

High Frequency Identification of Monetary Non-Neutrality

Emi Nakamura and Jón Steinsson*

Columbia University

October 25, 2015

Abstract

We present estimates of monetary non-neutrality based on evidence from high-frequency responses of nominal and real interest rates. Our identifying assumption is that unexpected changes in interest rates in a 30-minute window surrounding scheduled Federal Reserve announcements arises from news about monetary policy. At these times, nominal and real interest rates respond roughly one-for-one, several years out into the term structure, while the response of expected inflation is small. We use this evidence to estimate key parameters of a workhorse New Keynesian model. The implied degree of monetary non-neutrality is large. Moreover, we find evidence of a “Fed information effect”: FOMC announcements affect expectations not only about the evolution of monetary policy but also about future economic fundamentals.

Keywords: Real Interest Rates, Heteroskedasticity-based Estimation, Fed Information.

JEL Classification: E30, E40, E50

*We thank Matthieu Bellon, Vlad Bouchouev, Nicolas Crouzet, Stephane Dupraz, Jesse Garret, and Shaowen Luo, for excellent research assistance. We thank Michael Abrahams, Tobias Adrian, Richard K. Crump, Matthias Fleckenstein, Michael Fleming, Mark Gertler, Refet Gurkaynak, Peter Karadi, Hanno Lustig, Emanuel Moench, and Eric Swanson for generously sharing data and programs with us. We thank Marco Bassetto, Stephane Dupraz, Gauti Eggertsson, Mark Gertler, Refet Gurkaynak, Samuel Hanson, Sophocles Mavroeidis, Emanuel Moench, Serena Ng, Roberto Rigobon, David Romer, Christoph Rothe, Eric Swanson, Michael Woodford, Jonathan Wright and seminar participants at various institutions for valuable comments and discussions. We thank the National Science Foundation (grant SES-1056107) and the Columbia Business School Dean’s Office Summer Research Assistance Program for financial support.

1 Introduction

A fundamental question in macroeconomics is how monetary policy affects the economy. The key empirical challenge in answering this question is that most changes in interest rates happen for a reason. For example, the Fed might lower interest rates to counteract the effects of an adverse shock to the financial sector. In this case, the effect of the Fed’s actions are confounded by the financial shock, making it difficult to identify the effects of monetary policy. Two approaches used to overcome this endogeneity problem in the existing literature are structural vector autoregressions (e.g., Christiano, Eichenbaum, and Evans, 1999) and Romer and Romer’s (2004) approach of looking at the effects of changes in the intended federal funds rate that are orthogonal to the Fed’s information set as measured by its staff forecast. These approaches control for several important channels of potential endogeneity. The concern remains, however, that not all endogenous variation has been purged from these measures of monetary shocks.

An alternative approach—the one we pursue in this paper—is to focus on movements in bond prices in a narrow window around scheduled Federal Open Market Committee (FOMC) meetings. This high frequency identification approach was pioneered by Cook and Hahn (1989), Kuttner (2001), and Cochrane and Piazzesi (2002). It exploits the fact that monetary news is revealed in a lumpy fashion, with a disproportionate amount of monetary news revealed at the time of the eight regularly scheduled FOMC meetings each year.

What is appealing about the high frequency identification approach is how cleanly it is able to address the endogeneity concern. Our monetary shocks are constructed using unexpected changes in interest rates over a 30-minute window surrounding scheduled Federal Reserve announcements. All information that is public at the beginning of the 30-minute window will already be incorporated into financial markets, and will, therefore, not show up as spurious variation in the monetary shock. Such spurious variation is an important concern with VAR’s. Cochrane and Piazzesi (2002) argue, for example, that the interest rate change following the September 11, 2001 terrorist attacks is picked up as a monetary shock in conventional monthly and quarterly VAR-based identification strategies, though this is clearly an example of monetary policy responding to macroeconomic news occurring over that month. Even beyond the endogeneity issue, the predictions of VAR’s about the response of interest rates to monetary policy shocks are sensitive to model specification (what variables are included, how many lags, etc.), whereas the high frequency approach allows one to read this information directly out of bond market data.

Our measure of monetary shocks incorporates two important innovations developed in the high-

frequency identification literature. First, to capture the effects of forward guidance—i.e., announcements by the Fed that convey information about future changes in the Federal Funds Rate—we base our monetary policy indicator on a composite measure of changes in interest rates at different maturities spanning the first year of the term structure. Over the past 15 years, forward guidance has become an increasingly important tool in the conduct of monetary policy. This shift implies that the measure of monetary shocks used in most prior work—unexpected changes in the Federal Funds Rate—captures only a small fraction of monetary policy news associated with FOMC announcements (Gurkaynak, Sack, and Swanson, 2005). Second, we show that it is crucial to focus on a narrow 30-minute window (as opposed to a wider 1-day or 2-day window) to measure the effects of monetary policy. We use Rigobon’s (2003) heteroskedasticity-based estimation approach to show that estimates based on longer one-day windows are confounded by substantial “background noise,” leading to misleading results.

The monetary shocks we identify have large and persistent effects on both nominal and real interest rates. In fact, nominal and real interest rates respond roughly one-for-one several years out into the term structure. A monetary shock that raises the 2-year nominal yield on Treasuries by 110 basis points, raises the 2-year real TIPS yield by 106 basis points. The effect of this shock on the 2-year instantaneous real forward rate is 99 basis points. The impact of the shock then falls monotonically at longer horizons to 88 basis points at 3 years, 47 basis points at 5 years, and 12 basis point at 10 years. The effect of the monetary shock on the 5-year real forward rate is statistically significant, while its effect on the 10-year real forward rate is not.

Despite the large response of real interest rates to the monetary shock, the response of break-even inflation is essentially zero at horizons up to three years. At longer horizons, the response of break-even inflation becomes modestly, but significantly, negative. A tightening of monetary policy therefore eventually reduces inflation—as theory would predict. However, the response is small and occurs only after a long lag.

We use this high-frequency evidence to estimate the degree of monetary non-neutrality in the economy. We begin by employing the textbook, three-equation, New Keynesian model to develop intuition for how the high-frequency evidence can shed light on the degree of monetary non-neutrality. We show that several of the key parameters of this model are identified by the relative magnitude of the response of inflation and the response of real interest rates to a monetary shock. Intuitively, if the response of inflation is small when real interest rates move a substantial amount it must be that output responds little to real interest rates, inflation responds little to output, or both. Given

a reasonable elasticity of output to real interest rates, the small response of inflation to real interest rates is a gauge of the magnitude of price and wage adjustment frictions in the economy—the small response of inflation then indicates that nominal and real rigidities must be large. In addition to this, the fact that the response of inflation to the monetary shock is zero initially and builds over time suggests the presence of substantial inflation inertia—i.e., a lagged inflation term in the Phillips curve.

We build on these intuitions to estimate the workhorse business cycle model proposed by Christiano, Eichenbaum and Evans (2005, CEE) and further developed by Altig et al. (2011, ACEL). We estimate the parameters by simulated method of moments. Our empirical approach is analogous to the impulse response matching approach used by Rotemberg and Woodford (1997), CEE, and ACEL, except that we are estimating the parameters to fit our new high-frequency evidence on interest rate responses as opposed to impulse responses from a structural VAR. Our estimates imply that monetary non-neutrality is large. Output responds 3.8 times as much as inflation to a standard monetary shock for our estimates. This ratio is 3.3 for the parameters obtained by ACEL and 1.7 for the parameters obtained by CEE. On this metric, our estimates, thus, imply a similar amount of monetary non-neutrality as ACEL’s estimates, but substantially more than CEE’s. We also consider a modification of our baseline approach based on a recent hybrid high-frequency VAR approach proposed by Gertler and Karadi (2015) and show that it, too, yields similar estimates of the key parameters, despite relying on quite a different set of identifying assumptions. In particular, this alternative approach allows us to be agnostic about the extent to which monetary policy affects output through risk premium effects.

We extend our baseline model to allow for a “Fed information effect,” whereby FOMC announcements may affect private sector beliefs about the future evolution of exogenous economic fundamentals. We present evidence that private sector forecasts of future output growth from *Blue Chip Economic Indicators* rise when the Fed announces a surprise tightening of policy. This is the opposite of what one might expect from a model without the Fed information effect, but a natural response if the private sector responds to a Fed interest rate hike as a signal that the fundamentals of the economy are stronger than it previously believed.¹ When we calibrate our model to match these facts, we estimate somewhat less monetary non-neutrality than in our baseline case. Intuitively, the Fed information effect works in the opposite direction from the direct effect of a Fed interest rate cut:

¹Campbell et al. (2012) present similar evidence regarding the effect of surprise monetary shocks on *Blue Chip* expectations about unemployment. See also Romer and Romer (2000) and Faust, Swanson, and Wright (2004) for earlier empirical evidence regarding the effects of Fed information.

the interest rate cut stimulates private sector activity, but signals worse times ahead, making price increases less appealing. Hence, less price rigidities are required to explain why inflation responds so little to monetary shocks—lowering the degree of monetary non-neutrality in the model. In this case, the degree of real rigidity needed to fit the data is comparable to the degree of real rigidity assumed in the specific factor model of Woodford (2003, ch. 3).

An important question is whether some of the effects of monetary shocks on longer-term real interest rates we estimate reflect changes in risk premia as opposed to changes in expected future short-term real interest rates. We are certainly not arguing that risk premia do not play a role in interest rates more generally, but rather that, as Piazzesi and Swanson (2008) suggest, interest rate movements at the time of FOMC announcements are associated mainly with changes in expected interest rates. We evaluate this issue in several ways. First, we use direct measures of expectations from *Blue Chip Economic Indicators*. The *Blue Chip* data confirms the results of our baseline analysis that the real rate effects of monetary shocks are large, while the effect on expected inflation is small, though the estimates are much less precise. Second, we directly estimate the effects of our monetary shocks on risk premia using a state-of-the-art affine term structure model (Abrahams et al., 2015). While this model predicts that a large fraction of interest rate movements at other times *are* associated with risk premia, this is not the case for interest rate movements at the time of FOMC announcements. In other words, the expectations hypothesis of the term structure is a good approximation in response to our monetary shocks, even though it is not a good approximation unconditionally. This is what we need for our analysis to be valid. Third, we find little evidence that the interest rate effects we identify dissipate quickly after the announcement, as would be predicted by some models of liquidity premia.² Finally, we use Gertler and Karadi’s (2015) hybrid high-frequency VAR approach to construct estimates based only on the empirical responses of output and inflation—which do not require us to take a stand on the role of risk premia.

