individual surprises based on the securities' degree of substitutability. This specification of the supply shocks is motivated by the evidence on the importance of accounting for substitution effects across Treasuries when quantifying the impact of LSAPs, shown in D'Amico and King (2013).

The first step is to measure the size surprise for each maturity bucket, that is, the maturity distribution of the supply shock. For each asset purchase announcement, we first estimate investors' prevailing expectations of both its probability to occur, P, and its total size conditional to the program occurring, E(Q| the program occurs), from the PDS and/or market commentaries. Next, the associated vector of maturity-sector weights $E(\mathbf{w}_k)$ is calculated as follows.[14] For LSAP1, given the novelty of the Treasury purchase program, we simply assume that investors expected such purchases to be spread across all maturity sectors proportionally to the amount outstanding in each sector.[15] For all other announcements, we set the pre-announcement maturity weights to be identical to the actual purchase allocations observed under the immediately preceding program.[16] One exception is the MEP due to the expectations formation about purchase and sale sectors. For the purchase sectors, we take the weights associated with purchases allotted to the 6- to 30-year sector under the LSAP2 (the immediately preceding program) and renormalize them to sum to one.[17] For the sale sectors, under both the MEP and

[14] A detailed description of our computations for each of these variables is provided in Appendix A and B.
[15] Because of this somewhat arbitrary choice, in the robustness section we show an estimation of the local supply effect that is independent of the maturity weights' specification.
[16] For LSAP2, in particular, those maturity allocations were cited by market observers as the most likely weights given the similarity in the broader macroeconomic motivations to the previous programs.
[17] Most PDS respondents and market observers cited the 6- to 30-year sector as the most likely for purchases.

the MEP extension, sales are assumed to be concentrated in the 0- to 3-year bucket as indicated in the PDS responses. Finally, the expected maturity distribution of purchases/sales, $E(Q_k)$, is given by:[18]

$$E(Q_k) = P * E(Q|the\ program\ occurs) * E(w_k).$$

On the day of the FOMC, the Committee announces the actual size of the program, Q, and the Desk releases the actual intended weights for the distribution of purchases across *K* maturity buckets, $\mathbf{w}_k$, with k=1,..,K. The par value of the purchase/sale surprise for each maturity bucket, SQ_k, is then computed as the difference between the actual intended maturity distribution and the expected maturity distribution of the purchase amount:

$$SQ_k = Q * w_k - E(Q_k).$$

A positive surprise implies larger-than-expected purchases and a negative surprise implies smaller-than-expected purchases (or larger-than-expected sales in the case of MEPs) in the maturity-*k* sector.

Next, a measure of each security's purchase/sale surprise and related local-supply shock is estimated, with these two measures being identical only when the substitution effects across securities are ignored. We first assume that within a maturity bucket the purchase or sale surprise is allocated to each security *i* proportionally to its private or SOMA holdings, h_i and g_i, respectively:

$$s_i = \frac{SQ_k * h_i}{H_k}\ for\ purchases,$$

$$s_i = \frac{SQ_k * g_i}{G_k}\ for\ sales,$$

where $H_k = \sum_{j=1}^{N_k} h_j$ and $G_k = \sum_{j=1}^{N_k} g_j$ are the total private holdings and SOMA holdings, respectively, in the maturity bucket *k* containing N_k securities. This is because it is reasonable to assume that within a maturity bucket the amount of each security available for purchase depends on the amount left in the hands of the private investors, while the amount available for sale depends on the amount left in the SOMA portfolio.

[18] We assume that *Q* and *w* are independent.

To capture the substitution effect among nearby securities, for each security *i*, we define its local-supply shock (ls_i) as the weighted sum of its own supply shocks and the supply shocks of *J* securities having remaining maturities (τ_j) within a certain distance of security *i*'s maturity, and normalized by the weighted sum of the corresponding private holdings of those securities. We chose this specification because, as shown in D'Amico and King (2013), the price of each Treasury security reacts to its own purchases and those of securities with similar maturities. In particular:

$$ls_i = \frac{\sum_{j=1}^{J} s_j \delta_{ji}}{\sum_{j=1}^{J} \delta_{ji} h_j}. \qquad (12)$$

The weights (δ_{ji}) are a decreasing function of the maturity distance, which is defined over a variable window size, and are specified as follows:

$$\delta_{ji} = \left(1 - \frac{|\tau_j - \tau_i|}{\tau_i * \theta}\right) \mathbf{1}_{\{|\tau_j - \tau_i| \leq \theta * \tau_i\}},$$

with $\mathbf{1}_{\{.\}}$ being an indicator function, and the weight δ_{ji} equal 1 only for security *i*, symmetric, and approaching 0 as the maturity distance approaches $\theta * \tau_i$. In our baseline, we set θ equal to 0.5 to consider all securities within a maturity distance of 50 percent of security *i*'s maturity. This triangular kernel characterized by a variable window size should minimize biases due to the fact that the number of substitutes decreases as the maturity increases. In other words, as we move toward the long end of the yield curve it is important to increase the window size because the number of available securities becomes smaller. Further, in the robustness section, we show the sensitivity of our results to the choice of the window size and run an optimization routine to derive the best fitting value for θ.

Our empirical specification of ls_i is consistent with the theory sketched in Section 3 once we set (i) the change in each security's share $x_{\tau,t}$ that in equilibrium private investors must hold equal to the individual security's surprise:

$$\Delta x_{\tau_i} = s_i;$$

(ii) allow the number of supply risk factors to equal the number of securities outstanding, that is, *K*=*N*, as we move from supply surprises defined at maturity-bucket level to supply surprises defined at individual-security level to exploit the richness of our data; and (iii) assume that the sensitivity of each security's yield to its own and nearby securities' supply risk is determined by the relative maturity distance:

$$A_{q_k}(\tau_i) = \frac{\delta_{ji}}{\sum_{j=1}^{J} \delta_{ji} h_j}.$$

Finally, under these assumptions it is easy to see that ls_i is an empirical approximation of the unexpected variation in the market price of supply risk, $\lambda_{q_k,t}$:

$$ls_i = \frac{\sum_{j=1}^{J} s_j \delta_{ji}}{\sum_{j=1}^{J} \delta_{ji} h_j} \approx \left(\sum_{i=1}^{N_k} \Delta x_{\tau_i} A_{q_k}(\tau_i) \right),$$

where the N_k securities are defined over a variable moving window determined by $|\tau_j - \tau_i| \leq \theta * \tau_i$.

Figures 6 to 10 show a bar graph of the maturity bucket purchase/sale surprises for each announcement, where purchases are represented as positive numbers and sales are represented as negative numbers measured in billions of dollars. In particular, the yellow shaded portion of each bar measures the expected amount of purchases/sales and the blue shaded portion of each bar measures the unexpected amount of purchases/sales within each maturity bucket. In addition, Figure 11 illustrates an example of the computation of the local-supply shocks for the MEP. Maturity-bucket surprises, denoted by the bars, are normalized by the amount outstanding in each bucket, and the related local-supply shocks, denoted by the red dots, are calculated for each security using equation (12). Because there are very few substitutes for securities with maturities longer than ten years, the local-supply shocks at those maturities become very similar to the normalized maturity-sector surprises. Appendices A and B detail the computation of these shocks for each of the Fed's asset purchase announcements.

4.2 Measuring the duration-risk shock

The construction of our empirical proxy for the security-level duration-risk shocks proceeds in two steps. We first compute our measure of the aggregate exposure to duration risk and its unexpected variation around the Fed's announcements. Then, we move from the aggregate duration-risk shock to the individual security's duration-risk surprises.

In detail, as shown by V&V and already described in Section 3, the market price of interest-rate risk, $\lambda_{r,t} = \alpha\sigma_r \int_0^N x_{\tau,t} A_r(\tau) d\tau$, is mainly determined by the dollar value of the duration of the private-sector portfolio, whose discretized formulation is given by $\sum_{i=1}^{N} x_{\tau_i,t}\, A_r(\tau_i)$. In other words, since α and σ_r are constant across securities, in measuring the unexpected variation in $\lambda_{r,t}$, we can focus only on the changes in the dollar value of aggregate duration, which according to market convention is approximated by the amount of ten-year equivalents in billions of dollars left in the hands of private investors:

$$SDR = \sum_{k=1}^{K} \sum_{i=1}^{N_k} \frac{\Delta x_{d_i,t} * d_i}{d_{BM}} = \sum_{k=1}^{K} \sum_{i=1}^{N_k} \frac{s_i * d_i}{d_{BM}}, \quad (13)$$

where d_{i} is the Macaulay duration of security i, d_{BM} is the duration of the benchmark on-the-run 10-year Treasury, and $\Delta x_{d_i,t}$ is the change in the dollar value of the security's share, which in equilibrium has to equal the security's surprise s_i. Thus, equation (13) is measuring the surprise in the total duration risk (SDR) and is consistent with our basic theoretical framework when we set

$$A_r(\tau_i) = \frac{d_i}{d_{BM}},$$

as in this case SDR is an approximation of the unexpected change in the market price of duration risk

$$SDR \approx \left(\sum_{i=1}^{N} \Delta x_{t,\tau_i}\, A_r(\tau_i) \right).$$

A positive value implies that the Fed is removing from the market a larger-than-expected amount of duration risk.

Each individual security's duration-risk (*idr*) shock is then given by the product of the security's sensitivity to this risk factor and the aggregate duration-risk surprise (*SDR)*:

$$idr_i = \frac{(1 - \exp(-\kappa_r * \mathrm{d_i}))}{\kappa_r} * \mathrm{SDR},$$

where the specific functional form for the security's sensitivity is motivated by equation (13) in Greenwood and Vayanos (2014), with the short-rate mean-reversion parameter κ_r controlling the steepness of the concave function. We initially set κ_r equal to 0.2 based on the estimates of the Li and Wei (2013) model of the term-structure augmented with observed supply factors.[19] In the robustness section, we show the sensitivity of our results to the value of this parameter and run an optimization routine to derive the best fitting value for κ_r. Figure 12 provides an example of the computation of the individual duration risk shocks, denoted by the yellow dots, as function of duration in years.

4.3 Identification Strategy

As explained in the two previous sections, consistently with the existing theory (e.g., Vayanos and Vila 2009, Li and Wei 2013, and Greenwood and Vayanos 2014) we model the duration-risk effect with a concave and monotonic function, implying that this channel behaves like a slope risk factor. While, for the local-supply effect we use a non-parametric functional form that can generate multiple humps in different maturity sectors (D'Amico and King, 2013), implying that this channel's behavior is not necessarily consistent with any canonical yield curve factors (level, slope and curvature). Finally, in Section 6.1, we also show that the signaling channel behaves empirically like a level risk factor during the announcements under consideration. Therefore, we model its impact with the constant term in our cross-sectional specification.

Figures 13 to 17 show the Treasury yield reaction across different durations, measured from 15 minutes before the FOMC announcement to 4:00 p.m. of the next day (the green triangles), together with the plot of our measures for the underlying risk factors, that is, the individual duration-risk shocks (the red squares) and the local-supply shocks (the blue circles), for each of the program. It is quite evident that the two channels have very different functional forms across maturities: the duration-risk has a monotonic and concave shape (by construction) and the local-supply has indeed a multiple-hump shape. This implies that these two explanatory variables have a low correlation and therefore their impact in our

[19] Equation (13) in Greenwood and Vayanos (2013) is the same as V&V; and similarly to Li and Wei (2013), these authors also estimate κ_r to be 0.201, see their Appendix B.3 for more detail.

cross-section regression specification can be separately identified. This holds true also when we add the constant term to capture the signaling effect.

Overall, under our estimation strategy that exploits security-level variation across maturities/durations, the impact of these three channels can be separately identified as long as they have distinct functional forms, which seems to be the case based on the existing theory and empirical evidence.

