

The Signaling Channel for Federal Reserve Bond Purchases*

Michael D. Bauer[†], Glenn D. Rudebusch[‡]

First draft: September 14, 2011

This version: August 6, 2012

Abstract

Previous research has emphasized the portfolio balance effects of Federal Reserve bond purchases, in which a reduced bond supply lowers term premia. In contrast, we find that such purchases have important signaling effects that lower expected future short-term interest rates. Our evidence comes from a model-free analysis and from dynamic term structure models that decompose declines in yields following Fed announcements into changes in risk premia and expected short rates. To overcome problems in measuring term premia, we consider bias-corrected model estimation and restricted risk price estimation. We also characterize the estimation uncertainty regarding the relative importance of the signaling and portfolio balance channels.

Keywords: monetary policy, zero lower bound, QE, LSAP, signaling, portfolio balance, no arbitrage

JEL Classifications: E43, E52

*The views expressed herein are those of the authors and not necessarily shared by others at the Federal Reserve Bank of San Francisco or in the Federal Reserve System.

[†]Federal Reserve Bank of San Francisco, michael.bauer@sf.frb.org

[‡]Federal Reserve Bank of San Francisco, glenn.rudebusch@sf.frb.org

1 Introduction

During the recent financial crisis and ensuing deep recession, the Federal Reserve reduced its target for the federal funds rate—the traditional tool of U.S. monetary policy—essentially to the lower bound of zero. In the face of deteriorating economic conditions and with no scope for further cuts in short-term interest rates, the Fed initiated an unprecedented expansion of its balance sheet by purchasing large amounts of Treasury debt and federal agency securities of medium and long maturities.¹ Other central banks in comparable circumstances have taken broadly similar actions. Notably, the Bank of England also purchased longer-term debt during the financial crisis, and the Bank of Japan, when confronted over a decade ago with stagnation and near-zero short-term rates, purchased debt securities in its program of Quantitative Easing (QE).²

The goal of the Fed’s large-scale asset purchases (LSAPs) of bonds was to put downward pressure on longer-term yields in order to ease financial conditions and support economic growth. Using a variety of approaches, several studies have concluded that the Fed’s LSAP program was effective in lowering various interest rates below levels that otherwise would have prevailed (D’Amico and King, 2010; Gagnon et al., 2011; Hamilton and Wu, 2012a; Krishnamurthy and Vissing-Jorgensen, 2011). However, researchers do not yet fully understand the underlying mechanism and causes for the declines in long-term interest rates. Based on the usual decomposition of yields on safe long-term government bonds, there are two potential elements that central bank bond purchases could affect: the term premium and the average level of short-term interest rates over the maturity of the bond, also known as the risk-neutral rate. The term premium could have fallen because the Fed’s LSAPs reduced the amount of longer-term bonds in private-sector portfolios—which is loosely referred to as the *portfolio balance channel*. Alternatively, the LSAP announcements could have led market participants to revise down their expectations for future short-term interest rates, lengthening, for example, the expected period of a near-zero federal funds rate target. Such a *signaling channel* for LSAPs would reduce yields by lowering the average expected short-rate (or risk-neutral) component of long-term rates.

Much discussion of the financial market effects of the Fed’s bond purchases treats the portfolio balance channel as the key channel for that impact. For example, Chairman Ben Bernanke (2010) described the effects of the Fed’s bond purchases in this way:

¹The federal agency securities were debt or mortgage-backed securities that had explicit or implicit credit protection from the U.S. government.

²The Fed’s actions led to a larger central bank balance sheet and higher bank reserves much like the Bank of Japan’s QE; however, the Fed’s purchases were focused on longer-maturity assets.

I see the evidence as most favorable to the view that such purchases work primarily through the so-called portfolio balance channel, which holds that once short-term interest rates have reached zero, the Federal Reserve's purchases of longer-term securities affect financial conditions by changing the quantity and mix of financial assets held by the public. Specifically, the Fed's strategy relies on the presumption that different financial assets are not perfect substitutes in investors' portfolios, so that changes in the net supply of an asset available to investors affect its yield and those of broadly similar assets.

Along with central bank policy makers, researchers have also favored the portfolio balance channel in accounting for the effects of LSAPs. The most influential evidence supporting a portfolio balance channel has come from event studies that examine changes in asset prices following announcements of central bank bond purchases. Notably, Gagnon et al. (2011), henceforth GRRS, examine changes in the ten-year Treasury yield and Treasury yield term premium.³ They document that after eight key LSAP announcements, the ten-year yield fell by a total of 91 basis points (bps), while their measure of the ten-year term premium, which is based on the model of Kim and Wright (2005), fell by 71 bps. Based largely on this evidence, the authors argue that the Fed's LSAPs primarily lowered long-term rates through a portfolio balance channel that reduced term premia.

In this paper, we reexamine the notion that the signaling of lower future policy rates through LSAP announcements played a negligible role in lowering Treasury yields. First, we argue that the estimated contribution of policy expectations to decreases in long-term yield is likely a conservative measure of the importance of the signaling channel. For example, conventional monetary policy actions that signal lower future short rates tend to lower term premia as well. Therefore, assuming that all changes in term premia can be attributed to the portfolio balance channel is likely to underestimate the signaling effects of LSAPs.

We also provide model-free evidence suggesting that the Fed's actions lowered yields to a considerable extent by changing policy expectations about the future path of the federal funds rate. Under a market segmentation assumption that LSAPs primarily affected security-specific term premia in Treasury markets, changes after LSAP announcements in spreads between Treasury yields and money market and swap rates of comparable maturity illuminate the contribution of the portfolio balance channel. Joyce et al. (2011), for example, argue that increases in spreads between U.K. Treasury and swap yields following Bank of England QE announcements support a portfolio balance channel. In contrast, in the United States, we find

³Other event studies include Joyce et al. (2011), Neely (2010), Krishnamurthy and Vissing-Jorgensen (2011), and Swanson (2011).

that a large portion of the observed yield changes was also reflected in lower money market and swap rates. This suggests that the expectations component may make an important contribution to the declines in yields.

Our main contribution is to provide new model-based evidence that addresses two key statistical problems in decomposing the yield curve in previous studies—namely, small-sample bias and statistical uncertainty. We reconsider the GRRS results that are based on the Kim-Wright decompositions of yields into term premia and risk-neutral rates using a conventional arbitrage-free dynamic term structure model (DTSM). Although DTSMs are the workhorse model in empirical fixed income finance, they have been very difficult to estimate and are plagued by biased coefficient estimates as described by previous studies (e.g., Duffee and Stanton, 2004; Kim and Orphanides, 2005; and Bauer et al., forthcoming, henceforth BRW). Therefore, to get better measures of the term premium, we examine two alternative estimates of the DTSM. The first is obtained from a novel estimation procedure—following BRW—that directly adjusts for the small-sample bias in estimation of a maximally flexible DTSM. Since conventional biased DTSM estimates—like the Kim-Wright model that GRRS rely on—overstate the speed of mean reversion of the short rate, the model-implied forecast of the short rate is too close to the unconditional mean. Consequently, too much of the variation in forward rates is attributed to the term premium component. Intuitively then, conventional biased DTSM estimates understate the importance of the signaling channel. Indeed, we find that an LSAP event study using term premia obtained from DTSM estimates with reduced bias finds a larger role for the signaling channel. Our second estimation approach imposes restrictions on the risk pricing as in Bauer (2011). Intuitively, under restricted risk pricing, the cross-sectional interest rate dynamics, which are estimated very precisely, are being used to pin down the time series parameters. This reduces both small-sample bias and statistical uncertainty, so that short rate forecasts and term premium estimates are more reliable (Cochrane and Piazzesi, 2008; Joslin et al., 2010; Bauer, 2011). Here, too, we find a more substantial role for the signaling channel than is commonly acknowledged.

Importantly, we quantify the statistical uncertainty surrounding the DTSM-based estimates of the relative contributions of the portfolio balance and signaling channels. In particular, we take into account the parameter uncertainty that underlies estimates of the term premium and produce confidence intervals that reflect this estimation uncertainty. Our confidence intervals reveal that with a largely unrestricted DTSM, as is common in the literature, definitive conclusions about the relative importance of term premia and expectations effects of LSAPs are difficult. For the results based on unrestricted DTSMs, both of the extreme views of “only term premia” and “only expectations” effects are statistically plausible. How-

ever, under restrictions on the risk pricing in the DTSM, statistical uncertainty is reduced. Consequently, our decompositions of the LSAP effects using DTSM estimates under restricted risk prices not only point to a larger role of the signaling channel, but also allow much more precise inference about the respective contribution of signaling and portfolio balance. Taken together, our results indicate that an important effect of the LSAP announcements was to lower the market’s expectation of the future policy path, or, equivalently, to lengthen the expected duration of near-zero policy rates.

Our paper is most closely related to GRRS, since we also use a DTSM to decompose long-term Treasury yields in the context of an event study. Our results are not only quantitatively but also qualitatively different in that we show that the role of the signaling channel is not negligible, and in fact economically and statistically significant. The methodological differences that lead us to this conclusion are the use of alternative empirical DTSMs and, importantly, the construction of interval estimates. Another closely related paper is Krishnamurthy and Vissing-Jorgensen (2011, henceforth KVJ), which argues based on changes in money market futures and other model-free evidence, that signaling likely was an important channel through which LSAPs can affect safe and risky assets. Our new model-free results point to the same conclusion; furthermore, because rather strong auxiliary assumptions are needed for disentangling different LSAP channels without a formal model, we go beyond this model-free analysis.⁴

The paper is structured as follows. In Section 2, we describe the portfolio balance and signaling channels for LSAP effects on yields and discuss the event study methodology that we use to estimate the effects of the LSAPs. Section 3 presents model-free evidence on the importance of the signaling and portfolio balance channels. Section 4 describes the econometric problems with existing term premium estimates and outlines our two approaches for obtaining more appropriate decompositions of long rates. In Section 5, we present our model-based event study results. Section 6 concludes.

2 Identifying portfolio balance and signaling channels

Here we describe the two key channels through which LSAPs can affect interest rates and discuss how their respective importance can be quantified, albeit imperfectly, through an event study methodology.

⁴Christensen and Rudebusch (2012) also provide a model-based event study of the Fed’s LSAPs. We differ from their approach in that we use a new and different set of DTSM specifications, and importantly provide interval estimates for changes in policy expectations and term premia.

2.1 Portfolio balance channel

In the standard asset-pricing model, changes in the supply of long-term bonds do not affect bond prices. In particular, in a pricing model without frictions, bond premia are determined by the risk characteristics of bonds and the risk aversion of investors, both of which are unaffected by the quantity of bonds available to investors. In contrast, to explain the response of bond yields to central bank purchases of bonds, researchers have focused their attention exactly on the effect that a reduction in bond supply has on the risk premium that investors require for holding those securities. The key avenue proposed for this effect is the *portfolio balance channel*.⁵ As described by GRRS:

By purchasing a particular asset, a central bank reduces the amount of the security that the private sector holds, displacing some investors and reducing the holdings of others, while simultaneously increasing the amount of short-term, risk-free bank reserves held by the private sector. In order for investors to be willing to make those adjustments, the expected return on the purchased security has to fall. (p. 6)

The crucial departure from a frictionless model for the operation of a portfolio balance channel is that bonds of different maturities are not perfect substitutes. Instead, risk-averse arbitrageurs are limited in the market and there are “preferred-habitat” investors who have maturity-specific bond demands.⁶ In this setting, the maturity structure of outstanding debt can affect term premia.

The precise portfolio balance effect of purchases on term premia in different markets will vary depending on the interconnectedness of markets. To be concrete, consider the decomposition of the ten-year Treasury yield, y_t^{10} , into a risk-neutral component,⁷ YRN_t^{10} , and a term premium, YTP_t^{10} :

$$y_t^{10} = YRN_t^{10} + YTP_t^{10} \tag{1}$$

$$= YRN_t^{10} + YTP_{risk,t}^{10} + YTP_{instrument,t}^{10} \tag{2}$$

⁵Like most of the literature, we focus on the portfolio balance channel to account for term premia effects of LSAPs. Some recent papers have also discussed a liquidity/market functioning channel through which LSAPs could affect bond premia, including, for example, GRRS, Krishnamurthy and Vissing-Jorgensen (2011), and Joyce et al. (2010). This channel would seem most relevant for limited periods of market dislocation.

⁶Recent work on the theoretical underpinnings of the portfolio balance channel includes Vayanos and Vila (2009) and Hamilton and Wu (2012a).

⁷The risk-neutral yield equals the expected average risk-free rate over the lifetime of the bond under the real-world, or P, probability measure (plus a negligible convexity term). The risk-neutral yield is the interest rate that would prevail if all investors were risk neutral. It is not calculated under the risk-neutral, or Q, probability measure.

The term premium is further decomposed in equation (2) into a maturity-specific term premium, $YTP_{risk,t}^{10}$, that reflects the pricing of interest risk and an idiosyncratic instrument-specific term premium, $YTP_{instrument,t}^{10}$, that captures, for example, demand and supply imbalances for that particular security.⁸

Some researchers have focused on a *market segmentation* version of the portfolio balance channel (Joyce et al., 2011). Market segmentation between the government bond markets and other fixed income markets could reflect the specific needs of pension funds, other institutional investors, and foreign central banks to hold safe government bonds, and arbitrageurs that are institutionally constrained or simply too small in comparison to such huge demand flows. Changes in the bond supply then would have direct price effects through $YTP_{instrument,t}^{10}$ on the securities that were purchased. Because of market segmentation, the change in the price of a given security would depend on how much of that security was purchased. The effects through this type of portfolio balance channel on securities that were not purchased would be small. Notably, for the U.K., Joyce et al. (2011) find that the price effects on those securities purchased by the Bank of England were much larger than for other securities that were not purchased (e.g., swap contracts), which points to significant market segmentation.

