Does Quantitative Easing Affect Market Liquidity?

Jens H. E. Christensen†

and

James M. Gillan‡

Abstract

We argue that central bank large-scale asset purchases—commonly known as quantitative easing (QE)—can reduce priced frictions to trading through a liquidity channel that operates by temporarily changing the *shape* of the expected future price distributions of the targeted securities. For evidence we analyze how the Federal Reserve's second QE program that included purchases of Treasury inflation-protected securities (TIPS) affected a measure of liquidity premiums in TIPS yields and inflation swap rates. We find that, for the duration of the program, the liquidity premium measure averaged about 10 basis points lower than expected. This suggests that QE can improve market liquidity.

JEL Classification: E43, E52, E58, G12

Keywords: unconventional monetary policy, liquidity channel, financial market frictions, TIPS, inflation swaps

We thank participants at the Second International Conference on Sovereign Bond Markets, the 2016 Annual Meeting of the American Economic Association, the 9th Annual SoFiE Conference 2016, and the 2016 Annual Meeting of the Western Finance Association, in particular our discussants Paolo Pasquariello, Yuriy Kitsul, and Cynthia Wu, and seminar participants at the Federal Reserve Bank of San Francisco for helpful comments. Furthermore, we are grateful to Autria Christensen, Fred Furlong, Refet Gürkaynak, Jose Lopez, and Nikola Mirkov for many helpful comments and suggestions on previous drafts of the paper. Finally, we thank Lauren Ford for excellent research assistance. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Federal Reserve Bank of San Francisco or the Federal Reserve System.

†Corresponding author: Federal Reserve Bank of San Francisco, 101 Market Street MS 1130, San Francisco, CA 94105, USA; phone: 1-415-974-3115; e-mail: jens.christensen@sf.frb.org.

‡University of California at Berkeley, California, USA; e-mail: james.gillan@berkeley.edu.

This version: September 12, 2016.

1 Introduction

In response to the Great Recession induced by the global financial crisis of 2007-2009, the Federal Reserve quickly lowered its target policy rate—the overnight federal funds rate—effectively to its zero lower bound. Despite this stimulus, the outlook for economic growth remained grim and the threat of significant disinflation, if not outright deflation, was serious. As a consequence, the Fed began purchases of longer-term securities, also known as quantitative easing (QE), as part of its new unconventional monetary policy strategy aimed at pushing down longer-term yields and providing additional stimulus to the economy.

The success of the Fed's large-scale asset purchases in reducing Treasury yields and mortgage rates appears to be well established; see Gagnon et al. (2011), Krishnamurthy and Vissing-Jorgensen (2011), and Christensen and Rudebusch (2012) among others. These studies show that yields on longer-maturity Treasuries and other securities declined on days when the Fed announced it would increase its holdings of longer-term securities. Such announcement effects are thought to be related to the effects on market expectations about future monetary policy and declines in risk premiums on longer-term debt securities.

In this paper, we argue that, in addition to announcement effects, it is also possible for QE programs to affect yields by reducing priced frictions to trading as reflected in liquidity premiums through a liquidity channel.[1,2] This effect comes about because the operation of a QE program is tantamount to introducing into financial markets a large committed buyer who is averse to large asset price declines but does not mind price increases. For example, one repeatedly stated goal of the Fed's various asset purchases programs has been to put downward pressure on long-term interest rates or, equivalently, raise the prices of long-term bonds. As a consequence, for the duration of the program, the most severe downside risk of the targeted securities is effectively eliminated and the shape of their expected future price distributions is temporarily tweaked asymmetrically to the upside. We note that this tweak may not affect the first moment of the price distribution by much. However, since liquidity premiums represent investors' required compensation for assuming the risk of potentially having to liquidate long positions prematurely at significantly disadvantageous prices, the asymmetric twist to the asset price distributions of the targeted securities should reduce their liquidity premiums. By the same logic, liquidity premiums of securities *not* targeted by the QE program are not likely to be affected by the liquidity channel. Furthermore, while such liquidity effects in principle could extend beyond the operation of the QE program provided investors perceive future large declines in the prices of the targeted securities to be countered

[1]This paper represents a completion of the preliminary research described in Christensen and Gillan (2012).

[2]Gagnon et al. (2011) mention a liquidity, or market functioning, channel for the transmission of QE and stipulate a mechanism that shares similarities with the liquidity channel described in this paper, but they do not provide any empirical assessment of the importance of such a channel.

by additional central bank purchases, they are most likely to matter when the program is in operation.

The importance of the liquidity channel for a targeted class of securities could depend on several factors. First, the effect should be positively correlated with the amount purchased relative to the total market value of the targeted class of securities. Second, the intensity of the purchases, that is, the length of time it takes to acquire a given amount of the targeted securities, could play a role as well. The more intense the purchases are, the more loss absorbing capacity a QE program may provide in any given moment and the greater the reductions in the assets' downside risks are likely to be. Finally, the size of the liquidity premiums in the targeted securities should matter. Since such liquidity premiums are widely perceived to be small in the deep and liquid Treasury bond market, it may explain why the liquidity channel has gone unnoticed in the existing literature on the effects of QE.

For evidence on the liquidity channel we analyze how the Fed's second QE program (henceforth QE2), which started in November 2010 and concluded in June 2011, affected the priced frictions to trading in the market for Treasury inflation-protected securities (TIPS) and the related market for inflation swap contracts. The execution of the QE2 program provides an interesting natural experiment for studying liquidity effects in these two markets because the program included repeated purchases of large amounts of TIPS.

To further motivate the analysis and support the view that liquidity premium reductions from the QE2 TIPS purchases could exist and matter, we note that the existence of TIPS liquidity premiums is well established. Fleming and Krishnan (2012) report market characteristics of TIPS trading that indicate smaller trading volume, longer turnaround time, and wider bid-ask spreads than are normally observed in the nominal Treasury bond market (see also Campbell et al. 2009, Dudley et al. 2009, Gürkaynak et al. 2010, and Sack and Elsasser 2004). However, the degree to which they bias TIPS yields remains a topic of debate because attempts to estimate TIPS liquidity premiums directly have generated varying results.[3] Instead, to quantify the effects of the TIPS purchases on the functioning of the market for TIPS and the related market for inflation swaps, we use a novel measure that represents the sum of TIPS and inflation swap liquidity premiums.[4] The construction of the measure only relies on the law of one price and it provides a good proxy for the priced frictions to trading in these two markets independent of the purchase program's effect on market expectations for economic fundamentals. As such, the measure is well suited to capture the changes in TIPS and inflation swap liquidity premiums that we are interested in.

Although we view the primary channel of how QE affects market liquidity as going through

[3] Pflueger and Viceira (2013), D'Amico et al. (2014), Abrahams et al. (2015), and Andreasen et al. (2016) are among the studies that estimate TIPS liquidity premiums.

[4] As a derivative whose pricing is tied to TIPS, inflation swaps are even less liquid and contain their own liquidity premiums for that reason.

liquidity premiums, we note that other measures of market functioning could have been used. Kandrac and Schlusche (2013) analyze bid-ask spreads of regular Treasuries for evidence of any effects from the Treasury purchases during the various Fed QE programs, but do not find any significant results. Thus, they conclude that these purchases had no effect on the functioning of the Treasury bond market. In terms of the market for TIPS, the series of TIPS bid-ask spreads available to us do not appear to be reliable. Thus, we do not pursue an analysis similar to theirs. Fleming and Sporn (2013) study trading activity, quote incidences, and other indicators of market activity in the inflation swap market. We choose to focus on our liquidity premium measure because it quantifies the frictions to trading in the TIPS and inflation swap markets as a yield difference rather than as quantities.

Still, it remains the case that QE may reduce the frictions to trading in a broader sense. As a consequence, we explore the impact on TIPS trading volumes in our empirical analysis and find positive, but insignificant effects. However, we acknowledge that, in general, large-scale asset purchases such as the QE2 program have the potential to impair market functioning by reducing the amount of securities available for trading. Kandrac (2013, 2014) provide evidence of such negative effects on market functioning in the context of the Fed's purchases of mortgage-backed securities. In our case, though, the Fed's TIPS purchases during QE2 were not overly concentrated in any specific TIPS (as we document), and therefore there is little reason to suspect that this effect played any major role during the period under analysis, and our results are consistent with this view.

To estimate the effect of the TIPS purchases during QE2 on our liquidity premium measure, we use a selection on observables strategy to control for potentially confounding variation that may have occurred during the program. Our empirical strategy identifies the change in the mean liquidity premium during the program conditional on our set of controls and finds a significant effect of the program across various sample period definitions and control specifications. To recover the conditional mean, we estimate a linear regression with our liquidity measure and a set of lagged controls to flexibly allow for autoregressive persistence in the relationship. As an additional exercise, we estimate an event-study specification that recovers the change in the conditional mean relative to the 8 weeks prior to the launch of the program. Decomposing the estimate by week, we document a U-shaped pattern in the effect that peaks during the middle of the program at 20 basis points and returns to pre-QE2 period levels towards the end, suggestive of a short-lived cumulative effect of the TIPS purchases. As a final exercise, we employ a randomization inference strategy to test that our results are not an artifact of our model specification. To do this, we estimate changes in the conditional mean for periods of the same length during "placebo" QE2 programs to generate a distribution of estimates. The mean of these distributions is close to zero, and our estimate for the true program period is more negative than 90 percent of the placebo treatment periods. Taken together, we find that these results suggest that, for the duration of the QE2 program, the liquidity premium measure averaged about 10 basis points lower than expected, a reduction of almost 50 percent from pre-QE2 levels. We interpret these findings as indicating that part of the effect from QE programs derives from improvements in the market conditions for the targeted securities.

To assess whether the liquidity channel affects liquidity premiums of securities *not* targeted by the QE2 program, we repeat the analysis using credit spreads of AAA-rated industrial corporate bonds, an asset class that the Fed under normal circumstances is not allowed to acquire and hence could not possibly be expected to purchase. As the default risk of such highly rated bonds is negligible, their credit spreads mostly represent liquidity premiums.[5] Consistent with the theory of the liquidity channel, which emphasizes QE programs' effects on the shape of the price distributions for the targeted securities, we obtain no significant results in this exercise. Although not conclusive as we only consider one alternative asset class, we take this as evidence that the transmission of the liquidity channel is indeed limited to the purchased security classes.

In a recent paper, D'Amico and King (2013, henceforth DK) emphasize local supply effects as an important mechanism for QE to affect long-term interest rates. Under this local supply channel declines in the stock of government debt available for trading induced by QE purchases should push up bond prices (temporarily) due to preferred habitat behavior on the part of investors. DK find evidence of such instantaneous purchase effects in their analysis of the Treasury market response to the $300 billion of Treasury security purchases during the Fed's first QE program, which were announced on March 18, 2009, and concluded by October 30, 2009. They report an average decline in yields in the maturity segment purchased of 3.5 basis points on days when operations occurred. Meaning and Zhu (2011) repeat the analysis of DK for the purchases of regular Treasuries included in the QE2 program. They report that a typical QE2 purchase operation reduced Treasury yields by 4.7 basis points, while the cumulative stock effect of the entire program is estimated to be 20 basis points.

To analyze whether our results could be driven by local supply effects, we replicate the approach of DK to detect effects on individual TIPS prices from the TIPS purchases in the QE2 program. However, we fail to get any significant results, which suggests that local supply effects are not likely to be able to account for our findings.

To the best of our knowledge, this paper is the first to study the liquidity channel as a separate transmission mechanism for QE to affect long-term interest rates and to document that such liquidity effects are distinct from and more persistent than the local supply channel highlighted in the existing literature.

[5] See Collin-Dufresne et al. (2001) for evidence and a discussion of the weak link between corporate bond credit spreads and their default risk, frequently referred to as the credit spread puzzle, and Christensen (2008) for an overview of related research.

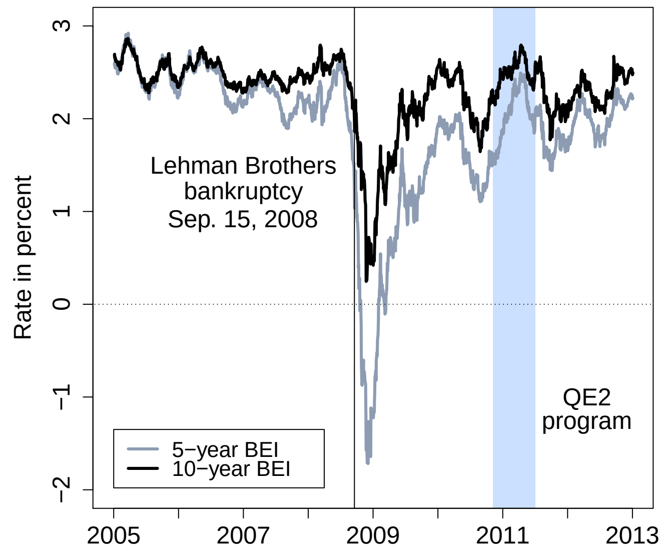

Figure 1: **TIPS Breakeven Inflation.**

In related research, De Pooter et al. (2015) analyze the government bond purchases performed by the European Central Bank (ECB) as part of its Securities Markets Programme (SMP) that operated from May 2010 to March 2012. To avoid assessing changes to expectations about monetary policy, they scrutinize the spreads between yields of targeted government bonds from the euro-area periphery and non-targeted German bund yields and control for credit risk using CDS rates. This leaves them with a measure of priced frictions that is very similar in concept to ours. In the empirical analysis, they report instantaneous effects on the order of 30-40 basis points from purchasing one percent of the outstanding market, while their results reveal longer lasting effects of 13 to 17 basis points. In light of our study, we interpret the difference of about 25 basis points between their two results as representing local supply effects of the nature discussed in DK, while the lasting effects of about 15 basis points can be taken to represent an estimate of the importance of the liquidity channel and are remarkably similar in magnitude to our results for the TIPS market. Eser and Schwaab (2016) also study instantaneous purchase effects from the ECB's SMP and report results of comparable magnitudes. In their own interpretation, reduced liquidity premiums appear to be the most important factor behind their findings.

Our findings could have important policy implications. First, for assessing the credibility of the Fed's price stability goal, it is a common practice to study the difference in yield between regular Treasury bonds and TIPS of the same maturity, known as breakeven inflation (BEI), which represent market-based measures of inflation compensation frequently mentioned in FOMC statements. Figure 1 shows daily five- and ten-year BEI since 2005, also highlighted is the operation of the QE2 program. During the period of its operation, BEI first experienced a sharp increase until the middle of the program followed by a notable downtick towards its end. Specifically, at the five-year maturity, BEI started at 1.51% on November 3, 2010, peaked at 2.49% on April 8, 2011, before retracing to 2.07% by the end of June 2011. At the ten-year maturity, BEI increased from 2.30% to 2.78% and fell back down to 2.59% between the same three dates. Based on our results, as much as one-third of the variation in BEI during this period could reflect effects arising from the QE2 TIPS purchases through the liquidity channel that, by definition, would have little to do with investors' inflation expectations or associated inflation risk premiums.[6] Thus, in determining how much the QE2 program helped boost investors' inflation expectations, it is crucial to account for the effects of the liquidity channel we unveil.