5. Estimation of local-supply and duration-risk effects

To estimate the effects of the different channels, we run a cross-section regression of yield changes on duration-risk shocks and local-supply shocks for each program announcement separately, and in addition we also consider all five announcements together and estimate a pooled data regression to improve the efficiency of our estimates. The regression specification is the following:[20]

$$\Delta y_i = \beta_0 + \beta_1 * ls_i + \beta_2 * idr_i + u_i,$$

where the yield changes are measured from 15 minutes before the FOMC announcement to 4:00 p.m. of the next day. We favor this event-window over a tighter time-window, often used for conventional monetary policy announcements, for various reasons. First, the novelty of the programs, as in the case of the first LSAP and MEP, implies that a longer amount of time is required to process the released information compared to conventional FOMC announcements. Second, as discussed in Section 2, in some cases (LSAP1 and Reinvestment) the New York Fed released the program details 30 minutes after the FOMC announcements (at 2:45pm), which also required some time to be fully processed. Third, the June 2012 FOMC announcement (MEP2) was followed by the Chairman Press briefing, which ended at about 3 p.m. and also contained relevant information that needed to be digested. Finally, the number of securities used in the regression ranged from 163 in LSAP1 to 245 in MEP2.

[20] It can be estimated by OLS due to the availability of a measure for the exogenous supply 'shocks.'

Table 1: Yield change regression results with variable window size, θ=0.5 and κ_r=0.2

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Pooled
	Two-day yield change regression					
Constant	0.466	-1.078	-2.982	3.169	0.367	0.629
	(0.52)	(-2.45)	(-4.65)	(3.32)	(0.91)	(1.92)
Duration risk shock	-3.000	-1.280	-0.952	-2.189	-0.399	-1.803
	(-22.07)	(-3.11)	(-1.97)	(-11.75)	(-3.60)	(-21.36)
Local supply shock	-0.385	-1.632	-1.210	-1.481	-0.480	-0.807
	(-12.51)	(-5.76)	(-25.16)	(-19.16)	(-12.00)	(-31.58)
R-squared	0.84	0.69	0.76	0.91	0.76	0.72
Observations	163	200	208	232	245	1048

Note: t-statistics in parenthesis.

As shown in Table 1, for each program announcement, the estimated coefficients on both the duration-risk shock and the local-supply shock are statistically significant and have the expected negative sign, except for the LSAP2 where the duration-risk coefficient is only marginally significant. As denoted by the R-squared, the total yield variation explained by the two shocks is in the range of 75 to 90 percent across the various programs. The remaining unexplained variation over this event-window can be due to other factors such as, for example, the liquidity and signaling effects. The liquidity effect, as indicated by Gagnon et al. (2011), can be due to the market perception that the presence of a large and consistent buyer like the Fed may enhance trading opportunities and therefore reduce the liquidity premium of some securities like off-the-run bonds. This effect should be more important during times of market stress, such as the first half of the LSAP1 period. However, the R-squared for LSAP1 is about 85%, which is quite large not only in absolute terms but also relative to the other programs, suggesting that even in this case the liquidity channel probably plays a marginal role. Regarding the signaling channel, we postpone its detailed discussion to Section 6.1, where we provide some evidence that its impact should be mostly captured by the constant term in the regression. Finally, the last column of Table 1 shows the results for the pooled data regression, which confirm that also in this case both coefficients are negative as well as statistically significant, and that, as expected, their magnitude is very close to the average of the estimates across the five events. The explained variation is 72 percent.

To provide an economic interpretation of these coefficients and to illustrate how the efficacy of this policy tool evolved over time, for each program, we compute the implied effects on the on-the-run 10-year yield (in basis points) from an unexpected $100-billion purchase announcement, which are reported in Table 2a. This is done using the coefficient estimates from each program and by assuming that the $100-billion surprise is distributed across the maturity sectors as in the actual announced program. As shown in the fourth column of Table 2a, the MEP is characterized by the largest impact on the 10-year yield, with a reduction of 13 basis points almost evenly split between the two channels. Further, LSAP1, the Reinvestment program, and LSAP2 have a very similar estimated effect of about 9 basis points, although in each program the relative importance of the two channels is quite different. Finally, the MEP extension is estimated to have reduced the same yield by about 4 basis points, of which 2.5 basis points are due to the local supply effect. On average, as indicated in the last column, using the individual coefficients, we find that a $100-billion purchase surprise translates in a 9-basis-point reduction in the 10-year yield, with each channel accounting for about half of the decline.

Table 2a: Implied effect on the 10-year yield from an unexpected $100B program using individual regression's coefficients

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Average
			Impact in basis points			
Total*	-8.9	-9.4	-9.2	-13.1	-3.7	-8.9
of which, duration risk	-7.6	-3.4	-2.5	-8.5	-1.5	-4.7
of which, local supply	-1.8	-5	-3.8	-7.8	-2.6	-4.2

*Includes the estimated constant term.

Furthermore, using both the actual surprises derived from the PDS, which are reported in the first column of Table 2b, and the implied effect per a $1-billion surprise, it is possible to compute the total impact of each program's announcement on the 10-year yield, shown in the second column of Table 2b.

It is interesting to note that the actual surprises are much smaller than the total amount announced, which is often used in many event studies to calculate the yield elasticity of a $1-billion purchase. As such, it is not surprising that these event studies often find a decreasing elasticity across consecutive Fed programs, as they do not properly capture that the surprise component is getting smaller over time, most likely due to the Fed's increased transparency about this monetary policy tool.

Table 2b: Implied effect on the 10-year yield from the actual surprises

	Surprise (billion)	10-year yield reaction (basis points)	
		Total	*per $1 billion surprise*
LSAP1 (3/18/2009)	418.2	-37.2	-0.089
Reinvestment (8/10/2010)	117.6	-11.0	-0.094
LSAP2 (11/3/2010)	109.5	-10.1	-0.092
MEP (9/21/2011)	114.0	-14.9	-0.131
MEP2 (6/20/2012)	148.3	-5.48	-0.037

In addition, to try to isolate the impact of the programs' design rather than the change in sensitivity to supply shocks over this sample period, we repeat the same exercise of Table 2a but using the pooled regression coefficients instead of the individual program coefficients. The reason why this should allow us to isolate the impact of the program's design is that we keep the estimated coefficients fixed across the different programs but we employ the program's actual maturity distribution of the supply shocks, which depends on the operational details released on the day of the announcement. The results of this exercise for a $100-billion surprise are reported in Table 3.

Table 3: Implied effect on the 10-year yield from an unexpected $100B program using pooled regression's coefficients

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Average
			Impact in basis points			
Total*	-7.7	-6.6	-6.5	-10.6	-10.6	-8.4
of which, duration risk	-4.6	-4.8	-4.7	-7	-6.9	-5.6
of which, local supply	-3.8	-2.4	-2.5	-4.3	-4.3	-3.5

*Includes the estimated constant term

In this case, with the exception of MEP2, the total impact on the 10-year yield from each program is slightly smaller in magnitude than those reported in Table 2a. However, it is very interesting to note the difference between purchase programs focused on the removal of quantities and purchase/sale programs focused on the removal of duration. In particular, while the first type of programs (LSAP1, the Reinvestment, and LSAP2) have a similar impact of about 7 basis points, the second type of programs (MEP and MEP2) have a larger impact of about 10.5 basis points, which is indeed identical across the two programs as they have very similar maturity distribution surprise (see Figure 9 and 10). Considered together, the results in Table 2 and 3 seem to indicate first, that the efficacy of this policy tool, measured by the effect on Treasury yields, has not been diminishing since the announcement of the first LSAP and second, that the design of the program can be as relevant as its size. In other words, asset purchases that remove both quantity and duration from the market, shifting the composition of the Fed's balance sheet toward longer-term maturities, seem more effective than those programs concentrating a larger amount of purchases in the 2- to 10-year maturity sector.

Finally, to better understand the relative importance of the two channels, in Table 4 we show the contribution of each channel to the total variation in yields. As pointed out earlier, the total variation explained by the two shocks is in the range of 75 to 90 percent across the various programs and for ease of comparison the individual total R-squared values are reported in the top row. The last two rows show the variation explained by each shock, which are obtained using the statistical method described in Kruskal (1987) to compute the relative importance of independent variables in multivariate regressions, where the variables used do not have an obvious ordering. His primary suggestion is to average relative importance over all ordering of the independent variables. Therefore, we first compute the proportion of variance in the dependent variable explained by the duration-risk channel and attribute the remaining explained variance to the local-supply channel. Then we switch the order of the independent variables and repeat the same exercise. The final relative contribution of each channel is obtained as average of the explained variances across the two exercises.

Specifically, except for LSAP2 where the local-supply shock accounts for almost the entire explained variation, this channel generally explains between 25 to 50 percent of the yields' reaction; while, the duration-risk shock, although almost irrelevant for LSAP2, on average, accounts for about 30 to 60 percent of the yield variation in the other four programs. Overall, in the pooled results, the two shocks seem to explain a similar portion of the Treasury yield responses to the program announcements, as the local-supply shock explains about 40 percent and the duration-risk shock explains about 30 percent of the total variation.

Table 4: Relative importance of the duration-risk and the local-supply channels

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Pooled
	Two-Day Yield Change Regression					
Total variation explained (R-squared)	0.84	0.69	0.76	0.91	0.76	0.72
of which, duration risk	0.58	0.33	0.01	0.41	0.32	0.29
of which, local supply	0.26	0.36	0.75	0.50	0.44	0.43

To also provide an illustration of the goodness-of-fit of our regressions, Figures 18 to 22 plot the predicted yield reactions, denoted by the black squares, versus the observed yield reactions, denoted by the green triangles. It is striking that these two shocks can capture quite well the highly nonlinear price reactions, replicating patterns very similar to those simulated by Vayanos and Vila (2009). This should not be surprising as our measures of local-supply and duration-risk surprises are trying to approximate the shocks to their theoretical factors and should therefore capture in observed data the same type of humps that their factors generate in simulated data. However, it is easy to note that the fitting errors tend to be larger for securities with higher durations. In the robustness section, we will show that a different choice of the parameters θ and κ_r can improve the fitting at the long-end of the yield curve. This improvement has an intuitive explanation: when optimally chosen, the values of the parameters tend to increase the relative contribution of the local-supply effect, which in turn is very important in explaining the size and shape of the yield reaction at very long maturities. In other words, a larger estimated coefficient for the local-supply channel allows capturing the wide humps observed at longer maturities, as, most likely, the scarcity of substitutes for securities with duration beyond 10 years magnifies the impact of a supply shock localized in these sectors. This explanation is consistent with the theoretical result of Gromb and Vayanos (2010) that in a simple model of cross-asset arbitrage, assets with higher idiosyncratic risk and fewer substitutes are more sensitive to demand shocks.

Overall, based on the reported estimates, there is no evidence that the announcement impact of these two channels has declined since the first LSAP. This result suggests that these two channels may be always operating and are not exceptional mechanisms elicited by the disruption of normal market functioning or the deterioration of market sentiment. Our results also indicate that, for each program, it is not only the unexpected component of the total size but also the unexpected component of the purchases' allocation to each maturity sector that matters, signifying the importance of the program's design and the accompanying communication strategy. However, like all event studies, our results depend on the assumption that Treasury yields responded only to the Fed's purchase announcement and not to other events in the chosen time interval, and do not provide evidence on the persistency of these supply effects.