Alternatively, markets for securities may be somewhat connected because of the presence of arbitrageurs. For example, GRRS have emphasized the case of investors that prefer a specific amount of duration risk along with a lack of maturity-indifferent arbitrageurs with sufficiently deep pockets. In this case, changes in the bond supply affect the aggregate amount of duration available in the market and the pricing of the associated interest rate risk term premia, $YTP_{risk,t}^{10}$. In this *duration removal* version of the portfolio balance channel, central bank purchases of even a few specific bonds can affect the risk pricing and term premia for a wide range of securities. Notably, in the absence of further frictions, all fixed income securities (e.g., swaps and Treasuries) of the same duration would be similarly affected. Furthermore, if the Fed were to remove a given amount of duration risk from the market by purchasing ten-year securities or by purchasing (a smaller amount of) 30-year securities, the effect through the duration removal version of the portfolio balance channel would be the same.

Thus, there are two ways in which bond purchases can affect term premia in Treasury yields: First, if markets for Treasuries and other assets (including Treasuries of varying maturity) are segmented, bond purchases can reduce Treasury-specific (or maturity-specific) premia. Second, by lowering aggregate duration risk, purchases can reduce term premia in all fixed-income securities.

⁸Also, any safety or liquidity premium, as discussed by KVJ, would be in this final term.

2.2 Signaling channel

The portfolio balance channel, which emphasizes the role of *quantities* of securities in asset pricing, runs counter to at least the past half century of mainstream frictionless finance theory. That theory, which is based on the presence of pervasive, deep-pocketed arbitrageurs, has no role for financial market segmentation or movements in idiosyncratic, security-specific term premia like $YTP_{instrument,t}^{10}$. Moreover, the duration removal version of the portfolio balance channel and its associated shifts in $YTP_{risk,t}^{10}$ would also be ignored in conventional models. In particular, the scale of the Fed’s LSAP program—\$1.725 trillion of debt securities—is arguably small relative to the size of bond portfolios. The U.S. fixed income market is on the order of \$30 trillion, and the global bond market—arguably, the relevant one—is several times larger. In addition, other assets, such as equities, also bear duration risk.

Instead, the traditional finance view of the Fed’s actions would focus on the new information provided to investors about the future path of short-term interest rates, that is, the potential *signaling channel* for central bank bond purchases to affect bond yields by changing the risk-neutral component of interest rates. In general, LSAP announcements may signal to market participants that the central bank has changed its views on current or future economic conditions. Alternatively, they may be thought to convey information about changes in the monetary policy reaction function or policy objectives, such as the inflation target. In such cases, investors may alter their expectations of the future path of the policy rate, perhaps by lengthening the expected period of near-zero short-term interest rates. According to such a signaling channel, announcements of LSAPs would lower the expectations component of long-term yields. In particular, throughout 2009 and 2010, investors were wondering how long the Fed would leave its policy interest rate unchanged at essentially zero. The language in the various FOMC statements in 2009 that economic conditions were “likely to warrant exceptionally low levels of the federal funds rate for an extended period,” provided some guidance, but the zero bound was terra incognita. In such a situation, the Fed’s unprecedented announcements of asset purchases with the goal of putting further downward pressure on yields might well have had an important signaling component, in the sense of conveying to market participants how bad the economic situation really was, and that extraordinarily easy monetary policy was going to remain in place for some time.

2.3 Event study methodology

The few studies to consider the relative contributions of the portfolio balance and signaling channels, specifically GRRS and KVJ for the U.S. and Joyce et al. (2011) for the U.K., have

used an event study methodology.⁹ This methodology focuses on changes in asset prices over tight windows around discrete events. We also employ such a methodology to assess the effects of LSAPs on fixed income markets.

In the portfolio balance channel described above, it is the quantity of asset purchases that affects prices; however, forward-looking investors will in fact react to *news* of future purchases. Therefore, because changes in the expected maturity structure of outstanding bonds are priced in immediately, credible *announcements* of future LSAPs can have the immediate effect of lowering the term premium component of long-term yields. In our event study, we focus on the eight LSAP announcements that GRRS include in their baseline event set, which are described in Table 1.

In calculating the yield responses to these announcements, there are two competing requirements for the size of the event window so that price changes reflect the effects of the announcements. First, the window should be large enough to encompass all of an announcement's effects. Second, the window should be short enough to exclude other events that might significantly affect asset prices. Following GRRS, we use one-day changes in market rates to estimate responses to the Fed's LSAP announcements.¹⁰ A one-day window appears to be a workable compromise. First, for large, highly liquid markets such as the Treasury bond market, and under the assumption of rational expectations, new information in the announcement about economic fundamentals should quickly be reflected in asset prices. Second, the LSAP announcements appear to be the dominant sources of news for fixed income markets on the days under consideration. On these announcement days, the majority of bond and money market movements appeared to be due to new information that markets received about the Fed's LSAP program.

On two of the LSAP event dates, the FOMC press release also contained direct statements about the path for the federal funds rate. On December 16, 2008, the FOMC decreased the target for the policy rate to a range from 0 to 1/4 percent, and indicated that it expected the target to remain there "for some time." On March 18, 2009, the FOMC changed the language about the expected duration of a near-zero policy rate to "for an extended period." Hence there were some conventional monetary policy actions, taking place at the same time as LSAP announcements. Our analysis will not be able to distinguish this direct signaling from the signaling effects through the LSAP announcements themselves. However, leaving out these

⁹GRRS also provide evidence on the portfolio balance channel from monthly time-series regressions of the Kim-Wright term premium on variables capturing macroeconomic conditions and aggregate uncertainty, as well as a measure of the supply of long-term Treasury securities. However, our experience with these regressions suggests the results are sensitive to specification (see also Rudebusch, 2007).

¹⁰Our results are robust to using the two-day change following announcements.

two dates from our event study analysis in fact increases the estimated relative contribution of the expectations component to the yield declines (see discussion below of Tables 6 and 7). Hence our empirical analysis is robust to this caveat.

Of course, if news about LSAPs is leaked or inferred prior to the official announcements, then the event study will underestimate the full effect of the LSAPs. The inability to account for important pre-announcement LSAP news makes us wary of analyzing later LSAP announcements after the eight examined. For example, expectations of a second round of asset purchases (QE2) were incrementally formed before official confirmation in fall 2010, which is a possible reason for why studies like KVJ find small effects on financial markets in their event study of QE2. For the events we consider, one can argue that markets mostly did not expect the Fed's purchases ahead of the announcements.¹¹

2.4 Changes in risk-neutral rates and the role of signaling

How can an event study distinguish between the portfolio balance and signaling channels? A simple conventional view would associate these two channels, respectively, with changes in term premia and risk neutral rates following LSAP announcements. However, there is an important complication in this empirical assessment: As a theoretical matter, the split between the portfolio balance and signaling channels is not the same as the decomposition of the long rate into expectations and risk premium components. In fact, because of second-round effects of the portfolio balance and signaling channels, estimated changes of risk-neutral rates are likely a lower bound for the contribution of signaling to changes in long-term interest rates.

To illustrate the mapping between the two channels and the long rate decomposition, first consider a scenario with just a portfolio balance channel and no signaling. In this case, LSAPs reduce term premia, which would act to boost future economic growth.¹² However, the improved economic outlook will also reduce the amount of conventional monetary policy stimulus needed because to achieve the optimal stance of monetary policy, the more policymakers add of one type of stimulus, the less they need to add of another. Thus, the operation of a portfolio balance channel would cause LSAPs to *increase* risk-neutral rates as well as reducing the term premium. In this case, we would measure higher policy expectations despite the absence of any direct signaling effects. The changes in risk-neutral rates following LSAP announcements will include both the direct signaling effects (presumably negative), as well as the indirect portfolio balance effects on future policy expectations (positive). Hence, this would mean that

¹¹On the issue of the surprise component of monetary policy announcements during the recent LSAP period see Wright (2011) and Rosa (2012).

¹²On this connection, see Rudebusch et al. (2007).

the true signaling effects on risk-neutral rates are likely larger than the estimated decreases in risk-neutral rates.

Conversely, consider the case with no portfolio balance effects but a signaling channel that operates because LSAP announcements contain news about easier monetary policy in the future. This news could take various forms, such as, (1) a longer period of near-zero policy rate, (2) lower risks around a little-changed but more certain policy path, (3) higher medium-term inflation and potentially lower real short-term interest rates, and (4) improved prospects for real activity, including diminished prospects for Depression-like outcomes. Taken together, it seems likely that this news, and the demonstration of the Fed’s commitment to act, would reduce the likelihood of future large drops in asset prices and hence lower the risk premia on financial assets. Indeed, although the effects of easier expected monetary policy on term premia could in general go either way, during the previous Fed easing cycle from 2001 to 2003, lower risk-neutral rates were accompanied by lower term premia. Table 2 shows changes in the actual, fitted, and risk-neutral ten-year yield, and in the corresponding yield term premium (according to the Kim-Wright model) for those days with FOMC announcements during 2001 to 2003 when the risk-neutral rate decreased.¹³ That is, on days on which the average expected future policy rate was revised downward by market participants—comparable to the potential signaling effects of LSAP announcements—the term premium usually fell as well. Over all such days, the cumulative change in the term premium was -21 bps, which has the same sign and more than half the magnitude of the cumulative change in the risk-neutral yield (-35 bps). Thus, during this episode, easing actions that lowered policy expectations at the same time lowered term premia. Arguably, the signaling effect of LSAPs on term premia would be even larger in the recent episode given the potential curtailment of extreme downside risk.

Both of these second-round effects work in the same direction of making the decomposition into changes in risk-neutral rates and term premia a downwardly biased estimate for the importance of the signaling channel. Therefore, the event study results should be considered conservative ones, with the true signaling effects likely larger than the estimated decreases in risk-neutral rates.

3 Model-free evidence

One possible approach to evaluate how an LSAP program affected financial markets is to consider model-free event-study evidence. A prominent example is the study by KVJ which

¹³The data for actual (fitted) yields and the Kim-Wright decomposition of yields are both available at <http://www.federalreserve.gov/econresdata/researchdata.htm> (accessed August 30, 2011). Similar qualitative conclusions are obtained when we use our preferred term premium measures described later.

attempts to disentangle different channels of LSAPs exclusively by studying different market rates, without using a model. In this section we do the same, focusing on just the portfolio balance and signaling channels. We use interest rate data on money market futures, overnight index swaps (OIS), and Treasury securities.

What can we learn about changes in policy expectations and risk premia from considering such interest rates without a formal model? Of course, these interest rates also contain a term premium and thus do not purely reflect the market’s expectations of future short rates. Hence we need auxiliary assumptions, and there are two kinds of plausible assumptions in this context. First, at short maturities, the term premium is likely small, because short-term investments do not have much duration risk. Thus, changes in near-term rates are plausibly driven by the expectations component. This argument can be used to interpret changes at the very short end of the term structure of interest rates, such as movements in near-term money market futures rates (see below) or in short-term yields (see, for example, GRRS, p. 24). Second, we can make assumptions related to market segmentation, which we now discuss in more detail.

3.1 Market segmentation

If markets are segmented to the extent that the portfolio balance effects of LSAPs operate mostly on instrument-specific premia, $YTP_{instrument,t}^n$, then the responses of futures and OIS rates mainly reflect the signaling effects of the announcements. Specifically, changes in the spreads between these interest rates and the rates on the purchased securities reflect portfolio balance effects on yield-specific term premia. For example, Joyce et al. (2011) assume that the Bank of England’s asset purchases only affect the term premium specific to gilts and neither the instrument-specific term premium in OIS rates (which were not part of the asset purchases) nor the general level of the term premium, $YTP_{risk,t}^n$. This market segmentation assumption enables them to draw inferences about the importance of signaling and portfolio balance purely from observed interest rates in OIS and bond markets: Movements in OIS rates reflect signaling effects, and movements in yield-OIS spreads reflect portfolio balance effects. They find that the responses of spreads are large, accounting for the majority of the responses of yields. This points to an important role for the portfolio balance channel in the U.K. It also indicates that the market segmentation assumption is plausible in their context, because a duration-removal story could not explain the differential effects on rates with similar risk characteristics.

Here we produce evidence similar to that of Joyce et al. (2011) for the U.S., considering both money market futures and OIS rates. We do not claim that the market segmentation as-

sumption is entirely plausible for the Treasury and OIS/futures markets, since these securities are close substitutes. To a reader that questions the effects on duration risk compensation and prefers the market segmentation story, the results below will be evidence about the importance of signaling and portfolio balance effects. More generally though, without the identifying assumption that changes in $YTP_{risk,t}^n$ are negligible, the changes in the spreads reflect changes in both YRN_t^n and $YTP_{risk,t}^n$, and thus constitute an upper bound for the magnitude of shifts in policy expectations.

3.2 Money market futures

Money market futures are bets on the future value of a short-term interest rate, and they are used by policymakers, academics, and practitioners to construct implied paths for future policy rates. Federal funds futures settle based on the federal funds rate, and contracts for maturities out to about six months are highly liquid. Eurodollar futures pay off according to the three-month London interbank offered rate (LIBOR), and the most liquid contracts have quarterly maturities out to about four years. While LIBOR and the fed funds rate do not always move in lockstep, these two types of futures contracts are typically used in combination to construct a policy path over all available horizons.

How has the futures-implied policy path has changed around LSAP dates? Figure 1 shows the futures-implied policy paths around the first five LSAP events, based on futures rates on the end of the previous day and on the end of the event day.¹⁴ On almost all days, the policy paths appear to have shifted down significantly at horizons of one year and longer in response to the LSAP announcements.¹⁵ Table 3 displays the changes at specific horizons on all eight LSAP event days. Also shown are total changes over all event days, as well as cumulative changes and standard deviations of daily changes over the LSAP period. At the short end, the path has shifted down by about 20-40 bps, while at longer horizons of one to three years the total decrease is around 50 bps. Because the decreases in short-term futures rates are arguably driven primarily by expectations, these results indicate that markets revised their near-term policy expectations downward around LSAP announcements by about 20-40 bps.¹⁶ Note that

¹⁴The policy paths are derived using federal funds futures contracts for the current quarter and two quarters beyond that. For longer horizons, we use Eurodollar futures, which are adjusted by the difference between the last quarter of the federal funds futures contracts and the overlapping Eurodollar contract. Beginning five months out, a constant term premium adjustment of 1bp per month of additional maturity is applied.

¹⁵The FOMC statement for January 28, 2009, contrary to the other announcements, actually caused sizable increases in yields and other market interest rates, as documented in GRRS and in our results below. Anecdotal evidence indicates that market participants were disappointed by the lack of concrete language regarding the possibility and timing of purchases of longer-dated Treasury securities.