More generally, for central banks in countries with somewhat illiquid sovereign bond markets (most euro-area countries likely belong in this category as suggested by the analysis in De Pooter et al. 2015), QE programs that target sovereign debt securities could be expected to reduce the liquidity premiums of those securities quite notably for the duration of the programs, which might be worthwhile to keep in mind when evaluating the effects of such QE programs. In this regard, we note that the TIPS market with a total outstanding notional of $1,078 billion as of the end of 2014 is quite comparable to the major European sovereign bond markets. Thus, our analysis could provide a useful reference point for understanding the effects of the liquidity channel in the European context.

Finally, since the Fed's TIPS purchases represented less than five percent of the TIPS market, our results suggest that even relatively modest QE programs could have sizable effects when the targeted security classes are illiquid.[7] Thus, the significance of the liquidity channel could matter for the design of QE programs; time frame, purchase pace, and targeted security classes are all decision variables that merit careful consideration under those circumstances.

The remainder of the paper is structured as follows. Section 2 discusses the channels of transmission of QE to long-term interest rates, paying special attention to the proposed liquidity channel. Section 3 details the execution of the TIPS purchases included in the QE2 program, while Section 4 describes the construction of the TIPS and inflation swap liquidity premium measure. Section 5 introduces the econometric model we use and presents our results. Section 6 concludes the paper and provides directions for future research. Appendices

[6] For the maximum effect of the liquidity channel on BEI to apply, it must be the case that there was no change in the liquidity premiums of inflation swaps in response to the QE2 TIPS purchases, a possibility our analysis allows for.

[7] The large effects on mortgage rates of the Fed's purchases of mortgage-backed securities during its first large-scale asset purchase program, which Krishnamurthy and Vissing-Jorgensen (2011) partly attribute to improved market functioning and reduced liquidity premiums, provide another example.

contain additional results, a description of our adaptation of DK's approach, and an extension of our analysis to the TIPS transactions included in the Fed's maturity extension program (MEP) that operated from September 2011 through the end of 2012.

2 Transmission Channels of QE to Long-Term Rates

In this section, we first give a theoretical overview of how to think about QE and its effects on the economy. We then describe the three main channels of transmission of QE discussed in the existing literature before we introduce the novel liquidity channel we highlight.

2.1 Theoretical Overview

Once a central bank has reduced its leading conventional policy rate to its effective lower bound, it may be forced to engage in QE to provide further monetary stimulus, if needed. This has been the reality facing several of the world's most prominent central banks in recent years.

This predicament raises several questions relevant to monetary policy. One key question is how QE works and affects the real economic outcomes such as employment and inflation that policymakers care about. This could matter for both design and management of QE programs and may ultimately have implications for how such programs should be wound down, a phase yet to be reached by any of the major central banks that have engaged in QE.[8]

The main mechanism for QE to affect the real economy is through its impact on long-term interest rates, which are key variables in determining many important economic decisions ranging from firm investment on the business side to home and auto purchases on the household side of the economy. Therefore, to understand how QE works, we need to study its transmission channels to long-term interest rates.

At its core, QE is merely a redistribution of assets in the economy as the total outstanding stock of assets in private hands is left unchanged. In the U.S. and the U.K. for example, QE has involved swapping medium- and long-term government-issued or government-backed securities for newly created reserves. In the aggregate, one could argue that not much has changed, in that one government-backed claim (Treasury securities) has been replaced by another (reserves that represent claims on the central bank, which is a branch of the government). Thus, the theoretical challenge in understanding the effects of QE is to identify conditions and circumstances under which this asset swap and the resulting change in private agents' portfolio compositions can have effects on long-term interest rates in equilibrium.

[8]The unwind of the small bond purchase program operated by the Swiss National Bank in 2009 and exited in 2010 is a rare exception. See Kettemann and Krogstrup (2014) for details and an analysis.

2.2 Signaling and Portfolio Balance Channels

The most straightforward way QE can affect long-term interest rates is by acting as a signaling device that changes agents' expectations about the future path for monetary policy. Christensen and Rudebusch (2012) and Bauer and Rudebusch (2014) are among the studies that emphasize the importance of the signaling channel for understanding the effects of QE.

Beyond potential signaling effects that would affect the yields of all securities, QE programs may change the supply of or demand for a given asset, which could affect its price and hence risk premium. Such effects are usually referred to as portfolio balance effects.[9] In recent research, Christensen and Krogstrup (2016a,b) introduce a distinction between supply-induced and reserve-induced portfolio balance effects.

Both types of portfolio balance effects share some common characteristics. Their existence requires market frictions or segmentation to matter so that a change in the relative market supply of an asset can have an impact on its relative price in equilibrium (a mechanism that is absent in standard models of the yield curve). To provide a theoretical justification for such effects, the seminal model introduced in Vayanos and Vila (2009) is a frequent reference. This model suggests that, when assets with otherwise near-identical risk and return characteristics are considered imperfect substitutes by some market participants (e.g., due to preferred habitat) and markets are segmented, a change in the relative market supply of an asset may affect its relative price (see also Tobin 1969).[10]

Most of the existing literature on the impact of QE on yields has focused on supply-induced portfolio balance effects where the central bank asset purchases are treated as a reduction in the market supply of the targeted assets.[11] Assuming unchanged investor demand, the prices of the targeted securities should go up or, equivalently, their yields go down.

The reserve-induced portfolio balance channel described in Christensen and Krogstrup (2016a,b) emphasizes instead the role of the reserves created by the central bank as part of any QE program. Provided the asset purchases are executed with nonbank financial market participants, the new reserves end up expanding banks' balance sheets with reserves on the asset side matched by increased deposits on the liability side. Since only banks can hold the reserves, this reduces the duration of their portfolios. Assuming banks had optimal portfolios before the central bank asset purchases, they increase their demand for long-term assets to counter the duration reduction, which pushes up asset prices.

As a third channel for QE to affect bond yields, DK highlight local supply effects as a

[9] This division into signaling and portfolio balance effects is a simplification. See Bauer and Rudebusch (2014) for a thorough discussion.

[10] See Hamilton and Wu (2012) and Greenwood and Vayanos (2014) for empirical applications of the Vayanos and Vila (2009) model to the U.S. Treasury market.

[11] Gagnon et al. (2011) and Joyce et al. (2011) are among the studies that emphasize this particular portfolio balance channel.

potentially important transmission mechanism. The principle behind this channel is that the individual purchase operations by the central bank represent small local reductions in the available supply of government debt. To a first order, such operations are small enough individually that they should not alter investors' preferences or portfolios. If so, the demand for government debt can be assumed constant around each purchase operation. Hence, any price effects from the purchases can be characterized as resulting from movements along demand curves.

2.3 The Liquidity Channel

The novel channel we propose in this paper is for QE to have effects on the liquidity premiums that investors demand to hold any security that is less than perfectly liquid.[12] To be specific, we think of the liquidity premium of a security as representing investors' required compensation for assuming the risk of potentially having to liquidate a long position in the security prematurely at a disadvantageous price, say, in a stressed market environment when market makers and arbitrageurs are severely capital constrained. We note that under normal circumstances the liquidity premium is determined as the outcome of a non-cooperative game between buyers and sellers and embeds their collective assessment of the net present value of the total sum of frictions to trading until maturity.

When a central bank launches a QE program, we argue that it is equivalent to introducing into financial markets a committed buyer with deep pockets and unusual preferences (from the perspective of a buyer). Specifically, we think of the central bank as averse to large asset price declines and does not mind (in fact, actually desires) asset price increases, and it will execute a trading strategy that attempts to ensure those outcomes. We stress that the aversion of the central bank to price declines is not tied to worries about the value of the acquired assets per se, but arises out of concerns that it could be viewed as a failure of its policy. This behavioral pattern effectively eliminates the most severe downside risk of the targeted securities. As a consequence, the shape of their expected future price distributions gets an asymmetric tweak to the upside in addition to any changes to the mean from the other QE transmission channels discussed above.

It then follows that the existence of a QE program changes the outcome of the game that determines the liquidity premium for the targeted securities. As rational agents, market participants are aware of the fact that, when confronted with disadvantageous prices, sellers can pursue the alternative strategy of submitting bids in the QE purchase auctions and sell targeted securities that way. As a result, sellers are less likely to be significantly squeezed

[12] A perfectly liquid security can be sold any time in arbitrarily small or large amounts at no trading costs (i.e., there is no bid-ask spread) and without affecting its price. A demand deposit is close to meeting these requirements if we abstract from the default risk of large deposits, which are not government guaranteed.

while the QE program is in operation, which makes all participants willing to accept a lower liquidity premium.

Furthermore, we note that these dynamics entail that the effects should taper off towards the end of a QE program as the number of remaining purchase auctions goes to zero; and once the program has ended, market participants are left playing their normal non-cooperative game. This suggests that effects tied to the liquidity channel could have a different dynamic profile than effects tied to the other transmission channels discussed earlier, in particular announcement effects are not necessarily material in size for the liquidity channel.

We stress that the liquidity channel is distinct from the insurance against macroeconomic tail risks that central bank asset purchases could potentially provide as discussed in Hattori et al. (2016). While the latter channel also affects the downside risk of assets, it is economy-wide in nature and would impact all asset classes instantaneously upon announcement thanks to the forward-looking behavior of investors, and we control for such announcement effects in our analysis as detailed below.

The importance of the liquidity channel for a given security class is likely to be determined by several factors. First, its effect should be positively correlated with the amount purchased relative to the total market value of the security class. Second, the intensity of the purchases, that is, the length of time it takes to purchase a given amount, could play a role as well. The more intense the purchases are, the greater is the ability of a given QE program to absorb negative liquidity shocks that force owners of targeted securities to sell and exert downward pressure on the securities' prices. As a consequence, the reduction in liquidity premiums should have a positive correlation with the purchase pace. Furthermore, the size of the liquidity premiums in the targeted security classes should matter. Since such liquidity premiums are widely perceived to be small in the deep and liquid Treasury bond market, it may explain why the liquidity channel has been overlooked in the existing literature.

In terms of the dynamic profile of the effect of the liquidity channel, we note that, since liquidity premiums reflect fears about the *future* resale value of securities, its effect is likely to taper off some time before the purchases are scheduled to end. In principle, though, its effect could extend beyond the operation of the QE program if investors perceive that undesirable price developments in the targeted securities would make the central bank return to the market. Still, it is clear that the liquidity effects could be expected to be strongest when the QE program is committed and in operation.

Finally, it is important to emphasize that, for the liquidity channel and the associated liquidity effects to exist, no portfolio balance effects are needed; only financial market frictions are required. Ultimately, the existence and importance of the liquidity channel may be tied to theories of limits to arbitrage capital with market makers and arbitrageurs; see Hu et al. (2013, henceforth HPW) for a discussion.[13] However, we leave it for future research to establish any such ties.

2.4 Identification of Liquidity Effects

In order to empirically identify effects on long-term interest rates arising through the liquidity channel, two criteria must be met. First, we need a QE program that is large, includes repeated repurchases of securities less liquid than Treasuries, and operates over a period long enough that the fears of forced resales implicit in the definition of liquidity premiums can be meaningfully affected by the purchases. Second, we must have a suitable measure of the priced frictions in the markets for the purchased securities.

The Fed's QE2 program meets these criteria. First, this program was large, operated over an eight-month period, and included repeated purchases of a significant amount of TIPS, which are widely perceived to be less liquid than Treasuries. Second, we devise a measure of the priced frictions in TIPS yields and inflation swap rates detailed in Section 4 that we use to detect evidence of the liquidity channel. Still, in trying to identify effects from the liquidity channel, we acknowledge that signaling, portfolio balance, and local supply effects could be operating as well.

In principle, effects of the signaling and portfolio balance channels should materialize immediately following the announcement of the QE2 program and not when it is implemented thanks to the rational, forward-looking behavior of investors. As a consequence, we look for effects related to the announcement of the program on November 3, 2010, but fail to detect any significant yield responses as documented in Appendix A. More likely, announcement effects tied to these channels materialized in the weeks and months ahead of the launch of the QE2 program as argued by Krishnamurthy and Vissing-Jorgensen (2011). Furthermore, the signaling and portfolio balance channels are thought to mainly affect the policy expectations and term premium components of bond prices, which should cancel out in the construction of our liquidity premium measure. Thus, neither of these channels are likely to be the drivers of our results. Also, it follows from this discussion that, in case there are unaccounted announcement effects, our results will be conservative and represent lower bound estimates of the importance of the liquidity channel.

To address the local supply channel, we replicate the analysis of DK in an attempt to identify local supply effects in individual TIPS prices, but fail to get any significant results as documented in Appendix B. However, this may not be as surprising as it could seem. First, we argue that their regressions suffer from misspecified time fixed effects. Second and more importantly, the mechanics of the liquidity channel suggest that its effects are not limited to

[13] Brunnermeier and Pedersen (2009) and Pasquariello (2015) provide examples of theoretical models where funding liquidity and informational frictions, respectively, may affect the workings of financial markets.

any specific security, but would apply to all securities at risk of being targeted by the QE program. For that reason these effects may go undetected in the type of analysis performed by DK that focuses on identifying local supply effects in individual security prices on purchase operation dates.

With signaling, portfolio balance, and local supply channels ruled out as important drivers of the variation in our measure of liquidity premiums in TIPS yields and inflation swap rates during the QE2 program, we turn our focus to the proposed liquidity channel. The remainder of the paper is dedicated to analyzing whether the TIPS purchases in the QE2 program had any effects on our liquidity premium measure consistent with this channel.

3 The TIPS Purchases in the QE2 Program

In this section, we provide a brief description of the Federal Reserve's QE2 program that included purchases of a sizable amount of TIPS.

The QE2 program was announced on November 3, 2010. In its statement, the Federal Open Market Committee (FOMC) said that the program would expand the Fed's balance sheet by $600 billion through Treasury security purchases over approximately an eight-month period. In addition, the FOMC had already decided in August 2010 to reinvest principal payments on its portfolio of agency debt and mortgage-backed securities in longer-term Treasury securities in order to maintain the size of the Fed's balance sheet, a policy that was maintained until September 2011.[14] As a consequence, the gross purchases of Treasury securities from November 3, 2010, until June 29, 2011, totaled nearly $750 billion, of which TIPS purchases represented about $26 billion. Since the total amount of marketable Treasury debt increased by $792 billion between the end of October 2010 and the end of June 2011, the Fed's Treasury purchases during this period nearly kept pace with the Treasury net issuance. In terms of TIPS, though, the net supply increased by $61 billion so that the Fed's purchases only represented an amount equal to 42 percent of the new supply. Thus, in the aggregate, the Fed's TIPS purchases did not come at the expense of private sector holdings.