5.1 Purchase versus sale price elasticity

In the next table, we try to address the following question: can we extrapolate our results to evaluate the potential impact of the 'exit strategy', that is, possible future sales of securities held in the SOMA portfolio for the purpose of tightening monetary policy? One possible way of addressing this question is to test whether the price elasticity in the case of purchases and sales is symmetric. In particular, exploiting the features of the MEP, during which both purchases and sales took place, we estimate different coefficients for the securities included in the purchase and sale sector, respectively. And since there is no particular reason to think that the duration risk coefficient should differ across these two sectors, we continue to run the regression for the full sample and simply use interactive dummies to estimate separate local-supply coefficients for securities purchased and sold under these programs. Further, to account for substitution effects, we extend to the 5-year maturity the threshold that divides the purchase and sale sectors, rather than limiting it to 3 years as indicated in the program operational details.[21]

Table 5: Regression results with different local-supply coefficients for sales and purchases

	MEP	LSAP2	Pooled
Constant	3.6707	-3.1313	-2.0322
	(2.58)	(-5.10)	(-5.917)
Duration risk shock	-2.2719	-0.6038	-1.3817
	(-9.77)	(-1.20)	(-14.80)
Local supply shock, ≤ 5 years	**-1.3964**		**-2.0869**
MEP sales	(-7.23)		(-22.56)
Local supply shock, > 5 years	**-1.4844**		-1.6162
MEP purchase	(-19.00)		(-22.79)
Local supply shock≤ 5 years		-2.006	**-2.0128**
LSAP2 purchases		(-10.80)	(-10.34)
Local supply shock> 5 years		-1.1531	-1.1708
LSAP2 purchases		(-24.06)	(-23.56)
R-squared	0.91	0.78	0.88

Note: t-statistics in parenthesis.

[21] In plotting yield responses against maturities following the MEP announcement, the 5-year maturity is where yields responses change from positive to negative values.

Looking at the numbers in bold in the first column of Table 5, it is striking how similar the estimated price elasticities for sales and purchases are. However, a 1% local-supply shock in the universe of outstanding Treasury securities with less than 5 years to maturity can be quite different in magnitude from a 1% shock in the universe of outstanding securities with more than 5 years to maturity. Therefore, for robustness we repeat the same experiment for the LSAP2 announcement, which, being the immediately preceding program, should not be characterized by significantly different amounts outstanding in the two sectors. Because LSAP2 included only purchase operations, the comparison of the local-supply coefficients across the two programs within the same maturity sector should just provide an idea of the stability of these estimates. As shown by the numbers in bold in the last column, the estimated local supply coefficient in the less-than-five-year sector is very similar across the two programs, suggesting that sales and purchases have similar elasticities.

6. Robustness

In this section we consider various exercises to test the robustness of our results.

6.1 Robustness to the signaling channel at the zero-lower bound.

In this section, we provide some evidence that the signaling channel, differently from the duration-risk and local-supply channels, mainly behaves like a level risk factor (that is, it affects the yields of different maturities by a similar amount), and therefore its impact in our cross-sectional specifications can be properly captured by the constant term, with the exception of the first LSAP. This is because among the FOMC announcements considered in this study, the March 2009 FOMC statement is the only one that contained a qualitative change to the forward rate guidance. In particular, the Committee indicated that "anticipates that economic conditions are likely to warrant exceptionally low levels of the federal funds rate for an extended period" rather than for "some time." Consequently, in this specific case, the LSAP program's impact can be partially confounded with that one from the change in forward guidance.

While forward rate guidance influences most directly investors' expectations of future short-term rates, the signaling channel of QE is an indirect effect of the Fed's balance sheet policy announcements. It arises from perceived new information that these announcements might relay about the state of the economy and the Fed's reaction function. Since this signaling channel, similarly to forward rate guidance, would mainly affect the expectation component of longer-term yields, and since the other two channels (as shown in Section 3) primarily affect the term premium, it is useful to decompose yields in these two components and analyze their changes during these announcements.[22]

To this purpose, we use daily estimates from three different term-structure models: Priebsch (2013) estimates a three-factor Gaussian *shadow-rate* model of the nominal term structure, which is the most reliable for the decomposition of shorter-term yields as it explicitly accounts for zero-lower bound (ZLB); Kim and Wright (2005) estimate a three-factor Gaussian model of the nominal term structure augmented with survey on interest rates; and D'Amico, Kim and Wei (2014) estimate a four-factor Gaussian model of the nominal and real term structure augmented with the inflation process and yields on inflation-indexed Treasury securities. As shown in Kim and Priebsch (2013), yield estimates generated by the Priebsch's model at longer horizons are very similar to those from Kim and Wright (2005) but can differ substantially at shorter-term horizons, which are the most affected by the ZLB.

Since we are interested in analyzing the impact of the signaling channel, we focus only on daily changes in the expectation component. As shown in Table 6, except for LSAP1, the changes in the 2-, 5-, and 10-year yields around the LSAP/MEP announcements are small in magnitude and very similar across maturities, suggesting that the signaling channel behaves *empirically* like a level risk factor during these announcements. This indicates that modelling its impact with the constant term in our cross-sectional estimation, although far from perfect, is a reasonable assumption. In fact, the estimated constant in our

[22] We favor this approach over the use of changes in the OIS rates, because the fixed-leg of OIS is equivalent to holding a long position on a bond with the same maturity of the OIS contract, therefore these rates contain a term premium very similar to those on Treasury securities. This suggests that at longer horizons they are an unreliable proxy for the expectation component of interest rates.

most robust specification (reported in Table 7) is quite close to the average change across maturities in the expectation component generated by the ZLB model, with an average magnitude of -1 basis point.[23]

Table 6: Daily Change in the Model-Implied Expectation Component of Treasury Yields

	Priebsch	ZLB		Kim-Wright		DKW	
	2-year	5-year	10-year	5-year	10-year	5-year	10-year
LSAP1	-15	-13	-10	-13	-16	-17	-18
Reinvestme	-1	-2	-2	-2	-3	-2	-3
LSAP2	-1	-2	-2	-1	-1	-1	-1
MEP	0	-2	-2	1	-2	1	-1
MEP2	1	1	1	1	1	2	1

Note: Changes are from COB of day before the FOMC to COB of the FOMC day.

The first LSAP, however, is a clear exception. As shown in the first row of the ZLB model, the change in the expectation component is fairly large and decreasing across maturities. This pattern is consistent with the results from all the models estimated in Bauer and Rudebusch (forthcoming) for the signaling channel of LSAP1, which provide evidence of a single-hump-shaped reaction in the expectation component of nominal Treasury yields, with the largest decline occurring around the 3-year horizon.[24] But, this is also the pattern that we would expect to observe when the impact on the expectation component is driven by a change to forward rate guidance rather than the QE signaling channel. Indeed, a single-hump-shaped yield reaction, with the largest decline at the 3- to 5-year maturities, has also been evident following the two FOMC announcements that introduced and extended the calendar-based forward rate guidance on August 9, 2011, and January 25, 2012, respectively.[25] Therefore, for the LSAP1

[23] The average one-day change in the expectation component across maturities for the Reinvestment program, LSAP2, MEP and MEP2 is -1.7, -1.7, -1.3, and 1, respectively. The average two-day change for the same announcements is -1.17, -1.7, -1.3, and 0.33, respectively.

[24] The magnitude of the decline is quite different, most likely because the models in Bauer and Rudebusch (forthcoming) do not account for the ZLB.

[25] The changes in the 2-, 5-, and 10-year Treasury yields from 15 minutes before the FOMC announcement to the end of day were: -7, -14, and -6 basis points for the August 2011 meeting; -2, -5, and -2 basis points for the January

announcement, it is extremely hard to disentangle the effect of forward rate guidance from the QE signaling channel.

The single-hump shape of the yield reaction to both changes in forward guidance and the QE signaling channel is distinct from the monotonic and concave shape of the duration-risk effect and the multiple-hump shape of the local-supply effect across maturities. This difference in the functional forms suggests a low correlation among the three channels' impact across yields. Therefore, although the regression specification for LSAP1 may suffer of an omitted variable problem because the constant term does not properly capture the signaling channel, the bias on the estimated coefficients should be small due to the low correlation between the omitted variable and the included explanatory variables. Consequently, for LSAP1, the estimates for the duration-risk and local-supply channels should be considered an upper bound of their effects.

6.2 Robustness to the choice of the parameters θ and κ_r

This section considers the robustness of our results to the choice of the parameter θ, which controls the window size in the computation of the individual local-supply shocks illustrated in Section 4.1, and to the choice of κ_r, which controls the steepness of the curve in the computation of the individual duration-risk shocks illustrated in Section 4.2. We compute optimal values for each parameter by jointly minimizing with respect to θ and κ_r the sum of squared residuals (or alternatively by maximizing the R-squared) in the pooled regression. The optimization surface is shown in Figure 23, where R-squared values are plotted as function of different values of θ and κ_r. It is possible to see that given the optimal value for θ, the R-squared is not very sensitive to the changes in κ_r. And on the other hand, for any choice of κ_r, the R-squared is quite sensitive to the changes in θ . Our results indicate that the optimal values of θ and κ_r are 0.769 and 0.095, respectively. This would suggest that if we use a broader concept of substitutability, that is, if the variable window size is extended to include all securities within a maturity distance of about

2012 meeting. The calendar-based forward guidance was further extended in the September 2012 statement, but at that meeting a new MBS purchase program was announced, making it harder to isolate the impact of the forward guidance extension.

77 percent of security *i*'s maturity, and if we choose a steeper concave function to measure the individual exposure to the aggregate duration risk, then we should be able to fit the yield reactions better. The regression results obtained using the optimal values for these parameters are reported in Table 7.

Table 7: Yield change regression results with variable window size, θ=0.769 and κ_r=0.095

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Pooled
			Two-day yield change regression			
Constant	-2.551	-1.399	-3.177	1.847	0.196	-0.624
	(-3.31)	(-4.27)	(-14.52)	(1.91)	(0.61)	(-2.44)
Duration risk shock	-1.375	1.638	0.065	-1.571	-0.314	-0.992
	(-17.37)	(6.97)	(0.56)	(-12.08)	(-5.01)	(-20.81)
Local supply shock	-0.680	-4.746	-2.003	-1.542	-0.539	-1.043
	(-18.61)	(-15.49)	(-68.93)	(-13.35)	(-11.69)	(-40.58)
R-squared	0.85	0.76	0.96	0.94	0.87	0.79
Observations	163	200	208	232	245	1048

Note: t-statistics in parenthesis. It is important to note that these t-statistics do not take into account the uncertainty about the parameters θ and κ_r.

Table 8: Relative importance of the duration-risk and local-supply channels, θ=0.769 and κ_r=0.095

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Pooled
			Two-Day Yield Change Regression			
Total variation explained (R-squared)	0.85	0.76	0.96	0.94	0.87	0.79
of which, duration risk	0.40	0.26	0.04	0.47	0.41	0.27
of which, local supply	0.48	0.49	0.92	0.48	0.46	0.52

As reported in Tables 7 and 8, when we use the optimized parameters, the adjusted R-squared for all individual programs and the pooled regression increase somewhat, and the variation explained by the local-supply channel becomes notably larger for LSAP1, LSAP2 and the Reinvestment program, as well as in the pooled specification. On the other hand, the variation explained by the duration-risk channel increases for the MEP and MEP2, which intuitively makes sense considering that these programs were designed to remove a significant amount of duration risk from the market and therefore should be characterized by a larger duration- risk effect on Treasury yields. This implies that, while the local-supply effect is very important in capturing these price reactions during the first three announcements, the duration-risk effect is crucial in approximating better the price reaction to the MEP and MEP2 announcements. As shown in figures 24 to 28, which plot the predicted yield reactions, denoted by the black squares, versus the observed yield reaction, denoted by the green triangles, using the optimized parameters, the fitting errors for longer-duration securities are substantially smaller than in the baseline case.

In addition, using these parameters' values, as shown in the second and third columns of Table 7, the estimated duration-risk coefficient becomes positive and significant in the case of the Reinvestment program and positive but not statistically significant for LSAP2. One possible explanation is that the optimal parameter values obtained using the pooled data specification are not very close to the optimal values for these two particular programs. In contrast, all the coefficients for the local-supply effect stay negative and significant, which seems to suggest that these estimates are more stable and less sensitive to the parameters' choice.