¹⁶One minor confounding factor is that on December 16, 2008, markets also were surprised by the target

this analysis is parallel to KVJ’s assessment of the importance of the signaling channel.

What about policy expectations at longer horizons? The last three columns of the table show the changes in the average futures-implied policy path over the next three years, the changes in the three-year yield, and the spread between the yield and the futures-implied rate.¹⁷ The futures-implied three-year yield declined by 43 bps, which corresponds to 54 percent of the decline in the yield. With the exception of March 2009, every LSAP announcement had a much larger effect on the futures-implied yield than on the Treasury yield. Under a market segmentation assumption, this evidence suggests that lower policy expectations accounted for more than half of the decrease in the three-year yield.

3.3 Overnight index swaps

In an overnight index swap (OIS), one party pays a fixed interest rate on the notional amount and receives the overnight rate, i.e., the federal funds rate, over the entire maturity period. Under absence of arbitrage, OIS rates reflect risk-adjusted expectations of the average policy rate over the horizon corresponding to the maturity of the swap. Intuitively, while futures are bets on the value of the short rate at a future point in time, OIS contracts are essentially bets on the average value of the short rate over a certain horizon.

Table 4 shows the results of an event study analysis of changes in OIS rates with maturities of two, five, and ten years, yields of the same maturities, and yield-OIS spreads. We consider the same set of event dates as before.¹⁸ The responses of yields to the Fed’s LSAP announcements are similar to the responses of OIS rates. For certain days and maturities, OIS rates respond even more strongly than yields, and at the ten-year maturity, the cumulative change of the OIS rate is larger than the yield change, which results in an increasing OIS spread. In those instances where the OIS spread significantly decreased, its relative contribution to the yield change is typically still much smaller than the contribution of the OIS rate change. The March 2009 announcement is the only one that significantly lowered spreads. On the other event days, yield-OIS spreads barely moved or increased, suggesting that large decreases in term premia are unlikely.

Clearly, yields and OIS rates moved very much in tandem in response to the LSAPs. Our evidence in this section is consistent with the finding of GRRS “that LSAPs had widespread

rate decision—expectations were for a new target of 25 bps, however the Federal Open Market Committee decided on a target range of 0-25 bps. Changes in short-term rates on this day reflect also reflect the effects of conventional monetary policy.

¹⁷Yields are zero-coupon yields from a smoothed yield curve data set constructed in Gürkaynak et al. (2007). See <http://www.federalreserve.gov/econresdata/researchdata.htm> (accessed July 29, 2011).

¹⁸OIS rates are taken from Bloomberg.

effects, beyond those on the securities targeted for purchase” (p. 20). Under a market segmentation identifying assumption, the evidence that OIS rates showed pronounced responses suggests an important contribution of lower policy expectations to the decreases in interest rates. Without such an assumption, it just indicates that instrument-specific premia in Treasuries did not move much around announcements.

Some readers might find our result unsurprising: Safe government bonds and swap contracts have similar risk characteristics, are likely to be close substitutes, and could therefore be expected a priori to respond similarly to policy actions. This of course simply amounts to not accepting the market segmentation assumption for these securities. However, there are two important points to keep in mind in response to this critique: First, the evidence for the U.K. has shown that yields and OIS rates do not necessarily need to respond similarly. For the case of the U.K., these instruments are not very close substitutes and there is considerable market segmentation, thus one might be inclined to find this plausible for the U.S. as well. Second, the same results hold for securities that are less close substitutes. Specifically, the evidence in KVJ as well as our own calculations using different data sources (results omitted) show that highly-rated corporate bonds responded about as much as Treasury yields to LSAPs.¹⁹ Clearly a Treasury bond and, say, a AA-rated corporate bond are not close substitutes, thus market segmentation is more plausible, and the fact that they respond in tandem is evidence that signaling played an important role.

However plausible one finds the necessary auxiliary assumptions, model-free analysis can only go so far. Thus, we now turn to model-based evidence to address whether Treasuries were affected by the LSAPs through downward shifts in the expected policy path and through shifts in a their term premium.

4 Term premium estimation

A theoretically rigorous decomposition of interest rates into expectations and term premium components requires a DTSM, which have generally proven difficult to estimate. Therefore, we consider several different model estimates to ensure robustness.

4.1 Econometric problems: bias and uncertainty

To estimate the term premium component in long-term interest rates, researchers typically resort to DTSMs. Such models simultaneously capture the cross section and time series

¹⁹Changes in default-risk premia do not account for this response, based on KVJ’s evidence that incorporates credit default swap data.

dynamics of interest rates, and impose absence of arbitrage, which ensures that the two are consistent with each other. Term premium estimates are obtained by forecasting the short rate using the estimated time series model, and subtracting the average short rate forecast (i.e., the risk-neutral rate) from the actual interest rate. The very high persistence of interest rates, however, causes major problems with estimating the time series dynamics. The parameter estimates typically suffer from small-sample bias and large statistical uncertainty, which makes the resulting estimated risk-neutral rates and term premia inherently unreliable.

The small-sample bias in conventional estimates of DTSMs stems from the fact that the largest root in autoregressive models for persistent time series is generally underestimated. Therefore the speed of mean reversion is overestimated, and the model-implied forecasts for longer horizons are too close to the unconditional mean of the process. Consequently, risk-neutral rates are too stable, and too much of the variation in long-term rates is attributed to the term premium component.²⁰ In the context of LSAP event studies, this bias works in the direction of attributing too large a share of changes in long-term interest rates to the term premium. Hence, the relative importance of the portfolio balance channel will be overestimated. Because of this concern, we conduct an event study using term premium estimates that correct for this bias.

Large statistical uncertainty underlies any estimate of the term premium, due to both specification and estimation uncertainty. The former reflects uncertainty about different plausible specifications of a DTSM, which might lead to quite different economic implications.²¹ We address this issue in a pragmatic way by presenting alternative estimates based on different specifications. Estimation uncertainty exists because the parameters governing the time series dynamics in a DTSM are estimated imprecisely, due to the high persistence of interest rates.²² Consequently, large statistical uncertainty underlies short rate forecasts and term premia calculated from such parameter estimates. Despite this fact, studies typically report only point estimates of term premia.²³ In our event study, we report interval estimates of changes in risk-neutral rates and of changes in the term premium.

²⁰This problem has been pointed out by Ball and Torous (1996) and discussed in subsequent studies including BRW.

²¹This issue has been highlighted, for example, by Rudebusch et al. (2007) and Bauer (2011).

²²The slow speed of mean reversion of interest rates makes it difficult to pin down the unconditional mean and the persistence of the estimated process. See, among others, Kim and Orphanides (2005).

²³Exceptions are the studies by Bauer (2011) and Joslin et al. (2010), who present measures of statistical uncertainty around estimated risk-neutral rates and term premia.

4.2 Alternative term premium estimates

We now briefly describe the alternative term premium estimates that we include in our event study. Details are provided in appendices. The data used in the estimation of our models consist of daily observations of interest rates from January 2, 1985, to December 30, 2009. We include T-bill rates at maturities of 3 and 6 months from the Federal Reserve H.15 release and zero-coupon yields at maturities of 1, 2, 3, 5, 7, and 10 years.

4.2.1 Kim-Wright

The term premium estimates used by GRRS are obtained from the model of Kim and Wright (2005). What distinguishes their model from an unrestricted, i.e., maximally flexible, affine Gaussian DTSM is the inclusion of survey-based short rate forecasts and some slight restrictions on the risk pricing. While Kim and Orphanides (2005) argue that incorporating additional information from surveys might help alleviate the problems with DTSM estimation, it is unclear to what extent bias and uncertainty are reduced. Survey expectations are problematic because on the one hand they are available only at low frequencies (monthly/quarterly), and on the other hand they might not represent rational forecasts of short rates (Piazzesi and Schneider, 2008). In terms of risk price restrictions, the model imposes only very few constraints, so the link between cross-sectional dynamics and time series dynamics is likely to be weak.

4.2.2 Ordinary least squares

As a benchmark, we estimate a maximally-flexible affine Gaussian DTSM. The risk factors correspond to the first three principal components of yields. We use the normalization of Joslin et al. (2011). The estimation is a two-step procedure: First, the parameters of the vector autoregression (VAR) for the risk factors are estimated using ordinary least squares (OLS). Second, we obtain estimates of the parameters governing the cross-sectional dynamics using the minimum-chi-square method of Hamilton and Wu (2012b). Because the model is exactly identified, these are also the maximum likelihood (ML) estimates. Details on the estimation can be found in Appendix B.1.

To account for the estimation uncertainty underlying the decompositions of long-term interest rates, we obtain bootstrap distributions of the VAR parameters. We can thus calculate risk-neutral rates and term premia for each bootstrap replication of the parameters, and calculate confidence intervals for all objects of interest. Details on the bootstrap procedure are provided in Appendix B.3.

4.2.3 Bias-corrected

One way to deal with the small-sample bias in DTSM estimates is to directly correct the estimates of the dynamic system for bias. Starting from the same model, we perform bias-corrected (BC) estimation of the VAR parameters in the first step and proceed with the second step of finding cross-sectional parameters as before. Our methodology, which closely parallels the one laid out in BRW, is detailed in Appendix B.2. We also obtain bootstrap replications of the VAR parameters.

The resulting estimates imply interest rate dynamics that are more persistent and short rate forecasts that revert to the unconditional mean much more slowly than is implied by the biased OLS estimates. Therefore, one would expect a larger contribution of the expectations component to changes in long-term rates around LSAP announcements. Because this estimation method only addresses the bias problem and not the uncertainty problem, confidence intervals cannot be expected to be any tighter than for OLS.

4.2.4 Restricted risk prices

The no-arbitrage restriction can be a powerful remedy for both the bias and the uncertainty problem if the risk pricing is restricted.²⁴ The intuition is that cross-sectional dynamics are precisely estimated and can help pin down the parameters governing the time series dynamics, reducing both bias and uncertainty in these parameters and leading to more reliable estimates of risk-neutral rates and term premia. There is a large set of possible restrictions on the risk pricing in DTSMs, and alternative restrictions may lead to different economic implications. To deal with these complications, we use a Bayesian framework parallel to the one suggested in Bauer (2011) for estimating our DTSM with restricted risk prices. This allows us to select those restrictions that are supported by the data and to deal with specification uncertainty by means of Bayesian model averaging. Another advantage is that interval estimates naturally fall out of the estimation procedure, because the Markov chain Monte Carlo (MCMC) sampler that we use for estimation, described in Appendix C.2, produces posterior distributions for any object of interest.

First, we estimate a maximally flexible model where risk price restrictions are absent using MCMC sampling. These estimates will be denoted by URP (*Unrestricted Risk Prices*). The point estimates of the model parameters are almost identical to OLS.²⁵ With regard to interval

²⁴This has been argued, for example, by Cochrane and Piazzesi (2008), Bauer (2011), and Joslin et al. (2010).

²⁵With uninformative priors, the Bayesian posterior parameter means are the same as the OLS/maximum likelihood estimates. In our case, differences between the two sets of point estimates, which could result from

estimation, there will however be some numerical differences, because the Bayesian credibility intervals (which we will for simplicity also call confidence intervals) for URP are conceptually different from the bootstrap confidence intervals for OLS. Because of potential differences between OLS and URP we include the URP estimates as a point of reference.

The estimates under *Restricted Risk Prices* will be denoted by RRP. To be clear, here parameters and the objects of interest such as term premium changes are estimated by means of Bayesian model averaging, since in this setting the MCMC sampler provides draws across model and parameter space. Because of the averaging over the set of restricted models, the inference takes into account both estimation and model uncertainty.

Because of the risk price restrictions, and in light of the results in Bauer (2011), one would expect a larger role for the expectations component in driving changes in long-term rates around LSAP announcements, as well as tighter confidence intervals around point estimates, i.e., more precise inference about the respective roles of the signaling and portfolio balance channels.

5 Changes in policy expectations and term premia

We now turn to model-based event study results to assess the effects of the Fed’s LSAP announcements on the term structure of interest rates. We decompose changes in Treasury yields around LSAP events into changes in risk-neutral rates, i.e., in policy expectations, and term premia using alternative DTSM estimation approaches.

5.1 Cumulative changes in long-term yields

Let us first consider cumulative changes in long-term Treasury yields over the LSAP events and how they are decomposed into expectations and risk premium components. The results are shown in Table 5. In addition to point estimates, we present 95%-confidence intervals for the changes in risk-neutral rates and premia. We decompose changes in the ten-year yield as in GRRS, and also include results for the five-year yield. Cumulatively over these eight days, the ten-year yield decreased by 89 bps, and the five-year yield decreased even more strongly by 97 bps.²⁶

The Kim-Wright decomposition of the change in the fitted ten-year yield of -102 bps results in a decrease in the risk-neutral yield (YRN) of 31 bps and a decrease in the yield term

the priors and from approximation error, turn out to be negligibly small.

²⁶GRRS consider the constant-maturity ten-year yield, which decreased by 91 bps, whereas we focus throughout on zero-coupon yields obtained from the GSW data set.

premium (YTP) of 71 bps. Notably, the cumulative change in the DTSM’s fitting error of -13 bps is contained in the term premium, which is calculated as the difference between the fitted yield and YRN. This is not made explicit in the GRRS study, and the authors compare the 71 bps decrease in the term premium to the 91 bps decrease in the actual (constant-maturity) ten-year yield. However, based on model-fitted results, the contribution of the term premium is not $\frac{-71}{-91} \approx 78\%$ but instead $\frac{-71}{-102} \approx 70\%$, with the risk-neutral component contributing 30% to the decrease. For the five-year yield, the relative contributions of expectations and term premium components are 32 percent and 68 percent, respectively.

The decomposition based on the OLS estimates leads to a slightly larger contribution of the expectations component than for the Kim-Wright decomposition, particularly for the five-year yield. For the ten-year yield, the contributions are 35 and 65 percent, respectively, and for the five-year yield they are 43 and 57 percent. The bootstrapped confidence intervals (CIs) reveal tremendous uncertainty attached to these point estimates. Based on these estimates, it is equally plausible that the entire yield change was driven by the term premium or by the expectations component. Similarly, these results suggest that the magnitude of the change in the Kim-Wright term premium is very uncertain.