The uniqueness of these TIPS purchases is evident in Figure 2(a), which shows the total book value of the Fed's TIPS holdings since 2008.[15] They increased the Fed's holdings by 52.8 percent and brought the total close to $75 billion.[16] Figure 2(b) shows the market share of individual TIPS held by the Fed at the beginning of the QE2 program and at its conclusion

[14] The Fed has all along reinvested principal payments on its portfolio of Treasury securities in Treasuries. Since September 2011, the Fed has been reinvesting principal payments on its portfolio of agency debt and mortgage-backed securities in agency mortgage-backed securities to support the housing market.

[15] The Fed has purchased TIPS outside the QE2 program, most notably during the MEP that ran from September 2011 through 2012. The effects of these TIPS transactions are analyzed separately in Appendix D.

[16] The slight decline in mid-April 2011 is due to a maturing five-year TIPS of which the Fed was holding $2.9 billion in principal and $327 million in accrued inflation compensation.

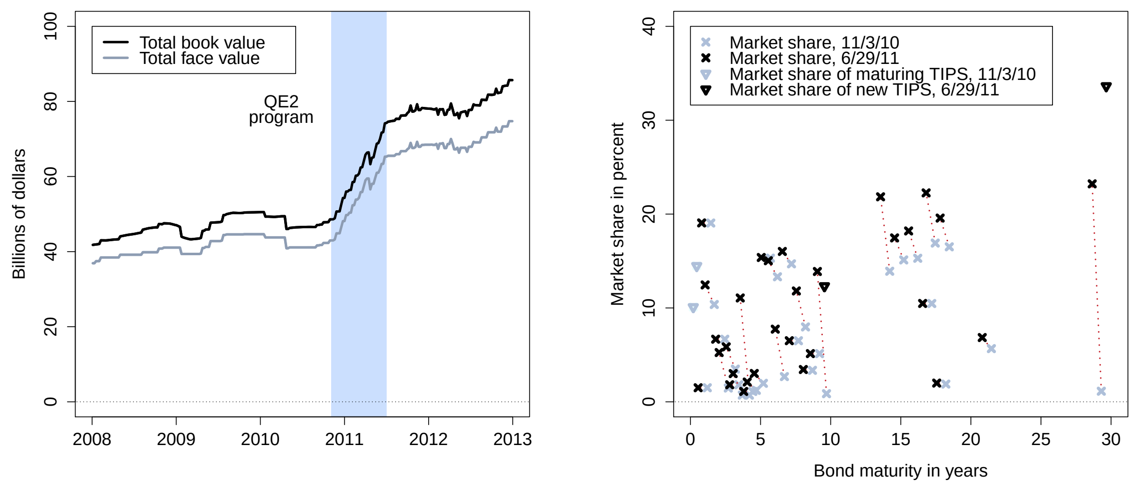

(a) Value of TIPS.

(b) Share of TIPS market.

Figure 2: **Fed's TIPS Holdings.**
Panel (a) shows the total book and face value of TIPS held in the Federal Reserve System's Open Market Account (SOMA). The difference between the two series reflects accrued inflation compensation. The data are weekly covering the period from January 2, 2008, to December 26, 2012. Panel (b) shows the market share of individual TIPS held by the Fed at the start of QE2 and at its conclusion with thin dashed red lines indicating the change in the shares held. Note that two TIPS held as of November 3, 2011, matured before the end of the program, and two new TIPS were issued during the program and acquired by the Fed.

with thin dashed red lines indicating the change for each TIPS. A total of three TIPS were issued during the QE2 program; the five-year 4/15/2016 TIPS issued on April 29, 2011, the ten-year 1/15/2021 TIPS issued on January 31, 2011, and the thirty-year 2/15/2041 TIPS issued on February 28, 2011. As of June 29, 2011, the Fed was only holding the two latter securities shown with black triangles in Figure 2(b). Note that the purchases were not heavily concentrated in any particular TIPS, and the Fed's TIPS holdings as a percentage of the stock of each security in general remained well below one-third.

The QE2 program was implemented with a very regular schedule. Once a month, the Fed publicly released a list of operation dates for the following 30-plus day period, indicating the relevant maturity range and expected purchase amount for each operation. There were 15 separate TIPS operation dates, fairly evenly distributed across time, each with a stated expected purchase amount of $1 billion to $2 billion. Table 1 lists the 15 operation dates, the total purchase amounts, and the weighted average maturity of the TIPS purchased. TIPS were the only type of security acquired on these dates, and the Fed did not buy any TIPS outside of those dates over the course of the program.[17] Furthermore, all outstanding TIPS with a

[17] Also, there were no TIPS auctions by the U.S. Treasury on any of the Fed's 15 TIPS operation dates. See Lou et al. (2013) for analysis of the effects of auctions in the regular nominal Treasury bond market.

QE2 TIPS purchase operation dates	TIPS purchases (mill.)	Weighted avg. maturity (years)
(1) Nov. 23, 2010	$1,821	9.43
(2) Dec. 8, 2010	$1,778	8.88
(3) Dec. 21, 2010	$1,725	16.09
(4) Jan. 4, 2011	$1,729	16.98
(5) Jan. 18, 2011	$1,812	14.64
(6) Feb. 1, 2011	$1,831	13.58
(7) Feb. 14, 2011	$1,589	14.16
(8) Mar. 4, 2011	$1,589	11.37
(9) Mar. 18, 2011	$1,653	17.77
(10) Mar. 29, 2011	$1,640	18.29
(11) Apr. 20, 2011	$1,729	23.17
(12) May 4, 2011	$1,679	13.62
(13) May 16, 2011	$1,660	20.49
(14) Jun. 7, 2011	$1,589	14.30
(15) Jun. 17, 2011	$2,129	5.98
Average	$1,730	14.58

Table 1: **QE2 TIPS Purchase Operations.**
The table reports the amount and weighted average maturity of TIPS purchased on the 15 TIPS operation dates during the QE2 program.

minimum of two years remaining to maturity were eligible for purchase on each operation date and, as shown in Figure 2(b), the Fed did purchase TIPS across the entire indicated maturity range. Thus, there does not appear to be a need to account for price movements of specific securities related to the release of the operation schedules. Also, market participants did not know in advance either the total amount to be purchased or the distribution of the purchases. However, since the actual purchase amounts all fall in the range from $1.589 billion to $2.129 billion, investors' perceived uncertainty about the total purchase amounts likely was lower than the width of the indicated range. Finally, the auction results containing this information were released a few minutes after each auction. As the auctions closed at 11:00 a.m. Eastern time, investors had sufficient time to process the information before the close of the market on each operation date. It is this structure of the execution of the TIPS purchases in the QE2 program that makes it a natural candidate for detecting local supply effects as we attempt in Appendix B.

4 A Measure of Liquidity Premiums in TIPS and Inflation Swaps

In this section, we describe the measure of liquidity premiums in TIPS yields and inflation swap rates that we use as a dependent variable in our empirical analysis.

Ideally, we would like to use a pure measure of liquidity premiums in TIPS yields in our analysis. However, empirically, it is very challenging to separate liquidity premiums from other factors that affect TIPS yields such as expectations for monetary policy and inflation. Instead, we combine the information in Treasury yields, TIPS yields, and inflation swap rates to get a handle on the size of the liquidity premiums in TIPS yields and inflation swap rates jointly as explained in the following.

To begin, note that, unlike regular Treasury securities that pay fixed coupons and a fixed nominal amount at maturity, TIPS deliver a real payoff because their principal and coupon payments are adjusted for inflation based on the consumer price index (CPI). The difference in yield between regular nominal, or non-indexed, Treasury bonds and TIPS of the same maturity is referred to as breakeven inflation, since it is the level of inflation that makes investments in indexed and non-indexed bonds equally profitable.

In an inflation swap contract, the owner of a long position pays a fixed premium in exchange for a floating payment equal to the change in the consumer price index used in the inflation indexation of TIPS. At inception, the fixed premium is set such that the contract has a value of zero.

Since the cash flows of TIPS and inflation swaps are adjusted with the same price index, economic theory implies a connection between their pricing. Specifically, in a frictionless world, the absence of arbitrage opportunities requires the inflation swap rate to equal BEI because buying one nominal discount bond today with a given maturity produces the same cash flow as buying one real discount bond of the same maturity and selling an inflation swap contract also of the same maturity. However, in reality, the trading of both TIPS and inflation swap contracts is impeded by frictions, such as wider bid-ask spreads and less liquidity relative to the market for regular nominal Treasury bonds. As a consequence, the difference between inflation swap rates and BEI will not be zero, but instead represents a measure of how far these markets are from the frictionless outcome described above.[18]

To map this to our data, we observe a set of nominal and real Treasury zero-coupon bond yields denoted $\widehat{y}_t^N(\tau)$ and $\widehat{y}_t^R(\tau)$, respectively, where τ is the number of years to maturity.

[18] Note that, due to collateral posting, the credit risk in inflation swap contracts is negligible and can be neglected for pricing purposes. Also, we assume the default risk of the U.S. government to be negligible, which is warranted for our sample that ends in June 2011 before the downgrade of U.S. Treasury debt in August 2011. However, even for this later period, which we consider in our analysis of the Fed's MEP described in Appendix D, any significant credit risk premium is not likely to bias our measure as it would presumably affect Treasury and TIPS yields in the same way, leaving BEI effectively unchanged.

Also, we observe a corresponding set of rates on zero-coupon inflation swap contracts denoted $\widehat{IS}_t(\tau)$. As noted above, these rates differ from the unobserved values that would prevail in a frictionless world without any obstacles to continuous trading denoted $y_t^N(\tau)$, $y_t^R(\tau)$, and $IS_t(\tau)$, respectively, with the theoretical relationship:

$$IS_t(\tau) = y_t^N(\tau) - y_t^R(\tau).$$

Now, we make three fundamental assumptions:

(1) The nominal Treasury yields we observe are very close to the unobservable frictionless nominal yields, that is, $\widehat{y}_t^N(\tau) = y_t^N(\tau)$ for all t and all relevant τ. Even if not exactly true (say, for example, during the financial crisis), this is not critical as the point is ultimately about the relative liquidity between securities that pay nominal and real yields.

(2) TIPS are no more liquid than nominal Treasury bonds. As a consequence, TIPS yields contain a time-varying liquidity premium denoted $\delta_t^R(\tau)$, which generates a wedge between the observed TIPS yields and their frictionless counterpart given by $\widehat{y}_t^R(\tau) = y_t^R(\tau) + \delta_t^R(\tau)$ with $\delta_t^R(\tau) \geq 0$ for all t and all relevant τ.

(3) Inflation swaps are no more liquid than nominal Treasury bonds. Hence, the observed inflation swap rates are also different from their frictionless counterpart with the difference given by $\widehat{IS}_t(\tau) = IS_t(\tau) + \delta_t^{IS}(\tau)$ and $\delta_t^{IS}(\tau) \geq 0$ for all t and all relevant τ.

In support of these assumptions, we note that market size, trading volume, and bid-ask spreads all indicate that regular Treasury securities are much more liquid than both TIPS and inflation swaps.[19] It then follows that the difference between observed inflation swap and BEI rates, which defines our liquidity premium measure, is given by

$$\begin{aligned}
LP_t(\tau) &\equiv \widehat{IS}_t(\tau) - \widehat{BEI}_t(\tau) \\
&= \widehat{IS}_t(\tau) - [\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)] \\
&= IS_t(\tau) + \delta_t^{IS}(\tau) - [y_t^N(\tau) - (y_t^R(\tau) + \delta_t^R(\tau))] \\
&= \delta_t^R(\tau) + \delta_t^{IS}(\tau) \geq 0.
\end{aligned}$$

This shows that $LP_t(\tau)$ is nonnegative and equal to the *sum* of liquidity premiums in TIPS yields and inflation swap rates. Hence, $LP_t(\tau)$ quantifies how far the observed market rates are from the frictionless outcome.

[19] Driessen et al. (2014) find statistically significant liquidity effects in both TIPS yields and inflation swap rates.

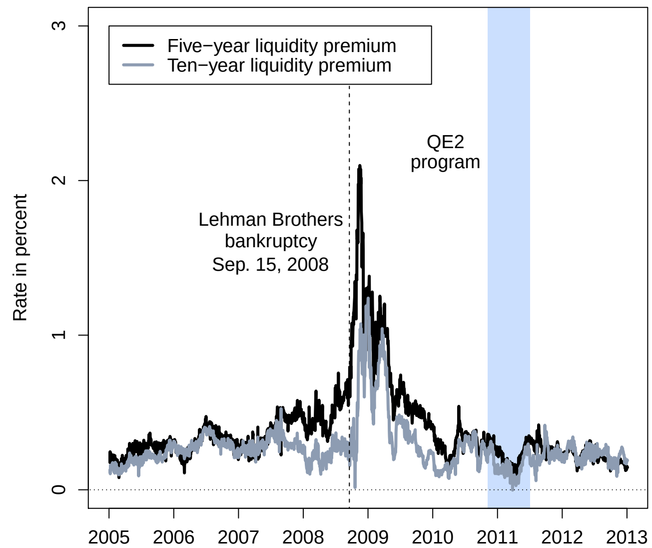

Figure 3: **Sum of Liquidity Premiums in TIPS and Inflation Swaps.**

4.1 Construction of the Liquidity Premium Measure

We use daily zero-coupon nominal and real Treasury bond yields as constructed by Gürkaynak et al. (2007, 2010) for our observed bond yields. For the inflation swap rates, we use daily quotes from Bloomberg. These rates are for zero-coupon inflation swap contracts, meaning they have no exchange of payment upon issuance and a single cash flow exchanged at maturity. The quoted rates represent the payment of the fixed leg at an annual rate, which we convert into continuously compounded rates using the formula $\widehat{IS}_t^c(\tau) = \ln(1 + \widehat{IS}_t(\tau))$ to make them comparable to the other interest rates. Bloomberg begins reporting quotes on inflation swap rates in early 2004, but the data are not densely populated until the end of the year. As a result, we begin the sample period on January 4, 2005, and end it on December 31, 2012, when the MEP was completed. Finally, we eliminate the few days where quotes are not available for all maturities, which leaves us with a sample of 1,977 observations.

Figure 3 shows $LP_t(\tau)$ at the five- and ten-year maturity. In the empirical analysis, we aim to quantify the liquidity effects of the QE2 TIPS purchases on the priced frictions in the markets for TIPS and inflation swaps as reflected in our liquidity premium measure. Importantly, in the construction of the measure, any effects of the QE2 program on bond investors' views of economic fundamentals, such as future monetary policy, inflation, and their implications for bond yields, will cancel out as they affect inflation swap rates and BEI of the same maturity equally. This is also the reason why the measure is likely affected to a minimum by signaling and portfolio balance effects.