Further, in Table 9, we report the implied effect on the 10-year yield in basis points from an unexpected $100 billion purchase announcement, obtained using the new estimated coefficients (with optimized values of θ and κ_r) for the individual programs. Under this parameter choice, the average impact of the local-supply channel is somewhat larger, although the overall impact of a $100-billion surprise remains about unchanged relative to those reported in Table 2a. Finally, similarly to the baseline regressions, we repeat the same exercise using the pooled specification coefficients, and the resulting effects on the 10-year yield are reported in Table 10. Also in this case, the total impacts are quite close to those implied by the baseline regressions and reported in Table 3. However, the total variation explained by the local-supply effect becomes relatively larger in the first three programs, and on average, each channel accounts for about half of the total yield decline.

It is also interesting to note that the estimated constant term, being slightly bigger in this specification, causes a wedge between the total impact and the sum two channels, with an average magnitude on -1 basis point, which as stressed in the previous section may be capturing the impact of the signaling channel. This also indicates that in our best specification (based on the higher R-squared), a $100-billion surprise translates in a 8-basis-point reduction in the 10-year Treasury yield due to the term-premium channels and a 1-basis-point decline in the same yield due to the signaling channel.

Table 9: Implied effect on 10-year yield from an unexpected $100B program using individual coefficients, θ=0.769 and κ_r=0.095

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Average
			Impact in basis points			
Total*	-10.8	-10.2	-9	-12.3	-3.3	-9.1
of which, duration risk	-5	6.2	0.2	-8.8	-1.7	-1.8
of which, local supply	-3.2	-15	-6	-5.4	-1.8	-6.3

*****Includes the estimated constant term.

Table 10: Implied effect on 10-year yield from an unexpected $100B program using pooled regression coefficients, θ=0.769 and κ_r=0.095

	LSAP1	Reinvestment	LSAP2	MEP	MEP2	Average
			Impact in basis points			
Total*	-9.1	-7.7	-7.4	-9.8	-9.6	-8.7
of which, duration risk	-3.6	-3.7	-3.7	-5.6	-5.5	-4.4
of which, local supply	-4.9	-3.3	-3.1	-3.6	-3.4	-3.7

*****Includes the estimated constant term.

Compared to previous studies, our total estimated effect on the 10-year Treasury yield in most specifications is quite similar to that obtained by Li and Wei (2013), as their results imply an average impact of about 7 basis points per $100 billion of 10-year equivalents; and considering that they do not explicitly account for the local supply channel, it is not surprising that in some specifications our estimates can be a bit larger in magnitude. On the other hand, our estimates of the local-supply effect in some specifications are just slightly smaller than those reported in D'Amico and King (2013), which imply that in the 10-year sector the total impact of $300 billion of purchases is about 15 basis points (i.e., about 5 basis points per $100 billion of purchases). However, in that study, the authors estimate the total stock effect from the day before the announcement of the first LSAP to the day of its last purchase operation; therefore, their estimates do not capture exclusively the announcement effect as in the case of this study.

6.3 Robustness to LSAP1 pre- and post-announcement maturity bucket weights

Considering the novelty of the first LSAP, that is, at the time there was not any empirical evidence on the maturity distribution of previous purchases, necessarily the choice of the pre-announcement maturity bucket weights has to be somewhat arbitrary. However, since in this instance the FOMC (2:15 p.m.) and the Desk (2:44 p.m.) announcements took place at two different times, we can use changes over this half-hour interval (between 2:15 and 2:44 p.m.) in market expectations about the maturity distribution of purchases to identify the local-supply and duration risk effects, without relying on any measure of the shocks.

In particular, given the pattern of the price reversal, shown in Figure 1, it seems safe to assume that most of the yield increase that took place between 2:44 p.m. and 4:00 p.m. was due to the local-supply effect, as it occurred only in the maturity sector that was excluded from the purchase program and seemed largely independent of the security's duration, judging from the parallel shift across different durations. Therefore, for each security in the excluded maturity sector we compute the yield increase between 2:44 and 4:00 p.m., and we find that these securities on average experienced a reversal of about 40 percent of the yield decline that had come soon after the FOMC announcement (i.e., from 2:15 to 2:44 p.m.), which in contrast should be driven by both channels. This implies that most likely about 40 percent of the total yield change was due to the local supply effect, which is larger than the variation explained by the local supply shock for LSAP1 in our baseline regression shown in Table 4.

In the second robustness exercise, also related to the choice of weights for the first program, we recomputed the surprise for each maturity bucket and consequently for each individual security under the assumption that market participants form expectations only about the total size of the program and are totally agnostic about the weights. Then, given the actual amount, we derive the total size shock and its distribution across the different maturity sectors based on the announced weights (that is, we use only the post-announcement weights). The regression results of this experiment are shown in the first column of Table 11, and since the new assumptions affect also the pooled regression results, the numbers reported in the second column are different from those shown in Table 1.

Table 11: Yield change regression results using alternative weights for LSAP1

	LSAP1	Pooled
Constant	0.5594	-1.269
	(0.683)	(-4.52)
Duration risk shock	-1.8041	-0.9901
	(-13.14)	(-13.12)
Local supply shock	-0.8031	-1.1183
	(-14.63)	(-40.20)
Observations	163	1048

Note: t-statistics in parenthesis.

Table 12: Relative importance of the duration-risk and local-supply channels using alternative weights for LSAP1

	LSAP1	Pooled
Total variation explained (R-squared)	0.87	0.79
of which, duration risk	0.42	0.25
of which, local supply	0.45	0.54

As shown in Table 12, with this weight specification, the R-squared for both LSAP1 and the pooled regressions improved slightly relative to our baseline regressions, and the variation explained by the local-supply shock is closer to the that one obtained by exploiting only the price reversal as just described above, which indicated that about 40 percent of the yield variation was due to the local-supply effect.

7. Conclusions

In summary, we find that, once the pre-announcement market expectations about the Fed's asset purchase announcements are carefully controlled for, both the duration-risk channel and the local-supply channel are always operating in our sample period, and are about equally important in explaining yield reactions of nominal Treasury securities to these announcements. This result suggests that it is not only the total size of the program but also its design that matters, as the latter determines the maturity composition of purchases/sales. It also signifies the importance of the FOMC's communication strategy, as it can strongly influence all three components—the size, the total dollar duration, and the location—of the supply shocks and the resulting Treasury yield responses.

Taken together, our findings indicate that these channels' impact did not decline across consecutive programs, implying that they may be key factors in the determination of Treasury securities prices and should be included in the widely-used asset pricing models. More importantly, these findings also confirm that the term premium component of interest rates and the channels impacting this component are important for the transmission mechanism of unconventional monetary policy, suggesting that they should be incorporated in traditional macro/monetary economic models, where usually monetary policy operates only through the expectation component of interest rates.

APPENDIX A: Computation of the Expected Total Size

The first LSAP (LSAP1) announcement in March 2009: In the PDS before the meeting, respondents indicated that they attached 49 percent probability to the Federal Reserve announcing purchases of long-term Treasury securities, 62 percent probability to an expansion of agency debt purchases, and 69 percent probability to an expansion of agency MBS purchases. The survey did not ask the sizes of the programs; we therefore examined the primary dealers' written comments to each answer for clues about the expected size of the purchases. Among the few who provided forecasts for the combined size of the program, none cited a number above $600 billion, which is far below the $1.15 trillion that was announced. A conservative measure of the surprise can therefore be calculated by multiplying the actual announced purchases of Treasury securities, agency debt, and agency MBS by the corresponding probability and sum up. This calculation implicitly treats all three types of securities as perfect substitutes.

Reinvestment policy announcement in August 2010: In the PDS, the respondents assigned 43 percent probability to the Fed starting to reinvest the principal payments from agency securities. Determining the expected size of the program requires us to choose a value for the expected cumulative agency MBS principal paydowns. For this we relied on the average across various projections available in market commentaries at the time, which estimated that roughly $200 billion would have been paid down over the next 6 months.

The second LSAP (LSAP2) announcement in November 2010: Primary Dealers assigned an average probability of 88 percent to the FOMC announcing at the upcoming meeting an expansion of its portfolio through additional asset purchases. They also on average indicated that the total size of the program would have been about $1 trillion over a 12-month horizon. We rescale the expected amount by the actual length of the announced program of about 7 months, assuming a constant monthly pace.

MEP announcement in September 2011: In the PDS, respondents assigned 73 percent probability to the FOMC announcing an increase in the average duration of SOMA holdings as an easing tool within 1 year. They also indicated that the program was expected to be about $376 billion over an almost 7-month horizon, with purchases occurring in the 7- to 30-year maturity sector and sales in the 0- to 3-year maturity sector. Given this detailed information, we are able to compute the surprise without relying on any strong assumption. In the statement from that meeting, the FOMC also announced it would begin reinvesting principal payments from its holdings of agency debt and agency MBS in agency MBS instead of Treasuries, which reportedly came as a surprise to the market. The effect of this surprise was likely to be small and we ignore it in the current analysis.

MEP extension announcement in June 2012: Primary dealers assigned about 54 percent probability to the FOMC easing at the upcoming meeting utilizing the size or composition of its portfolio. To estimate the size of the extension, investors appeared to look at the amount of Treasury securities in SOMA holdings with maturities that were expected to fall below 3 years over the next few months, which over a 6-month period totaled about $205 billion at the time. It is therefore safe to conjecture that, had they expected the MEP to be extended in its original form for 6 months, the expected size of the program was likely $205 billion, compared to the actual announced amount of $267 billion. Also in this case, we do not need to rely on strong assumptions and simply calculate the expected size of the program by multiplying the perceived probability of an MEP extension by the expected size of the extension conditional on the program being announced.

APPENDIX B: Computation of the Maturity Distribution of Purchases

The post-announcement maturity weights are calculated as follows. The Desk's technical notes following the LSAP1 and the reinvestment program announcements were very similar and specified only that the purchases would be concentrated in the 2- to 10-year sector. For LSAP1, we therefore assume that investors recalibrated their post-announcements weights in the following way: they understood the language to mean about 80% of purchases would occur in the 2- to 10-year sector, and weighted securities in this sector proportionally to their amounts outstanding as a fraction of the total outstanding in the 2- to 10-year sector only. Similarly, we assign the remaining 20% to the securities outside the eligible maturity sector proportionally to the amounts outstanding in the excluded maturity ranges. For the reinvestment program announcement, we assume that, based on their previous experience with purchases allocation under LSAP1, they continued to assign weights identical to those associated with the actual LSAP1 purchases; which implies that the pre- and post-announcements weights are identical for this announcement. Starting with LSAP2, the Desk would release statements simultaneously with the FOMC announcements outlining the intended weights for each of the maturity buckets, which are used as the post-announcement weights for those programs. The following tables show the pre- and post-announcements weights for each program.

LSAP1

	1.5-2.5 years	2.5-4 years	4-5.5 years	5.5-7 years	7-10 years	10-17 years	17-30 years
LSAP1 expected*	17.4%	20%	18%	8.5%	17.8%	8%	10.3%
LSAP1 announced**	9.7%	25%	22.3%	10.6%	22.1%	4.5%	5.8%

* The expected weights are assumed to be proportional to the percentage outstanding in each bucket.
** The actual weights have been recalibrated to reflect the Desk announcement specifying that purchases were going to be concentrated in the 2- to 10-year sector, which we interpret as 80% allocation to the 2- to 10-year sector.

Reinvestment Program

	1.5-2.5 years	2.5-4 years	4-5.5 years	5.5-7 years	7-10 years	10-17 years	17-30 years
Reinvestment expected*	8.9%	24.5%	19.7%	20.3%	11.6%	8.2%	6.2%
Reinvestment announced**	8.9%	24.5%	19.7%	20.3%	11.6%	8.2%	6.2%

* The expected weights are assumed to be equal to the actual purchases' maturity distribution observed for LSAP1.
** As at that time there was no announcement about the weights, the announced weights are assumed to be equal to the expected, as following both LSAP1 and Reinvestment program announcements the Desk specified that purchases were going to be concentrated in the 2- to 10-year sector.