The BC estimates imply a larger role for the expectations component, which now accounts for about 50 percent of the yield change, both at the five-year and ten-year maturity. The CIs are even wider than for the OLS estimates. Addressing the bias problem in term premium estimation via direct bias correction increases the estimated contribution of the signaling channel, but the inference is still very imprecise, since the uncertainty problem remains.

The last two decompositions are for the URP and RRP estimates. The URP point estimates are almost identical to the OLS results and indicate that both components contributed to the decrease in yields.²⁷ The URP confidence intervals, which are conceptually different as mentioned above, are slightly narrower than the OLS ones. However, there still is considerable statistical uncertainty: The contribution of risk-neutral rates could plausibly be anywhere between $\frac{-7}{-94} \approx 7\%$ and $\frac{-71}{-94} \approx 76\%$. With restricted risk prices, the point estimates for the five-year yield closely correspond to the BC results, with a contribution of expectations that is slightly larger than the contribution of the term premium. The split between changes in expectations and premia here is 52 and 48 percent. For the ten-year yield, the RRP decomposition also attributes more, if only by a little, to the expectations component than the Kim-Wright and OLS results—with an expectation and term premium split of 38 and 62 percent. Importantly, the confidence intervals around the RRP estimates are much

²⁷Slight differences are due to the fact that the decompositions for URP are posterior means of the object of interest, whereas for OLS the decompositions are calculated at the point estimates of the parameters.

tighter than for unrestricted DTSM estimates. The intervals clearly indicate that both the expectations and term premium components have played an important role in lowering yields. For the ten-year yield, the relative contribution of risk-neutral rates is estimated to be between $\frac{-29}{-94} \approx 30\%$ and $\frac{-53}{-94} \approx 56\%$.

5.2 Shifts in the forward curve and policy expectations

To understand these decompositions of yield changes and to get a more comprehensive perspective of the effects of the LSAP announcements on the term structure, it is useful to look at forward rates and the expected policy path in Figures 2 and 3. Based on our four alternative DTSM estimates, the figures show the cumulative change over the LSAP event days in instantaneous forward rates out to ten years maturity, as well as cumulative changes in expected policy rates with 95%-confidence intervals.

The shift in forward rates, shown as a solid line, is common to all four decompositions because fitted rates are essentially identical across DTSM estimates. The shift is hump-shaped, with the largest decrease, about -110 bps, occurring at a horizon of three years. At the short end, the change is about -45 bps for the six-month horizon, and about -80 bps for the twelve-month horizon. At the long end, forward rates decreased by approximately 80 bps. The decreases at the short end are particularly interesting, because the size of the term premium is presumably small at short horizons. Based on this argument, most of the drop in the six-month forward rate and a significant portion of the drop in the one-year rate would be attributed to a lowering of policy expectations. This is confirmed by our model-based decompositions.

Figure 2 contrasts the OLS (left panel) and BC results (right panel). The decompositions at the short end are very similar, with essentially all of the decrease in the six-month rate and a sizable fraction of the decrease in other near-term rates attributed to the expectations component. The difference between OLS and BC is most evident in the decompositions of changes in long-term rates with horizons of five to ten years. The OLS estimates imply a rather small contribution for the expectations component, whereas the BC estimates attribute around half of the decrease in forward rates to lower expectations. The very large estimation uncertainty underlying these decompositions is also apparent. For either decomposition, at horizons longer than five years, the forward rate curve and the zero line are both within the confidence bands for the changes in expectations. Neither the “all expectations” hypothesis—that these forward rates decreased solely because of lower policy expectations—nor the “all term premia” hypothesis—that expectations did not change and only term premia drove long rates lower—can be rejected.

Figure 3 shows the decompositions resulting from the URP (left panel) and RRP estimates (right panel). Again, the improved decomposition in the right panel leads to a larger role for expectations. The main difference between the two panels is that under restricted risk prices a larger share of the decrease in short- and medium-term forward rates is attributed to lower expectations, whereas decompositions of changes in long-term forward rates are rather similar. Thus, the economic implications for changes in term premia are somewhat different under our BC and RRP estimates. These differences reinforce the need to include more than one set of estimates to draw robust conclusions.

Figure 3 also shows how imposing risk price restrictions greatly increases the precision of inference. In the left panel, the confidence bands around the estimated downward shift in expectations are quite large. In the right panel, the RRP confidence bands are comparably tight, and our conclusions about the role of expectations are a lot more precise. In a maximally-flexible DTSM, the estimation uncertainty is so large that we cannot really be sure about the relative contribution of changes in policy expectations. However, plausible restrictions on risk prices lead to the conclusion that both components, expectations as well as premia, played an important role for lowering rates around LSAP events.

5.3 Day-by-day results

To drill down further into the shifts in the term structure, Tables 6 and 7 show the decompositions of ten-year and five-year yield changes on each of the eight event days. In the top panels of each table, we compare the Kim-Wright decompositions of daily changes to the OLS and BC results. In the bottom panels, we compare Kim-Wright to the URP and RRP results. In the bottom three rows of each panel, we show total changes over the event days (which correspond to the point estimates in Table 5), as well as cumulative changes and standard deviations of daily changes over the LSAP period.

The tables show in detail how the event days differ from each other. The first three days, in 2008, show very similar decreases in yields and decompositions. In contrast, as discussed above, rates increased on January 28, 2009, because market participants were disappointed by the lack of concrete announcements of Treasury purchases. On March 18, 2009, the most dramatic decrease occurred, with the long-term yield falling by half a percentage point. This announcement seems to have had the largest impact on term premia. The last three days showed only minor movements, which when compared to the standard deviations of daily changes are not significant.²⁸

²⁸As noted above, the December 16, 2008, and the March 18, 2009, FOMC statements also contained direct signaling of future interest rate policy. However, excluding these two dates does not weaken our overall results.

The typical pattern is that the estimated contribution of risk-neutral rates to the changes in yields is larger for BC/RRP than for OLS/URP. Notably, the RRP decompositions always have the same signs as the Kim-Wright decompositions. The OLS and BC decompositions, on the other hand, differ from Kim-Wright and RRP in that they imply decreases in the risk-neutral yield on every day, due to the downward movement of the short-end of the term structure.

5.4 Summary of model-based evidence

Previous findings in GRRS were based on the Kim-Wright decomposition of long-term rates and seemed to show a large contribution of term premium changes. In addition to the caveat that the decrease in the estimated term premium also included a sizable pricing error component, there are two other important reasons why these results need to be taken with a large grain of salt. First, in terms of point estimates, the decomposition of rate changes based on alternative DTSM estimates imply a larger contribution of the expectations component to rate changes around LSAP announcements than the Kim-Wright decomposition. And second, putting confidence intervals around the estimated changes in risk-neutral rates and term premia reveals that large changes in policy expectations around LSAP announcements are consistent with the data. Increasing the precision by restricting the risk pricing of the DTSM leads to a statistically significant role for both the expectations component and the term premium component in lowering yields.

In terms of quantitative conclusions, one would take away from the GRRS study that only $1 - \frac{71}{91} \approx 22\%$ of the cumulative decrease in the ten-year yield around LSAP events was due to changing policy expectations. Our model estimates and the resulting confidence intervals, however, suggest that this number is too low, and that the true contribution of policy expectations to lower long-term Treasury yields is more likely to be around 40-50 percent.

6 Conclusion

In this paper, we have challenged the common wisdom that the Fed's LSAP program mostly worked through a portfolio balance channel. Evidence from different sources, both model-free and based on DTSM estimates, points to a larger role of the signaling channel than previous studies have acknowledged.²⁹ Our results suggest that changes in the expectations component of long-term interest rates were both economically and statistically significant. Furthermore,

²⁹Similar evidence using additional alternative measures of term premia, constructed from an arbitrage-free Nelson-Siegel DTSM, is provided in Christensen and Rudebusch (2012).

we argue that because of second-round effects of signaling and portfolio balance, the relative contribution of expectations to changes in interest rates are conservative estimates of the importance of the signaling channel.

Therefore, it appears that the Fed affected long rates not only by changing the risk premium in long-term interest rates, but also to an important extent by altering market expectations of the future path of monetary policy. The plausible interpretation is that, through announcing and implementing LSAPs, the FOMC signaled to market participants that it would maintain an easy stance for monetary policy for a longer time than previously anticipated. This result raises the question: If the FOMC wants to move interest rate expectations, why doesn't it simply communicate its intentions directly to the public? Of course, central banks have long been reluctant to directly reveal their views on likely future policy actions (see Rudebusch and Williams, 2008). This reluctance arises from the belief that financial markets would tend to interpret any central bank statements about the likely future path of policy as commitments to future action, as opposed to projections based on existing information and subject to considerable change. Thus, central banks have in the past given only indirect hints or used coded language about future interest rate inclinations. Since the start of the financial crisis, the FOMC has been more forthcoming and provided direct signals; however, bond purchases may provide some advantage as an additional reinforcing indirect signaling device about future interest rates.

The effectiveness of LSAPs will typically be judged based on whether they lowered various borrowing rates and not only government bond yields.³⁰ After all, private borrowing rates—corporate bond rates, bank and loan rates, and, importantly, mortgage rates—are the most relevant interest rates for the transmission of monetary policy. While we study only Treasury yields in this paper, our results have a close connection to the question whether LSAPs lowered effective lending rates: Signaling effects will lower rates in all fixed income markets, because all interest rates depend on the expected future path of policy rates. Portfolio balance effects, on the other hand, are not guaranteed to affect various markets in a similar fashion. Our finding that signaling was important during QE1 is consistent with the widespread effects of LSAPs that other studies have found. In this way, our paper explains how these policy actions have been successfully in lowering the interest rates most relevant for consumption and investment.

As directions for future research, one important issue is to account for the zero lower bound on the nominal short-term interest rate. Affine DTSMs ignore this restriction, and incorporating it might lead to slightly different results. However, the most promising models that ensure

³⁰For non-Treasury markets, LSAPs can improve market functioning and liquidity when these markets are under distress (such as the agency MBS market in 2008). In this way, LSAPs can in some circumstances be more effective than forward guidance alone.

a non-negative short rate, such as the shadow rate models estimated in Kim and Singleton (2012), lack analytical bond pricing formulas and are computationally too expensive.³¹ There is much work to be done about the ZLB constraint. Another interesting avenue for exploration is to augment our event study approach with information about the quantity of outstanding Treasury debt (actual or announced), which can be incorporated into DTSMs (see Li and Wei, 2012), or with other additional risk factors (such as market-based uncertainty measures, or higher order yield-curve factors). Finally, there is a need for measures of LSAP expectations, particularly for the analysis of subsequent programs of the Federal Reserve, which have to some extent been anticipated.³²

References

- Ang, Andrew, Jean Boivin, and Sen Dong**, “Monetary Policy Shifts and the Term Structure,” *Review of Economic Studies*, 2011, 78 (2), 429–457.
- , **Sen Dong, and Monika Piazzesi**, “No-Arbitrage Taylor Rules,” NBER Working Paper 13448, National Bureau of Economic Research September 2007.
- Ball, Clifford A. and Walter N. Torous**, “Unit roots and the estimation of interest rate dynamics,” *Journal of Empirical Finance*, 1996, 3 (2), 215–238.
- Bartlett, M. S.**, “A comment on D. V. Lindley’s statistical paradox,” *Biometrika*, 1957, 44 (3-4), 533–534.
- Bauer, Michael D.**, “Bayesian Estimation of Dynamic Term Structure Models under Restrictions on Risk Pricing,” Working Paper 2011-03, Federal Reserve Bank of San Francisco November 2011.
- , **Glenn D. Rudebusch, and Jing Cynthia Wu**, “Correcting Estimation Bias in Dynamic Term Structure Models,” *Journal of Business and Economic Statistics*, July 2012, 30 (3), 454–467.
- Bernanke, Ben**, “The Economic Outlook and Monetary Policy,” speech at Jackson Hole, Wyoming August 2010.

³¹For a model with simulation-based bond pricing, the cost of *evaluating* the likelihood function can easily amount to many minutes or even hours. This cost multiplies for a daily bond price model.

³²Some approaches exist, for example Wright (2011) and Rosa (2012). However, these measures for expectations rely either on market interest rates themselves or on qualitative judgement and discrete categories. Market or dealer surveys about LSAP expectations (probabilities or quantities) are not publicly available.

- Carlin, Bradley P. and Siddhartha Chib**, “Bayesian Model Choice via Markov Chain Monte Carlo Methods,” *Journal of the Royal Statistical Society. Series B (Methodological)*, September 1995, *57* (3), 473–484.
- Chib, S. and B. Ergashev**, “Analysis of Multifactor Affine Yield Curve Models,” *Journal of the American Statistical Association*, 2009, *104* (488), 1324–1337.
- Christensen, Jens H. E. and Glenn D. Rudebusch**, “The Response of Interest Rates to U.S. and U.K. Quantitative Easing,” Working paper 2012-06, Federal Reserve Bank of San Francisco 2012.
- Christensen, Jens H.E., Francis X. Diebold, and Glenn D. Rudebusch**, “The Affine Arbitrage-Free Class of Nelson-Siegel Term Structure Models,” *Journal of Econometrics*, September 2011, *164* (1), 4–20.
- Cochrane, John H. and Monika Piazzesi**, “Decomposing the Yield Curve,” unpublished manuscript 2008.
- Cowles, Mary Kathryn and Bradley P. Carlin**, “Markov Chain Monte Carlo Convergence Diagnostics: A Comparative Review,” *Journal of the American Statistical Association*, 1996, *91*, 883–904.
- D’Amico, Stefania and Thomas B. King**, “Flow and Stock Effects of Large-Scale Treasury Purchases,” Finance and Economics Discussion Series 2010-52, Federal Reserve Board of Governors February 2010.
- Dellaportas, Petros, Jonathan J. Forster, and Ioannis Ntzoufras**, “On Bayesian model and variable selection using MCMC,” *Statistics and Computing*, January 2002, *12* (1), 27–36.
- Duffee, Gregory R. and Richard H. Stanton**, “Estimation of Dynamic Term Structure Models,” working paper, Haas School of Business Mar 2004.
- Gagnon, Joseph, Matthew Raskin, Julie Remache, and Brian Sack**, “The Financial Market Effects of the Federal Reserve’s Large-Scale Asset Purchases,” *International Journal of Central Banking*, March 2011, *7* (1), 3–43.
- Gürkaynak, Refet S., Brian Sack, and Jonathan H. Wright**, “The U.S. Treasury yield curve: 1961 to the present,” *Journal of Monetary Economics*, 2007, *54* (8), 2291–2304.