We note that other model-free measures of priced frictions could have been used. One such alternative measure can be constructed from asset swap spreads as described in Pflueger and Viceira (2013). In Appendix C, we demonstrate that, in theory, this measure should be closely correlated with our liquidity premium measure and we make a brief comparison of the two and find them to be highly positively correlated. This suggests that the two measures indeed appear to contain the same information.

5 Empirical Results

Figure 3 provides visually suggestive evidence that the combined liquidity premiums in the TIPS and inflation swap markets dipped during the QE2 program. The challenge in identifying the causal effect of quantitative easing programs on our measure of liquidity is that there are many potential confounding factors that could also affect liquidity and be incorrectly attributed to the program. Our empirical strategy relies on the identification assumption that we can adequately control for confounding factors with other observable measures of liquidity. We employ several strategies to quantify the impact of the program and perform several robustness checks detailed in the following.

5.1 The Econometric Model

Throughout this section, the baseline assumption is that the sum of liquidity premiums in the TIPS and inflation swaps market is determined by the econometric model:

$$LP_t(\tau) = \alpha + \delta D_t^{QE2} + \sum_{s=0}^{T} X'_{t-s}(\tau)\beta + \varepsilon_t, \quad (1)$$

where $LP_t(\tau)$ is the sum of the liquidity premiums in the TIPS market in week t for maturity τ and D_t^{QE2} is an indicator that the QE2 program was in effect during week t.[20] We model the liquidity premium as a linear function of a constant α, a vector of exogenous measures of market liquidity $X_t(\tau)$ in week t for maturity range τ, and a stochastic residual ε_t. The specification flexibly allows for T lags of the exogenous controls to account for autoregressive persistence across periods. The coefficient of interest in equation (1) is δ, which measures the change in the mean liquidity premium while the QE2 program was in operation.

[20]We define the QE2 period as beginning the week the program was announced (November 3, 2010) until the week the program was completed (June 30, 2011).

5.2 Identification

Identification of δ in equation (1) relies on a selection-on-observables strategy. That is, to recover the change in the conditional mean liquidity measure due to the QE2 program necessitates the assumption that observations outside of the QE2 program serve as good controls for variation associated with other observable measures of market liquidity. Further, we must assume that our control measures of liquidity are exogenous and unaffected by the program, which is a valid assumption under the null hypothesis that the program had no effect on market liquidity. Formally, OLS estimation of equation (1) recovers the average treatment effect of the program if $\mathbb{E}[\varepsilon_t|X_t(\tau)] = 0$.

Given that our identification relies on a selection-on-observables strategy, it is essential that our control variables capture potential confounding factors. The control measures we include in our analysis are the VIX options-implied volatility index, the HPW measure of market illiquidity, off-the-run Treasury par-yield spreads, and inflation swap bid-ask spreads. Note that the maturities of the last two variables are chosen to correspond to the maturity of $LP_t(\tau)$. We briefly discuss the motivation for the inclusion of each control variable in the following:

VIX – The VIX represents near-term uncertainty about the general stock market as reflected in options on the Standard & Poor's 500 stock price index and is widely used as a gauge of investor fear and risk aversion. The motivation for including this variable is that elevated economic uncertainty would imply increased uncertainty about the future resale price of any security and therefore could cause liquidity premiums that represent investors' guard against such uncertainty to go up.

HPW Illiquidity Measure – The second control variable we include is a market illiquidity measure introduced in the recent paper by HPW.[21] Their analysis suggests that this measure is a priced risk factor across several financial markets, which we interpret to imply that it represents an economy-wide illiquidity measure that should affect all financial markets.

On-The-Run Treasury Par-Yield Spreads – The third set of control variables are spreads that are typically associated with liquidity in the Treasury market. We use the yield difference between seasoned (off-the-run) Treasury securities and the most recently issued (on-the-run) Treasury security of the same maturity.[22] For each maturity segment in the Treasury yield curve, the on-the-run security is typically the most traded security and therefore penalized the least in terms of liquidity premiums.

[21] The data are publicly available at Jun Pan's website: https://sites.google.com/site/junpan2/publications.

[22] We do not construct on-the-run spreads for the TIPS market since Christensen et al. (2012) show that such spreads have been significantly biased in the years following the peak of the financial crisis due to the value of the deflation protection option embedded in the TIPS contract.

IS Bid-Ask Spreads – To account for liquidity effects in the inflation swap market, we include the bid-ask spreads on inflation swap contracts as reported by Bloomberg for the inflation swap contracts with five- and ten-year maturities. The microstructure frictions that such spreads represent could potentially account for part of the variation in our liquidity premium measure and we want to control for that effect. The bid-ask spreads of the inflation swap contracts exhibit reasonable time variation at levels consistent with transaction costs in the inflation swap market reported by Fleckenstein et al. (2014) based on conversations with traders.[23]

Since the last two variables are measured as yield spreads, we refer to them jointly as the spread series in the following.

5.3 Change in the Conditional Mean

Our first strategy is to estimate the average effect of the QE2 program on our liquidity premium series through OLS estimation of equation (1). Table 2 reports estimates for different sample periods and specifications. Panels A and B report estimates of the change in the conditional mean of the 5-year and 10-year sum of liquidity premiums, respectively. The columns are grouped by the sample used for estimation with columns (1)-(3) corresponding to the full sample. Columns (4)-(6) are estimated on a sample that excludes weeks associated with the financial crisis (1/1/2008 to 6/30/2009), while columns (7)-(9) additionally exclude all weeks associated with other QE programs (QE1 and the MEP). Within each sample grouping, we report three specifications with an increasing number of control measures. Columns (1), (4), and (7) are the same specification and only include an indicator for the QE2 period. Columns (2), (5), and (8) add the VIX volatility measure and the HPW measure as controls. Columns (3), (6), and (9) add the on-the run spreads for Treasuries and the bid-ask spreads for inflation swaps (both measured at maturities corresponding to the maturity of the dependent variable). All regressions include 8 lags of control variables so that $T = 8$ in equation (1), but estimates are not sensitive to increases in the number of lags. Standard errors reported are estimated using the Newey-West (1987) heteroskedasticity and autocorrelation robust estimator with a 26-period (half-year) maximum lag.

[23]We do not include bid-ask spreads for the TIPS from Bloomberg since they are implausibly flat before the spring of 2011, leading us to question the validity of the data. Haubrich et al. (2012) report bid-ask spreads for ten-year TIPS, which are higher than the Bloomberg data, in particular around the peak of the financial crisis in the fall of 2008 and early 2009. Unfortunately, their series ends in May 2010 and cannot be used for our analysis.

	Full Sample			No Financial Crisis			No Financial Crisis or other QE Programs		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
				Panel A: Sum of 5-year Liquidity Premiums					
$\hat{\delta}$ (QE2)	-17.8^{***}	-8.1^{***}	-10.2^{***}	-7.9^{**}	-8.6^{***}	-10.6^{***}	-8.6^{***}	-8.9^{***}	-7.8^{***}
	(6.8)	(2.4)	(2.6)	(3.2)	(2.3)	(2.6)	(2.9)	(2.4)	(2.4)
Adj. R^2	0.03	0.86	0.88	0.05	0.48	0.54	0.10	0.35	0.37
				Panel B: Sum of 10-year Liquidity Premiums					
$\hat{\delta}$ (QE2)	-12.5^{***}	-7.9^{***}	-10.6^{***}	-8.7^{***}	-8.8^{***}	-6.5^{***}	-8.5^{***}	-9.2^{***}	-5.5^{**}
	(3.7)	(2.6)	(3.2)	(2.2)	(2.1)	(2.4)	(2.5)	(1.8)	(2.6)
Adj. R^2	0.04	0.64	0.67	0.11	0.11	0.14	0.13	0.27	0.37
Controls:									
+ VIX & HPW		✓	✓		✓	✓		✓	✓
+ Spreads			✓			✓			✓
N	417	409	409	339	331	331	254	246	246

Table 2: **Estimates of the change in Conditional Mean during QE2**

Panels A and B report OLS estimates of δ from specification (1) with the 5-year and 10-year sums of the liquidity premiums in TIPS and inflation swaps as the dependent variables. Standard errors are reported in parenthesis and are estimated using the Newey and West (1987) heteroskedasticity and autocorrelation robust estimator. Each column reports a specification for a given sample period. Columns (1)-(3) report results for the full sample (1/1/2005 to 12/28/2012), columns (4)-(6) report results for the full sample excluding weeks associated with the financial crisis (1/1/2008 to 6/30/2009), and columns (7)-(9) report results for the full sample excluding the financial crisis and all other periods associated with QE programs (QE1 and the MEP). Columns (1), (4), and (7) are the same specification and only include an indicator for the QE2 period. Columns (2), (5), and (8) add the VIX volatility measure and the HPW illiquidity measure. Columns (3), (6), and (9) add the off-the-run/on-the run Treasury and bid-ask spreads for inflation swaps of the corresponding maturity. All control variables include 8 lags. Stars indicate levels of statistical significance corresponding to p-values less than 0.05 (*), 0.01 (**), and 0.001 (***).

The estimates of the change in the conditional mean recovered with our econometric model are fairly robust to the inclusion of the control variables and the selection of the sample. Furthermore, they are generally significantly different from zero. Controlling for other measures of liquidity our estimates range from -5 to -10 basis points for the 5-year and 10-year liquidity premiums, depending on the sample period and measures included. Notably a negative change is apparent in all specifications, without controlling for any unobserved measures of liquidity. This could suggest our model could be picking up an anomalous period of lower liquidity premiums that happened to coincide with the launch of the QE2 program. We address this concern below using a randomization inference strategy.

We stress that there are several reasons why we must be very cautious in interpreting our estimates as causal. If there are additional omitted factors that our control variables do not capture that are correlated with the QE2 program, our estimate of δ will be biased. In order to address these concerns we include a diverse set of liquidity measures. Given the substantial values of the adjusted R^2 achieved by including these measures of market liquidity, we believe the likelihood that confounding factors are materially biasing our estimates is small.

Secondly, we might be concerned that the inclusion of certain control measures may result in a "bad controls" problem if the QE2 TIPS purchases caused changes in the control variables. While it is possible that the Fed's participation in financial markets could have broad implications, we argue the VIX and HPW measures capture market liquidity that is unlikely to be significantly affected by the Fed's QE2 program. The more direct controls from the Treasury and IS markets may suffer from a stronger bad controls problem. In this case, estimates from (2), (5), and (8) would be preferable to those in (3), (6), and (9).

5.4 Event-Study Analysis

To test the ability of the controls to account for confounding variation, we perform an event-study analysis by estimating the same regressions as in regressions (2), (5), and (8) in Table 2 including a set of indicators for weeks relative to the QE2 announcement. These are recovered with the specification

$$LP_t(\tau) = \alpha + \sum_{s=-8}^{36} \delta_s \mathbf{1}[t - t_0 = s] + \sum_{s=0}^{8} X'_{t-s}(\tau)\beta + \varepsilon_t \tag{2}$$

where t_0 denotes the week before QE2 was announced. We normalize coefficients to the week before the announcement and plot $\delta_s - \delta_0, \forall s \in \{-8, -7, \ldots, 36\}$ in Figure 4. The black line indicates point estimates and the dark shaded area are 95 percent confidence intervals with the same Newey-West standard errors as in Table 2. The dark line indicates week zero, the week before November 3, 2010, when QE2 was announced. The dotted lines and lighter shaded

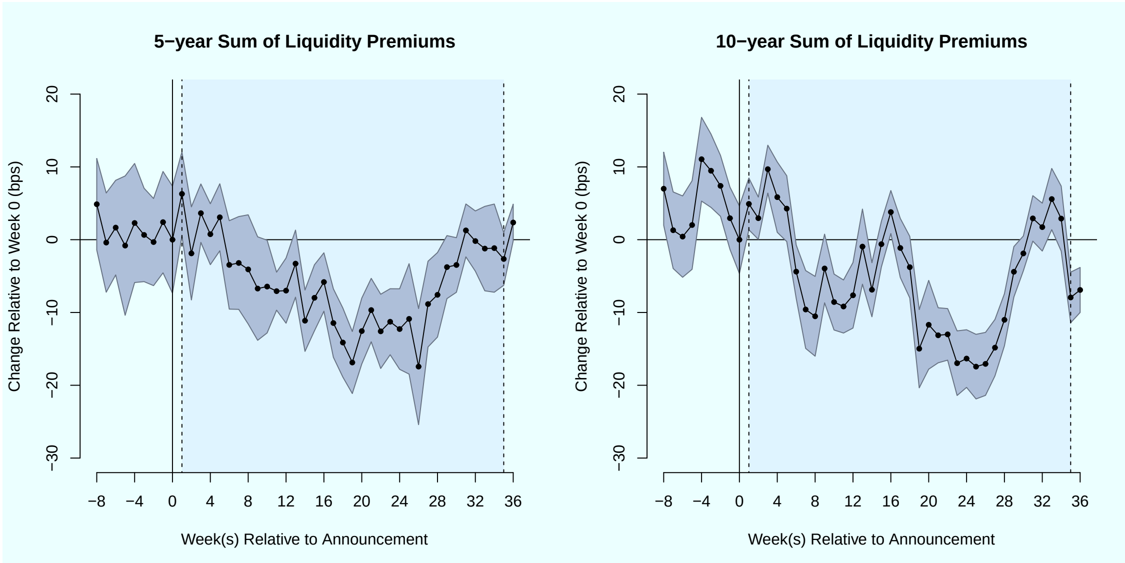

Figure 4: **Effect on Sum of Liquidity Premiums following QE2 Announcement.** The figure shows estimates for the change in the conditional mean during QE2 estimated using specification (8) in Table 2 decomposed by week relative to the start of the program. The left (right) panel shows the distribution for the 5-year (10-year) sum of liquidity premiums in the TIPS and inflation swap markets. The solid line indicates the week prior to the first week of purchases and the shaded area indicates the time during which the Federal Reserve was actively purchasing TIPS with the dotted lines indicating the first and last week of purchases. Standard errors are reported in parenthesis and are estimated using the Newey and West (1987) heteroskedasticity and autocorrelation robust estimator.

area indicate the duration of the QE2 program. Due to our normalization, the coefficient for week zero is zero by construction so that we may read the value of -7 basis points in week 8 of the left panel of Figure 4 as indicating the liquidity premium in that week was 7 basis points lower than the week before the announcement after conditioning on other measures of liquidity.