The two most related empirical papers to our paper are Hanson and Stein (2015) and Gertler and Karadi (2015). As we discuss in section 3, we make different identifying assumptions than Hanson and Stein and use a different definition of the monetary shock, and come to quite different conclusions about the long-run effects of monetary policy. There are also very substantial methodological differences between our work and that of Gertler and Karadi (2015). They rely on a VAR to estimate the dynamic effects of monetary policy. Our identification approach is entirely VAR-free. Many researchers are skeptical of the ability of VARs to overcome endogeneity concerns. Gertler

²Hanson and Stein (2015) present a behavioral model in which “search for yield” generates significant risk premium effects of monetary shocks that dissipate over time.

and Karadi (2015) are able to avoid the timing assumptions typically made when VARs are used to identify the effects of monetary policy. But they are subject to the concern that the lag structure of the VAR is not able to capture all available relevant information that monetary policy actions are based on (recall the 9/11 example we discuss above). We think it is valuable to develop an alternative identification approach that is entirely unreliant on these assumptions. In addition, Gertler and Karadi do not study the implications of their estimates for the degree of monetary non-neutrality.

There are a number of other papers that use the high frequency identification approach to estimate the effects of monetary shocks, though none attempt to use this evidence to study the extent of monetary non-neutrality. Wright (2012) studies the effects of unconventional monetary policy on interest rates (both nominal and real) during the recent period over which short-term nominal interest rates have been at their zero lower bound. Gagnon et al. (2010), Krishnamurthy and Vissing-Jorgensen (2011), and Rosa (2012) use high frequency identification methods to study the effect of large-scale asset purchases by the Federal Reserve since the 2008 financial crisis. Gilchrist, Lopez-Salido, and Zakrajsek (2015) use a high frequency approach to compare the effects of conventional monetary policy to those of the unconventional measures employed after the Federal Funds Rate hit the zero lower bound.³

The paper proceeds as follows. Section 2 describes the data we use in our analysis. Section 3 presents our main empirical results regarding the response nominal and real interest rates and TIPS break-even inflation to monetary policy shocks. Section 4 shows what structural parameters our empirical evidence provides information on in the context of a textbook New Keynesian model and quantitatively assesses the degree of monetary non-neutrality implied by our empirical evidence by estimating the CEE/ACEL model using simulated method of moments. This section also presents alternative estimates based on Gertler and Karadi’s (2015) hybrid high-frequency-VAR estimation approach. Finally, section 5 modifies our estimation strategy to allow for a Fed information effect. Section 6 concludes.

2 Data

To construct our measure of monetary shocks, we use tick-by-tick data on Federal Funds futures and Eurodollar futures from the CME Group (owner of the Chicago Board of Trade and Chicago

³Also, Beechey and Wright (2009) analyze the effect of unexpected movements in the near-term Federal Funds rate at the time of FOMC announcements on nominal and real 5-year and 10-year yields and the five-to-ten year forward for the sample period February 17th 2004 to June 13th 2008. Their results are similar to ours for the 5-year and 10-year yields.

Mercantile Exchange). These data can be used to estimate changes in expectations about the Federal Funds Rate at different horizons after an FOMC announcement (see appendix A). The tick-by-tick data we have for Federal Funds futures and Eurodollar futures is for the sample period 1995-2012. For the period since 2012 we use data on changes in the prices of the same five interest rate futures over the same 30-minute windows around FOMC announcements that was graciously shared with us by Refet Gurkaynak.

We obtain the dates and times of FOMC meetings up to 2004 from the appendix to Gurkaynak, Sack, and Swanson (2005). We obtain the dates of the remaining FOMC meetings from the Federal Reserve Board website at <http://www.federalreserve.gov/monetarypolicy/fomccalendars.htm>. For the latter period, we verified the exact times of the FOMC announcements using the first news article about the FOMC announcement on *Bloomberg*. We cross-referenced these dates and times with data we obtained from Refet Gurkaynak and in a few cases used the time stamp from his database.

To measure the effects of our monetary shocks on interest rates, we also use several other daily interest rate series. To measure movements in Treasuries at horizons of 1 year or more, we use daily data on zero-coupon nominal treasury yields and instantaneous forward rates constructed by Gurkaynak, Sack, and Swanson (2007). These data are available on the Fed’s website at <http://www.federalreserve.gov/pubs/feds/2006/200628/200628abs.html>. We also use the yields on 3M and 6M Treasury bills. We retrieve these from the Federal Reserve Board’s H.15 data release.

To measure movements in real interest rates, we use zero-coupon yields and instantaneous forward rates constructed by Gurkaynak, Sack, and Wright (2010) using data from the TIPS market. These data are available on the Fed’s website at <http://www.federalreserve.gov/pubs/feds/2008/200805/200805abs.html>. TIPS are “inflation protected” because the coupon and principal payments are multiplied by the ratio of the reference CPI on the date of maturity to the reference CPI on the date of issue.⁴ The reference CPI for a given month is a moving average of the CPI two and three months prior to that month, to allow for the fact that the Bureau of Labor Statistics publishes these data with a lag.

TIPS were first issued in 1997 and were initially sold at maturities of 5, 10 and 30 years, but only the 10-year bonds have been issued systematically throughout the sample period. Other maturities have been issued more sporadically. While liquidity in the TIPS market was initially poor, TIPS now represent a substantial fraction of outstanding Treasury securities. We start our analysis in

⁴This holds unless cumulative inflation is negative, in which case no adjustment is made for the principle payment.

2000 to avoid relying on data from the period when TIPS liquidity was limited. For 2- and 3-year yields and forwards we start our analysis in 2004. Gurkaynak, Sack, and Wright (2010) only report zero-coupon yields for these maturities from 2004 onward. One reason is that to accurately estimate zero-coupon yields at this maturity it is necessary to wait until longer maturity TIPS issued several years earlier have maturities in this range. To facilitate direct comparisons between nominal and real interest rates, we restrict our sample period for the corresponding 2- and 3-year nominal yields and forwards to the same time period.

We use a daily decomposition of nominal and real interest rate movements into risk-neutral expected future rates and risk premia obtained from Abrahams, Adrian, Crump, and Moench (2015). We also use a daily decomposition of the nominal term structure into risk-neutral expected future rates and risk premia based on the model of Kim and Wright (2005) that is available on the Fed’s website at <http://www.federalreserve.gov/pubs/feds/2005/200533/200533abs.html>. We use data on expectations of future nominal interest rates, inflation and output growth from the *Blue Chip Economic Indicators*. *Blue Chip* carries out a survey during the first few days of every month soliciting forecasts of these variables for up to the next 8 quarters. We use monthly data on industrial production, the consumer price index, one-year nominal Treasury yields, the Federal Funds Rate, and the Gilchrist and Zakrajsek (2012) excess bond premium, as well as high frequency data on the change in the expected Federal Funds rate three months 3-month ahead around the time of FOMC announcements obtained from Mark Gertler and Peter Karadi. We use data on inflation swaps from *Bloomberg*. Finally, we use data on the level of the S&P500 stock price index obtained from *Yahoo Finance*.

3 Empirical Analysis

Our goal in this section is to identify the effect of the monetary policy news contained in scheduled FOMC announcements on nominal and real interest rates of different maturities. Specifically, we estimate

$$\Delta s_t = \alpha + \gamma \Delta i_t + \epsilon_t, \tag{1}$$

where Δs_t is the change in an outcome variable of interest (e.g., the yield on a five year zero-coupon Treasury bond), Δi_t is a measure of the monetary policy news revealed in the FOMC announcement, ϵ_t is an error term, and α and γ are parameters. The parameter of interest is γ , which measures the effect of the FOMC announcement on Δs_t relative to its effect on the policy indicator Δi_t .

To identify a pure monetary policy shock, we consider the change in our policy indicator (Δi_t) in a 30-minute window around scheduled FOMC announcements.⁵ The idea is that changes in the policy indicator in these 30-minute windows are dominated by the information about future monetary policy contained in the FOMC announcement. Under the assumption that this is true, we can simply estimate equation (1) by ordinary least squares. We also present results for a heteroskedasticity based estimation approach (Rigobon, 2003; Rigobon and Sack, 2004) which is based on a weaker identifying assumption to verify that our baseline identifying assumption is reasonable. In our baseline analysis, we focus on only scheduled FOMC announcements, since unscheduled meetings may occur in reaction to other contemporaneous shocks.

The policy indicator we use is a composite measure of changes in interest rates at different maturities spanning the first year of the term structure. Until recently, most authors used unanticipated changes in the Fed Funds Rate (or closely related changes in very short term interest rates) as their policy indicator. The key advantage of our measure is that it captures the effects of “forward guidance.” Forward guidance refers to announcements by the Fed that convey information about future changes in the Federal Funds Rate. Over the past 15 years, the Federal Reserve has made greater and greater use of such forward guidance. In fact, changes in the Federal Funds Rate have often been largely anticipated by markets once they occur. Gurkaynak, Sack, and Swanson (2005) convincingly argue that unanticipated changes in the Fed Funds Rate capture only a small fraction of the monetary policy news associated with FOMC announcements in recent years (see also, Campbell et al., 2012).

The specific composite measure we use as our policy indicator is the first principle component of the unanticipated change over the 30-minute windows discussed above in the following five interest rates: the Federal Funds rate immediately following the FOMC meeting, the expected Federal Funds rate immediately following the next FOMC meeting, and expected 3-month Eurodollar interest rates at horizons of two, three and four quarters. We refer to this policy indicator as the “policy news shock.” We use data on Fed Funds futures and Eurodollar futures to measure changes in market expectations about future interest rates at the time of FOMC announcements. The scale of the policy news shock is arbitrary. For convenience, we rescale it such that its effect on the 1-year nominal Treasury yield is equal to one. Appendix A provides details about the construction of the policy news shock.⁶

⁵Specifically, we calculate the monetary shock using a 30-minute window from 10 minutes before the FOMC announcement to 20 minutes after it.