- Hamilton, James D. and Jing Cynthia Wu**, “The Effectiveness of Alternative Monetary Policy Tools in a Zero Lower Bound Environment,” *Journal of Money, Credit and Banking*, 2012, *44*, 3–46.
- **and –**, “Identification and estimation of Gaussian affine term structure models,” *Journal of Econometrics*, 2012, *168* (2), 315–331.
- Joslin, Scott, Kenneth J. Singleton, and Haoxiang Zhu**, “A New Perspective on Gaussian Dynamic Term Structure Models,” *Review of Financial Studies*, 2011, *24* (3), 926–970.
- **, Marcel Pribsch, and Kenneth J. Singleton**, “Risk Premiums in Dynamic Term Structure Models with Unspanned Macro Risks,” working paper September 2010.
- Joyce, Michael, Ana Lasasosa, Ibrahim Stevens, and Matthew Tong**, “The Financial Market Impact of Quantitative Easing in the United Kingdom,” *International Journal of Central Banking*, 2011, *7* (3), 113–161.
- Kilian, Lutz**, “Small-sample confidence intervals for impulse response functions,” *Review of Economics and Statistics*, 1998, *80* (2), 218–230.
- Kim, Don H. and Athanasios Orphanides**, “Term Structure Estimation with Survey Data on Interest Rate Forecasts,” *Computing in Economics and Finance* 2005 474, Society for Computational Economics November 2005.
- **and Jonathan H. Wright**, “An arbitrage-free three-factor term structure model and the recent behavior of long-term yields and distant-horizon forward rates,” *Finance and Economics Discussion Series* 2005-33, Board of Governors of the Federal Reserve System (U.S.) 2005.
- **and Kenneth J. Singleton**, “Term Structure Models and the Zero Bound: An Empirical Investigation of Japanese Yields,” *Journal of Econometrics*, September 2012, *170* (1), 32–49.
- Krishnamurthy, Arvind and Annette Vissing-Jorgensen**, “The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy,” *Brookings Papers on Economic Activity*, Fall 2011, pp. 215–265.
- Li, Canlin and Min Wei**, “Term Structure Modelling with Supply Factors and the Federal Reserve’s Large Scale Asset Purchase Programs,” *Finance and Economics Discussion Series* 2012-37, Federal Reserve Board May 2012.

- Neely, Christopher J.**, “The Large-Scale Asset Purchases Had Large International Effects,” Working Paper Series 2010-018C, Federal Reserve Bank of St. Louis July 2010.
- Piazzesi, Monika and Martin Schneider**, “Bond positions, expectations, and the yield curve,” Working Paper 2008-02, Federal Reserve Bank of Atlanta January 2008.
- Rosa, Carlo**, “How “Unconventional” Are Large-Scale Asset Purchases? The Impact of Monetary Policy on Asset Prices,” Staff Report No. 560, Federal Reserve Bank of New York May 2012.
- Rudebusch, Glenn**, “Commentary on “Cracking the Conundrum”,” *Brookings Papers on Economic Activity*, 2007, 38 (1), 317–325.
- Rudebusch, Glenn D. and John C. Williams**, “Revealing the Secrets of the Temple: The Value of Publishing Central Bank Interest Rate Projections,” in John Y. Campbell, ed., *Asset Prices and Monetary Policy*, University of Chicago Press, 2008, pp. 247–284.
- , **Brian P. Sack, and Eric T. Swanson**, “Macroeconomic Implications of Changes in the Term Premium,” *Federal Reserve Bank of St. Louis Review*, July/August 2007, 89 (4), 241–269.
- Swanson, Eric T.**, “Lets Twist Again: A High-Frequency Event-Study Analysis of Operation Twist and Its Implications for QE2,” *Brookings Papers on Economic Activity*, Spring 2011, pp. 151–188.
- Vayanos, Dimitri and Jean-Luc Vila**, “A Preferred-Habitat Model of the Term Structure of Interest Rates,” NBER Working Paper 15487, National Bureau of Economic Research November 2009.
- Wright, Jonathan H.**, “What does Monetary Policy do to Long-Term Interest Rates at the Zero Lower Bound?,” working paper June 2011.

Appendices

A Model specification

We use a discrete-time affine Gaussian DTSM. A vector of N pricing factors, X_t , follows a first-order Gaussian VAR:

$$X_{t+1} = \mu + \Phi X_t + \Sigma \varepsilon_{t+1}, \quad (3)$$

where $\varepsilon_t \stackrel{iid}{\sim} N(0, I_N)$ and Σ is lower triangular. The short rate, r_t , is an affine function of the pricing factors:

$$r_t = \delta_0 + \delta_1' X_t. \quad (4)$$

The stochastic discount factor (SDF) is of the form

$$-\log(M_{t+1}) = r_t + \frac{1}{2} \lambda_t' \lambda_t + \lambda_t' \varepsilon_{t+1},$$

where the N -dimensional vector of risk prices is affine in the pricing factors,

$$\Sigma \lambda_t = \lambda_0 + \lambda_1 X_t,$$

for N -vector λ_0 and $N \times N$ matrix λ_1 . Under these assumptions X_t follows a first-order Gaussian VAR under the risk-neutral pricing measure \mathbb{Q} ,

$$X_{t+1} = \mu^{\mathbb{Q}} + \Phi^{\mathbb{Q}} X_t + \Sigma \varepsilon_{t+1}^{\mathbb{Q}}, \quad (5)$$

and the prices of risk determine how VAR parameters under the objective measure and the \mathbb{Q} measure are related:

$$\mu^{\mathbb{Q}} = \mu - \lambda_0 \quad \Phi^{\mathbb{Q}} = \Phi - \lambda_1. \quad (6)$$

Furthermore bond prices are exponentially affine functions of the pricing factors:

$$P_t^m = e^{\mathcal{A}_m + \mathcal{B}_m X_t},$$

and the loadings $\mathcal{A}_m = \mathcal{A}_m(\mu^{\mathbb{Q}}, \Phi^{\mathbb{Q}}, \delta_0, \delta_1, \Sigma)$ and $\mathcal{B}_m = \mathcal{B}_m(\Phi^{\mathbb{Q}}, \delta_1)$ follow the recursions

$$\begin{aligned} \mathcal{A}_{m+1} &= \mathcal{A}_m + (\mu^{\mathbb{Q}})' \mathcal{B}_m + \frac{1}{2} \mathcal{B}_m' \Sigma \Sigma' \mathcal{B}_m - \delta_0 \\ \mathcal{B}_{m+1} &= (\Phi^{\mathbb{Q}})' \mathcal{B}_m - \delta_1 \end{aligned}$$

with starting values $\mathcal{A}_0 = 0$ and $\mathcal{B}_0 = 0$. Model-implied yields are determined by $y_t^m = -m^{-1} \log P_t^m = A_m + B_m X_t$, with $A_m = -m^{-1} \mathcal{A}_m$ and $B_m = -m^{-1} \mathcal{B}_m$. Risk-neutral yields, the yields that would prevail if investors were risk-neutral, can be calculated using

$$\tilde{y}_t^m = \tilde{A}_m + \tilde{B}_m X_t, \quad \tilde{A}_m = -m^{-1} \mathcal{A}_m(\mu, \Phi, \delta_0, \delta_1, \Sigma), \quad \tilde{B}_m = -m^{-1} \mathcal{B}_m(\Phi, \delta_1).$$

Risk-neutral yields reflect policy expectations over the life of the bond, $m^{-1} \sum_{h=0}^{m-1} E_t r_{t+h}$, plus a convexity term. The yield term premium is defined as the difference between actual and risk-neutral yields, $ytp_t^m = y_t^m - \tilde{y}_t^m$.

Denote by \hat{Y}_t the vector of observed yields on day t . The number of observed yield maturities is $J = 8$. We take the risk factors X_t to be the first $N = 3$ principal components of observed yields. That is, if W denotes the $N \times J$ matrix with rows corresponding to the first three eigenvectors of the covariance matrix of \hat{Y}_t , we have $X_t = W\hat{Y}_t$.

We parameterize the model using the canonical form of Joslin et al. (2011). Thus, the free parameters of the model are $r_\infty^Q = E^Q(r_t)$, the risk-neutral long-run mean of the short rate, λ^Q , the eigenvalues of Φ^Q , and the VAR parameters μ , Φ , and Σ . For the canonical model this leaves $1 + 3 + 3 + 9 + 6 = 22$ parameters to be estimated, apart from the measurement error specification. To see how μ^Q , Φ^Q , δ_0 , and δ_1 are calculated from $(W, \lambda^Q, r_\infty^Q, \Sigma)$ refer to Proposition 2 in Joslin et al. (2011).

B Frequentist estimation

B.1 Ordinary least squares

First we use OLS to obtain the VAR parameters in equation (3). The mean-reversion matrix Φ is estimated using a demeaned specification without intercept, and then the intercept vector is calculated as $\mu = (I_N - \Phi)\bar{X}$, where \bar{X} is the unconditional sample mean vector. The innovation covariance matrix is estimated from the residuals in the usual way. Denote these OLS estimates by $\hat{\mu}$, $\hat{\Phi}$ and $\hat{\Omega}$.

We obtain estimates of the cross-sectional parameters r_∞^Q and λ^Q using the approach of Hamilton and Wu (2010, henceforth HW). As cross-sectional measurements, Y_t^2 in HW's notation, we use the fourth principal component of yields. Write the corresponding eigenvector as the row vector W_2 , then we have $Y_t^2 = W_2\hat{Y}_t$. The reduced-form equations in the first step of the HW approach are the VAR for $Y_t^1 = X_t$ and the single measurement equation, which we write as

$$Y_t^2 = a + bY_t^1 + u_t, \quad (7)$$

for scalar a and row vector b , where u_t is a measurement error. The reduced-form parameters are $(\mu, \Phi, \Omega, a, b, \sigma_u^2)$, where $\sigma_u^2 = Var(u_t)$. The second step of the HW approach is to find the structural parameters that result in a close match for the reduced-form parameters, to be found by minimizing a chi-square distance statistic. A simplification is possible because we have exact identification, where the number of reduced-form parameters equals the number of structural parameters. Because the chi-square distance of the HW's second step reaches exactly zero, the weighting matrix is irrelevant and the problem separates into simpler, separate analytical and numerical steps, particularly simple in our case. The parameters for the VAR for Y_t^1 are directly available, namely $(\hat{\mu}, \hat{\Phi}, \hat{\Omega})$, because these parameters are both reduced-form and structural parameters. The parameters for the cross-sectional equation, a and b are found by choosing r_∞^Q and λ^Q so that the distance between the least squares estimates, (\hat{a}, \hat{b}) , and the model-implied values (W_2A_m, W_2B_m) is small. Here the J -vector A_m and the $J \times N$ matrix B_m contain the model-implied yield loadings. In addition to a dependence on Ω , B_m is determined

only by λ^Q , and A_m depends both on r_∞^Q and λ^Q . Therefore we can first search over values for λ^Q to minimize the distance between \hat{b} and $W_2 B_m$ – we use the Euclidean norm as the distance metric – and then pick r_∞^Q to minimize the distance between \hat{a} and $W_2 A_m$. Denote the resulting estimates by \hat{r}_∞^Q and $\hat{\lambda}^Q$.

Because OLS does most of the work in this estimation procedure, it is very fast even for a daily model. We have 6245 observations and the estimation takes only seconds.

The table shows the OLS estimates in the left column. The estimated intercept and the risk-neutral mean are scaled up by $100n$, where $n = 252$ is the number of periods (business days) per year. Thus these numbers correspond to annualized percentage points.

The estimated persistence is high: The largest eigenvalue of $\hat{\Phi}$, .999484, is close to one. The half life calculated from $\hat{\Phi}$ of the level factor in response to a level shock is 4.6 years.

	OLS			BC		
$\mu \cdot 100n$	-0.0276	0.0022	0.0076	-0.0223	0.0046	0.0073
Φ	0.9995	-0.0004	0.0251	0.9998	0.0000	0.0249
	-0.0004	0.9982	-0.0168	-0.0003	0.9986	-0.0167
	-0.0001	-0.0001	0.9876	-0.0001	-0.0002	0.9883
λ	0.999484	0.998266	0.987565	0.999770	0.998824	0.988035
$r_\infty^Q \cdot 100n$	12.37			12.38		
λ^Q	0.999774	0.998069	0.994425	0.999774	0.998069	0.994425

Note: Parameter estimates from frequentist estimation, obtained using OLS and BC. λ are the eigenvalues of Φ , λ^Q are the eigenvalues of Φ^Q .

B.2 Bias-corrected estimation

The intuition for our bias-corrected estimation procedure is to find parameters for the VAR that yield a median of the OLS estimator equal to the OLS estimates from the data. We use the indirect inference estimator detailed in Bauer et al. (2012). A residual bootstrap is used for every attempted value of Φ to generate data and find the median of the OLS estimator. In successive iterations, the attempted parameter values are adjusted using an updating scheme based on stochastic approximation, until the median of the OLS estimator on the generated data is sufficiently close to $\hat{\Phi}$. Denote the resulting estimate by $\tilde{\Phi}^{unr}$, indicating the unrestricted bias-corrected estimate.