LSAP2	1.5-2.5 years	2.5-4 years	4-5.5 years	5.5-7 years	7-10 years	10-17 years	17-30 years
LSAP2 expected	9%	24%	20%	20%	11%	9%	7%
LSAP2 announced	5%	20%	20%	23%	23%	2%	4%

Notes: The expected LSAP2 maturity distribution is equal to the Reinvestment program actual maturity distribution of the final purchases except for rounding.

MEP	6-8 years	8-10 years	10-20 years	20-30 years
MEP expected	44.2%	44.2%	3.8%	7.7%
MEP announced	32%	32%	4%	29%

Notes: The expected weights for the MEP have been obtained by redistributing the percentage amount previously purchased (and as such observed by the market) in the 1.5- to 6- year sector over the 6- to 30-year sector proportionally to the percentage of the sector.

MEP Extension	6-8 years	8-10 years	10-20 years	20-30 years
MEP extension expected	32%	32%	4%	29%
MEP extension announced	32%	32%	4%	29%

Note: The maturity distribution expected for the MEP extension is assumed to be equal to that one announced for the original MEP.

References

Amihud, Y., Mendelson, H., 1991. Liquidity, maturity, and the yields on US Treasury securities. Journal of Finance 46, 1411–1425.

Andres, J., Lopez-Salido, J. D., Nelson, E., 2004. Tobin's imperfect asset substitution in optimizing general equilibrium. Journal of Money, Credit and Banking 36, 665–690.

Banerjee, R., Daros, S., Latto, D., McLaren, N., 2012. Using changes in auction maturity sectors to help identify the impact of QE on gilt yields. Quarterly Bulletin 2012 Q2, Bank of England.

Bauer, M. D., Rudebusch G.D. (Forthcoming). The signaling channel for Federal Reserve Bond Purchases. *International Journal of Central Banking.*

Bauer, M. D., Rudebusch G.D. 2014. Monetary Policy Expectations at the Zero Lower Bound. Working Paper 2013-18, Federal Reserve of San Francisco.

Bernanke, Ben S., 2013. "Communication and Monetary Policy." Remarks before the National Economists Club, Washington, D.C., November 19.

Bernanke, B. S., Reinhart, V. R., Sack, B. P., 2004. Monetary policy alternatives at the zero bound: an empirical assessment. Brookings Papers on Economic Activity 2, 1–78.

Brandt, M.W., Kavajecz, K. A., 2004. Price discovery in the US Treasury market: the impact of order flow and liquidity on the yield curve. Journal of Finance 59, 2623–2654.

Cochrane, J. H., 2008. Comments on 'Bond supply and excess bond returns' by Robin Greenwood and Dimitri Vayanos. May 12, presentation at the Initiative in Global Financial Markets Beyond Liquidity Conference, University of Chicago.

Cox, J. C., Ingersoll, J. E., Ross, S. A., 1985. A theory of the term structure of interest rates. Econometrica 53, 385–407.

Culbertson, J. M., 1957. The term structure of interest rates. Quarterly Journal of Economics 71, 485–517.

D'Amico, S., English, W., López-Salido, D., Nelson, E., 2012. The Federal Reserve's Large-Scale Asset Purchase Programs: Rationale and Effects. Economic Journal 122, November, 415-446.

D'Amico, S., King, T.B., 2013. Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply. Journal of Financial Economics 108, 425-448.

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

Engen, E., Hubbard, R. G., 2005. Federal government debt and interest rates. In: Gertler. M., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2004. MIT Press, Cambridge, MA, pp. 83-160.

Gabaix, X., Krishnamurthy, A., Vigneron, O., 2007. Limits of arbitrage: theory and evidence from the mortgage-backed securities market. Journal of Finance 62, 557–595.

Gagnon, J., Raskin, M., Remache, J., Sack, B., 2011. Large-scale asset purchases by the Federal Reserve: did they work? International Journal of Central Banking 7 (1): 3-43.

Garleanu, N., Pedersen, L. H., Poteshman, A. M., 2009. Demand-based option pricing. Review of Financial Studies 22, 4259–4299.

Greenwood, R., 2005. Short- and long-term demand curves for stocks: theory and evidence on the dynamics of arbitrage. Journal of Financial Economics 75, 607–649.

Greenwood, R., Vayanos, D., 2010a. Price pressure in the government bond market. American Economic Review 100, 585–590.

Greenwood, R., Vayanos, D., 2014. Bond supply and excess bond returns. Review of Financial Studies 27, 663-713.

Gromb, D., Vayanos, D., 2010. Limits of arbitrage: the state of the theory. Annual Review of Financial Economics 2, 251–275.

Gurkaynak, R., Wright, J., 2012. Macroeconomics and the term structure. Journal of Economic Literature 50, 331–367.

Hamilton, J., Wu, J., 2012. The effectiveness of alternative monetary policy tools in a zero lower bound environment. Journal of Money, Credit and Banking 44, 3–46.

Han, B., Longstaff, F. A., Merrill, C., 2007. The US Treasury buyback auctions: the cost of retiring illiquid bonds. Journal of Finance 62, 2673–2693.

Holland, T. E., 1969. 'Operation Twist' and the movement of interest rates and related economic time series. International Economic Review 10, 260–265.

Hu, X., Pan, J., J., Wang., 2012. Noise as information for illiquidity. NBER Working Paper 16468.

Jordan, B. D., Jordan, S. D., 1997. Special repo rates: an empirical analysis. Journal of Finance 52, 2051–2072.

Joyce, M. A. S., Tong, M., 2012. QE and the Gilt market: a disaggregated Analysis. Economic Journal 122, November, 348-384.

Kaminska, I., Vayanos, D., Zinna, G., 2011. Preferred-habitat investors and the US term structure of real rates. Bank of England Working Paper 435.

Kim, H. Kim, Priebsch, Marcel A., 2013. Estimation of Multi-Factor Shadow-Rate Term Structure Models. Working Paper. Mimeo.

King, T. B., 2013. A portfolio-balance approach to the nominal term structure. Federal Reserve Bank of Chicago Working Paper 2013-18.

Kohn, Donald L., 2009. "Monetary Policy Research and the Financial Crisis: Strengths and Shortcomings." Speech at the Conference on Key Developments in Monetary Economics, Washington, D.C., October 9.

Krishnamurthy, A., 2002. The bond/old-bond spread. Journal of Financial Economics 66, 463–506.

Krishnamurthy, A., Vissing-Jorgensen, A., 2011. The effects of quantitative easing on interest rates: channels and implications for policy. Brookings Papers on Economic Activity 2, 215-87.

Krishnamurthy, A., Vissing-Jorgensen, A., 2012. The aggregate demand for Treasury debt. Journal of Political Economy 120 (1), 233-67.

Kruskal, William, 1987. Relative importance by averaging over orderings. The American Statistician, February, vol. 41, 6-10.

Kuttner, K., 2006. Can central banks target bond prices? NBER Working Paper 12454.

Li, C., Wei, M., 2013. Term Structure Modelling with Supply Factors and the Federal Reserve's Large Scale Asset Purchase Programs. International Journal of Central Banking, March 2013.

Longstaff, F. A., 2004. The flight-to-liquidity premium in US Treasury bond prices. Journal of Business 77, 511–526.

Lou, D., Yan, H., Zhang, J., 2010. Anticipated and repeated shocks in liquid markets. Unpublished working paper. London School of Economics and Yale University.

Meaning, J., Zhu, F., 2011. The impact of recent central bank asset purchase programs. BIS Quarterly Review, December, 73–83.

Modigliani, F., Sutch, R., 1966. Innovations in interest rate policy. American Economic Review 56, 178–197.

Musto, D., Nini, G., Schwarz, K., 2011. Notes on bonds: liquidity at all costs in the great recession. Unpublished working paper, University of Pennsylvania.

Priebsch, Marcel A., 2013. Computing Arbitrage-Free Yields in Multi-Factor Gaussian Shadow-Rate Term Structure Models. Finance and Economics Discussion Series 2013-63. Board of Governors of the Federal Reserve System.

Rosa, Carlo, 2012. How "unconventional" are large-scale asset purchases? The impact of monetary policy on asset prices. Federal Reserve Bank of New York Staff Reports, 560, May.

Ross, M., 1966. 'Operation Twist': a mistaken policy? Journal of Political Economy 74, 195–199.

Svensson, L. E. O., 1994. Estimating and interpreting forward interest rates: Sweden 1992–1994. 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. Brookings Papers on Economic Activity 42, 151–188.

Vayanos, D., Vila, J-L., 2009. A preferred-habitat model of the term structure of interest rates. NBER Working Paper 15487.

Williams, John L., 2011. "Unconventional Monetary Policy: Lessons from the Past Three Years." Presentation to the Swiss National Bank Research Conference, Zurich, September 23.

Wright, Jonathan H., 2012. What does Monetary Policy do to Long-term Interest Rates at the Zero Lower Bound? Economic Journal 122, 447-466.

Yellen, Janet, 2011. "Unconventional Monetary Policy and Central Bank Communications." Speech at the U.S. Monetary Policy Forum, New York, February 25.