⁶Our policy news shock variable is closely related to the “path factor” considered by Gurkaynak, Sack, and Swanson

3.1 Baseline Estimates

Table 1 presents our baseline estimates of monetary shocks on nominal and real interest rates and break-even inflation. Each estimate in the table comes from a separate OLS regression of the form discussed above (equation (1)). In each case the independent variable is the policy news shock measured over a 30-minute window around an FOMC announcement, while the change in the dependent variable is measured over a one-day window.⁷

The first column of Table 1 presents the effects of the policy news shock on nominal Treasury yields and forwards. Recall that the policy news shock is scaled such that the effect on the one-year Treasury yield is 100 basis points. Looking across different maturities, we see that the effect of the shock is somewhat smaller for shorter maturities, peaks at 110 basis points for the 2-year yield and then declines monotonically to 38 basis points for the 10-year yield. Since longer-term yields reflect expectations about the average short-term interest rate over the life of the long bond, it is easier to interpret the time-path of the response of instantaneous forward rates. Abstracting from risk premia, the effect on the 2-year instantaneous forward rate (say) is the effect on the expected short-term interest rate that the market expects to prevail in 2 years time. The impact of our policy news shock on forward rates is also monotonically declining in maturity from 114 basis points at 2-years to -8 basis points at 10-years. We show below that the negative effect on the 10-year nominal forward rate reflects a decline in break-even inflation at long horizons.

The second column of Table 1 presents the effects of the policy news shock on real interest rates measured using TIPS. While the policy news shock affects nominal rates by construction, this is not the case for real interest rates. In neoclassical models of the economy, the Fed controls the nominal interest rate but has no impact on real interest rates. We estimate the impact of our policy news shock on the 2-year real yield to be 106 basis points, and the impact on the 3-year real yield to be 102 basis points. Again, the time-path of effects is easier to interpret by viewing estimates for instantaneous forward rates. The effect of the shock on the 2-year real forward rate is 99 basis points. It falls monotonically at longer horizons to 88 basis points at 3 years, 47 basis points at 5 years, and 12 basis point at 10 years (which is not statistically significantly different from zero).

(2005). The five interest rate futures that we use to construct our policy news shock are the same five futures as Gurkaynak, Sack, and Swanson (2005) use. They motivate the choice of these particular futures by liquidity considerations. They advocate the use of two principle components to characterize the monetary policy news at the time of FOMC announcements—a “target factor” and a “path factor.” We focus on a single factor for simplicity. See also Barakchian and Crowe (2010).

⁷The longer window for the dependent variable adds noise to the regression without biasing the coefficient of interest.

Evidently, monetary policy shocks can affect real interest rates for substantial amounts of time. However, in the long-run, the effect of monetary policy shocks on real interest rates is zero as theory would predict.

The third column of Table 1 presents the effect of the policy news shock on break-even inflation as measured by the difference between nominal Treasury rates and TIPS rates. The first several rows provide estimates based on bond yields, which indicate that the response of break-even inflation is small. The shorter horizon estimates are actually slightly positive but then become negative at longer horizons. None of these estimates are statistically significantly different from zero. Again, it is helpful to consider instantaneous forward break-even inflation rates to get estimates of break-even inflation at points in time in the future. The response of break-even inflation implied by the 2 year forwards is slightly positive, though statistically insignificant. The response is negative at longer horizons: for maturities of 3, 5 and 10 years, the effect is -6, -21 and -20 basis points, respectively. It is only the responses at 5 and 10 years that are statistically significantly different from zero. Our evidence thus points to break-even inflation responding modestly and quite gradually to monetary shocks that have a substantial effect on real interest rates.

Table 1 presents results for a sample period from January 1st 2000 to March 19th 2014, except that we drop the period spanning the height of the financial crisis in the second half of 2008 and the first half of 2009.⁸ We choose to drop the height of the financial crisis because numerous well-documented asset pricing anomalies arose during this crisis period, and we wish to avoid the concern that our results are driven by these anomalies. However, similar results obtain for the full sample including the crisis, as well as a more restrictive data sample ending in 2007, and for a sample that also includes unscheduled FOMC meetings (see Table A.1).

3.2 Background Noise in Interest Rates

A concern regarding the estimation approach we describe above is that other non-monetary news might affect our monetary policy indicator during the window we consider around FOMC announcements. If this is the case, it will contaminate our measure of monetary shocks. This concern looms much larger if one considers longer event windows than our baseline 30-minute window. It has been common in the literature on high frequency identification of monetary policy to consider a one- or two-day window around FOMC announcements (e.g., Kuttner, 2001; Cochrane and Piazzesi, 2002;

⁸The sample period for 2- and 3-year yields and forwards is somewhat shorter (it starts in 2004) because of data limitations (see section 2 for details).

Hanson and Stein, 2015). In these cases, the identifying assumption being made is that no other shocks affect the policy indicator during these one or two days. Especially when the policy indicator is based on interest rates several quarters or years into the term structure—as has recently become common to capture the effects of forward guidance—the assumption that no other shocks affect interest rates over one or two days is a strong assumption. Interest rates of those maturities fluctuate substantially on non-FOMC days, suggesting that other shocks than FOMC announcements affect these interest rates on FOMC days. There is no way of knowing whether these other shocks are monetary shocks or non-monetary shocks.

To assess the severity of this problem, Table 2 compares estimates of equation (1) based on OLS regressions to estimates based on a heteroskedasticity-based estimation approach developed by Rigobon (2003) and Rigobon and Sack (2004). We do this both for a 30-minute window and for a 1-day window. The heteroskedasticity-based estimator is described in detail in Appendix B. It allows for “background” noise in interest rates arising from other shocks during the event windows being considered. The idea is to compare movements in interest rates during event windows around FOMC announcements to other equally long and otherwise similar event windows that do not contain an FOMC announcement. The identifying assumption is that the variance of monetary shocks increases at the time of FOMC announcements, while the variance of other shocks (the background noise) is unchanged.

The top panel of Table 2 compares estimates based on OLS to those based on the heteroskedasticity-based estimator (Rigobon estimator) for a subset of the assets we consider in Table 1 when the event window is 30-minutes as in our baseline analysis. The difference between the two estimators is very small, both for the point estimates and the confidence intervals.⁹ This result indicates that there is in fact very little background noise in interest rates over a 30-minute window around FOMC announcements. In this case, the OLS identifying assumption—that only monetary shocks occur within the 30-minute window—yields a point estimate and confidence intervals that are close to correct. Table A.2 presents a full set of results based on the Rigobon estimator and a 30-minute window. This table confirms that OLS yields very similar results to the Rigobon estimator for all the assets we consider.

⁹The confidence intervals for the Rigobon estimator in Table 2 are constructed using a procedure that is robust to inference problems that arise when the amount of background noise is large enough that there is a significant probability that the difference in the variance of the policy indicator between the sample of FOMC announcements and the “control” sample is close to zero. In this case, the conventional bootstrap approach to constructing confidence intervals will yield inaccurate results. Appendix C describes the method we use to construct confidence intervals in detail. We thank Sophocles Mavroeidis for suggesting this approach to us.

In contrast, the problem of background noise is quite important when the event window being used to construct our policy news shocks is one day. The second panel of Table 2 compares estimates based on OLS to those based on the Rigobon estimator for policy news shocks constructed using a one-day window. In this case, the difference between the OLS and Rigobon estimates are substantial. The point estimates in some cases differ by dozens of basis points and have different signs in three of the six cases considered. However, the most striking difference is that the confidence intervals that OLS yields are much narrower than those generated using the Rigobon method. According to OLS, the effects on the 5-year nominal and real forwards are highly statistically significant, while the Rigobon estimator indicates that these effects are far from being significant. Clearly, the approach of using OLS with a 1-day window massively overstates the true statistical precision of the estimates. This indicates that there is a large amount of background noise in the interest rates used to construct the policy news shock over a one day window.

These differences are even larger when a longer-term interest rate is used as the policy indicator that proxies for the size of monetary shocks. The third panel of Table 2 compares results based on OLS to those based on the Rigobon estimator when the policy indicator is the change in the two-year nominal yield over a one day window. Again, the confidence intervals are much wider using the Rigobon estimator than OLS. In fact, here we report 90% confidence intervals for the Rigobon estimator since the 95% confidence intervals are in some cases infinite (i.e., we were unable to find any value of the parameter of interest that could be rejected at that significance level).

An important substantive difference arises between the OLS and Rigobon estimates in the case of the 10-year real forward rate when the 2-year nominal yield is used as the policy indicator. Here, OLS estimation yields a statistically significant effect of the monetary shock on forward rates at even a 10-year horizon. This result is emphasized by Hanson and Stein (2015). However, the Rigobon estimator with appropriately constructed confidence intervals reveals that this result is statistically insignificant. Our baseline estimation approach using a 30-minute window and the policy news shock as the proxy for monetary shocks yields a point estimate that is small and statistically insignificant.¹⁰

¹⁰Hanson and Stein (2015) also present an estimator based on instrumenting the 2-day change in the 2-year rate with the change in the two-year rate during a 60-minute window around the FOMC announcement. This yields similar results. Since this procedure is not subject to the concerns raised above, it suggests that there are other sources of difference between our results and those of Hanson and Stein than econometric issues. One possible source of difference is that we use different monetary shock indicators. Their policy indicator (the change in the 2-year yield) is further out in the term structure and may be more sensitive to risk premia. As we discuss in section 3.3, our measure of monetary shocks is uncorrelated with the risk premia implied by the affine term structure model of Abrahams et al. (2015), whereas Hanson and Stein’s monetary shocks are associated with substantial movements in risk premia. The difference could also arise from the fact that Hanson and Stein focus on a 2-day change in long-term real forwards; which could yield different results if the response of long-term bonds to monetary shocks is inertial.