In working with daily data, where the persistence is extremely high, our bias-corrected estimation procedure can lead to estimates for Φ with eigenvalues that are either greater than one or below but extremely close to one. This is unsatisfactory because it implies VAR dynamics that are either explosive or display mean reversion that is so slow as to be unnoticeable. Therefore we impose a restriction on our bias-corrected estimates, ensuring that the largest eigenvalue does not exceed *the largest eigenvalue under the pricing measure*. This seems to us a useful and intuitively appealing restriction, since from a finance perspective the far-ahead real-world expectations (under the physical measure) should not be more variable than the far-ahead risk-neutral expectations (under Q).³³ To obtain our bias-corrected estimate of Φ , we thus shrink $\tilde{\Phi}^{unr}$ toward $\hat{\Phi}$ using the adjustment procedure of Kilian (1998), until its largest

³³This intuition is also built into other models in the DTSM literature, such as Christensen et al. (2011)

eigenvalue is smaller, in absolute value, than the largest eigenvalue of $\hat{\Phi}$. The final adjusted bias-corrected estimate is denoted by $\tilde{\Phi}$.

Based on our estimate $\tilde{\Phi}$, we calculate the intercept $\tilde{\mu}$ and the innovation covariance matrix $\tilde{\Omega}$, as well as the cross-sectional parameters \tilde{r}_{∞}^Q and $\tilde{\lambda}^Q$ in analogous fashion as for OLS.

B.3 Bootstrap

To infer changes in risk-neutral rates and term premia, we construct a bootstrap distribution for the parameters of the DTSM. The focus is on the VAR parameters, since these crucially affect the characteristics of risk-neutral rates and premia. Because the cross-sectional parameters are estimated very precisely and re-estimating them on each bootstrap sample would be computationally costly, we only produce bootstrap distributions for Φ , μ , and Ω . As is evident from the estimation results, different values of the VAR parameters essentially have no effect on the estimated values for the cross-sectional parameters, so this simplification is completely innocuous.

By definition of the BC estimates, if we generate bootstrap samples (indexed by $b = 1, \dots, B$) using $\tilde{\Phi}^{unr}$, the OLS estimator has a median equal to $\hat{\Phi}$. The realizations of the OLS estimator on these samples thus provide a bootstrap distribution around $\hat{\Phi}$, which is conveniently obtained as a by-product of the bias correction procedure. We denote these bootstrap values by $\hat{\Phi}_b$.

To obtain a bootstrap distribution around the BC estimates $\tilde{\Phi}$, we shift the OLS bootstrap distribution by the estimated bias. That is, we set $\tilde{\Phi}_b = \hat{\Phi}_b + \tilde{\Phi} - \hat{\Phi}$, with the result that the values of $\tilde{\Phi}_b$ are centered around $\tilde{\Phi}$.

To ensure that the resulting VAR dynamics are stationary for every bootstrap replication, we again apply a stationarity adjustment similar to the one suggested by Kilian (1998). For the BC bootstrap replications, we shrink non-stationary values of $\tilde{\Phi}_b$ toward $\tilde{\Phi}$. We also apply such a stationarity adjustment if values of $\hat{\Phi}_b$ have non-stationary roots, in that case shrinking toward $\hat{\Phi}$. These stationarity adjustments have no impact on the median.

For each value of $\hat{\Phi}_b$ and $\tilde{\Phi}_b$ we calculate the corresponding estimates of μ and Ω as described earlier.

In terms of computing time, these bootstrap distributions are very quick to obtain. They naturally fall out of the bias-corrected estimation procedure. The only time-consuming task is the stationarity adjustment, which, however, has manageable computational cost.

Having available bootstrap distributions for the VAR parameters allows us to obtain bootstrap distributions for every object of interest, for example for the ten-year risk-neutral rate at a specific point in time, or for the cumulative changes in the ten-year yield term premium over a set of days. While our methodology is in some respects *ad hoc*, it has the unique advantage of enabling us to account in a relatively straightforward and computationally efficient way for the underlying estimation uncertainty of our inference about policy expectations and term premia.

where the largest Q-eigenvalue is unity and the VAR is stationary, or Joslin et al. (2010) where the largest eigenvalues under the two measures are restricted to be equal.

C Bayesian estimation

We employ Markov chain Monte Carlo (MCMC) methods to perform Bayesian estimation. Specifically, we obtain a sample from the joint posterior distribution of the model parameters using a block-wise Metropolis-Hastings (MH) algorithm. Other papers that have used MCMC methods for estimation of DTSMs include Ang et al. (2007), Ang et al. (2011) and Chib and Ergashev (2009). Our methodology is closely related to the one in Bauer (2011).

First we estimate the canonical model, and then, in a second step, we estimate over-identified models with zero restrictions on elements of λ_0 and λ_1 . For this purpose it is convenient to parameterize the model in terms of $(\lambda_0, \lambda_1, \Omega, r_\infty^Q, \lambda^Q)$.

The prior for the elements of λ_0 and λ_1 is independent normal, with mean zero and standard deviation .01. This prior cannot be too diffuse because that would affect the model selection exercise in the direction of favoring parsimonious models (the Lindley-Bartlett paradox; see Bartlett, 1957). In light of the magnitude of the frequentist estimates that we have obtained, this prior is not overly informative.

The priors for Ω and r_∞^Q are taken to be completely uninformative. The elements of λ^Q are a priori assumed to be independent, uniformly distributed over the unit interval.

For the measurement equations, we deviate slightly from our previous specification and simply take all J yields individually as the measurements, as in Joslin et al. (2011). The measurement errors are assumed to have equal variance, denoted by σ_u^2 . Notably, there are only $J - N$ independent linear combinations of these measurement errors, because N linear combinations of yields, namely the first three principal components, are priced perfectly by the model. We specify the prior for σ_u^2 to be uninformative.

C.1 Maximally-flexible model

Denote the parameters of the model as $\theta = (\lambda_0, \lambda_1, \Omega, r_\infty^Q, \lambda^Q, \sigma_u^2)$. There are five blocks of parameters which we draw successively in our MCMC algorithm.

The likelihood of the data factors into the likelihood of the risk factors, denoted by $P(X|\theta)$, and the cross-sectional likelihood, written as $P(Y|X, \theta)$ – X stands for all observations of X_t and Y stands for the data, i.e., all observations of \hat{Y}_t . The factor likelihood function is simply the conditional likelihood function of a Gaussian VAR.³⁴ It depends on the VAR parameters, which in this parameterization are determined by $(\lambda_0, \lambda_1, \Omega, r_\infty^Q, \lambda^Q)$. The cross-sectional likelihood function depends on $(\Omega, r_\infty^Q, \lambda^Q, \sigma_u^2)$. Thus we have

$$\begin{aligned} P(Y|\theta) &= P(X|\theta) \cdot P(Y|X, \theta) \\ &= P(X|\lambda_0, \lambda_1, \Omega, r_\infty^Q, \lambda^Q) \cdot P(Y|X, \Omega, r_\infty^Q, \lambda^Q, \sigma_u^2). \end{aligned}$$

The sampling algorithm allows us to draw from the joint posterior distribution

$$P(\theta|Y) \propto P(Y|\theta) \cdot P(\theta),$$

where $P(\theta)$ denotes the joint prior over all model parameters, despite the fact that this dis-

³⁴We always condition on the first observation.

tribution is only known up to a normalizing constant. This, of course, is the underlying idea of essentially all MCMC algorithms employed in Bayesian statistics.

As starting values of the chain, we use OLS estimates for μ , Φ , and Ω , the sample mean of all yields for r_∞^Q , the eigenvalues of $\hat{\Phi}$ for λ^Q , and a tenth of the standard deviation of all yields for σ_u (since yield pricing errors have smaller variance than yields).

We run the sampler for 50,000 iterations. We discard the first half as a burn-in sample and then take every 50'th iteration of the remaining sample. This constitutes our MCMC sample, which approximately comes from the joint posterior distribution of the parameters.

To ensure that the MCMC chain has converged, we closely inspect trace plots and make sure that our starting values have no impact on the results. In addition, we calculate convergence diagnostics of the type reviewed in Cowles and Carlin (1996).

C.1.1 Drawing (λ_0, λ_1)

Every element of λ_0 and λ_1 is drawn independently, iterating through them in random order, using a random walk (RW) MH step. For the conditional posterior distribution of these parameters we have

$$\begin{aligned} P(\lambda_0, \lambda_1 | \theta_-, X, Y) &\propto P(Y|\theta, X)P(X|\theta)P(\theta) \\ &\propto P(X|\theta)P(\theta), \end{aligned}$$

where θ_- denotes all parameters except for λ_0 and λ_1 . The second line follows because the likelihood of the data for given risk-neutral dynamics does not depend on the prices of risk, as noted earlier. For each parameter, we use a univariate random walk proposal with t_2 -distributed innovations that are multiplied by scale factors to tune the acceptance probabilities to be in the range of 20-50 percent. After obtaining the candidate draw, the restriction that the physical dynamics are non-explosive is checked, and the draw is rejected if the restriction is violated. Otherwise the acceptance probability for the draw is calculated as the minimum of one and the ratio of the factor likelihood times the ratio of the priors for the new draw relative to the old draw.

C.1.2 Drawing Ω

For the conditional posterior of Ω we have

$$P(\Omega | \theta_-, X, Y) \propto P(Y|\theta, X)P(X|\theta)P(\theta)$$

where θ_- denotes all parameters except Ω . Since we need successive draws of Ω to be close to each other—otherwise the acceptance probabilities will be too small—independence Metropolis is not an option. Element-wise RW MH does not work particularly well either. A better alternative in terms of efficiency and mixing properties is to draw the entire matrix Ω in one step. We choose a proposal density for Ω that is Inverse-Wishart (IW) with mean equal to the value of the previous draw and scale adjusted to tune the acceptance probability, which

is equal to

$$\alpha(\Omega^{(g-1)}, \Omega^{(g)}) = \min \left\{ \frac{P(X|\Omega^{(g)}, \theta_-)P(\Omega^{(g)}, \theta_-)q(\Omega^{(g)}, \Omega^{(g-1)})}{P(X|\Omega^{(g-1)}, \theta_-)P(\Omega^{(g-1)}, \theta_-)q(\Omega^{(g-1)}, \Omega^{(g)})}, 1 \right\},$$

where g is the iteration. Here $q(A, B)$ denotes the transition density, which in this case is the density of an IW distribution with mean A . The ratio of priors is equal to one since we assume an uninformative prior, unless the draw would imply nonstationary VAR dynamics, in which case the prior ratio is zero. The reason that some draws of Ω can imply nonstationary VAR dynamics is that in our normalization, the value of Ω matters for the mapping from r_∞^Q and λ^Q into μ^Q and Φ^Q , which together with λ_0 and λ_1 determine the VAR parameters.

C.1.3 Drawing r_∞^Q

Both factor likelihood and cross-sectional likelihood depend on r_∞^Q , thus

$$P(r_\infty^Q|\theta_-, X, Y) \propto P(Y|\theta, X)P(X|\theta)P(\theta),$$

where θ_- denotes all parameters except r_∞^Q . We use an RW MH step, with proposal innovations from a t-distribution with two degrees of freedom, multiplied by a scaling parameter to tune the acceptance probabilities. The ratio of priors is equal to one, because we have an uninformative prior, if the implied VAR dynamics are stationary and zero otherwise, in which case the prior ratio is zero. The acceptance probability is equal to the minimum of one and the product of prior ratio, the ratio of cross-sectional likelihoods, and the ratio of factor likelihoods.

C.1.4 Drawing λ^Q

Again both likelihoods depend on this parameter, so we have

$$P(\lambda^Q|\theta_-, X, Y) \propto P(Y|\theta, X)P(X|\theta)P(\theta),$$

where θ_- denotes all parameters except λ^Q . We draw all three elements in one step, using an RW proposal with independent t-distributed innovations, each with two degrees of freedom and multiplied to tune acceptance probabilities. The prior ratio is one if all three proposed values are within the unit interval and the implied VAR dynamics are stationary, and zero otherwise. We implement the requirement that the three elements of λ^Q are in descending order by rejecting draws that would change this ordering. Again the acceptance probability is equal to the minimum of one and the product of prior ratio, the ratio of cross-sectional likelihoods, and the ratio of factor likelihoods.

C.1.5 Drawing σ_u^2

In this block the conditional posterior distribution of σ_u^2 is known in close form. The problem of drawing this error variance corresponds to drawing the error variance of a pooled regression. The condition posterior distribution is inverse gamma, because an uninformative prior on this parameter is conjugate.

C.2 Restricted risk prices

We closely follow the methodology laid out in Bauer (2011), where Gibbs variable selection (Dellaportas et al., 2002) is applied to the context of DTSM estimation. Let λ denote a vector stacking all elements of λ_0 and λ_1 . For the purpose of model selection, we introduce a vector of indicator variables, γ , that describes which risk price parameters, i.e., which elements of λ , are restricted to zero. The parameters of the model are now $(\gamma, \theta) = (\gamma, \lambda, \Omega, r_\infty^Q, \lambda^Q, \sigma_u^2)$. The goal of course is to sample from the joint posterior

$$P(\gamma, \theta|Y) \propto P(Y|\gamma, \theta)P(\theta|\gamma)P(\gamma).$$

The likelihood $P(Y|\gamma, \theta)$ is the product of factor likelihood and cross-sectional likelihood, as before. The difference is that here it is evaluated by treating those elements of λ as zero for which the corresponding element in γ is zero. The priors for the parameters conditional on the model indicator $P(\theta|\gamma)$ are specified as before. The prior for the model indicators $P(\gamma)$ is such that all elements are independent Bernoulli random variables with .5 prior probability.

The parameters $\Omega, r_\infty^Q, \lambda^Q$, and σ_u^2 are drawn exactly as in the estimation algorithm for the URP model. What is different here is we sample the vector indicating the model specification, γ , and the parameter vector γ , which all models have in common.

For each iteration g of the MCMC sampler, we draw the block (γ, λ) by drawing pairs (γ_i, λ_i) , going through the $N + N^2 = 12$ risk price parameters in random order.