Figure 1. LSAP 1 Announcement

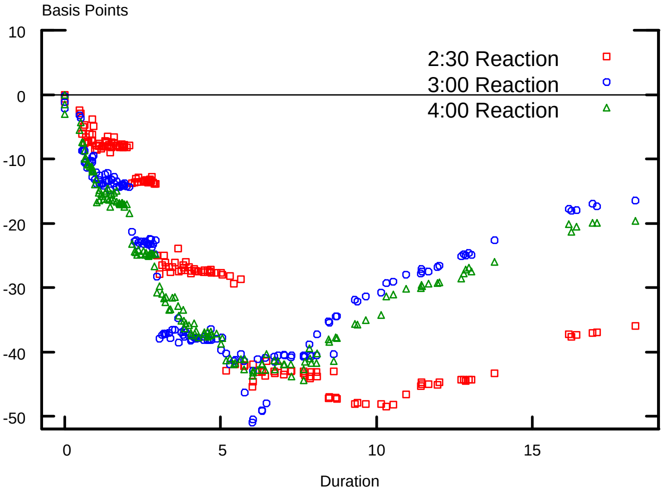

Figure 2. Reinvestment Announcement

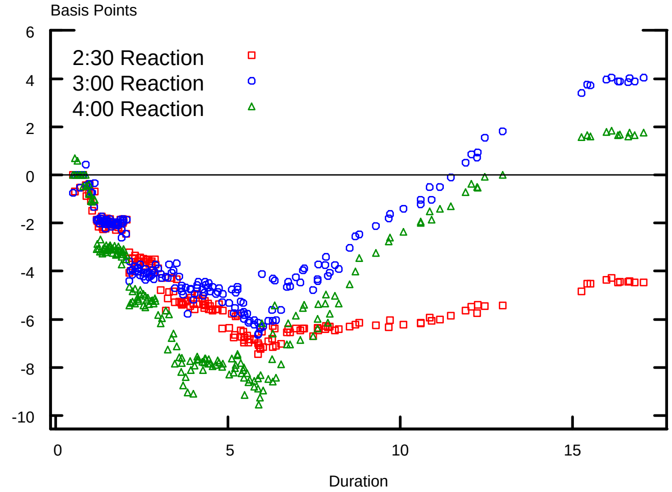

Figure 3. LSAP 2 Announcement

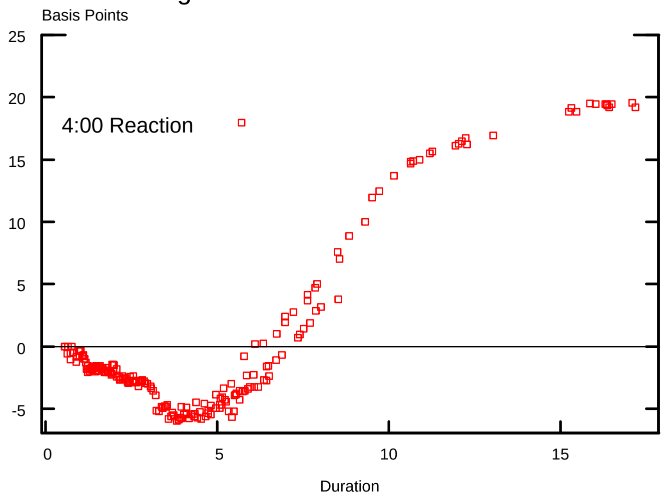

Figure 4. MEP Announcement

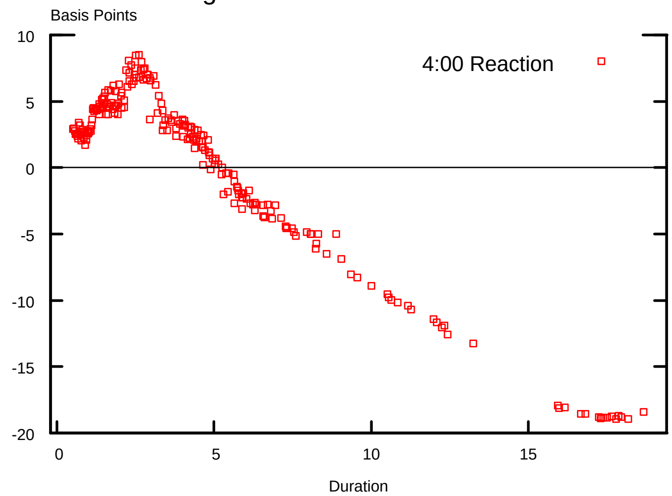

Figure 5. MEP Extension Announcement

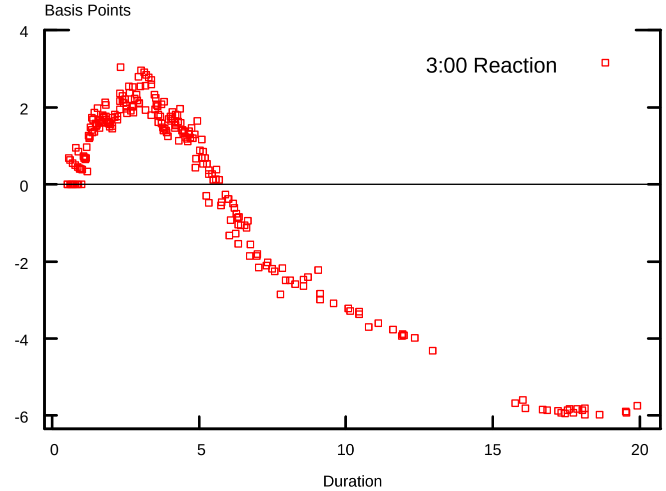

Figure 6. LSAP1 Announcement

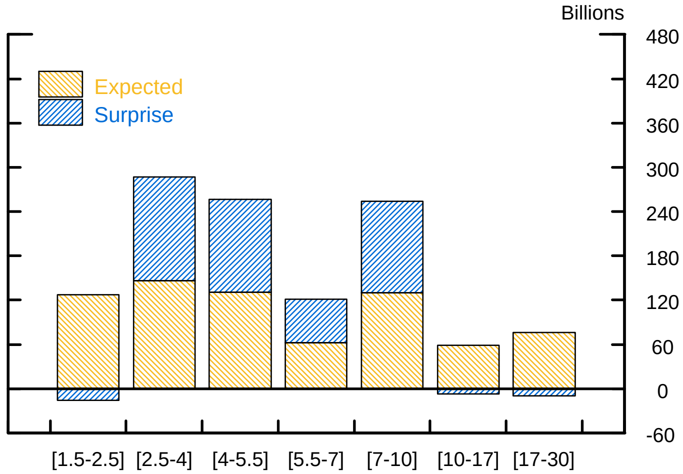

Figure 7. Reinvestment Announcement

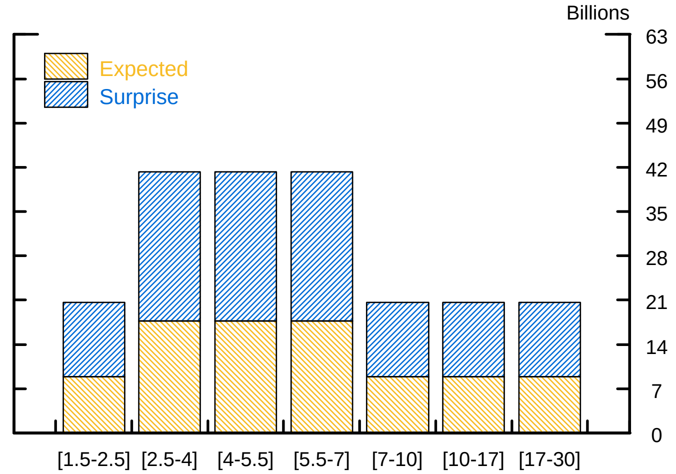

Figure 8. LSAP2 Announcement

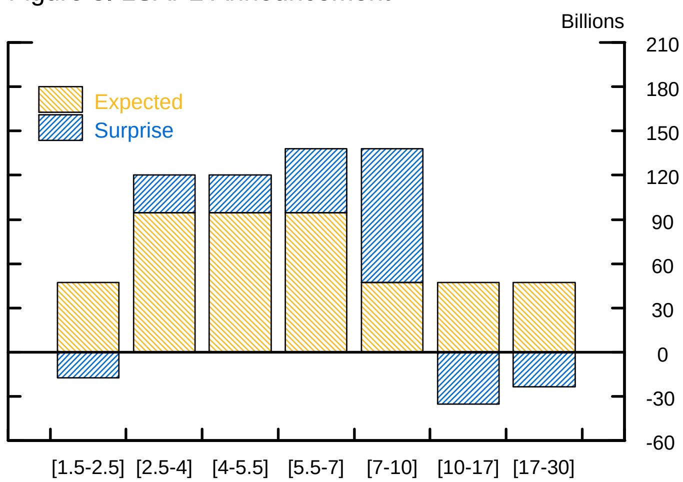

Figure 9. MEP Announcement

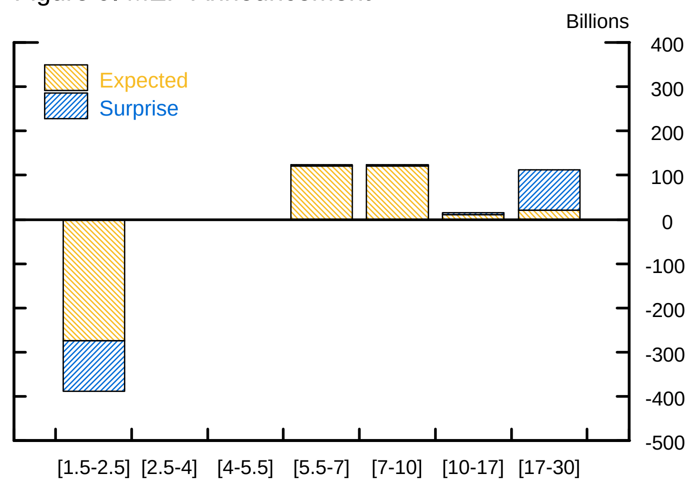

Figure 10. MEP Extension Announcement

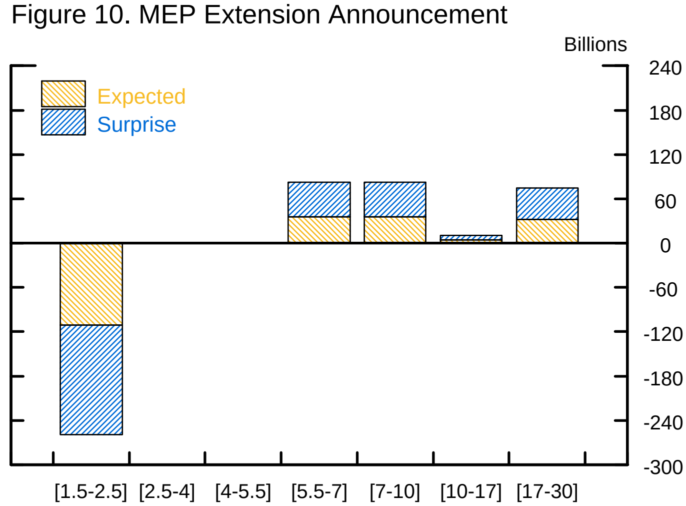

Figure 11

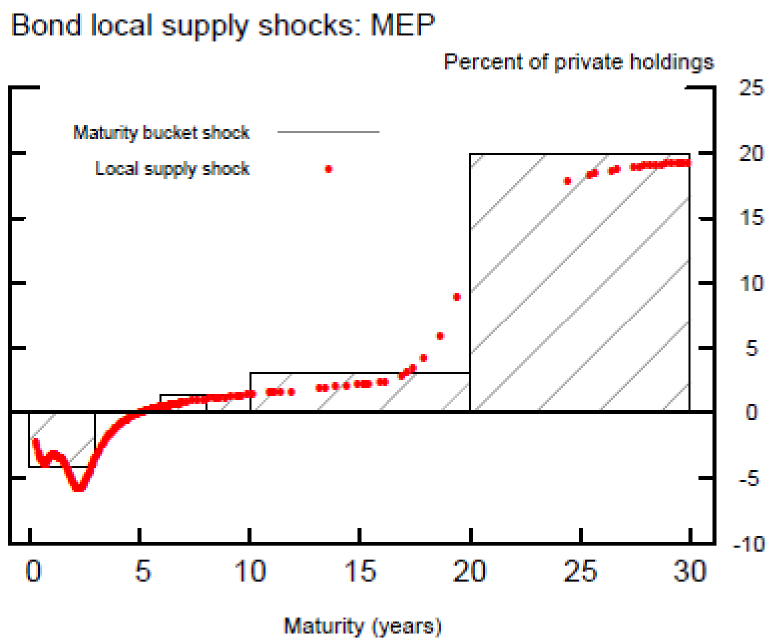

Figure 12

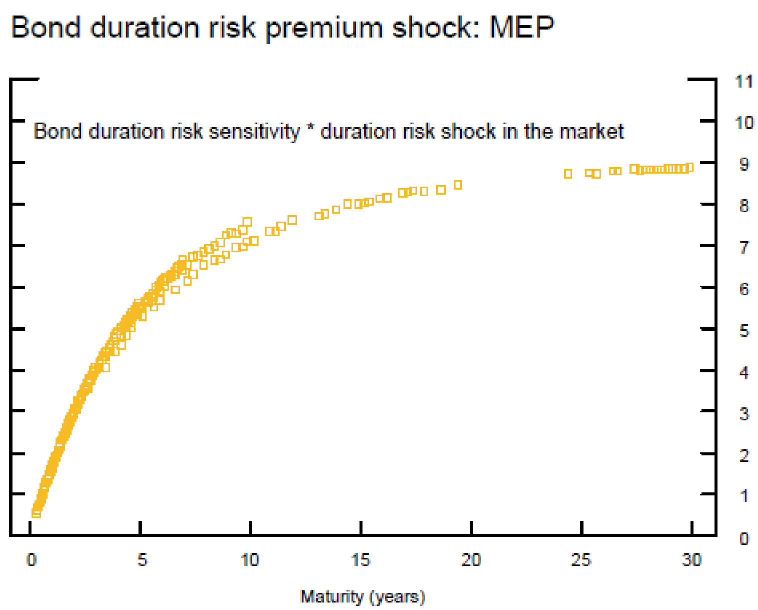

Note: Bond duration risk sensitivity is [1 - exp(- γ * d)]/ γ where d is bond duration and the parameter γ is set as 0.2 according to Li and Wei (2012).