3.3 Risk Premium or Expected Future Short-Term Rates?

An important question when interpreting our results is to what extent the movements in long-term interest rates we identify reflect movements in risk premia as opposed to changes in expected future short-term interest rates. A great deal of evidence suggests that changes in risk premia do play an important role in driving movements in long-term interest rates. Yet, for our analysis, the key question is not whether risk premia matter in general, but rather how important they are in explaining the abrupt changes in interest rates that occur in the narrow windows around FOMC announcements that we focus on.¹¹

We present three sets of results that indicate that risk premium effects are not driving our empirical results: 1) the impact of our policy news shock on direct measures of expectations from the *Blue Chip Economic Indicators*; 2) the impact of our policy news shock on risk-neutral expected short rates from a state-of-the-art affine term structure model; and 3) the impact of our policy news shock on interest rates over longer event windows than in our baseline results. Section 4.3 presents additional evidence based on Gertler and Karadi’s hybrid high-frequency-VAR approach.

Let us begin with our analysis of the *Blue Chip* forecast data. *Blue Chip* surveys professional forecasters on their beliefs about macroeconomic variables over the next two years in the first few days of every month. From this survey, it is possible to obtain direct measures of expectations that are not contaminated by risk premium effects. We use expectations about future values of the 3-month T-Bill rate as our measure of short-term nominal interest rate expectations and expectations about changes in the GDP deflator as our measure of expectations about inflation (and the difference between the two as our measure of expectations about short-term real rates).

We estimate the impact of monetary shocks on expectations by running regressions of the change from one month to the next in expectations regarding a particular forecast horizon on any policy news shock that occurs over the month except for those that occur in the first week (because we do not know whether these occurred before or after the survey response). Unfortunately, *Blue Chip* asks respondents only about the current and subsequent calendar year on a monthly basis, so fewer observations are available for longer-term expectations, leading to larger standard errors.¹² The

¹¹Piazzesi and Swanson (2008) show that Fed Funds futures have excess returns over the Federal Funds rate and that these excess returns vary counter-cyclically at business cycle frequencies. However, they argue that high frequency changes in Fed Funds futures are likely to be valid measures of changes in expectations about future Federal Funds rates since they difference out risk premia that vary primarily at lower frequencies.

¹²For example, towards the end of each year, forecasters are only asked about their beliefs a little more than 1-year in advance; while in the first quarter they are asked about their beliefs for almost the next full 2-years. *Blue Chip* also asks for longer-term inflation forecasts, but only twice a year (March and October) implying that there is too much noise for our event study analysis.

sample period for this analysis is January 1995 to January 2012, except that we exclude the apex of the 2008-2009 financial crisis as we do in our baseline analysis.

Table 3 presents the results from this analysis. The table shows that the policy news shock has a persistent impact on expected short-term interest rates, both nominal and real. The interest rate effects are somewhat larger than in our baseline results, but rather noisily estimated. The effect on expected inflation is small and statistically insignificant at all horizons except that it is marginally significantly negative at 2 quarters. The much larger standard errors in Table 3 arise from the fact that the *Blue Chip* variables are available only at a monthly as opposed to a daily frequency. Overall, these estimates appear consistent with our baseline findings that monetary shocks have large effects on expected short-term nominal and real rates.

Our second approach is to regress estimates of changes in expected future short rates from a state-of-the-art affine term structure model on our monetary policy shocks. Abrahams et al. (2015) employ an affine term structure model to decompose changes in both nominal and real interest rates at different maturities into changes in risk-neutral expected future short rates and changes in risk premia.¹³ Table 4 presents results based on their decomposition. The response of model-implied risk-neutral interest rates to our policy news shock is very similar to the response of raw interest rates in our baseline results. This piece of evidence, thus, points to our monetary shocks having large effects on future short-term nominal and real rates and small effects on expected inflation (even smaller than in our baseline results).

It is important to stress that the Abrahams et al. (2015) model by no means rules out the potential importance of risk premium effects. In fact, risk premia for long-term bonds are large and volatile in this model. Moreover, Abrahams et al. (2015) estimate large effects of monetary shocks on risk premia for Hanson and Stein’s (2015) measure of monetary shocks (the 2-day change in the 2-year nominal yield around FOMC announcements). Our analysis shows that the effect of monetary shocks on long-term real interest rates (and long-term real term premia) is much smaller for our monetary shocks measure. This is consistent with our finding in section 3.2 that the 2-day window includes too much “background noise” for an OLS regression to accurately assess the effects of monetary shocks. Our results suggest that, while risk premia explain a large fraction of the variation in long-term interest rates in general, the movements in interest rates at the time of FOMC announcements are largely due to changes in expected future short rates.

¹³What we refer to as the risk premia here is the difference between raw interest rate changes and changes in model-implied risk neutral interest rates. Abrahams et al. (2015) further decompose this difference into a term premium, a liquidity premium, and a model error term.

Our third approach to gauging the role of risk premia in our results is to consider longer event windows for the outcome variables of interest. If the effects we estimate are due to transient liquidity premia (as Hanson and Stein (2015) hypothesize), we should see smaller effects over longer event windows as the temporary liquidity premia reverse themselves. Table 5 presents the effects of our policy news shock on nominal and real interest rates over event windows of 1, 5, 10, 20, 60, 125, and 250 trading days.¹⁴ While the estimates become very noisy as the event window becomes larger, there is little evidence that the effects on interest rates tend to dissipate over time. Indeed, in most cases, the point estimates appear to grow over time (though, again, the standard errors are extremely large).

We also consider an alternative market-based measure of inflation expectations based on inflation swap data.¹⁵ The sample period for this analysis is limited by the availability of swaps data to begin in January 1st 2005. Unfortunately, due to the short sample available to us, the results are extremely noisy, and are therefore not particularly informative. As in our baseline analysis, there is no evidence of large negative responses in inflation to our policy news shock (as would arise in a model with flexible prices). Indeed the estimates from this approach (which are compared to our baseline results in Table A.3) suggest a somewhat larger “price puzzle”—i.e., positive inflation response—at shorter horizons, though this is statistically insignificant.

4 Evidence on Monetary Non-Neutrality

How much monetary non-neutrality does our high frequency evidence on interest rates imply? In this section, we answer this question through the lens of a workhorse New Keynesian model. We estimate key parameters of such a model to match the high frequency evidence from section 3. This follows in the tradition of work by Rotemberg and Woodford (1997), Christiano, Eichenbaum, and Evans (2005) and others that estimates the parameters of models of this kind to match responses of output, inflation, and other variables to monetary shocks identified using structural vector autoregressions (VARs). For comparison, we also estimate the same workhorse New Keynesian model to match the response from a hybrid high-frequency VAR identification approach developed by Gertler and

¹⁴In all cases, the policy news shock is measured over a 30-minute event window. We only vary the length of the event window for the dependent variables.

¹⁵An inflation swap is a financial instrument designed to help investors hedge inflation risk. As is standard for swaps, nothing is exchanged when an inflation swap is first executed. However, at the maturity date of the swap, the counterparties exchange $R_t^x - \Pi_t$, where R_t^x is the x -year inflation swap rate and Π_t is the reference inflation over that period. If agents were risk neutral, therefore, R_t would be expected inflation over the x year period. See Fleckenstein, Longstaff, and Lustig (2014) for an analysis of the differences between break-even inflation from TIPS and inflation swaps.

Karadi (2015). But before moving to the full workhorse New Keynesian model, it is useful to develop intuition regarding what structural parameters our evidence helps identify using a textbook, three-equation, New Keynesian model.

4.1 Intuition in a Textbook New Keynesian Model

4.1.1 The Behavior of Households and Firms

Consider a setting in which the behavior of households and firms can be described by the following Euler equation and Phillips curve:

$$\hat{x}_t = E_t \hat{x}_{t+1} - \sigma(\hat{i}_t - E_t \hat{\pi}_{t+1} - \hat{r}_t^n), \quad (2)$$

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa \zeta \hat{x}_t. \quad (3)$$

Hatted variables denote percentage deviations from steady state. The variable $\hat{x} = \hat{y}_t - \hat{y}_t^n$ denotes the “output gap”—the difference between actual output $\hat{y}_t$ and the “natural” level of output $\hat{y}_t^n$ that would prevail if prices were flexible, $\hat{\pi}_t$ denotes inflation, $\hat{i}_t$ denotes the gross return on a one-period, risk-free, nominal bond, and $\hat{r}_t^n$ denotes the “natural rate of interest.” Both the natural rate of output and the natural rate of interest are functions of exogenous shocks to tastes and technology. Appendix D presents a detailed derivation of these equations from primitive assumptions about tastes and technology. Woodford (2003) and Galí (2008) present textbook treatments.

The Euler equation (2) is common to both Real Business Cycle and New Keynesian models, and describes how household consumption responds to movements in real interest rates. The parameter σ in the Euler equation denotes the intertemporal elasticity of substitution. The Phillips curve is fundamental to the New Keynesian paradigm. It describes how inflation responds to deviations of output from the natural rate of output. We have split the slope of the Phillips curve into two parameters κ and ζ to emphasize that sluggish price adjustment in the model arises from the combination of two forces: nominal rigidity—i.e., infrequent prices changes—and coordination failure among price setters often referred to as “real rigidity”—i.e., the fact that firms respond incompletely to shocks even when they do change their prices because other firms have yet to respond.

4.1.2 Monetary Policy

To build intuition, we assume that the monetary authority sets interest rates according to the following simple rule:

$$\hat{i}_t - E_t \hat{\pi}_{t+1} = \bar{r}_t + \phi_\pi \hat{\pi}_t, \quad (4)$$

where $\bar{r}_t = r_t^n + \epsilon_t$. Here, the monetary authority varies the real interest rate in such a way as to track the natural real rate of interest r_t^n . However, it does this with some error ϵ_t . The monetary authority also systematically varies the real interest rate with the rate of inflation in line with the well known Taylor principle.

4.1.3 Intuition for Identification

In this simple model, it is straightforward to show how our evidence on the response of the real interest rate and expected inflation to monetary shocks identifies key parameters relating to the extent of monetary non-neutrality. Assuming that monetary shocks have no effect on output in the long run, we can solve the Euler equation—equation (2)—forward and get that the response of the output gap to a monetary shock is,

$$\hat{x}_t = -\sigma \sum_{j=0}^{\infty} E_t \hat{r}_{t+j} = -\sigma \hat{r}_t^\ell. \quad (5)$$

where $\hat{r}_{t+j}$ denotes the response of the short-term real interest rate at time $t + j$ —i.e., $\hat{r}_{t+j} = \hat{r}_{t+j} - E_{t+j} \hat{r}_{t+j+1}$ —and $\hat{r}_t^\ell$ denotes the response of the long-run real interest rate.