We might think that the weeks immediately prior to the QE2 announcement are the best control periods since the economic landscape would be the most similar. If the weeks prior to the announcement are in fact good controls, we should observe estimates statistically insignificant from the pre-announcement week zero. We observe this in the 5-year maturity specification clearly with every period being statistically indistinguishable from zero. The 10-year liquidity measure is less robust with only 4 of 8 being insignificantly different, but does not have a clear trend.

We estimate the impact of the QE2 program in an additional way by comparing the residual variation in the immediate pre-program period to that during the program period.

	5-year Sum of Liquidity Premiums	10-year Sum of Liquidity Premiums
$\hat{\delta}^{PRE}$	-1.4	0.9
(9-week pre-period)	(3.5)	(2.5)
$\hat{\delta}^{QE2}$	-9.0^{***}	-9.2^{***}
(QE2 period)	(2.7)	(1.8)
$\hat{\delta}^{QE2} - \hat{\delta}^{PRE}$	-7.6^{**}	-10.0^{***}
(change)	(2.5)	(2.7)
Observations	246	246
Adjusted R^2	0.35	0.27

Table 3: **Event-study estimates of the change in Conditional Mean during QE2**
The top panel reports OLS estimates of from specification (3) for the 5-year and 10-year sums of the liquidity premiums in TIPS and inflation swaps in the left and right columns, respectively. Standard errors are reported in parenthesis and are estimated using the Newey and West (1987) heteroskedasticity and autocorrelation robust estimator. We exclude the financial crisis and all other periods associated with QE programs (QE1 and the MEP) from the sample. We control for the VIX volatility measure and the HPW illiquidity measure with 8 lags in each variable. We also report the difference in the estimated δ^{PRE} and δ^{QE2} coefficients with the associated standard errors. Stars indicate levels of statistical significance corresponding to p-values less than 0.05 (*), 0.01 (**), and 0.001 (***).

To do so, we estimate the following regression:

$$LP_t(\tau) = \alpha + \delta_{PRE} D_t^{PRE} + \delta_{QE2} D_t^{QE2} + \sum_{s=0}^{8} X'_{t-s}(\tau)\beta + \varepsilon_t, \quad (3)$$

where $D_t^{PRE} = 1$ in the 9 weeks prior to the announcement and is zero otherwise. The rest of the variables are identical to the definition in equation (1). Table 3 reports the results of estimating equation (3) using OLS and the same specification as columns (2), (5), and (8) in Table 2 where we exclude the financial crisis and other QE program periods. The columns each report estimates for the two maturity ranges and also perform the statistical test that the mean in the pre-period is different from the mean during the program. Both differences are close to the estimate reported in Table 2 indicating a decline of between 7 and 10 basis points.

That said, Figure 4 puts the difference between the actual realization of the liquidity premium and the counterfactual path into sharper focus for the duration of the QE2 program. Interestingly, in both maturity ranges the measure declines during the first third of the program and then increases back to its level at the program start in a fairly symmetric fashion, suggesting financial market participants may have repeatedly priced the liquidity premiums

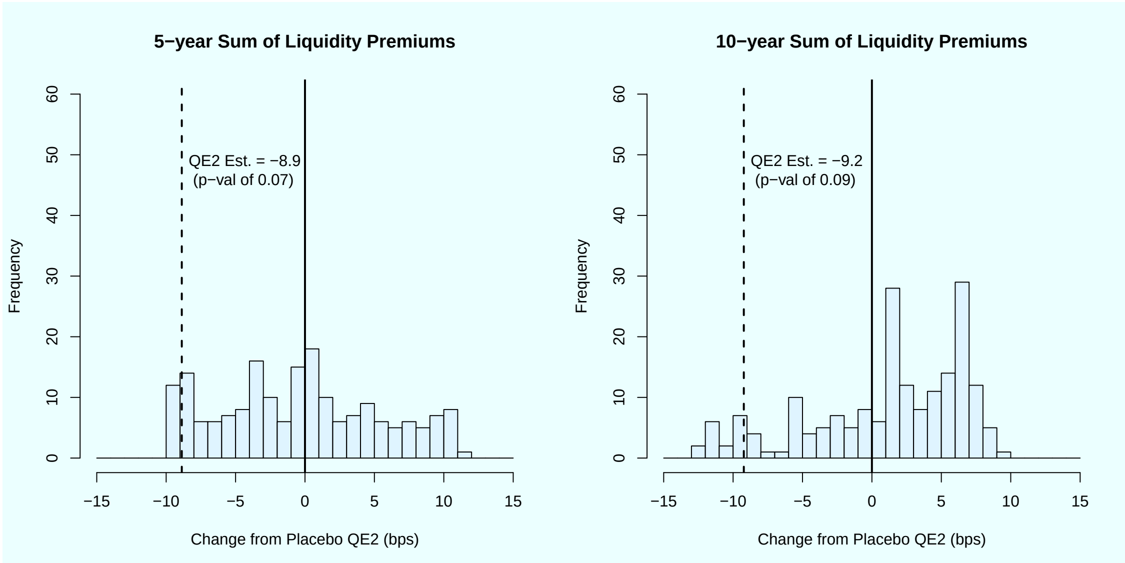

Figure 5: **Results from Randomization Inference Exercise.**
The figure shows the distribution of estimates from specification (8) in Table 2 for the change in the conditional mean for "placebo" periods of the same length as QE2. The left (right) panel shows the distribution for the 5-year (10-year) sum of liquidity premiums in the TIPS and inflation swap markets. Each distribution is composed of 188 estimates with the estimate for the true program period shown by the dotted line.

of TIPS and inflation swaps lower for the first half of the program before gradually returning to pre-program levels.[24]

5.5 Robustness of Specification

Given the clear change in the conditional mean without the inclusion of any control measures, we explore the robustness of our estimates to specification bias. That is, the period in which the Federal Reserve enacted QE2 may have been anomalously more negative than other periods. To test this hypothesis, we perform a randomization inference exercise that randomly assigns a "placebo" treatment period of equal length as QE2 to the sample and recovers an estimate of the effect of these random programs. Specifically, for one placebo estimate we replace the D_t^{QE2} in equation (1) with an indicator equal to one for the consecutive 34-week period beginning at the start of the sample and estimate $\hat{\delta}$. Under the null hypothesis that the effect of the program was zero, we can then construct a p-value that indicates how likely it would have been to observe our QE2 estimates by chance alone.

We perform the exercise using our preferred specification that excludes the financial crisis and other QE program periods from the sample and includes the VIX and HPW measures

[24] Coroneo (2016) performs a counterfactual analysis of the effect of the QE2 TIPS purchases on estimated average TIPS liquidity premiums and report positive results consistent with ours.

with 8 lags. Given our exclusion of various periods, we run a simulation for each period where a sequential set of weeks equal to the length of the QE2 program exists. This results in 188 simulated programs.

Figure 5 shows the results of the randomization inference exercise by plotting the distribution of $\hat{\delta}$'s with each observation representing a placebo estimate. The left panel plots the results for the 5-year measure and the right panel the results for the 10-year measure. The distributions are roughly centered around zero and is consistent with the null hypothesis with means of -0.4 and 0.8 basis points for the 5-year and 10-year measures, respectively. Our estimates for the actual program are indicated by vertical dotted lines in each panel. They lie at the left tail of the distributions and are more negative than 93 percent of the estimates for the 5-year measure and 91 percent of the estimates of the 10-year measure. This indicates that if a program of zero effect had randomly been assigned to the sample period, it would be fairly unlikely to observe the large negative estimates we recover. We believe this provides evidence that it is unlikely that model error is driving our results.

5.6 Effect of QE2 on Other Outcomes

In this section, we estimate the effect of QE2 on other measures in order to better characterize the full effect of the program and the underlying mechanism. Specifically, we estimate the effect of QE2 on TIPS trading volumes and AAA-rated U.S. industrial corporate bond credit spreads using the same econometric model as in equation (1) except replacing the dependent variable with the new outcomes of interest.

For the TIPS trading volume exercise, the dependent variable in the TIPS trading volume analysis is the weekly average of the daily trading volume in the secondary market for TIPS as reported by the Federal Reserve Bank of New York and shown in Figure 6(b).[25] We use the eight-week moving average to smooth out short-term volatility. Figure 6(b) provides suggestive evidence that there was an increase during the program, as compared to the Treasury volume which remained fairly flat. However, as with our identification strategy above, we control for other sources of variation in liquidity to measure the effect of the program.

We also study whether the effects from the liquidity channel extend beyond the targeted securities, which in the case of QE2 were Treasuries and TIPS. We do this by examining highly rated industrial corporate bond credit spreads. If the effect of the program is local, our hypothesis is that the liquidity channel should have no effect as the Fed did not buy any corporate bonds and financial market participants knew this. The dependent variable in this exercise is the excess yield of AAA-rated U.S. industrial corporate bonds over comparable Treasury yields.[26] In choosing the maturity, we face a trade-off. On one side, we would

[25] The trading volume data are available at: http://www.newyorkfed.org/markets/statrel.html.
[26] The corporate bond yield data are from Bloomberg; see Christensen et al. (2014) for details.

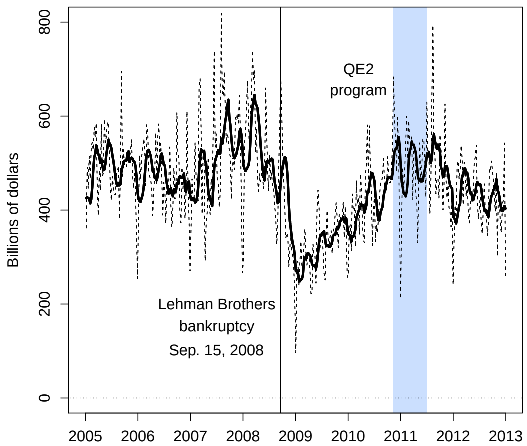

(a) Treasury trading volume.

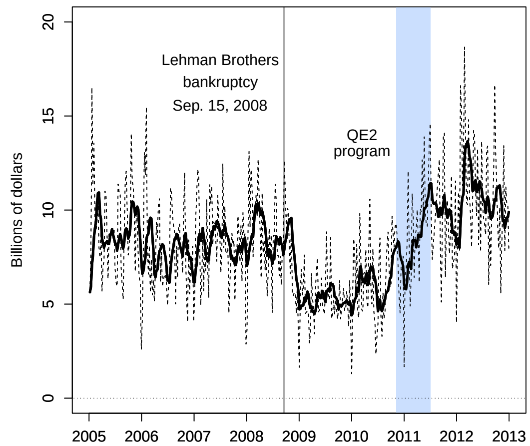

(b) TIPS trading volume.

Figure 6: **Treasury and TIPS Trading Volumes.**
Panel (a) shows the weekly average of daily trading volume in the secondary market for Treasury coupon bonds (dashed black line) and the smoothed eight-week moving average (solid black line). Panel (b) shows the weekly average of daily trading volume in the secondary market for TIPS (dashed black line) and the smoothed eight-week moving average (solid black line).

ideally like to match the maturity of our liquidity premium measure to be consistent with the previous analysis. However, the credit risk of even AAA-rated industrial bond issuers cannot be deemed negligible at a five- to ten-year horizon. On the other hand, if we focus on very short-term debt where the credit risk is entirely negligible, we are far from the desired maturity range. We believe using the two-year credit spread strikes a reasonable balance, and the results are not sensitive to this particular choice. As the credit risk component of such highly rated shorter-term corporate bond yields is minimal, the yield spread largely reflects the premium bond investors require for being exposed to the lower trading volume and larger bid-ask spreads in the corporate bond market vis-à-vis the liquid Treasury bond market.

Table 4 reports the results of the estimates of the change in conditional mean during the program for samples that do not include the financial crisis and those that also exclude other QE programs. The estimates are positive for the TIPS trading volume, but insignificant for all specifications. The AAA credit spreads appear to have a negative change during the course of the program, but are also insignificant. These findings suggest that it is priced frictions and expectations thereof more so than actual liquidity that is affected by the operation of QE programs.

	TIPS Trading Volume				AAA Credit Spread			
	No Financial Crisis		No Financial Crisis or other QE Programs		No Financial Crisis		No Financial Crisis or other QE Programs	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Panel A: Estimates of Conditional Mean							
$\hat{\delta}$ (QE2)	0.44	0.30	0.91	0.25	−4.23	−1.09	−5.17	−2.36
	(0.63)	(0.38)	(0.64)	(0.23)	(2.70)	(2.43)	(2.89)	(2.68)
Controls:								
+ VIX & HPW	✓	✓	✓	✓	✓	✓	✓	✓
+ Spreads		✓		✓		✓		✓
Adj. R^2	0.45	0.60	0.30	0.72	0.36	0.37	0.38	0.40
N	331	331	246	246	289	289	246	246

Table 4: **QE2 Response of TIPS Trading Volume and AAA Credit Spreads**
The table reports the results of regressing the eight-week moving average of the weekly TIPS trading volume series and the excess yield of AAA-rated U.S. industrial corporate bonds over comparable Treasury yields on specification (1) with the VIX and HPW measures as control variables with 8 lags. Columns (1)-(4) report results for the trading volume and (5)-(8) results for the AAA credit spread. The sample used in estimation is shown at the top of the columns with (1), (2), (5), and (6) corresponding to the sample without the financial crisis and (3), (4), (7), and (8) corresponding to the sample that also excludes other QE programs. The AAA credit spread regressions are also limited by the fact that the series we use is only reported through March 9, 2012. We report Newey-West standard errors with 26 lags.

6 Conclusion

In this paper, we argue that one channel of transmission for central bank large-scale asset purchases to long-term interest rates comes about through a reduction of the priced frictions in the targeted security classes.

For evidence we analyze the effects the TIPS purchases included in the Fed's QE2 program had on the functioning of the market for TIPS and the related market for inflation swaps. To quantify the frictions in the markets for these two types of financial claims, we use a model-independent measure of the sum of liquidity premiums in TIPS yields and inflation swap rates constructed from the difference between inflation swap rates and BEI. This measure is ideal for our purposes as it is unaffected by how the QE2 program and its implementation might have changed investors' expectations for economic fundamentals such as inflation and monetary policy.

Our results from regressions with a switch in the conditional mean and a counterfactual analysis both suggest that the TIPS purchases reduced liquidity premiums in the markets for TIPS and inflation swaps. Specifically, our counterfactual analysis indicates that the purchases persistently depressed the liquidity premium measure by an average of about 10 basis points for the duration of the QE2 program from what we would otherwise have expected it to be. In our view, this represents a considerable reduction. Furthermore, and critical to our interpretation, the liquidity premium effects dissipated towards the end of the QE2 TIPS purchases. This leads us to conclude that one benefit of QE programs is to improve financial market functioning by reducing liquidity premiums through a liquidity channel. However, our results also show the limitation of such liquidity effects in that they appear to only be sustained as long as QE purchases are ongoing and expected to continue. Furthermore, such liquidity effects appear, indeed, to be limited to the targeted securities as we find no significant effects on the liquidity premiums of AAA-rated U.S. industrial corporate bonds.