C.2.1 Drawing λ_i

For each pair we first draw $\lambda_i^{(g)}$ conditional on $\gamma_i^{(g-1)}$ and all other parameters. If the parameter is currently included (unrestricted), i.e., if $\gamma_i = 1$, we draw from the conditional posterior. If the parameter is currently restricted to zero ($\gamma_i = 0$) the data is not informative about the parameter and we draw from a so-called pseudo-prior (Carlin and Chib, 1995; Dellaportas et al., 2002). That is,

$$P(\lambda_i|\lambda_{-i}, \gamma_i = 1, \gamma_{-i}, \theta_-, X, Y) \propto P(X|\theta, \gamma)P(\lambda_i|\gamma_i = 1) \quad (8)$$

$$P(\lambda_i|\lambda_{-i}, \gamma_i = 0, \gamma_{-i}, \theta_-, X, Y) \propto P(\lambda_i|\gamma_i = 0), \quad (9)$$

where θ_- denotes all parameters in θ other than λ , and λ_{-i} (γ_{-i}) contains all elements of λ (γ) other than λ_i (γ_i).³⁵ We assume prior conditional independence of the elements of λ given γ , and the prior for each price of risk parameter, $P(\lambda_i|\gamma_i = 1)$, is taken to be standard normal. The conditional posterior in equation (8) is not known analytically and we use an RW MH step to obtain the draws, with a fat-tailed RW proposal and scaling factor as before. For the pseudo-prior $P(\lambda_i|\gamma_i = 0)$ we use a normal distribution, with moments corresponding to the marginal posterior moments from our estimation of the URP model.

³⁵These conditional distributions parallel the ones in equations (9) and (10) of Dellaportas et al. (2002).

C.2.2 Drawing γ_i

When we get to the second element of the pair, the indicator γ_i , the conditional posterior distribution is known and we can directly sample from it without the MH step. It is Bernoulli, and the success probability is easily calculated based on the ratio:

$$q = \frac{P(\gamma_i = 1 | \gamma_{-i}, \theta, X, Y)}{P(\gamma_i = 0 | \gamma_{-i}, \theta, X, Y)} = \frac{P(X | \gamma_i = 1, \gamma_{-i}, \theta) P(\lambda_i | \gamma_i = 1) P(\gamma_i = 1, \gamma_{-i})}{P(X | \gamma_i = 0, \gamma_{-i}, \theta) P(\lambda_i | \gamma_i = 0) P(\gamma_i = 0, \gamma_{-i})}. \quad (10)$$

The first factor in the numerator and the denominator is the factor likelihood. The second factor in the numerator is the parameter prior, and in the denominator it is the pseudo-prior. The third factor cancels out, since we use an independent, uninformative prior with prior inclusion probability of each element of 0.5, putting equal weight on $\gamma_i = 1$ and $\gamma_i = 0$. The conditional posterior probability for drawing $\gamma_i = 1$ is given by $q/(q+1)$.³⁶

C.2.3 Bayesian model averaging

As output from the MCMC algorithm, we have available a sample that comes approximately from the joint posterior distribution of (γ, θ) . When we want to calculate the posterior distribution of any object of interest, such as for the value of the ten-year term premium on a certain day, we simply calculate it for every iteration of the MCMC sample. In each iteration that we use from this sample – as before we discard the first half and then only use every 50'th iteration – different elements might be restricted to zero. By effectively sampling across models and parameter values we are taking into account model uncertainty in our posterior inference. This technique is called Bayesian model averaging: the model specification is effectively averaged out, and the inference is not conditional on a specific model but instead takes into account model uncertainty.

³⁶A subtlety, which is ignored in the above notation, is that the joint prior $P(\gamma, \theta)$ imposes that the physical dynamics resulting from any choice of γ and λ_1 can never be explosive. This is easily implemented in the algorithm: If including a previously excluded element would lead to explosive dynamics then we simply do not include it, i.e., we set $\gamma_i = 0$, and vice versa.

Table 1: LSAP announcements

Date	Announcement	Description
25 November 2008	initial LSAP announcement	Federal Reserve announces purchases of up to \$100 billion in agency debt and up to \$500 billion in agency MBS.
1 December 2008	Chairman's speech	Chairman states that the Federal Reserve "could purchase longer-term Treasury securities [...] in substantial quantities."
16 December 2008	FOMC statement	Statement indicates that the FOMC is considering expanding purchases of agency securities and initiating purchases of Treasury securities.
28 January 2009	FOMC statement	Statement indicates that the FOMC "is prepared to purchase longer-term Treasury securities."
18 March 2009	FOMC statement	Statement announces purchases "up to an additional \$750 billion of agency [MBS]," \$100 billion in agency debt, and \$300 billion in Treasury securities.
12 August 2009	FOMC statement	Statement drops "up to" language and announces slowing pace for purchases of Treasury securities.
23 September 2009	FOMC statement	Statement drops "up to" language for purchases of agency MBS and announces gradual slowing pace for purchases of agency debt and MBS.
4 November 2009	FOMC statement	Statement declares that the FOMC would purchase "around \$175 billion of agency debt."

Table 2: Easing actions and term premium changes, 2001-2003

Date	Change in FFR target	Change in 10y yield		
		Actual	YRN	YTP
01/31/2001	-50	-4	-3	0
03/20/2001	-50	-3	-2	-1
04/18/2001	-50	-6	-5	-1
08/21/2001	-25	-3	-2	-1
10/02/2001	-50	-2	-2	1
11/06/2001	-50	-2	-3	1
12/11/2001	-25	-3	-2	-2
05/07/2002	0	0	-1	0
06/26/2002	0	-12	-4	-7
08/13/2002	0	-9	-4	-5
09/24/2002	0	-1	-1	-1
11/06/2002	-50	-3	-3	1
05/06/2003	0	-8	-3	-6
Cumulative	-350	-56	-35	-21

Note: Changes, in basis points, in the fed funds rate (FFR) target, actual ten-year yield, and the Kim-Wright estimated risk-neutral yield and yield term premium, on days with FOMC meetings during the 2001-2003 easing cycle that also had a decline in the risk-neutral yield. Changes in YRN and YTP do not always sum up to actual yield changes because the DTSM does not fit yields perfectly.

Table 3: Changes in futures-implied policy paths around LSAP announcements

Date	1m	6m	1y	2y	3y	avg. 3y	3y yld.	diff.
11/25/2008	-5	-6	-10	-13	-22	-12	-18	-7
12/1/2008	1	-4	-7	-18	-21	-11	-16	-5
12/16/2008	-17	-16	-12	-11	-16	-12	-13	-1
1/28/2009	0	0	5	11	15	7	8	0
3/18/2009	-1	-4	-11	-10	-11	-8	-35	-27
8/12/2009	-1	-6	-8	-3	-1	-4	-1	3
9/23/2009	0	-3	-5	-6	-2	-4	-4	0
11/4/2009	0	-2	-1	1	5	1	0	-1
Total	-23	-40	-49	-49	-53	-43	-80	-37
Cum. changes	-33	-27	28	107	122	62	24	-38
Std. dev.	1	2	5	8	9	6	7	4

Note: Changes, in basis points, of futures-implied policy paths at fixed horizons. Paths are linearly interpolated if no futures contract is available for required horizon. The last three columns show the change of the average policy path over the next three years, the change in the three-year zero coupon yield, and the difference between the yield change and the change in the average policy path. The bottom two rows show the cumulative changes and standard deviations of daily changes over the period 11/24/08 to 12/30/09.

Table 4: Changes in yields, OIS rates, and spreads around LSAP announcements

Date	OIS rates			yields			yield-OIS		
	2y	5y	10y	2y	5y	10y	2y	5y	10y
11/25/2008	-14	-25	-28	-14	-22	-21	-1	2	7
12/1/2008	-13	-21	-19	-12	-21	-22	1	-1	-2
12/16/2008	-15	-29	-32	-11	-16	-17	5	12	14
1/28/2009	6	11	14	5	10	12	-1	-1	-2
3/18/2009	-12	-27	-38	-26	-47	-52	-14	-20	-14
8/12/2009	-1	-2	1	-1	1	6	0	3	5
9/23/2009	-5	-6	-5	-4	-4	-2	1	3	3
11/4/2009	-3	1	5	-1	3	7	2	2	2
Total	-58	-97	-102	-65	-97	-89	-7	0	14
Cum. changes	-8	19	59	2	31	16	10	11	-43
Std. dev.	5	8	10	6	8	9	3	3	4

Note: Changes, in basis points, in OIS rates, zero-coupon yields, and yield-OIS spreads around LSAP announcements. The bottom two rows show the cumulative changes and standard deviations of daily changes over the period 11/24/08 to 12/30/09.

Table 5: Decomposition of LSAP effect on long-term yields

	ten-year yield			five-year yield		
	yield	YRN	YTP	yield	YRN	YTP
actual	-89			-97		
Kim-Wright	-102	-31	-71	-94	-30	-64
OLS	-93	-33	-60	-93	-40	-53
OLS UB		-90	-3		-85	-9
OLS LB		9	-102		0	-94
BC	-93	-46	-47	-93	-48	-46
BC UB		-141	48		-112	19
BC LB		0	-93		-3	-90
URP	-94	-31	-62	-93	-39	-53
URP UB		-71	-23		-69	-24
URP LB		-7	-86		-14	-78
RRP	-94	-36	-58	-93	-48	-44
RRP UB		-53	-40		-59	-33
RRP LB		-29	-65		-41	-51

Note: Alternative decompositions of yield changes, in basis points, on announcement days. The first line shows actual yield changes, the following lines show changes in fitted yields, risk-neutral yields (YRN) and yield term premia (YTP) for alternative DTSM estimates. Also shown are upper bounds (UB) and lower bounds (LB) for the change in the term premium, based on bootstrap confidence intervals (for OLS and BC) or quantiles of posterior distributions (for URP and RRP).

Table 6: Ten-year yield, decompositions of day-by-day changes

Date	act.	Kim-Wright			OLS			BC		
		yld.	YRN	YTP	yld.	YRN	YTP	yld.	YRN	YTP
11/25/2008	-21	-24	-7	-17	-23	-6	-17	-23	-8	-15
12/1/2008	-22	-24	-7	-17	-22	-5	-17	-22	-7	-15
12/16/2008	-17	-18	-7	-12	-17	-5	-13	-17	-6	-11
1/28/2009	12	12	3	9	13	-2	15	13	-2	15
3/18/2009	-52	-56	-16	-40	-53	-7	-46	-53	-10	-43
8/12/2009	6	4	1	3	5	-3	8	5	-4	8
9/23/2009	-2	-2	-1	-1	-2	-3	1	-2	-4	3
11/4/2009	7	7	2	5	7	-3	10	7	-4	11
Total	-89	-102	-31	-71	-93	-33	-60	-93	-46	-47
Cum. changes	16	24	-7	31	30	-10	40	30	-12	42
Std. dev.	9	9	3	7	9	4	9	9	5	9

Date	act.	Kim-Wright			URP			RRP		
		yld.	YRN	YTP	yld.	YRN	YTP	yld.	YRN	YTP
11/25/2008	-21	-24	-7	-17	-23	-6	-17	-23	-9	-14
12/1/2008	-22	-24	-7	-17	-22	-6	-17	-22	-9	-14
12/16/2008	-17	-18	-7	-12	-17	-5	-13	-17	-7	-10
1/28/2009	12	12	3	9	13	-1	14	13	5	8
3/18/2009	-52	-56	-16	-40	-54	-9	-44	-54	-21	-32
8/12/2009	6	4	1	3	5	-2	7	5	2	3
9/23/2009	-2	-2	-1	-1	-2	-3	1	-2	-1	-1
11/4/2009	7	7	2	5	7	-2	9	7	2	4
Total	-89	-102	-31	-71	-94	-34	-60	-94	-37	-56
Cum. changes	16	24	-7	31	30	-7	37	30	10	20
Std. dev.	9	9	3	7	9	3	8	9	4	6

Note: Decompositions of yield changes, in basis points, on each LSAP announcement day. The first column shows actual yield changes, the following columns show changes in fitted yields, risk-neutral yields (YRN) and yield term premia (YTP) for alternative DTSM estimates. The bottom three rows show the total changes over all events, as well as cumulative changes and standard deviations of daily changes over the period 11/24/08 to 12/30/09.