Similarly, we can solve forward the Phillips curve—equation (3)—and get that the response of inflation to a monetary shock is

$$\hat{\pi}_t = \kappa \zeta \sum_{j=0}^{\infty} \beta^j E_t \hat{x}_{t+j}. \quad (6)$$

Combining equations (5) and (6), we get a relationship between the response of inflation and the real interest rates:

$$\hat{\pi}_t = -\kappa \zeta \sigma \sum_{j=0}^{\infty} \beta^j E_t \hat{r}_{t+j}^\ell. \quad (7)$$

We wish to draw two main conclusions from equation (8). First, the relative size of the response of inflation and real interest rates to a monetary shock pins down $\kappa \zeta \sigma$. In section 3, we estimate the response of expected inflation and real interest rates. Our evidence thus sheds light on $\kappa \zeta \sigma$. A small response of expected inflation relative to the magnitude of the real interest rate response implies a small value of $\kappa \zeta \sigma$ —i.e., a large amount of nominal and real rigidities, a small value of the intertemporal elasticity of substitution, or both.

Second, the dynamics of the response of expected inflation to a monetary shock are informative about the degree of inflation inertia in the economy. Equation (8) shows clearly that (almost) irrespective of the values of the parameters of the model, inflation should fall more in the short run

than in the long run in response to a positive shock to real interest rates (since positive real interest rate terms “fall out” of the infinite sum on the right hand side of equation (8) as time passes).¹⁶

The implications of the textbook model for inflation persistence are illustrated in Figure 1 for particular values of the structural parameters. Figure 2 presents our estimated response of TIPS break-even inflation and nominal and real interest rates in the form of a figure for ease of comparison with the results from the model. In sharp contrast with the predictions of equation (8) the inflation response we estimate in the data is initially small but builds over time. Our estimated responses, thus, point towards substantial inflation inertia in the economy that the simple model described above cannot capture. Such inflation inertia is incorporated into the workhorse business cycle model that we consider below.

An important question is whether this argument continues to hold even if the monetary shock leads to a shift in the long-run inflation target of the central bank (and therefore the long-run inflation rate). In this case, equation (7) becomes

$$\hat{\pi}_t = -\kappa\zeta\sigma \sum_{j=0}^{\infty} \beta^j E_t \hat{r}_{t+j}^{\ell} + \hat{\pi}_{\infty}, \quad (8)$$

where $\hat{\pi}_{\infty}$ denotes the change in the long-run inflation rate. The simple monetary policy we consider above implies that $\pi_{\infty} = 0$. Even if this term is non-zero, however, it is important to recognize that it affects inflation in *every* period after the shock. Hence, it would not change the slope of the response of expected inflation. The extra term *does* have the potential to lead to a larger response of inflation to a monetary shock than in our baseline model. Empirically, however, the response of expected inflation to the monetary shock already appears to be very low. Adding this feature to the model would further increase the degree of rigidities we estimate in the data, and therefore, the degree of monetary non-neutrality. In what follows, we assume monetary policy rules for which $\pi_{\infty} = 0$.

4.2 Estimating the CEE/ACEL Model with High Frequency Data

We build on these intuitions to estimate the workhorse medium-scale business cycle model proposed by Christiano, Eichenbaum, and Evans (2005, henceforth CEE) and further developed by Altig et al. (2011 henceforth ACEL). Relative to the simple model above, this model incorporates additional

¹⁶ The exception to this is if the persistence of the monetary policy shock is sufficiently high (more persistent than β). In this case, the fact that the terms further out in the sum are getting closer to the present as time passes will lead the response of inflation to grow over time. Our estimated policy news shock is far less persistent than it would need to be to generate this effect.

features that have been shown to be important in generating realistic business cycles, and is therefore more suitable for quantitative analysis.

CEE and ACEL present detailed descriptions of their model. We refrain from repeating this material here. Rather, we only discuss the elements of the model that are most relevant for our analysis. ACEL develop a version of this model in which capital is firm specific. They show that this version of the model is equivalent to the homogeneous capital version of the model analyzed in CEE up to a linear approximation (though with different parameter interpretations, as we discuss below). We therefore refer to this model as the CEE/ACEL model.

4.2.1 Estimation Approach

We estimate the model by simulated method of moments. The moments we use in our estimation are the responses of 2, 3, 5, and 10-year nominal and real yields and the responses of 2, 3, 5, and 10-year instantaneous nominal and real forward rates to our policy news shock. We minimize the sum of the squared difference between the moments in the data and the model. So as not to have to estimate the size of the shock, we scale the responses from the model in such a way that they perfectly match the response of the 3Y real forward rate.

We construct confidence intervals by bootstrapping. Our bootstrap procedure is to re-sample the data with replacement, estimate the empirical moments on the re-sampled data, and then estimate the structural parameters using a loss function based on the estimated empirical moments for the re-sampled data. We repeat this procedure 500 times and report the 2.5% and 97.5% quantiles of the statistics of interest. Importantly, this procedure for constructing the confidence intervals captures the statistical uncertainty associated with our empirical estimates in Table 1.

We estimate five structural parameters of the model. Two of these describe the dynamics of the monetary shock; two relate to the response of inflation to output; and one relates to the response of output to the real interest rate. We fix all other parameters equal to their estimated values in CEE. The primary reason that we do not estimate a larger set of parameters is that our empirical evidence provides us with information about certain aspects of the CEE/ACEL model—namely the response interest rates and inflation to a monetary shock—but not all aspects.

CEE show that the linearized first-order condition for investment in their model may be solved forward to yield

$$\hat{\lambda}_t = \hat{\lambda}_{t-1} + \frac{1}{k_I} \sum_{j=0}^{\infty} \beta^j E_{t-1} \hat{p}_{k,t+j}, \quad (9)$$

where $\hat{\lambda}_t$ denotes investment and $\hat{p}_{k,t}$ is the shadow value of a unit of installed capital. From equation (9), we see that $1/k_I$ is the elasticity of investment with respect to a 1 percent temporary increase in the current price of installed capital. The parameter k_I , thus plays a key role in determining the response of output to changes in real interest rates in the CEE/ACEL model. We estimate k_I .

The two key parameters governing the response of inflation to variation in output in the homogeneous capital version of the CEE/ACEL model are ξ_p and ξ_w . These parameters govern the frequency of price change and the frequency of wage change. Specifically, the frequency of price change is $1 - \xi_p$ and the frequency of wage change is $1 - \xi_w$. We estimate these two parameters. ACEL show that the homogeneous capital version of the model with a particular value for ξ_p yields the same aggregate dynamics as the firm-specific capital version of the model with a much lower value of ξ_p . The reason for this is that firm-specific capital is a powerful source of real rigidity that dramatically lowers the slope of the price Phillips curve in the model for any given values of ξ_p .

The baseline specification of monetary policy in CEE/ACEL is a rule for money growth. However, CEE show that their model behaves very similarly if they instead specify monetary policy as the following rule for the nominal interest rate

$$i_t = \rho i_{t-1} + (1 - \rho)(\phi_\pi E_{t-1} \Pi_{t+1} + \phi_y y_t) + \bar{i}_t, \quad (10)$$

where $\rho = 0.8$, $\phi_\pi = 1.5$, $\phi_y = 0.1$, and $\bar{i}_t$ is i.i.d. We specify monetary policy by this rule and follow CEE in setting $\phi_\pi = 1.5$ and $\phi_y = 0.1$. However, to capture the “hump-shaped” pattern of real interest rates we identify following the monetary shocks, we allow $\bar{i}_t$ to follow the AR(1) process $\bar{i}_t = \nu \bar{i}_{t-1} + \epsilon_t$ and estimate both ρ and ν . As we show in Figure 3, this specification of monetary policy is able to capture very well the nature of the monetary shocks we estimate in the data, which combine a small contemporaneous shocks with a much larger, highly persistent forward guidance shock.

CEE/ACEL assume that firms that do not have an opportunity to reoptimize their prices index their prices to past inflation. Likewise, CEE/ACEL assume that unions that do not have an opportunity to reoptimize their wages index their wages to past wage inflation. CEE/ACEL, thus, build into their model the high degree of price and wage inflation inertia that we argue above is essential in fitting the delayed response of inflation to monetary shocks we estimate in section 3.

4.2.2 Estimates of Monetary Non-Neutrality

Our primary interest is the extent of monetary non-neutrality implied by our high frequency evidence. We measure the degree of monetary non-neutrality as the ratio of the cumulative impulse

response of output to the cumulative impulse response of inflation after a monetary shock.¹⁷ Intuitively, we are defining monetary non-neutrality as the size of the output response to a monetary shock relative to the response of the price level. If monetary shocks lead to large movements in the price level relative to output, our measure of monetary non-neutrality is small. If however, output moves a great deal relative to the price level in response to monetary shock, our measure is large.

Panel A of Table 6 reports this measure of monetary non-neutrality for our baseline estimation of the CEE/ACEL model as well as for CEE and ACEL’s original estimates. Our baseline estimates imply that output responds 3.8 times as much as inflation to the monetary shock. This degree of monetary non-neutrality is similar to that implied by ACEL’s estimates (3.3), but substantially larger than that implied by CEE’s estimates (1.7). Evidently, our estimates imply a very substantial amount of monetary non-neutrality, comparable to ACEL, but statistically significantly more than CEE.

Figure 3 presents the response of nominal and real interest rates and inflation to our monetary policy shock in the model with the parameter estimates we get from our structural estimation procedure. Comparing these responses to those in Figure 2 and the numbers in Table 1, we see that the model fits the data quite well. The shock to the path for nominal interest rates captures well the shock we estimate in the data: a small contemporaneous response followed by a very persistent, hump-shaped response (i.e., a great deal of forward guidance). The response of inflation is very small initially and then gradually increases. The response of real interest rates is close to identical to the response of nominal interest rates out to about 3 years. At longer horizons, the response of nominal interest rates falls below the response of real interest rates.