Figure 13. LSAP 1 Announcement

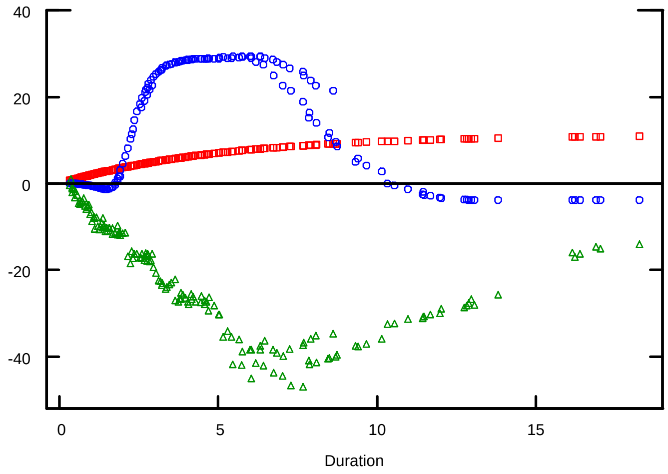

Figure 14. Reinvestment Announcement

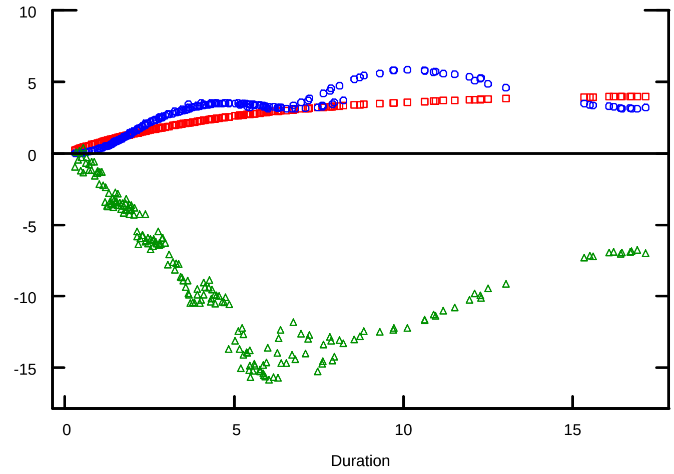

Figure 15. LSAP2 Announcement

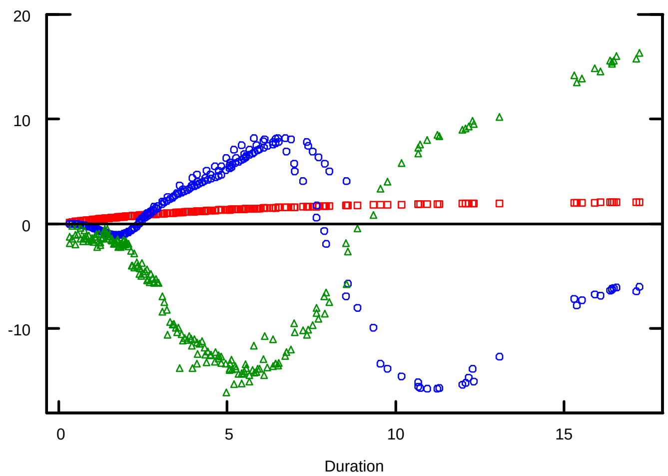

Figure 16. MEP Announcement

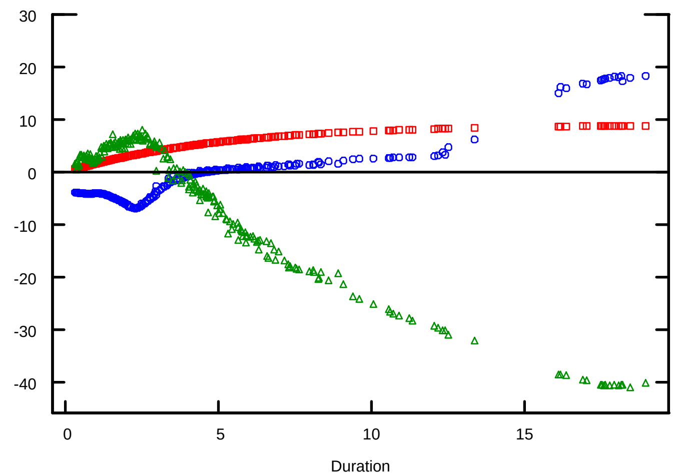

Figure 17. MEP Extension Announcement

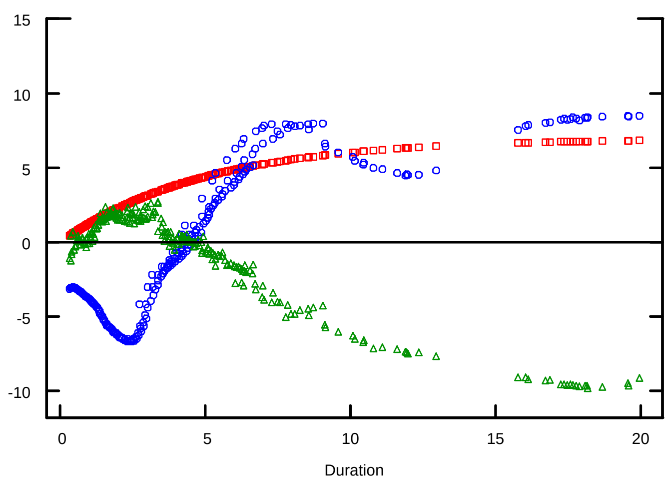

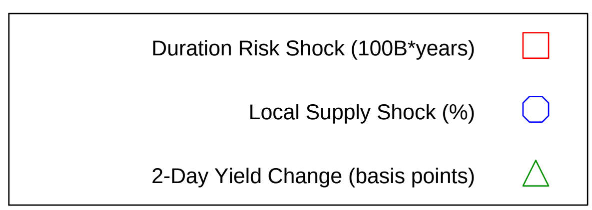

Figure 18. LSAP 1 Announcement

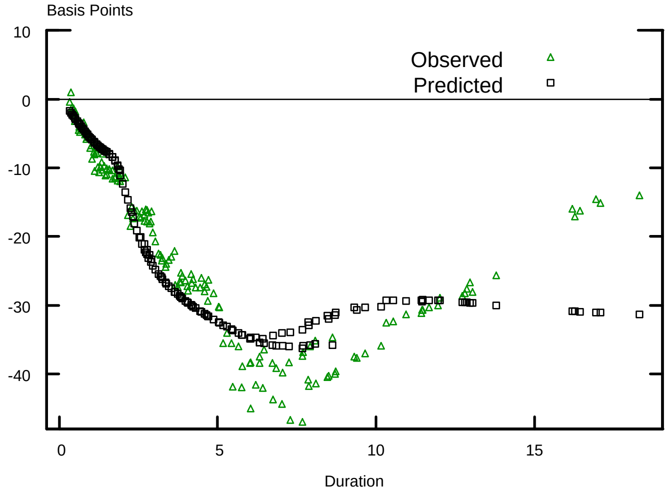

Figure 19. Reinvestment Announcement

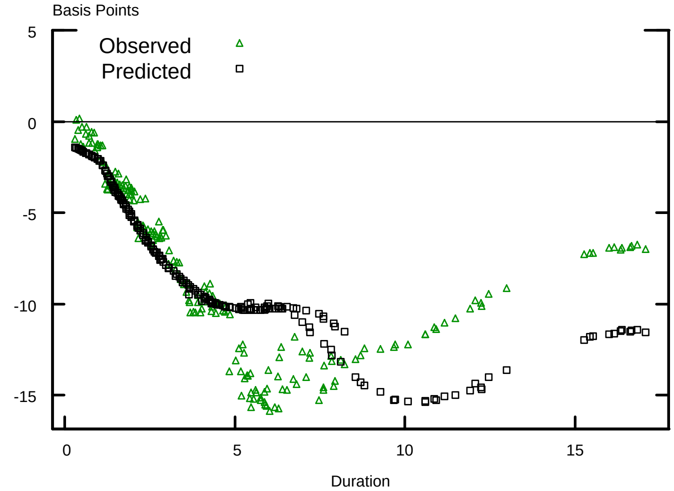

Figure 20. LSAP 2 Announcement