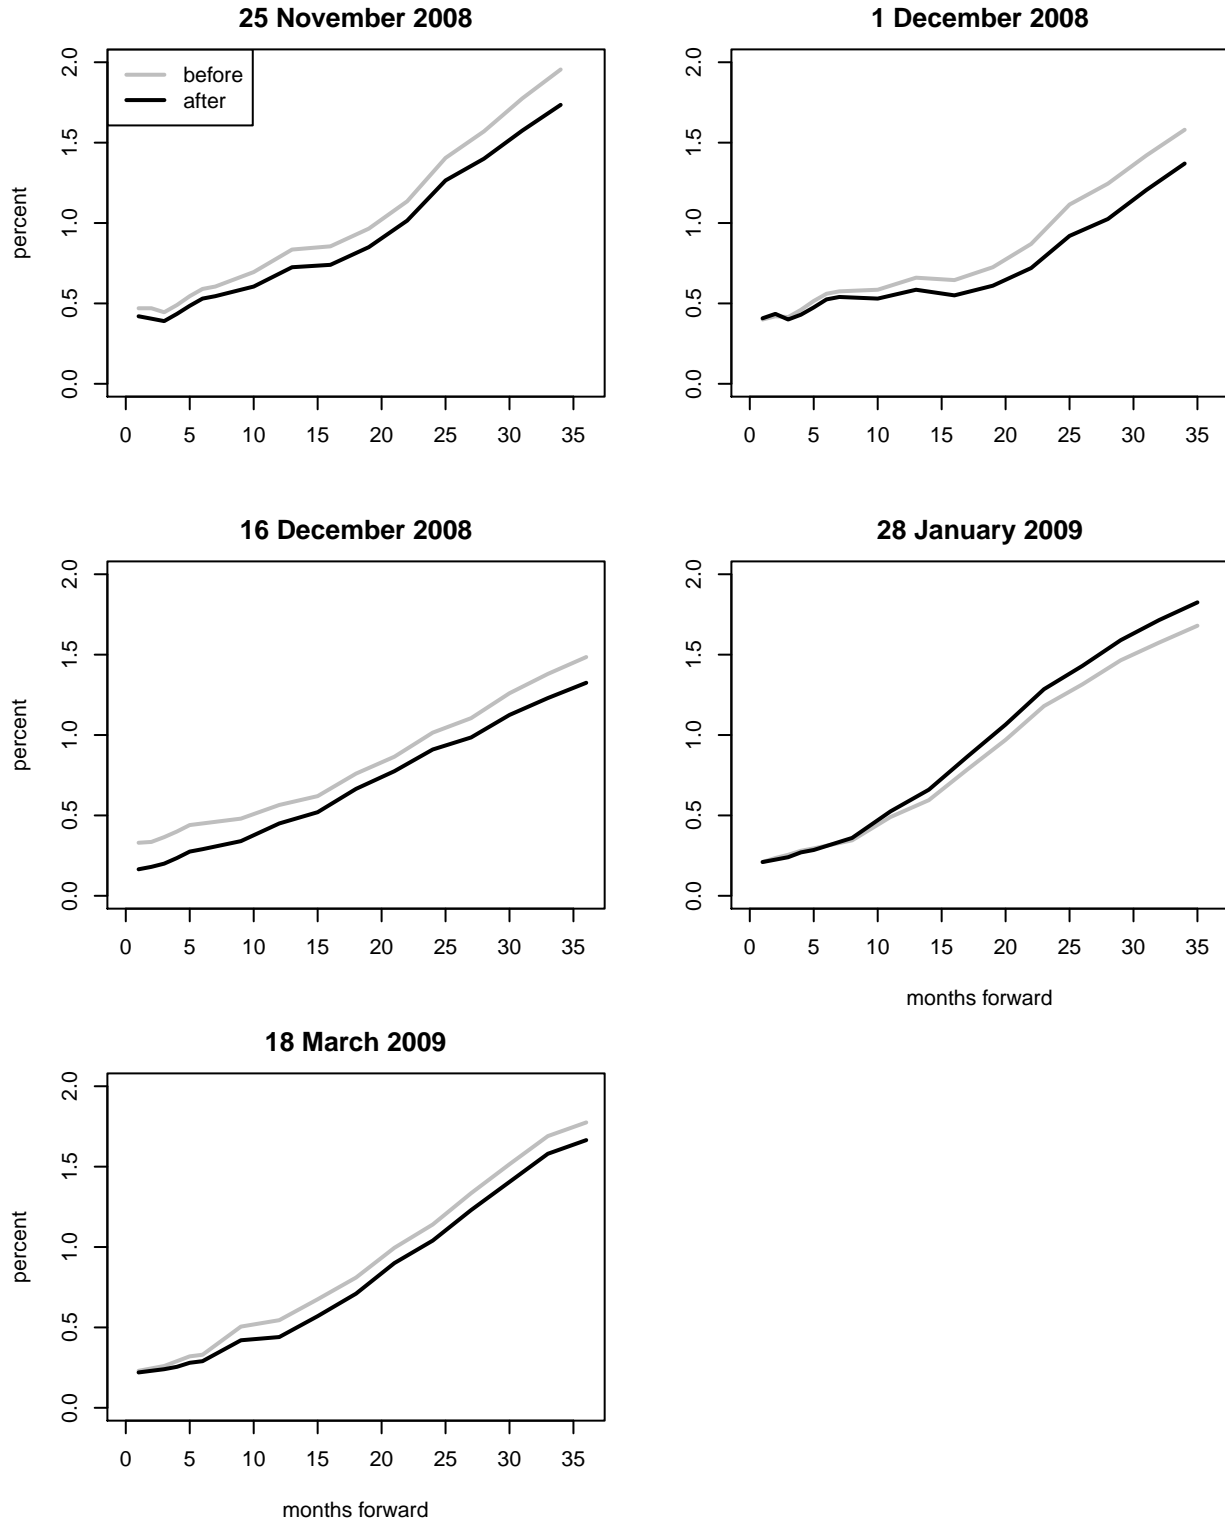
Table 7: Five-year yield, decompositions of day-by-day changes

Date	act.	Kim-Wright			OLS			BC		
		yld.	YRN	YTP	yld.	YRN	YTP	yld.	YRN	YTP
11/25/2008	-22	-22	-7	-15	-21	-7	-15	-21	-8	-13
12/1/2008	-21	-21	-6	-15	-21	-6	-15	-21	-7	-14
12/16/2008	-16	-16	-6	-10	-16	-5	-11	-16	-6	-10
1/28/2009	10	9	3	7	9	-2	12	9	-3	12
3/18/2009	-47	-47	-13	-34	-46	-8	-39	-46	-9	-37
8/12/2009	1	2	0	2	2	-4	6	2	-4	7
9/23/2009	-4	-3	-1	-2	-3	-4	1	-3	-5	1
11/4/2009	3	4	1	3	4	-4	8	4	-5	8
Total	-97	-94	-30	-64	-93	-40	-53	-93	-48	-46
Cum. changes	31	20	-10	29	19	-14	33	19	-16	35
Std. dev.	8	8	3	6	8	5	7	8	5	7

Date	act.	Kim-Wright			URP			RRP		
		yld.	YRN	YTP	yld.	YRN	YTP	yld.	YRN	YTP
11/25/2008	-22	-22	-7	-15	-21	-7	-14	-21	-11	-10
12/1/2008	-21	-21	-6	-15	-21	-6	-14	-21	-10	-10
12/16/2008	-16	-16	-6	-10	-16	-6	-11	-16	-9	-8
1/28/2009	10	9	3	7	9	-1	10	9	5	5
3/18/2009	-47	-47	-13	-34	-46	-10	-36	-46	-24	-22
8/12/2009	1	2	0	2	2	-3	5	2	1	1
9/23/2009	-4	-3	-1	-2	-3	-4	0	-3	-2	-1
11/4/2009	3	4	1	3	4	-3	7	4	1	2
Total	-97	-94	-30	-64	-93	-40	-53	-93	-49	-44
Cum. changes	31	20	-10	29	19	-11	30	19	7	24
Std. dev.	8	8	3	6	8	4	7	8	4	5

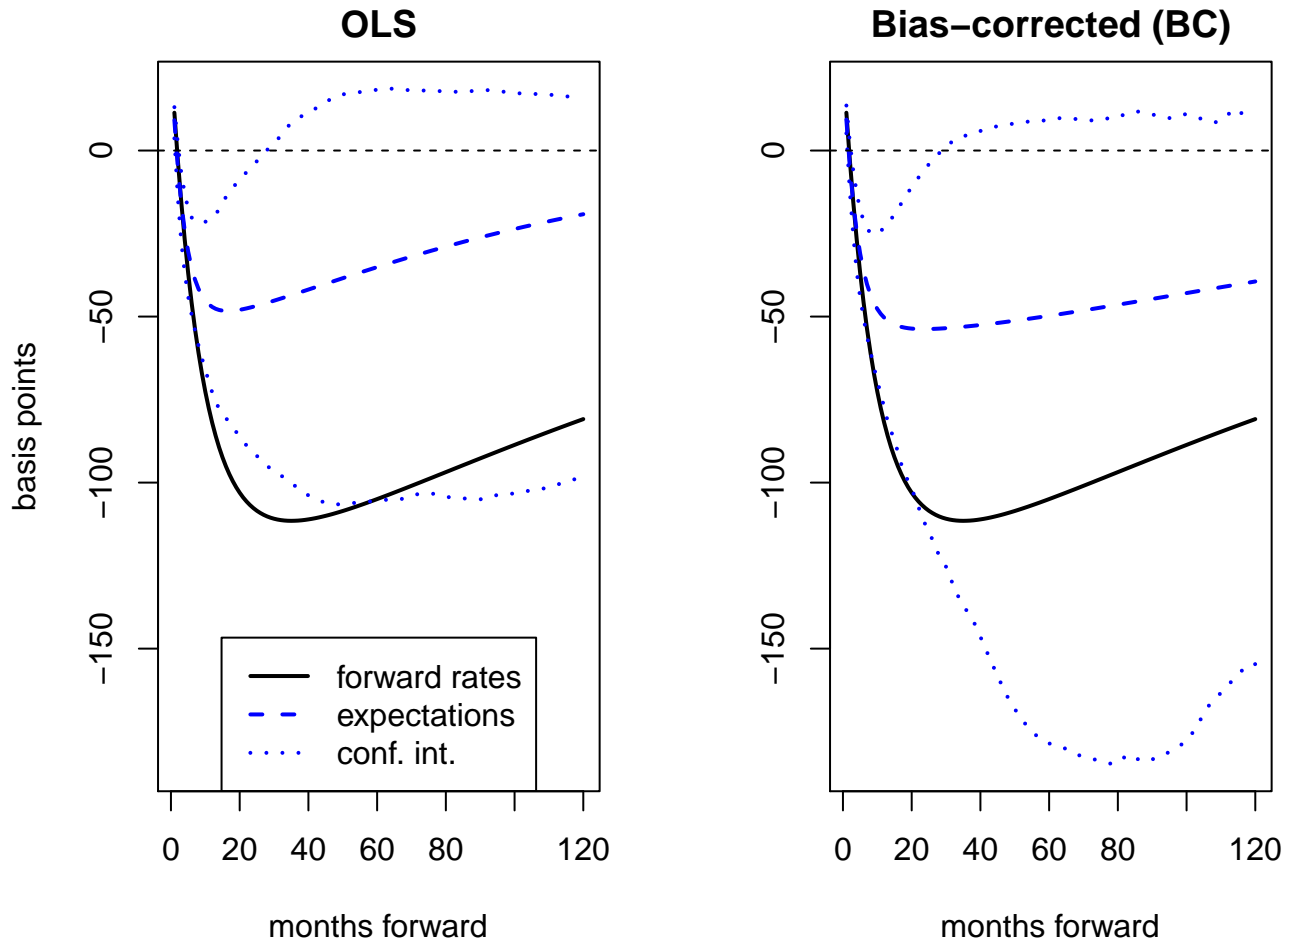
Note: See Table 6.

Figure 1: Shifts of futures-implied policy paths around key LSAP dates



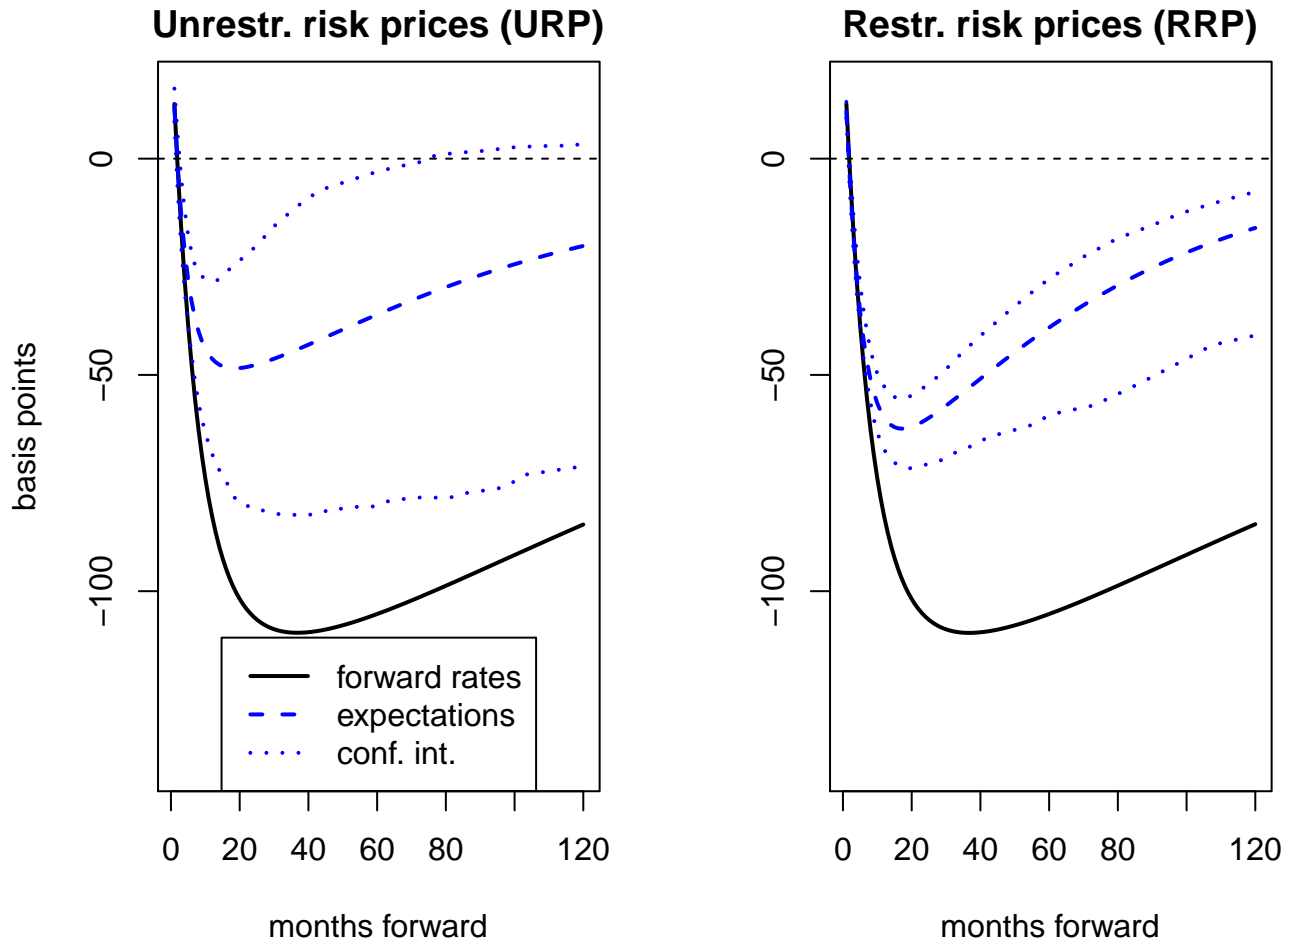
Note: Policy paths before and after five key LSAP announcements that are implied by market rates of federal funds futures and Eurodollar futures. For details on calculation, refer to main text.

Figure 2: Shift of forward curve and policy path: OLS vs. BC



Note: Cumulative changes, in basis points, on announcement days in fitted forward rates (solid line) and policy expectations (dashed line) together with 95%-confidence intervals for changes in expectations (dotted lines). Left panel shows decomposition based on OLS estimates, right panel for BC estimates.

Figure 3: Shift of forward curve and policy path: URP vs. RRP



Note: Cumulative changes, in basis points, on announcement days in fitted forward rates (solid line) and policy expectations (dashed line) together with 95%-confidence intervals for changes in expectations (dotted lines). Left panel shows decomposition based on URP estimates, right panel for RRP estimates.