Table 7 presents our individual parameter estimates. We estimate $\rho = 0.96$ and $\nu = 0.74$. These estimates allow us match the hump-shaped response of interest rates to the policy news shock. The remaining three parameters are not precisely estimated. This reflects the fact that they all contribute to a sluggish response of prices to movements in real interest rates. They do so in slightly different ways—which is why the model is formally identified—but these differences are not large enough to yield sharp inference for each parameter separately. This is illustrated in Figure 4, which presents a scatter-plot of the joint sampling distribution of ξ_p and ξ_w that we estimate. The figure shows clearly that low values of ξ_w are accompanied by very high values of ξ_p and vice versa. Our

¹⁷More specifically, we calculate $\sum_{j=0}^{500} |\hat{y}_{t+j}|/4$ and $\sum_{j=0}^{500} |\hat{\pi}_{t+j}|$ and take the ratio. For output, we divide by four because a 1% higher level of output in all quarters of a year is equivalent to a 1% higher level of annual output for the year as a whole. For inflation, there is no need to divide by four, because a 1% higher inflation rate in all quarters of a year is equivalent to a 4% higher inflation rate on an annual basis.

results, thus, provide strong evidence for a large amount of nominal and real rigidities, but they provide little guidance on whether the source of these rigidities is wage rigidity or price rigidity.¹⁸ The loss function in our estimation favors a large value of the investment adjustment cost parameter k_I . However, the loss function is very flat for values of k_I larger than 20 and the lower end of the confidence interval for k_I is as small as 0.7. We therefore restrict k_I to be less than 25. A value of $k_I = 25$ implies that a 1% permanent increase in the price of installed capital leads to a 4% increase in investment.¹⁹

4.3 A Hybrid High-Frequency-VAR Identification Approach

We now consider a hybrid high-frequency-VAR (HF-VAR) identification approach developed by Gertler and Karadi (2015, henceforth GK). According to this approach, high-frequency measures of monetary policy surprises are used as “external instruments” in a VAR to identify the contemporaneous effects of monetary policy. This approach allows us to estimate the key parameters of the model while remaining agnostic about the extent to which the interest rate effects we estimate arise from risk premium effects.

The key difference between the HF-VAR approach and our baseline estimation approach is that, the HF-VAR approach uses the high-frequency identification scheme only to identify the contemporaneous effect of monetary shocks. The dynamic effects of the monetary shock—i.e., how the effects of the shock play out over time—are calculated using a VAR. This makes it possible to identify directly the effect of monetary shocks on a wider set of variables. In particular, it can be used to estimate the dynamic effects of monetary shocks on output (for which the contemporaneous effect as measured by the high-frequency identification scheme is estimated to be zero). These estimated responses can then be used to re-estimate the CEE/ACEL model. However, the ability to estimate the response of monetary shocks to more variables comes at the cost of stronger identifying assumptions. The accuracy of this approach relies on the accuracy of the VAR in capturing the dynamics of the key variables.

We replicate the HF-VAR results presented in Figure 1 of GK. The VAR is estimated at a

¹⁸Our finding that nominal and real rigidities are large is in line with direct GMM estimates of the New Keynesian Phillips curve. Mavroeidis, Plagborg-Møller, and Stock (2014) survey this literature and, using a common data set, run a huge number of a priori reasonable specifications which span different choices made in various papers in the literature. They find that values of the slope coefficient in these Phillips curves vary substantially across specifications and are symmetrically dispersed around a value of zero.

¹⁹This lines up well with existing micro-evidence. Using variation in the price of capital associated with tax changes, Cummins, Hassett, and Hubbard (1994) estimate an elasticity of investment with respect to a permanent change in the price of capital of 6.6.

monthly frequency with four series: logarithm of industrial production, the logarithm of the CPI, the one-year nominal treasury yield, and the Gilchrist-Zakrajsek excess bond premium (Gilchrist and Zakrajsek, 2012). The VAR has 12 lags. The one-year nominal treasury yield is the policy indicator and the external instrument used to identify monetary policy shocks is the change in the 3-month ahead Federal Funds future during the 30-minute window around FOMC announcements.²⁰ We use the same sample period as GK: for the high frequency estimation, the sample period is 1991:1-2012:6 but excluding 2008:7-2009:6, while for the estimation of the VAR, it is 1979:7-2012:6. GK provide a detailed description of the exact procedure used (see also Mertens and Ravn (2013) and Stock and Watson (2012)).

We use the resulting impulse responses to re-estimate the CEE/ACEL model by simulated methods of moments. First, we convert the monthly impulse responses into quarterly impulse responses.²¹ Following CEE, we use the first 25 quarters of the impulse responses in our estimation. We weight all moments equally. Following Mertens and Ravn (2013) and GK, we construct confidence intervals using an application of the recursive-design wild bootstrap with 500 iterations (Goncalves and Kilian, 2004).

Panel B of Table 6 presents estimates of our measure of the degree of monetary non-neutrality based on the HF-VAR estimation approach. The first row presents our baseline HF-VAR estimates, and the two subsequent rows present alternative procedures that we discuss below. The baseline HF-VAR estimates use impulse responses for output, inflation, and the 1-year nominal Treasury yield as moments and estimates the same set of structural parameters of the CEE/ACEL model that we estimate in our benchmark high-frequency estimation. For this case, the HF-VAR estimation approach yields a similar estimate of monetary non-neutrality as our baseline estimation (3.4 versus 3.8 for our baseline). These results are also very similar to those for ACEL’s estimates, but somewhat higher than CEE’s.

The estimates of the structural parameters of the CEE/ACEL model for this HF-VAR estimation are presented in the second column of Table 7. Given the results about monetary non-neutrality, it is not surprising that the HF-VAR approach yields similar estimates of the parameters to our baseline approach. In particular, the HF-VAR estimation yields quite similar estimates of the price

²⁰GK consider a variety of possible instruments for the one-year treasury yield and show that the 30-minute change in the 3-month ahead Federal Funds future provides the most powerful instrument among the alternative instruments they consider.

²¹For industrial production, the one-year treasury yield, and the excess bond premium series, we simply take an average over the quarter. For inflation, we calculate the response on impact as the response of the CPI in the second month after the shock and the response in quarter h as the response of the CPI in month $3h + 2$ less the response of the CPI in month $3h - 1$.

and wage rigidity parameters to our baseline estimation. Both parameters are estimated with low precision, reflecting the difficulty of separately identifying price and wage rigidity.

The two differences that do arise are: First, the point estimate of the investment elasticity parameter k_I is smaller using the HF-VAR approach, though it is not estimated with much statistical precision either here or in our baseline case. And, second, the AR(1) parameter of the monetary shock ν is estimated to be smaller in the HF-VAR case, i.e, the HF-VAR approach does not capture the hump-shaped response of interest rates that we estimate in our baseline approach. The well-known downward bias in AR(1) coefficients in small samples may lead the VAR to underestimate the true persistence of the monetary shock; a view that is supported by evidence presented in GK regarding the evolution of Blue Chip expectations about interest rates after these monetary shocks.

An advantage of the HF-VAR approach is that it is possible to estimate the degree of monetary non-neutrality without taking a stand on whether the real interest rate response we observe to a monetary shock arises from changes in expected future short-term interest rates or changes in risk premia. The second last row of Table 6 presents the results of redoing the structural estimation using only the output and inflation impulse responses. This approach allows one to remain agnostic about whether the real interest rate response we observe to a monetary shock arises from changes in expected future short-term interest rates or changes in risk premia because in either scenario, the higher real interest rate reduces output and this feeds through to inflation via the Phillips curve. The second last row of Table 6 presents the results of redoing the structural estimation using only the output and inflation impulse responses. This approach yields almost identical estimates of monetary non-neutrality to our baseline methodology.

One difference between GK's results and our results is that GK argue that monetary shocks have a substantial effect on risk primia. However, GK show that their VAR generates a substantially less persistent reaction of the Federal Funds rate than Blue Chip data suggest. This difference could arise either because the Blue Chip forecasters overreact to monetary shocks or because of misspecification in the VAR. Our results based on the affine term structure model of Abrahams et al. (2015) supports the later interpretation. But Table 6 shows that, in any case, the implications for monetary non-neutrality are the same. An increase in the risk premium reduces output much like an increase in expected real rates. Despite the reduction in output, there is almost no effect on inflation, suggesting a very flat Phillips curve. This yields a large estimate of monetary non-neutrality, consistent with our baseline results.

Finally, the last row of Table 6 present estimates for a case where we estimate the habit parameter

in the utility function in addition to the other five structural parameters. The estimation of this parameter relies on the impulse response for output, for which the habit parameter is important. This is an advantage of the HF-VAR approach, since the output impulse response is not directly estimated in our baseline approach. This approach also yields similar conclusions regarding our three metrics of monetary non-neutrality. This is true even though we estimate a large value of 0.95 for the habit parameter, quite a bit larger than CEE (0.65) and ACEL (0.76).

5 Allowing for a Fed Information Effect

In the analysis above, we have taken the conventional view that FOMC announcements convey information only about future monetary policy. This view may seem reasonable given that the FOMC has access to the same data as the private sector, with minor exceptions.²² However, the Fed does employ a legion of talented economists whose primary role is to process all the information being released about the economy. This may imply that the FOMC has an informational advantage over the private sector when it comes to data processing. Romer and Romer (2000) argue that monetary policy actions by the Fed reveal information to the public that is useful for forecasting inflation and that this informational advantage is due to superior information processing.²³

Here, again, the textbook New Keynesian model is useful for building intuition. The solved-forward Euler equation

$$\hat{x}_t = -\sigma \sum_{j=0}^{\infty} E_t(\hat{i}_{t+j} - \hat{\pi}_{t+j+1} - \hat{r}_{t+j}^n). \quad (11)$$

shows that the output gap is determined by the current and expected future values of the “interest rate gap”—the difference between the real interest rate $\hat{i}_{t+j} - E_{t+j}\hat{\pi}_{t+j+1}$ and the natural rate of interest $\hat{r}_{t+j}^n$. Recall that the natural rate of interest is the real interest rate that would prevail if prices (and wages) were perfectly flexible. In the simple model laid out in appendix D, the natural rate of interest is determined by expected future productivity growth as well as preference shocks. In richer models, other shocks—such as shocks to the financial sector and household borrowing limits—will affect the natural rate of interest.