In an attempt to identify local supply effects in individual TIPS prices from the QE2 TIPS purchases, we adapted the approach of DK. However, our analysis did not yield any significant results. The interpretation we offer for this finding is that the liquidity effects we document are general in nature and not tied to any specific TIPS, which would make them go undetected in the analysis of DK. Clearly, a better understanding of the connection between the liquidity effects we document and potential local supply effects would be desirable, but we leave it for future research. Furthermore, we find that reductions in the priced frictions are not accompanied by improvements in actual liquidity as measured by TIPS trading volumes. Thus, more research is needed to better understand the nature and mechanics of the liquidity channel we unveil in this paper.

Our findings could have important policy implications. On a practical note, for understanding BEI and the underlying variation in investors' inflation expectations during the QE2 program, it is crucial to account for the effects we document associated with the liquidity channel. Second, for central banks in countries with somewhat illiquid sovereign bond markets, QE programs targeting sovereign debt could be quite effective in lowering liquidity premiums in addition to any benefits arising from other transmission channels. More generally, it appears that even relatively modest QE programs could have large effects if the targeted security classes are illiquid. Thus, the significance of the liquidity channel could matter for the design of QE programs; time frame, purchase pace, and targeted security classes would all be variables that could make a meaningful difference for the effectiveness of a QE program provided liquidity premiums in the targeted securities are of nontrivial magnitude.

As a final thought, we note that the QE program launched by the European Central Bank in January 2015 could provide rich cross-country data for studying questions related to the liquidity channel highlighted in this paper thanks to the length, size, and implementation of that particular QE program. We encourage others to undertake this research in the future.

Appendix A: Event Study of QE2 Announcement Effects

Response		Maturity					
		5-year	6-year	7-year	8-year	9-year	10-year
Nominal yields	Nov. 2, 2010	122	159	195	227	256	282
	Nov. 3, 2010	118	156	192	227	258	286
	Change	**-4**	**-3**	**-2**	**0**	**2**	**4**
TIPS yields	Nov. 2, 2010	-28	-9	10	27	41	54
	Nov. 3, 2010	-33	-12	8	26	42	56
	Change	**-5**	**-4**	**-2**	**-1**	**0**	**2**
TIPS BEI rates	Nov. 2, 2010	150	168	185	201	215	227
	Nov. 3, 2010	151	168	185	201	216	230
	Change	**1**	**0**	**0**	**1**	**1**	**2**
Inflation swap rates	Nov. 2, 2010	183	199	216	228	238	251
	Nov. 3, 2010	185	199	215	229	237	248
	Change	**2**	**0**	**-1**	**1**	**-1**	**-3**
LP measure	Nov. 2, 2010	32	32	31	28	23	23
	Nov. 3, 2010	34	31	30	28	21	18
	Change	**2**	**0**	**-1**	**0**	**-2**	**-6**

Table 5: **Market Response to QE2 Announcement.**
The table reports the one-day response of nominal Treasury yields, real TIPS yields, TIPS breakeven inflation, inflation swap rates, and the LP measure at six maturities to the announcement of QE2 on November 3, 2010. All numbers are measured in basis points and reported in continuously compounded equivalents. The Treasury and TIPS yields are from Gürkaynak et al. (2007, 2010), while the inflation swap rates are from Bloomberg.

Table 5 summarizes the market reaction to the announcement of the QE2 program on November 3, 2010, using a one-day event window. The key observation is the rather muted response of medium- and long-term Treasury and TIPS yields. Importantly for our analysis, this converts into an even more muted response of TIPS breakeven inflation and inflation swap rates that leave the liquidity premium measure little affected. To put the reported yield changes into perspective, we note that the standard deviation of daily changes in the measure over the period from January 4, 2005, to November 2, 2010, is 5.4 basis points at the five-year maturity and declines monotonically with maturity reaching 4.0 basis points at the ten-year maturity. We take this as evidence that there are no statistically significant effects related to the announcement of the QE2 program that need to be accounted for.

Appendix B: Local Supply Effects in the TIPS Market

In this appendix, we describe our adaptation of DK's analysis to attempt to identify local supply effects in the TIPS market.

First, we introduce notation and define the fundamental statistical objects, which are as follows:

(i). N is the total number of TIPS in existence during the QE2 program.

(ii). $O^n(t)$ equals the notional amount of security n outstanding at t, $n \in \{1, \ldots, N\}$.

(iii). $Q^n(t)$ equals the dollar amount of security n purchased at t, $n \in \{1, \ldots, N\}$.

(iv). $R^n(t) = \frac{P^n(t) - P^n(t-1)}{P^n(t-1)}$ is the daily percentage price change of security n at time t, $n \in \{1, \ldots, N\}$.

(v). T^n is the maturity date of security n, $n \in \{1, \ldots, N\}$.

The second step is to calculate the variables used in the subsequent regressions. Similar to DK, for each security n, we define buckets of substitutes, but limit the number to three buckets due to the smaller number of TIPS trading relative to the number of securities in the market for regular Treasuries.

The first bucket is denoted $S_0(n)$ and only contains security n. For this bucket, two variables are defined:

(i). $O_0^n(t) = O^n(t)$ is the notional amount of security n outstanding.

(ii). $Q_0^n(t) = Q^n(t)$ is the amount of security n purchased at time t.

The second bucket is denoted $S_1(n)$ and contains all securities with maturities within two years of the maturity of security n, that is, $S_1(n) = \{m : |T^m - T^n| \leq 2\}$. Following DK we refer to these securities as the near substitutes for security n.

Finally, the third bucket is denoted $S_2(n)$ and contains all securities with a difference in maturity of more than two years relative to the maturity of security n, that is, $S_2(n) = \{m : |T^m - T^n| > 2\}$. Again, using language similar to DK, we refer to these securities as the far substitutes for security n.

Related to the last two buckets, the following variables are defined:

(i). $O_i^n(t) = \sum_{m \in S_i(n)} O^m(t)$ is the notional amount outstanding of bucket i substitutes for security n at time t, $i \in \{1, 2\}$.

(ii). $Q_i^n(t) = \sum_{m \in S_i(n)} Q^m(t)$ is the amount of bucket i substitutes for security n purchased at time t, $i \in \{1, 2\}$.

As in DK, we use normalized variables in the regressions:

(i). $q_0^n(t) = \frac{Q_0^n(t)}{O_0^n(t) + O_1^n(t)}$ is the amount of security n purchased at time t relative to the notional amount outstanding of security n itself and its near substitutes.

(ii). $q_i^n(t) = \frac{Q_i^n(t)}{O_0^n(t) + O_1^n(t)}$ is the amount of bucket i substitutes for security n purchased at time t relative to the notional amount outstanding of security n itself and its near substitutes, $i \in \{1, 2\}$.

Finally, similar to DK, we run regressions of the daily percentage price change of each TIPS security on a set of variables:

$$R^n(t) = \gamma_0 q_0^n(t) + \gamma_1 q_1^n(t) + \gamma_2 q_2^n(t) + \delta(t) + \alpha^n + \varepsilon^n(t),$$

where

- γ_0 is security n's price elasticity to own purchases,
- γ_1 is security n's price elasticity to purchases of near substitutes,
- γ_2 is security n's price elasticity to purchases of far substitutes,
- $\delta(t)$ represents time fixed effects, and
- α^n represents security fixed effects.

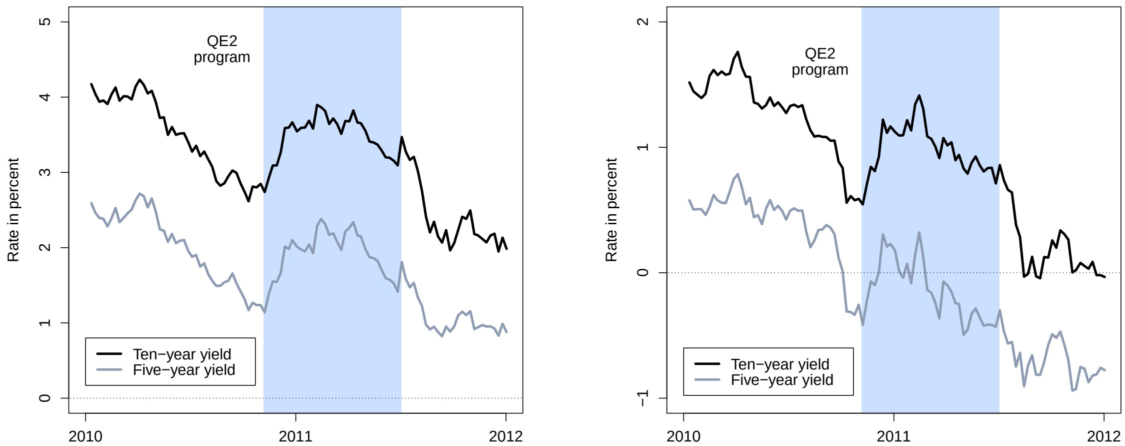

(a) Treasury yields.

(b) TIPS yields.

Figure 7: **Treasury and TIPS Yields.**
Panel (a) illustrates the five- and ten-year Treasury yields from the Gürkaynak et al. (2007) database over the 2010-2011 period. Panel (b) illustrates the five- and ten-year TIPS yields from the Gürkaynak et al. (2010) database over the same period.

The results presented in the main text suggest that the QE2 TIPS purchases led to a sustained reduction in the frictions to trading in the markets for TIPS and inflation swaps. However, the exact channel through which the effects came about is not identified. At face value, the purchases could have lowered liquidity premiums in both markets. Alternatively, if there are local supply effects from the purchases, this would tend to push down TIPS yields, while nominal yields and inflation swap rates presumably would be unaffected in that case. As a consequence, BEI would widen leading to a decline in our liquidity premium measure. In this appendix, to shed light on this latter alternative channel, we attempt to estimate any direct effects on TIPS prices from the QE2 TIPS purchases by replicating the approach of DK.

Assuming the purchased securities are held for a considerable period of time, QE purchases are effectively equivalent to a reduction in the available stock of the targeted securities. The empirical question is whether fluctuations in the supply of government debt should affect yields. Under the expectations hypothesis and in standard term structure models, such supply effects are ruled out. However, models with imperfect asset substitutability or preferred-habitat investors allow for local supply effects on bond yields (see DK for a detailed discussion). Still, as is evident from Figure 7, which shows the changes in the five- and ten-year Treasury and TIPS yields around the time of the QE2 program, the naked eye is a poor guide for detecting such supply effects as both nominal and real yields increased on net during the QE2 program, but the latter less than the former causing BEI to widen as well. Thus, a statistical model is needed to tease out any effects from the asset purchases against this backdrop of generally rising yields. By using security-level data one might hope to be able to identify local supply effects and how they vary across securities with different maturities and liquidity characteristics. To do so, we replicate the approach of DK as briefly summarized in the following. However, we note up front that, unlike the analysis so far, the key element in their approach is to control appropriately for changes in expectations about monetary policy and other economic fundamentals that may affect TIPS prices independent of QE2. Below we will discuss the complications this may entail.

To begin, we follow DK and conduct the regressions in price changes. Second, we drop all TIPS with

Purchases	All TIPS	<10 years to maturity	> 10 years to maturity
Own	-0.023	0.080	-0.035
	(-0.83)	(0.950)	(-1.990)
Near substitutes	-0.068	-0.068	-0.036
(maturity w/in 2 years of own)	(-1.470)	(-0.910)	(-1.100)
Far substitutes	0.008	0.001	0.004
(maturity more than 2 years from own)	(0.560)	(0.030)	(0.460)
# Obs.	427	284	143
# CUSIPs	30	20	10
Adjusted R^2	0.733	0.762	0.953

Table 6: **Flow Effects on Day of Purchase.**
The table reports the results of regressions of the flow effects from the QE2 TIPS purchases as described in the text. The first column reports the results of using all available TIPS with more than two years to maturity, while the following two columns report the result of splitting that sample into one subsample for TIPS with between two and ten years to maturity, and one subsample for TIPS with more than ten years to maturity. T-statistics are reported in parentheses. Asterisks * and ** indicate significance at the 5 percent and 1 percent levels, respectively.

less than two years remaining to maturity at the beginning of the QE2 program because TIPS near maturity have rather erratic price behavior due to the seasonality and general unpredictability of shocks to the headline consumer price index.[27] Third, unlike DK, we only have three maturity buckets related to each security, namely (1) the security itself, (2) the near substitutes with maturities within two years of that of the security, and (3) the far substitutes whose maturities are more than two years from that of the security.

Next, we run regressions of the daily percentage price change of each TIPS security n, denoted $R^n(t)$, on a set of variables:

$$R^n(t) = \gamma_0 q_0^n(t) + \gamma_1 q_1^n(t) + \gamma_2 q_2^n(t) + \delta(t) + \alpha^n + \varepsilon^n(t), \quad (4)$$

where $q_0^n(t)$ represents the normalized amount purchased of security n itself, $q_1^n(t)$ is the normalized amount purchased of near substitutes for security n, while $q_2^n(t)$ is the normalized amount purchased of far substitutes for security n. Thus, the corresponding coefficients can be interpreted as elasticities where γ_0 is security n's price elasticity to own purchases, γ_1 is its price elasticity to purchases of near substitutes, and γ_2 is its price elasticity to purchases of far substitutes. Finally, $\delta(t)$ and α^n represent time and security fixed effects, respectively.

Table 6 reports the regression results for the full sample using all available TIPS with more than two years to maturity as well as the results from two subsamples, one for TIPS with between two and ten years to maturity, the other for TIPS with more than ten years to maturity.[28] We note that all estimated purchase elasticities are insignificant and frequently do not even have the right sign. In short, we are not able to detect any local supply effects in TIPS prices directly.

Various explanations could account for this outcome. First, as emphasized by DK, according to the theory of local supply effects in bond markets (see Vayanos and Vila 2009), they are more likely to matter when liquidity and market functioning is poor, that is, when the arbitrageurs who trade away profit opportunities along the yield curve are capital constrained and are taking on only the most profitable trades, not necessarily

[27] For similar reasons, TIPS with less than two years to maturity are discarded in the construction of the Gürkaynak et al. (2010) TIPS yield curve.

[28] We split the sample around the ten-year maturity point as there is a discrete jump in TIPS outstanding with remaining maturity above ten years, as can also be seen in Figure 2(b).

all available arbitrages. As noted in Figure 3, our measure of TIPS and inflation swap market functioning had reached pre-crisis levels well before the announcement of the QE2 program. Thus, it is indeed possible that market functioning could have been restored and local supply effects would be small for that reason.[29] In addition, we think that there are issues with the specification of the time fixed effects represented by $\delta(t)$. This specification provides a poor proxy for changes in the shape of the yield curve on purchase dates. For example, a level shift in the TIPS yield curve will affect the prices of long-maturity TIPS in a very different way than the prices of short-maturity TIPS.[30] By contrast, the time fixed effect imposes an identical price response across all TIPS. Furthermore, the bias from this misspecification might be more severe in our case than in the analysis of DK for two reasons. First, our pool of TIPS is smaller and more heterogeneous than their sample of regular Treasuries that is dominated by securities with three to ten years remaining to maturity.[31] Second, the limited number of purchase dates in our analysis could matter as well since it allows for less averaging of any errors induced by the misspecified time fixed effects.