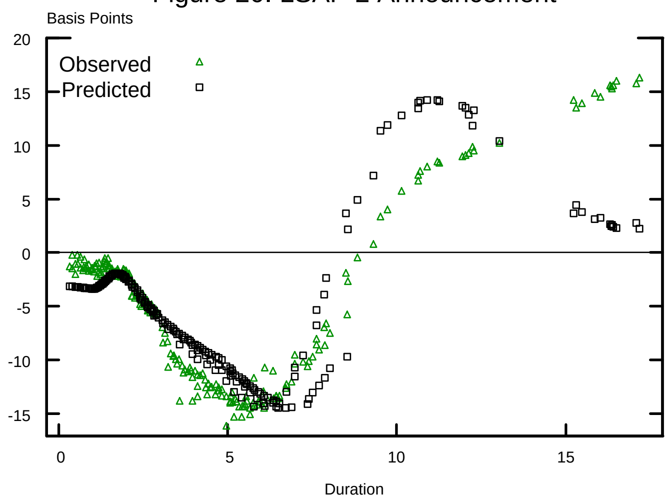

Figure 21. MEP Announcement

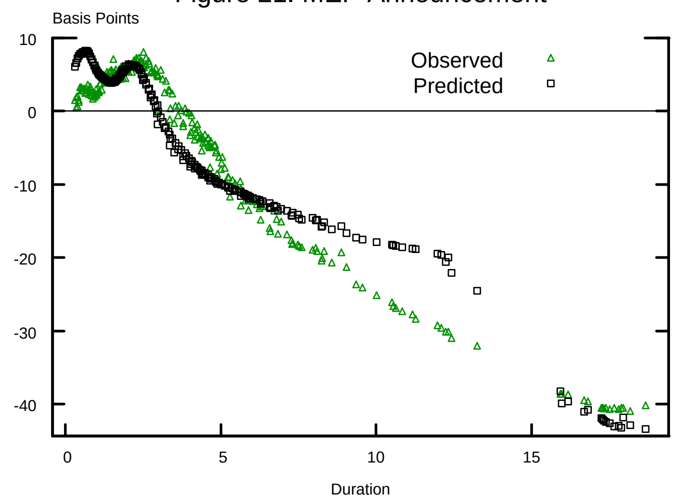

Figure 22. MEP Extension Announcement

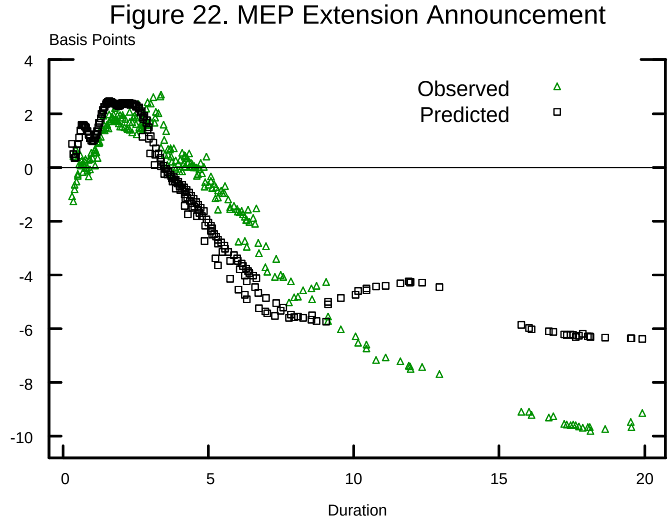

Figure 23. Parameter Optimization Surface

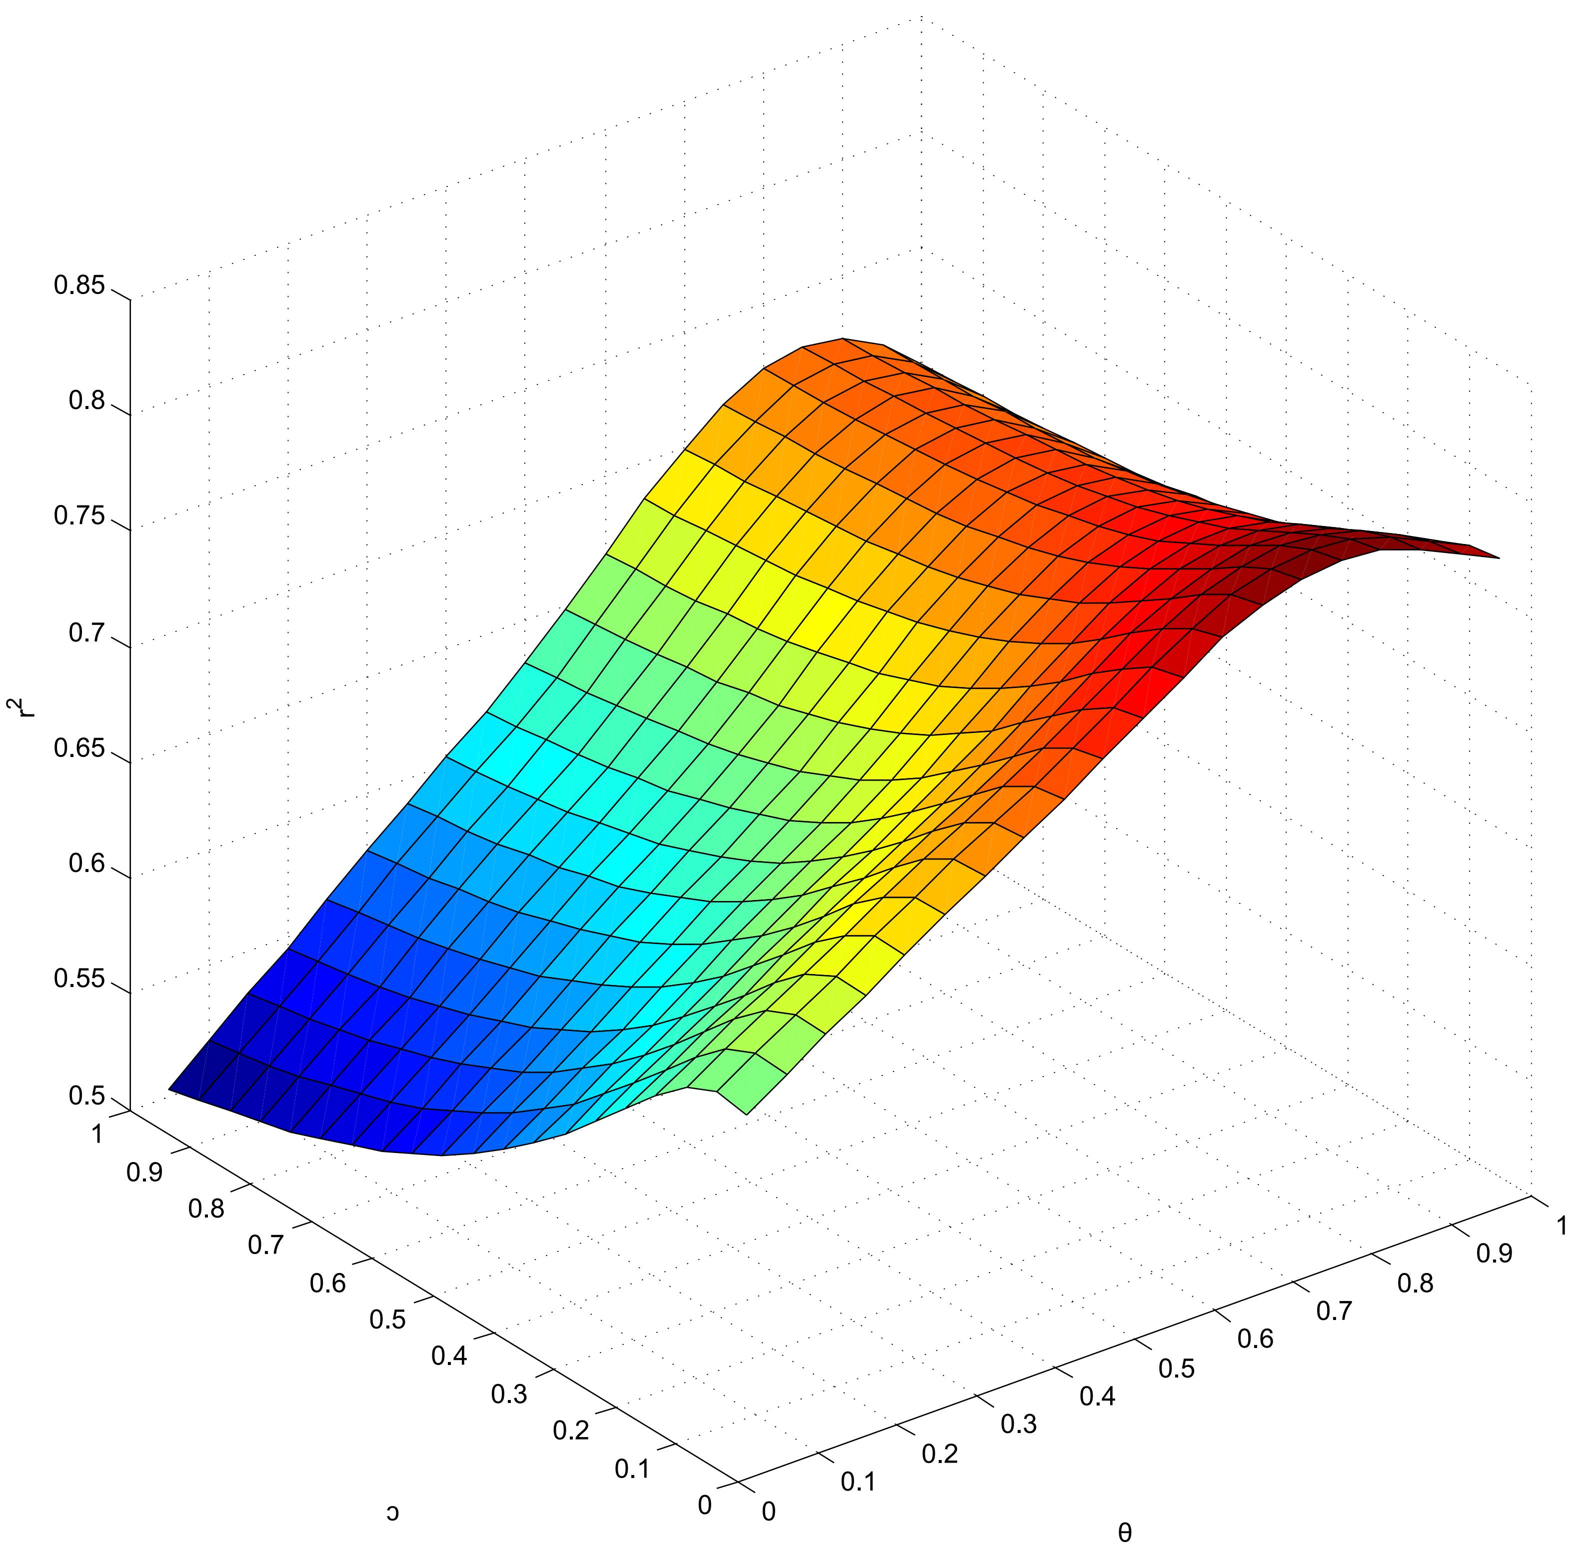

Figure 24. LSAP 1 Announcement

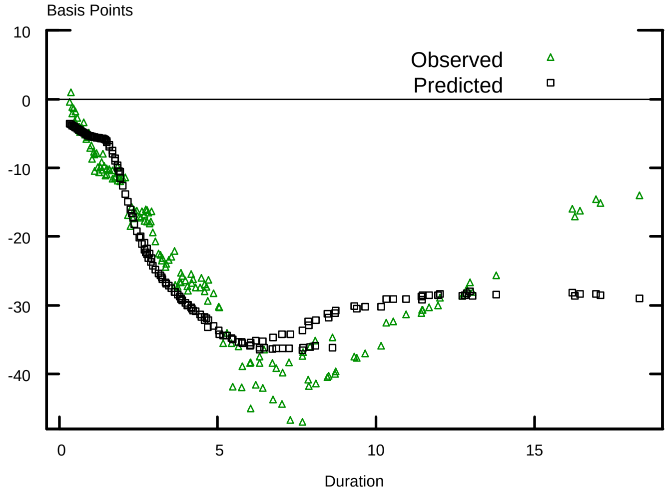

Figure 25. Reinvestment Announcement

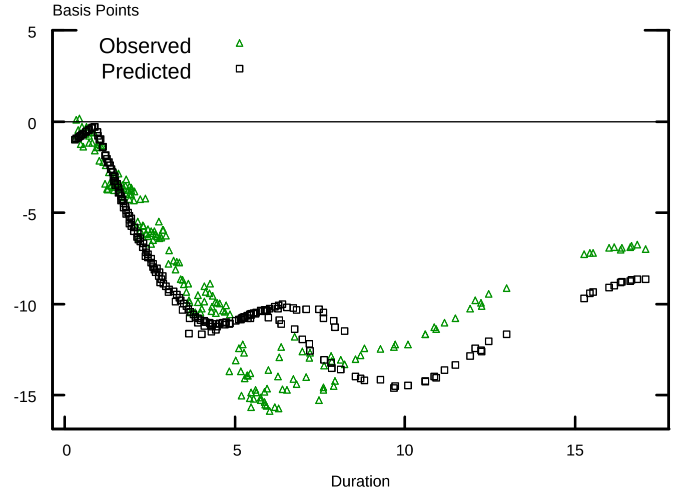

Figure 26. LSAP 2 Announcement

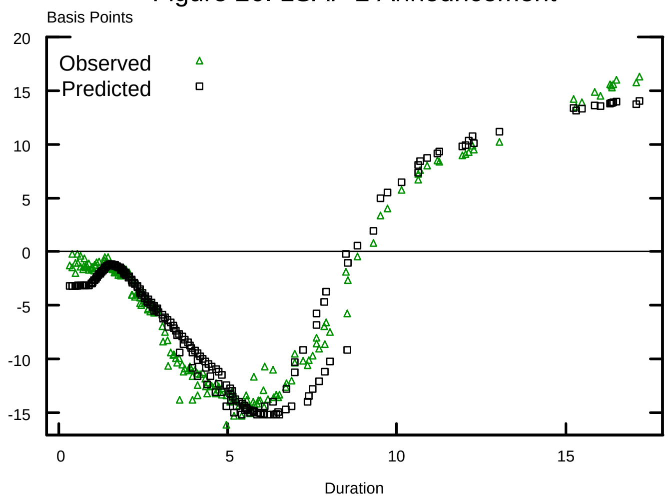

Figure 27. MEP Announcement

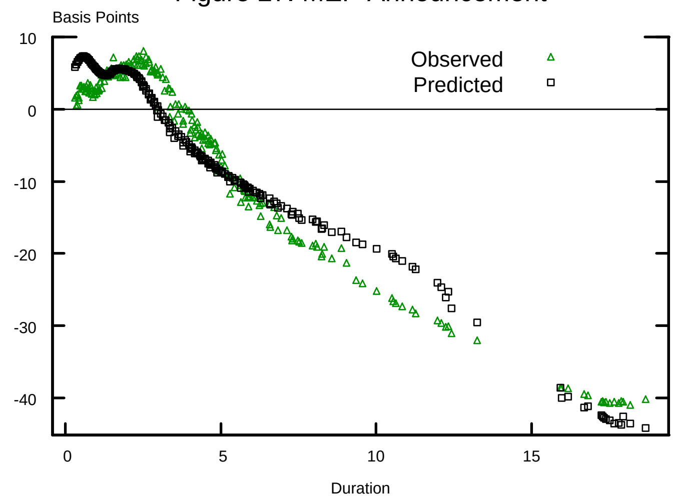

Figure 28. MEP Extension Announcement