²²The FOMC has some advance knowledge of industrial production data since the Federal Reserve produces these data. It also collects anecdotal information on current economic conditions from reports submitted by bank directors and through interviews with business contacts, economists, and market experts. This information is subsequently published in reports commonly known as the Beige Book.

²³Faust, Swanson, and Wright (2004) argue that Romer and Romer’s results do not hold up for a more recent sample period and are sensitive to using the unexpected component of the change in the Federal Funds rate as the monetary surprise as opposed to the entire change in the Federal Funds rate.

If the Fed is expected to be able to maintain a zero interest rate gap, the output gap will be zero today and in the future. This, furthermore, implies that inflation will be zero today and in the future—see equation (6). Varying the real interest rate so as to perfectly track the natural rate of interest, therefore, constitutes optimal monetary policy in this simple model. From this perspective, it is natural to think of the Fed’s announcements as potentially conveying information about current and future values of the natural rate of interest.

Table 8 presents evidence that FOMC announcements may in fact convey information about current and future values of the natural rate of interest. The table reports the response of expectations of output growth from *Blue Chip* to our policy news shock. If the policy news shock only conveyed information about future monetary policy, expectations about output growth should fall (since we are looking at an increase in interest rates). In fact, expectations about output growth rise. One way to interpret this evidence is that whenever the FOMC surprises the markets by indicating that it will tighten policy more than the markets thought, the private sector infers that the FOMC is more optimistic about the economy than it had thought and it responds by raising its own expectations about output growth.²⁴

To fit this additional piece of evidence we now abandon the conventional view of monetary shocks, and assume, instead, that FOMC announcements convey information both about future monetary policy and about current and future exogenous shocks such as productivity growth.²⁵ For simplicity, we do this within the context of the textbook model presented in section 4.1 augmented in two ways. First, to be able to capture inflation inertia, we adopt the price setting assumptions of CEE/ACEL. These assumptions give rise to a hybrid Phillips curve which implies that current inflation is influenced by past inflation in addition to deviations of future marginal cost from its natural rate:

$$\hat{\pi}_t = \hat{\pi}_{t-1} + \kappa \sum_{j=0}^{\infty} \beta^j E_t \hat{m}c_{t+j}, \quad (12)$$

where $\hat{m}c_t$ denotes deviations of marginal cost from its natural rate. Second, we allow for external habit formation in consumption. This implies that the output gap is influenced by its past value in addition to future interest rate gaps:

$$\hat{x}_t = b\hat{x}_{t-1} - (1-b)\sigma \sum_{j=0}^{\infty} E_t (\hat{i}_{t+j} - \hat{\pi}_{t+j+1} - \hat{r}_{t+j}^n). \quad (13)$$

²⁴Campbell et al. (2012) present similar evidence regarding the effect of surprise monetary shocks on *Blue Chip* expectations about unemployment.

²⁵This alternative view about the information content of Fed announcements is closely related to the notion of endogenous monetary policy actions in Ellingsen and Soderstrom (2001).

We set the habit parameter $b = 0.65$ —the value estimated by CEE.

To capture the notion that surprise policy tightening by the FOMC leads the private sector to revise its expectation about current and future values of exogenous shocks, we assume that FOMC announcements lead to changes to expectations about current and future values of the natural rate of interest $\Delta E_t \hat{r}_{t+j}^n$ that are proportional to the change in expectations about current and future monetary policy, $\Delta E_t \bar{r}_{t+j}$, i.e., $\Delta E_t \hat{r}_{t+j}^n = \psi \Delta E_t \bar{r}_{t+j}$.²⁶ Intuitively, rather than assuming that the entire increase in expectations about future real interest rates is an increase *relative* to future values of the natural rate of interest, we assume that a fraction ψ is an increase in private sector expectations about current and future natural rates. This implies that only a fraction $1 - \psi$ of the increase in expected real interest rates translates into an increase in the interest rate gap that drives the output gap and inflation in the model. To capture the hump-shaped dynamics of nominal interest rates that we estimate in section 3, we specify monetary policy by equation (4) with an AR(2) shock.²⁷ In addition, we assume that the shock to expectations about the current value of the natural rate of output is proportional to the shock to expectations about the current monetary policy with the same factor of proportionality, i.e., $\Delta E_t \hat{y}_t^n = \psi \Delta E_t \bar{r}_t$.²⁸ For simplicity, we think of the increases in natural rates as arising from good news about productivity growth. In this case given our external habit assumption, $r_t^n = (\sigma^{-1}/(1-b))E_t \Delta y_{t+1}^n - (\sigma^{-1}b/(1-b))\Delta y_t^n$.

We calibrate the model to match the evidence on expected output from Table 8 and the evidence on interest rates and expected inflation from Table 1. We set the autoregressive roots of the monetary policy shock to $\rho_1 = 0.94$ and $\rho_2 = 0.70$. The response of expected output in the model depends on the degree to which shocks to real interest rates are a shock to the natural rate of interest, which is parameterized by ψ in our model. We choose $\psi = 0.8$ to roughly match the response of expected output. The larger is the values of ψ , the larger will be the response of expected output growth.

The response of expected inflation in the model is highly sensitive to the degree of real rigidity. The degree of real rigidity in the model is, in turn, highly sensitive to the elasticity of substitution

²⁶Here ΔE_t denotes the change in expectations in the 30-minute window around the FOMC announcement.

²⁷We captured these dynamics in section 4.2 using an inertial policy rule, but this structure complicates the modeling of the information effect on the natural rate. Ultimately, this distinction is unimportant: the effects of monetary shocks depend only on the realized dynamics of nominal and real interest rates, as opposed to the underlying feedback rule. We have estimated the CEE/ACEL model using both monetary policies, with very similar results, and use the estimated roots from the simpler policy in this section. Moreover, Rudebusch (2006) presents both narrative and statistical evidence suggesting that a non-inertial monetary policy rule—such as equation (4)—is a better description of policy than more standard inertial policy rules such as equation (10).

²⁸Here we assume that the FOMC meeting occurs at the beginning of the period, before the value of $\hat{y}_t^n$ is revealed to the agents. In reality, uncertainty persists about output in period t until well after period t , due to heterogeneous information. We abstract from this.

between different products θ (since this influences the degree to which marginal costs are sensitive to a firm’s demand). We choose $\theta = 10$ to roughly match the response of expected inflation. We choose standard values for all other parameters.²⁹

The resulting fit of the model is shown in Figure 5. Panel A presents the response of interest rates and expected inflation to the FOMC announcement, while Panel B presents the response of expectations about output growth and the output gap. The response of expected output growth is positive because the private sector revises upward their expectations about future productivity growth. However, to match the fact that expected inflation falls in response to the announcement, we must assume that the announcement changes beliefs about future real interest rates by more than it changes beliefs about future natural rates—i.e. $\psi < 1$. This implies that expectations about the output *gap* become negative.

The degree of real rigidity needed to match the response of inflation is substantially smaller than under the conventional interpretation of monetary policy shocks ($\psi = 0$). In that case, a value of $\theta = 400$ is needed to roughly match the response of inflation in Figure 5. The reason we are able to match the empirical responses of interest rates and expected inflation with a smaller amount of real rigidity is that the shock to the interest rate gap is smaller since the change in real interest rates arises partly from a change in beliefs about the natural rate of interest. Nevertheless, the degree of real rigidity assumed in this calibration is substantial. It is similar to the degree of real rigidity in the specific factor model discussed in Woodford (2003, ch. 3). That model was designed to generate a large amount of real rigidity.

Table 9 presents one additional piece of evidence that sheds light on the information content of FOMC announcements. This is the response of stock prices to FOMC announcements. Intuitively, a pure tightening of monetary policy leads stock prices to fall (higher discount rates and lower output), while good news about future fundamentals can raise stock prices (if higher future cash-flows outweigh higher future discount rates). In the data, we estimate that the S&P500 index falls by 6.5% in response to a policy news shock that raises the 2-year nominal forward by 1%.³⁰ This estimate is rather noisy, with a standard error of 3.9%.

Table 9 also presents the response of stock prices to our monetary policy shock in the model.³¹ In

²⁹We set the subjective discount factor to $\beta = 0.99$, the elasticity of intertemporal substitution to $\sigma = 0.5$, the Frisch elasticity of labor supply to $\eta = 1$, the curvature of the production function to $a = 2/3$, and we assume that firms change prices on average once a year ($\alpha = 0.75$).

³⁰Earlier work by Bernanke and Kuttner (2005) and Rigobon and Sack (2004) finds large responses of the stock market to surprise movements in the Federal Funds rate.

³¹For simplicity, we model stocks as an unlevered claim to the consumption stream in the economy as is common in the asset pricing literature.

the calibration of our model where monetary policy announcements convey information about both future monetary policy and future exogenous economic fundamentals, stock prices fall by 12.8% in response to the FOMC announcement. If monetary policy only conveys information about monetary policy, stock prices fall by 23.4%. The response of stock prices in the data is thus another indicator that favors the view that monetary policy conveys information to the public about future exogenous fundamentals.

6 Conclusion

We use a high-frequency identification approach to estimate the extent of monetary non-neutrality. Our evidence suggests small effects of monetary shocks on expected inflation, despite large effects on real interest rates several years into the term structure. Our empirical analysis underscores the importance of focusing on narrow windows around Fed announcements as a way of measuring monetary shocks—we show that analysis based on wider windows can yield spurious results. Our findings do not appear to derive from effects of monetary shocks on risk premia.

We use these reduced-form estimates to estimate a workhorse monetary business cycle model. Our results suggest that the Phillips curve is quite flat and that there is a large amount of inflation persistence—implying substantial nominal and real rigidities in the economy and a large amount of monetary non-neutrality. We also present evidence suggesting that the Fed influences private sector expectations about future exogenous fundamentals.