To summarize, we believe there are compelling reasons why we are not able to identify any purchase effects on individual TIPS prices from the QE2 TIPS purchases using the approach of DK, despite the clear results we obtain when we analyze the effects on our TIPS and inflation swap liquidity measure. However, we stress that there is not necessarily a contradiction between the two sets of results. One key difference is that our approach based on the liquidity measure is unaffected by changes in expectations about economic fundamentals, unlike the method used by DK which could be severely biased by them. Furthermore, our results suggest that the QE2 TIPS purchase operations led to a reduction in the general frictions to trading in the market for TIPS and the related market for inflation swaps that may not be tied to any specific TIPS. Finally, the liquidity effects we detect are persistent and not limited to a few days around each TIPS purchase operation. Hence, they may go undetected in the approach used by DK that relies on day-to-day variation for its identification.

[29] Using an approach similar to DK, Kandrac and Schlusche (2013) analyze the effects of Treasury securities purchases on Treasury bond prices in all the Fed's QE programs. They find that effects do appear to fade in the later programs.

[30] Figure 7(b) shows that the TIPS curve did experience several level shifts during the QE2 program.

[31] The closer securities are in terms of maturity, the smaller is the room for error from the misspecification of the time fixed effects.

Appendix C: Comparison to Asset Swap Spreads

In this appendix, we elaborate on the construction of our liquidity premium measure described in Section 4 and explain how it relates to the asset swap spreads used in Pflueger and Viceira (2013), which represents an alternative measure of the priced frictions in the TIPS market relative to the Treasury market.

To begin, we believe that assumptions (1) and (2) in Section 4 are uncontroversial because the market for U.S. Treasury bonds is one of the most liquid fixed-income markets. In comparison, TIPS are widely considered to be less liquid. However, assumption (3) that the observed inflation swap rates are above their ideal frictionless rate is less obvious and merits elaboration.

We justify this assumption based on previous research on the mechanics of hedging activity in the inflation swap market. Campbell et al. (2009) suggest that the observed inflation swap rate should be marked up from the unobserved frictionless rate due to the financing and transaction costs of replicating cash flows in related asset swap markets. In practice, there are two main strategies for generating CPI-linked floating cash flows. The first is to buy the TIPS with the desired maturity. This requires funding and implies receiving cash flows on all coupon dates of that security, which investors may not find desirable. The alternative is to enter into a zero-coupon inflation swap of the desired maturity. As its value is zero at inception, there are no funding costs in a zero-coupon inflation swap and investors should be willing to pay a premium for this convenience, which explains why the inflation swap rate can be above BEI in equilibrium. However, the size of the inflation swap rate markup is primarily determined from the supply side. The counterparty to the inflation swap (typically a hedge fund or investment bank) hedges the CPI-linked cash flows by going long in TIPS and short in nominal Treasury bonds through the asset swap market. Thus, the markup represents the compensation the counterparty requires for assuming the liquidity risk of multiple transactions to hedge the contract.

This hedging activity creates a connection between our measure of TIPS and inflation swap liquidity premiums and asset swap rates. In an asset swap, the party who has a long position in the contract pays LIBOR plus a spread while receiving the cash flow of a specific bond without exchange of the principal. In an inflation swap, the party who has a short position in the contract typically generates CPI-linked cash flows by making the following set of transactions at time t to hedge the assumed risk:

- A short position in the τ-year zero-coupon inflation swap struck at $\widehat{IS}_t(\tau)$, that is, the investor will receive $\widehat{IS}_t(\tau)\tau$ at maturity in return for delivering the net change in the price level $\frac{CPI(t+\tau)}{CPI(t)} - 1$.
- A long asset swap position for the τ-year zero-coupon TIPS, that is, agree to paying [LIBOR+$\beta_t^R(\tau)$]τ to receive the fixed accrued coupon $\widehat{y}_t^R(\tau)\tau$ and the accrued inflation compensation $\frac{CPI(t+\tau)}{CPI(t)} - 1$.[32]
- A short asset swap position in the τ-year zero-coupon Treasury bond, that is, agree to paying the nominal Treasury yield $\widehat{y}_t^N(\tau)\tau$ to receive [LIBOR+$\beta_t^N(\tau)$]τ.

Here, $\beta_t^N(\tau)$ and $\beta_t^R(\tau)$ denote the asset swap spreads for the nominal Treasuries and TIPS, respectively. As all transactions involve swaps on zero-coupon assets, there is no outlay upon inception because they all have zero net value and payments are only exchanged at maturity. Table 7 summarizes the outlays and receipts from this set of transactions at maturity. The net receipt to the party who has a short position in the inflation swap is

$$[\widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)) + \beta_t^N(\tau) - \beta_t^R(\tau)]\tau \geq 0. \quad (5)$$

Note that this strategy is really a hedge as the value on the left-hand side of equation (5) is deterministic and

[32] Here, we are neglecting the value of the deflation protection in the TIPS in that the actual payment on the TIPS asset swap is $\max[\frac{CPI(t+\tau)}{CPI(t)} - 1, 0]$. We thank Xiaopeng Zhang for pointing this out. Thus, the calculations are accurate provided the value of the deflation protection for the particular TIPS in the asset swap is negligible. If not, they can be corrected by calculating its value in a way similar to the one described in Appendix D.

Receipts	$\widehat{IS}_t(\tau)\tau+\widehat{y}_t^R(\tau)\tau+\frac{CPI(t+\tau)}{CPI(t)}$-1 +[LIBOR+$\beta_t^N(\tau)$]$\tau$
Payments	$\widehat{y}_t^N(\tau)\tau$ + $\frac{CPI(t+\tau)}{CPI(t)}$-1 +[LIBOR+$\beta_t^R(\tau)$]$\tau$
Net receipts	$[\widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau)) + \beta_t^N(\tau) - \beta_t^R(\tau)]\tau$

Table 7: **Cash Flow of Investment Strategy that Hedges a Short Position in a Zero-Coupon Inflation Swap Contract.**
Illustration of the cash flows involved in the investment strategy that hedges a short position in an inflation swap. It involves: (1) The inflation swap position itself, (2) a long asset swap position in the τ-year zero-coupon TIPS, and (3) a short asset swap position in the τ-year zero-coupon Treasury bond.

set at the inception of the contract. For the leveraged investor to be willing to participate in the inflation swap market this value must be nonnegative as indicated. Since we define our liquidity premium measure as $LP_t(\tau) = \widehat{IS}_t(\tau) - (\widehat{y}_t^N(\tau) - \widehat{y}_t^R(\tau))$, the inequality in equation (5) can also be written as

$$LP_t(\tau) + \beta_t^N(\tau) - \beta_t^R(\tau) \geq 0. \tag{6}$$

Campbell et al. (2009) note that, normally, the asset swap spreads are negative and more so for the nominal Treasuries, that is,[33]

$$\beta_t^N(\tau) < \beta_t^R(\tau) \leq 0. \tag{7}$$

Under competitive circumstances (zero cost of entry to the inflation swap market, etc.), we expect equation (6) to hold with equality. Using the inequality in equation (7), we can then rewrite equation (6) as

$$LP_t(\tau) = \beta_t^R(\tau) - \beta_t^N(\tau) > 0. \tag{8}$$

Thus, our liquidity premium measure should equal the difference between the TIPS and nominal Treasury asset swap spreads and be strictly positive. Of course, this is an idealized calculation based on zero-coupon bonds, but the difference between asset swap spreads on TIPS and regular Treasuries should still provide a good approximation to our measure.

In Figure 8, we compare the ten-year TIPS-Treasury asset swap spread used in Pflueger and Viceira (2013) to our ten-year liquidity premium measure.[34] We find that the two measures are highly positively correlated and of approximately the same magnitude as theory would suggest. Unfortunately, this asset swap spread series is only monthly and therefore not useful for our empirical analysis. Still, we note that the series exhibit a temporary drop during the QE2 program with a peak trough in April 2011 that matches well with the variation in our liquidity premium series.

Finally, we stress that, without additional information, any combination of nonnegative $\delta_t^R(\tau)$ and $\delta_t^{IS}(\tau)$ that satisfies the condition $LP_t(\tau) = \delta_t^R(\tau) + \delta_t^{IS}(\tau)$ is admissible and cannot be legitimately excluded ex ante. Also, we underscore that our construction is valid for any sample of nominal Treasury and real TIPS yields as long as our three key assumptions are satisfied by the data. This observation implies that the size and shape of our measure depend on the underlying pool of bonds, the method used in the yield curve construction, etc., but that its interpretation remains valid despite such differences.

[33] During the financial crisis of 2008 and 2009, $\beta_t^R(\tau)$ turned positive, but the relative relationship between $\beta_t^N(\tau)$ and $\beta_t^R(\tau)$ remained as indicated by the first inequality.

[34] We thank Carolin Pflueger for sharing their data.

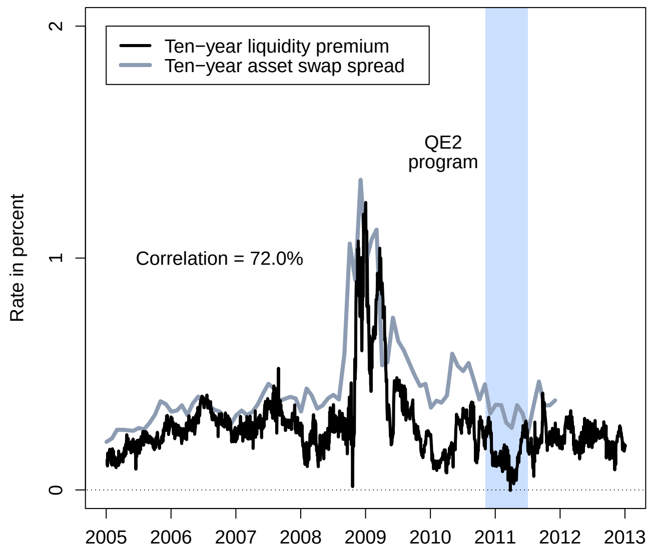

Figure 8: **Comparison of Liquidity Premium Measure and Asset Swap Spreads.** Illustration of the ten-year TIPS and inflation swap liquidity premium measure, which is daily covering the period from January 4, 2005, to December 31, 2012, a total of 1,977 observations. Also shown is the difference between ten-year TIPS and Treasury asset swap spreads used in Pflueger and Viceira (2013), which is a monthly series covering the period January 2005 to December 2011, a total of 84 observations.

Appendix D: The TIPS Purchases and Sales during the MEP

In this appendix, we provide a brief description of the Federal Reserve's maturity extension program (MEP) that included purchases and sales of a sizable amount of TIPS.

The MEP program was announced on September 21, 2011. At first, it was intended to run through June 2012 and involve buying $400 billion of Treasury securities with more than 6 years to maturity financed by selling a similar amount of Treasury securities with less than 3 years to maturity. At the June 2012 FOMC meeting it was decided to continue the MEP through the end of 2012 at which point it would total more than $600 billion in purchases and sales of securities. Similar to the QE2 program, the MEP involved transactions in TIPS the effects of which we briefly detail and analyze below.

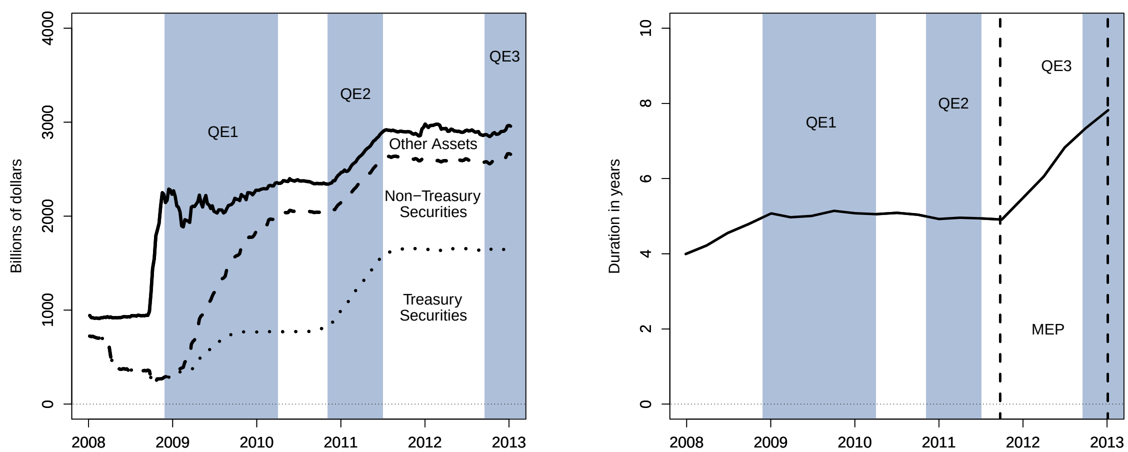

(a) Fed's assets. (b) Duration of Fed's Treasury securities.

Figure 9: **The Fed's Assets and the Duration of Its Treasury Securities.**

For a start, though, Figure 9 shows how the Fed's asset holdings have changed since 2008. We note that the first asset purchase program (QE1) consisted of a modest expansion of its Treasury securities holdings combined with substantial purchases of mortgage-backed securities (MBS). During the QE2 program it was only the Treasury holdings that increased, while the MEP analyzed in this appendix barely changed the size of the Fed's balance sheet. However, obviously, it did achieve the intended goal of increasing the average maturity of the Fed's securities holdings. This is illustrated in Figure 9(b), which shows the change in the average duration of the Fed's nominal Treasury securities since 2008.[35] The weighted average duration increased from about five years to almost eight years over the course of the MEP.

Like the QE2 program, the MEP was implemented with a fairly regular schedule. Once a month, the Fed publicly released a list of operation dates for the following 30-plus day period, indicating the relevant maturity range and expected purchase and sale amount for each operation.[36] There were 15 separate TIPS purchase operation dates, effectively once a month, each with a stated expected purchase amount of $1 billion to $2 billion. TIPS were the only type of asset purchased on these dates. In addition, there were 10 separate TIPS

[35]The durations are calculated based on real-time quarterly estimation of the shadow-rate term structure model analyzed in Christensen et al. (2015). This model respects the zero lower bound for yields, which has been a prominent characteristic of the Treasury yield curve since 2009.

[36]The information can be found at http://www.newyorkfed.org/markets/tot_operation_schedule.html.

MEP TIPS purchase operation dates	TIPS purchases (Mill.)	Weighted avg. maturity (years)
(1) Oct. 5, 2011	$1,861	22.77
(2) Nov. 3, 2011	$1,916	25.62
(3) Dec. 12, 2011	$1,872	25.02
(4) Jan. 10, 2012	$1,905	28.56
(5) Feb. 10, 2012	$1,926	26.98
(6) Mar. 14, 2012	$1,272	27.53
(7) Apr. 3, 2012	$1,765	19.01
(8) May 9, 2012	$1,565	15.44
(9) Jun. 15, 2012	$1,730	16.29
(10) Jul. 10, 2012	$1,809	21.08
(11) Aug. 9, 2012	$1,947	24.58
(12) Sep. 10, 2012	$1,979	26.77
(13) Oct. 11, 2012	$1,819	24.67
(14) Nov. 9, 2012	$1,939	25.49
(15) Dec. 11, 2012	$1,829	23.01
Average	$1,809	23.52

Table 8: **The MEP TIPS Purchase Operation Dates.**
The table reports the amount and weighted average maturity of TIPS purchased on the 15 TIPS purchase operation dates during the MEP. The TIPS purchase amounts are reported in millions of dollars, while the weighted average maturities are measured in years.

MEP TIPS sale operation dates	TIPS sales (Mill.)	Weighted avg. maturity (years)
(1) Oct. 17, 2011	$1,456	2.17
(2) Nov. 9, 2011	$1,376	1.20
(3) Dec. 7, 2011	$1,353	0.74
(4) Jan. 5, 2012	$1,367	0.73
(5) Feb. 7, 2012	$1,407	1.68
(6) Mar. 5, 2012	$1,415	1.66
(7) Apr. 9, 2012	$1,289	0.37
(8) May 2, 2012	$1,427	2.15
(9) Jun. 11, 2012	$1,146	2.24
(10) Oct. 19, 2012	$1,198	2.87
Average	$1,343	1.58

Table 9: **The MEP TIPS Sale Operation Dates.**
The table reports the amount and weighted average maturity of TIPS sales on the 10 TIPS sale operation dates during the MEP. The TIPS sale amounts are reported in millions of dollars, while the weighted average maturities are measured in years.

sale operation dates distributed with sale operations once a month from October 2011 to June 2012 plus a final sale operation in mid-October 2012. These all had indicated expected sale amounts of $1 billion to $1.5 billion.

One complicating factor in analyzing the MEP relative to the QE2 program is that not all TIPS were eligible in each operation. The sales were targeting TIPS with less than 3 years to maturity,[37] while the purchases were targeted at TIPS with more than 6 years to maturity. However, given that this would remain true throughout the operation of the MEP, this should show up as an announcement effect when the MEP was first introduced in September 2011, but not change from day to day during the implementation of the program. Thus, we proceed with an analysis similar to the one we used to analyze the effects of the QE2 program.

Table 8 lists the 15 TIPS purchase operation dates during the MEP, while Table 9 reports the corresponding statistics for the 10 TIPS sale operation dates. The MEP TIPS purchases totaled $27.1 billion, all of which involved TIPS with more than 6 years to maturity. The MEP TIPS sales totaled $13.4 billion and only included TIPS with less than 3.5 years to maturity. Thus, the net TIPS purchases in the MEP were $13.7 billion stretched out over a 15-month period. In comparison, the QE2 program involved almost twice the amount of net TIPS purchases and took less than half the time to implement.[38] Hence, based on these statistics, the QE2 program can be viewed as four times more intense than the MEP in terms of the operations related to the TIPS market. All else equal, this would suggest that liquidity channel effects from the MEP could be expected to be substantially smaller than the effects we reported for the QE2 program.

[37] Except for the sale of $572 million of a 3.24-year TIPS on October 19, 2012, all TIPS sold during the MEP had maturities less than 3 years.

[38] The QE2 TIPS purchases ran from November 23, 2010, to June 17, 2011, a 206-day period, while the MEP TIPS operations were implemented from October 5, 2011, to December 11, 2012, a 433-day period.

Appendix E: Fed TIPS Purchases outside QE2 and the MEP

TIPS purchase operation dates	TIPS purchases (mill.)	Weighted avg. maturity (years)
(1) Apr. 16, 2009	\$1,619	11.74
(2) May 26, 2009	\$1,562	2.16
(3) Jul. 16, 2009	\$1,525	3.30
(4) Aug. 30, 2010	\$398	10.20
(5) Sep. 28, 2010	\$655	10.86
(6) Oct. 10, 2010	\$788	10.74
Average	\$1,091	8.17

Table 10: **TIPS Purchase Operation Dates outside QE2 and the MEP.**
The table reports the amount and weighted average maturity of TIPS purchased on the six TIPS purchase operation dates outside QE2 and the MEP. The TIPS purchase amounts are reported in millions of dollars, while the weighted average maturities are measured in years.

Table 10 contains information for the six TIPS purchase operations that were included as part of the Treasury securities purchases in the QE1 program in 2009 and during the re-investment program that was initiated in the months before the announcement of the QE2 program. The total amount purchased was \$6.1 billion. The three TIPS purchase operations in 2009 were close in size to the ones during QE2, but six weeks apart instead of biweekly. On the other hand, the three TIPS purchase operations in the fall of 2010 had a frequency not that different from the schedule operated during QE2, but the purchased amounts were about one-third of the purchase amounts during QE2. Thus, in both cases, the intensity of the TIPS purchases was but a fraction of that experienced during the QE2 program and for that reason we choose not to analyze these TIPS purchases further.

References

Abrahams, Michael, Tobias Adrian, Richard K. Crump, and Emanuel Moench, 2015, "Decomposing Real and Nominal Yield Curves," Federal Reserve Bank of New York Staff Reports #570.

Andreasen, Martin M., Jens H. E. Christensen, Kevin Cook, and Simon Riddell, 2016, "The TIPS Liquidity Premium," Manuscript, Federal Reserve Bank of San Francisco.

Bauer, Michael D. and Glenn D. Rudebusch, 2014, "The Signaling Channel for Federal Reserve Bond Purchases," *International Journal of Central Banking*, Vol. 10, No. 3, 233-289.

Brunnermeier, Markus K. and Lasse Heje Pedersen, 2009, "Market Liquidity and Funding Liquidity," *Review of Financial Studies*, Vol. 22, No. 6, 2201-2238.

Campbell, John Y., Robert J. Shiller, and Luis M. Viceira, 2009, "Understanding Inflation-Indexed Bond Markets," *Brookings Papers on Economic Activity*, Spring, 79-120.

Cheridito, Patrick, Damir Filipović, and Robert L. Kimmel, 2007, "Market Price of Risk Specifications for Affine Models: Theory and Evidence," *Journal of Financial Economics*, Vol. 83, No. 1, 123-170.

Christensen, Jens H. E., 2008, "The Corporate Bond Credit Spread Puzzle," Federal Reserve Bank of San Francisco, *Economic Letter* 2008-10.

Christensen, Jens H. E. and James M. Gillan, 2012, "Do Fed TIPS Purchases Affect Market Liquidity?" Federal Reserve Bank of San Francisco, *Economic Letter* 2012-07.

Christensen, Jens H. E. and Signe Krogstrup, 2016a, "A Portfolio Model of Quantitative Easing," Working Paper 2016-12, Federal Reserve Bank of San Francisco.

Christensen, Jens H. E. and Signe Krogstrup, 2016b, "Transmission of Quantitative Easing: The Role of Central Bank Reserves," Working Paper 2014-39, Federal Reserve Bank of San Francisco.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2012, "Extracting Deflation Probability Forecasts from Treasury Yields," *International Journal of Central Banking*, Vol. 8, No. 4, 21-60.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2014, "Do Central Bank Liquidity Facilities Affect Interbank Lending Rates?" *Journal of Business and Economic Statistics*, Vol. 32, No. 1, 136-151.

Christensen, Jens H. E., Jose A. Lopez, and Glenn D. Rudebusch, 2015, "A Probability-Based Stress Test of Federal Reserve Assets and Income," *Journal of Monetary Economics*, Vol. 73, 26-43.

Christensen, Jens H. E. and Glenn D. Rudebusch, 2012, "The Response of Interest Rates to U.S. and U.K. Quantitative Easing," *Economic Journal*, Vol. 122, F385-F414.

Collin-Dufresne, Pierre, Robert S. Goldstein, and J. Spencer Martin, 2001, "The Determinants of Credit Spread Changes," *Journal of Finance*, Vol. 56, No. 6, 2177-2207.

Coroneo, Laura, 2016, "TIPS Liquidity Premium and Quantitative Easing," Manuscript, University of York, Department of Economics and Related Studies.

D'Amico, Stefania, Don H. Kim, and Min Wei, 2014, "Tips from TIPS: The Informational Content of Treasury Inflation-Protected Security Prices," Finance and Economics Discussion Series 2014-24, Board of Governors of the Federal Reserve System.

D'Amico, Stefania and Thomas B. King, 2013, "The Flow and Stock Effects of Large-Scale Treasury Purchases: Evidence on the Importance of Local Supply," *Journal of Financial Economics*, Vol. 108, No. 2, 275-564.

De Pooter, Michiel, Robert F. Martin, and Seth Pruitt, 2015, "The Liquidity Effects of Official Bond Market Intervention," Manuscript. Federal Reserve Board of Governors.

Driessen, Joost, Theo E. Nijman, and Zorka Simon, 2014, "The Missing Piece of the Puzzle: Liquidity Premiums in Inflation-Indexed Markets," Manuscript. Department of Finance, Tilburg University.

Dudley, William C., Jennifer Roush, and Michelle Steinberg Ezer, 2009, "The Case for TIPS: An Examination of the Costs and Benefits," *Federal Reserve Bank of New York Economic Policy Review*, Vol. 15, No. 1, 1-17.

Eser, Fabian and Bernd Schwaab, 2016, "Evaluating the Impact of Unconventional Monetary Policy Measures: Empirical Evidence from the ECB's Securities Markets Programme," *Journal of Financial Economics*, Vol. 119, No. 1, 147-167.

Fleckenstein, Mathias, Francis A. Longstaff, and Hanno Lustig, 2014, "The TIPS-Treasury Bond Puzzle," *Journal of Finance*, Vol. 69, No. 5, 2151-2197.

Fleming, Michael J. and Neel Krishnan, 2012, "The Microstructure of the TIPS Market," *Federal Reserve Bank of New York Economic Policy Review*, Vol. 18, No. 1, 27-45.

Fleming, Michael J. and John R. Sporn, 2013, "Trading Activity and Price Transparency in the Inflation Swap Market," *Federal Reserve Bank of New York Economic Policy Review*, Vol. 19, No. 1, 45-57.

Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack, 2011, "Large-Scale Asset Purchases by the Federal Reserve: Did They Work?" *International Journal of Central Banking*, Vol. 7, No. 1, 3-43.

Greenwood, Robin and Dimitri Vayanos, 2014, "Bond Supply and Excess Bond Returns," *Review of Financial Studies*, Vol. 27, No. 3, 663-713.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2007, "The U.S. Treasury Yield Curve: 1961 to the Present," *Journal of Monetary Economics*, Vol. 54, No. 8, 2291-2304.

Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright, 2010, "The TIPS Yield Curve and Inflation Compensation," *American Economic Journal: Macroeconomics*, Vol. 2, No. 1, 70-92.

Hamilton, James D. and Jing Wu, 2012, "The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment," *Journal of Money, Credit, and Banking*, Vol. 44, No. 1 (Supplement), pp. 3-46.

Hattori, Masazumi, Andreas Schrimpf, and Vladyslav Sushko, 2016, "The Response of Tail Risk Perceptions to Unconventional Monetary Policy," *American Economic Journal: Macroeconomics*, forthcoming.

Haubrich, Joseph, George Pennacchi, and Peter Ritchken, 2012, "Inflation Expectations, Real Rates, and Risk Premia: Evidence from Inflation Swaps," *Review of Financial Studies*, Vol. 25, No. 5, 1588-1629.

Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, "Noise as Information for Illiquidity," *Journal of Finance*, Vol. 68, No. 6, 2341-2382.

Joyce, Michael A. S., Ana Lasaosa, Ibrahim Stevens, and Matthew Tong, 2011, "The Financial Market Impact of Quantitative Easing in the United Kingdom," *International Journal of Central Banking*, Vol. 7, No. 3, 113-161.

Kandrac, John, 2013, "Have Federal Reserve MBS Purchases Affected Market Functioning?," *Economics Letters*, Vol. 121, 188-191.

Kandrac, John, 2014, "The Costs of Quantitative Easing: Liquidity and Market Functioning Effects of Federal Reserve MBS Purchases," Manuscript, Board of Governors of the Federal Reserve System.

Kandrac, John and Bernd Schlusche, 2013, "Flow Effects of Large-Scale Asset Purchases," *Economics Letters*, Vol. 121, 330-335.

Kettemann, Andreas and Signe Krogstrup, 2014, "Portfolio Balance Effects of the Swiss National Bank's Bond Purchase Program," *Journal of Macroeconomics*, Vol. 40, 132-149.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen, 2011, "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity*, Fall 2011, 215-265.

Li, Canlin and Min Wei, 2013, "Term Structure Modeling with Supply Factors and the Federal Reserve's Large-Scale Asset Purchase Programs," *International Journal of Central Banking*, Vol. 9, No. 1, 3-39.

Lou, Dong, Yan Hongjun, and Jinfan Zhang, 2013, "Anticipated and Repeated Shocks in Liquid Markets," *Review of Financial Studies*, Vol. 26, No. 8, 1890-1912.

Meaning, Jack and Feng Zhu, 2011, "The Impact of Recent Central Bank Asset Purchase Programmes," *BIS Quarterly Review*, December, 73-83.

Newey, Whitney K. and Kenneth D. West, 1987, "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 1987, vol. 55, No. 3, 703-708.

Pasquariello, Paolo, 2015, "Government Intervention and Arbitrage," Manuscript, Department of Finance, Ross School of Business, University of Michigan.

Pflueger, Carolin E. and Luis M. Viceira, 2013, "Return Predictability in the Treasury Market: Real Rates, Inflation, and Liquidity," Manuscript, Harvard Business School.

Sack, Brian and Robert Elsasser, 2004, "Treasury Inflation-Indexed Debt: A Review of the U.S. Experience," *Federal Reserve Bank of New York Economic Policy Review*, Vol. 10, No. 1, 47-63.

Tobin, James, 1969, "A General Equilibrium Approach to Monetary Theory," *Journal of Money, Credit and Banking*, Vol. 1, No. 1, 15-29.

Vayanos, Dimitri and Jean-Luc Vila, 2009, "A Preferred-Habitat Model of the Term Structure of Interest Rates," NBER Working Paper 15487.