A Construction of the Policy News Shock

The policy news shock is constructed as the first principle component of the change in five interest rates. The first of these is the change in market expectations of the Federal Funds Rate over the remainder of the month in which the FOMC meeting occurs. To construct this variable we use data on the price of the Federal Funds Futures contract for the month in question. The Federal Funds futures contract for a particular month (say April 2004) trades at price p and pays off $100 - \bar{r}$ where $\bar{r}$ is the average of the effective Federal Funds Rate over the month.³² To construct the change in expectations for the remainder of the month, we must adjust for the fact that a part of the month has already elapsed when the FOMC meeting occurs. Suppose the month in question has m_0 days and the FOMC meeting occurs on day d_0 . Let $f_{t-\Delta t}^1$ denote the price of the current month's Federal Funds Rate futures contract immediately before the FOMC announcement and f_t^1 the price of this contract immediately following the FOMC announcement. Let r_{-1} denote the average Federal Funds Rate during the month up until the point of the FOMC announcement and r_0 the average Federal Funds Rate for the remainder of the month. Then

$$f_{t-\Delta t}^1 = \frac{d_0}{m_0} r_{-1} + \frac{m_0 - d_0}{m_0} E_{t-\Delta t} r_0,$$

$$f_t^1 = \frac{d_0}{m_0} r_{-1} + \frac{m_0 - d_0}{m_0} E_t r_0.$$

As a result

$$E_t r_0 - E_{t-\Delta t} r_0 = \frac{m_0}{m_0 - d_0} (f_t^1 - f_{t-\Delta t}^1).$$

When the FOMC meeting occurs on a day when there are 7 days or less remaining in a month, we instead use the change in the price of next month's Fed Funds Futures contract. This avoids multiplying $f_t^1 - f_{t-\Delta t}^1$ by a very large factor.

The second variable used in constructing the policy news shock is the change in the expected Federal Funds Rate at the time of the next scheduled FOMC meeting. Similar issues arise in constructing this variable as with the variable described above. Let m_1 denote the number of days in the month in which the next scheduled FOMC meeting occurs and let d_1 denote the day of the meeting. The next scheduled FOMC meeting may occur in the next month or as late as 3 months after the current meeting. Let $f_{t-\Delta t}^n$ denote the price of the Federal Funds Rate futures contract

³²Fed Funds futures have been traded since 1988. The effective Federal Funds Rate is the rate that is quoted by the Federal Reserve Bank of New York on every business day. See the Chicago Board of Trade Reference guide [http://www.jamesgoulding.com/Research_II/FedFundsFutures/FedFunds\(FuturesReferenceGuide\).pdf](http://www.jamesgoulding.com/Research_II/FedFundsFutures/FedFunds(FuturesReferenceGuide).pdf) for a detailed description of Fed futures contracts. On a trading day in March (say), the April Federal Funds futures contract is labeled as 2nd expiration nearby and also as 1st beginning nearby, in reference to the month over which $\bar{r}$ is computed.

for the month of the next scheduled FOMC meeting immediately before the FOMC announcement and f_t^n the price of this contract immediately following the FOMC announcement. Let r_1 denote the Federal Funds Rate after the next scheduled FOMC meeting. Analogous calculations to what we present above yield

$$E_t r_1 - E_{t-\Delta t} r_1 = \frac{m_1}{m_1 - d_1} \left[(f_t^n - f_{t-\Delta t}^n) - \frac{d_1}{m_1} (E_t r_0 - E_{t-\Delta t} r_0) \right].$$

As with the first variable, if the next scheduled FOMC meeting occurs on a day when there are 7 days or less remaining in a month, we instead use the change in the price of next month's Fed Funds Futures contract.

The last three variables used are the change in the price of three Eurodollar futures at the time of the FOMC announcements. A Eurodollar futures contract expiring in a particular quarter (say 2nd quarter 2004) is an agreement to exchange, on the second London business day before the third Wednesday of the last month of the quarter (typically a Monday near the 15th of the month), the price of the contract p for 100 minus the then current three-month US dollar BBA LIBOR interest rate. The contract thus provides market-based expectations of the three month nominal interest rate on the expiration date.³³ We make use of Eurodollar futures at horizons of n quarters in the future for $n = 2, 3, 4$ or, more precisely, the expiration date of the " n quarter" Eurodollar future is between $n - 1$ and n quarters in the future at any given point in time.

We approximate the change in these variables over a 30-minute window around FOMC by taking the difference between the price in the last trade that occurred more than 10 minutes before the FOMC announcement and the first trade that occurred more than 20 minutes after the FOMC announcement. On control days in the analysis using the heteroskedasticity based estimation approach, we take the last trade before 2:05pm and the first trade after 2:35pm (since FOMC announcements tend to occur at 2:15pm). On some days (most often control days), trading is quite sparse and there sometimes is no trade before 2:05 or after 2:35. To limit the size of the windows we consider, we only consider trades on the trading day in question and until noon the next day. If we do not find eligible trades to construct the price change we are interested in within this window, we set the price change to zero (i.e., we interpret no trading as no price change).

³³Eurodollar futures began trading in the early 1980's. See the CME Group Eurodollar futures reference guide <http://www.cmegroup.com/trading/interest-rates/files/eurodollar-futures-reference-guide.pdf> for more details about how Eurodollar futures are defined.

B Rigobon’s Heteroskedasticity-Based Estimator

Table 2 presents results from a heteroskedasticity-based estimator of the type developed by Rigobon (2003) and Rigobon and Sack (2004). The empirical model we consider in this analysis is the following. Let ϵ_t denote a pure monetary shock and suppose that movements in the policy indicator Δi_t we observe in the data is governed both by monetary and non-monetary shocks:

$$\Delta i_t = \alpha_i + \epsilon_t + \eta_t, \quad (14)$$

where η_t is a vector of all other shocks that affect Δi_t . Here α_i and β_i are constants and we normalize the impact of ϵ_t on Δi_t to one. We wish to estimate the effects of the monetary shock ϵ_t on an outcome variable Δs_t . This variable is also affected by both the monetary and non-monetary shocks:

$$\Delta s_t = \alpha_s + \gamma \epsilon_t + \beta_s \eta_t. \quad (15)$$

The parameter of interest is γ , which should be interpreted as the impact of the pure monetary shock ϵ_t on Δs_t relative to its impact on Δi_t .

Our identifying assumption is that the variance of monetary shocks increases at the time of FOMC announcements, while the variance of other shocks is unchanged. Define $R1$ as a sample of narrow time intervals around FOMC announcements, and define $R2$ as a sample of equally narrow time intervals that do not contain FOMC announcements but are comparable on other dimensions (e.g., same time of day, same day of week, etc.). We refer to $R1$ as our “treatment” sample and $R2$ as our “control” sample. Our identifying assumption is that $\sigma_{\epsilon,R1} > \sigma_{\epsilon,R2}$, while $\sigma_{\eta,R1} = \sigma_{\eta,R2}$.

Let Ω_{Ri} denote the variance-covariance matrix of $[\Delta i_t, \Delta s_t]$ in regime Ri . Then Ω_{Ri} is given by

$$\Omega_{Ri} = \begin{bmatrix} \sigma_{\epsilon,Ri}^2 + \sum_j \beta_{i,j}^2 \sigma_{\eta,j}^2 & \gamma \sigma_{\epsilon,Ri}^2 + \sum_j \beta_{i,j} \beta_{s,j} \sigma_{\eta,j}^2 \\ \gamma \sigma_{\epsilon,Ri}^2 + \sum_j \beta_{i,j} \beta_{s,j} \sigma_{\eta,j}^2 & \gamma^2 \sigma_{\epsilon,Ri}^2 + \sum_j \beta_{s,j}^2 \sigma_{\eta,j}^2 \end{bmatrix},$$

where j indexes the elements of η_t . Notice that

$$\Delta \Omega = \Omega_{R1} - \Omega_{R2} = (\sigma_{\epsilon,R1}^2 - \sigma_{\epsilon,R2}^2) \begin{bmatrix} 1 & \gamma \\ \gamma & \gamma^2 \end{bmatrix}.$$

Thus,

$$\gamma = \frac{\Delta \Omega_{12}}{\Delta \Omega_{11}} = \frac{\text{cov}_{R1}(\Delta i_t, \Delta s_t) - \text{cov}_{R2}(\Delta i_t, \Delta s_t)}{\text{var}_{R1}(\Delta i_t) - \text{var}_{R2}(\Delta i_t)}. \quad (16)$$

This is the estimator we use to construct the results in Table 2 and Table A.2. Notice that if we set the variance of the “background noise” η_t to zero, then the heteroskedasticity-based estimator,

equation (16) reduces to the coefficient from an OLS regression of Δs_t on Δi_t . Intuitively, the full heteroskedasticity-based estimator can be thought of as the simple OLS estimator, adjusted for the “normal” covariance between Δs_t and Δi_t and the “normal” variance of Δi_t .

C Weak Instruments Robust Confidence Intervals

The confidence intervals in Table 2 are constructed using a more sophisticated bootstrap procedure than is conventional. The reason is that the conventional bootstrap approach to constructing confidence intervals yields inaccurate results in the case when there is a significant probability that the difference in the variance of Δi_t between the treatment and control sample is close to zero.³⁴ Figure A.1 illustrates that this is the case for the 1-day window estimation but not the 30-minute window. The problem is essentially one of weak instruments. Rigobon and Sack (2004) show that the estimator in equation (16) can be formulated as an IV regression. When the difference in the variance of Δi_t between the treatment and control sample is small, the instrument in this formulation is weak, leading to biased point estimates and confidence intervals.

In Table 2, we, therefore, employ a weak-instruments robust approach to constructing confidence intervals. The approach we employ is a test inversion approach. A 95% confidence interval for our parameter of interest γ can be constructed by performing a hypothesis test for all possible hypothetical true values of γ and including those values that are not rejected by the test in the confidence interval. The test statistic we use is

$$g(\gamma) = \Delta \text{cov}(\Delta i_t, \Delta s_t) - \gamma \Delta \text{var}(\Delta i_t), \quad (17)$$

where Δcov and Δvar denote the difference between the covariance and variance, respectively, in the treatment and control samples. Intuitively, $g(\gamma) = 0$ at the true value of γ . We estimate the distribution of $g(\gamma)$ for each hypothetical value of γ and include in our confidence interval values of γ for which $g(\gamma) = 0$ cannot be rejected. Figure A.2 plots the 2.5%, 50% and 97.5% quantiles of the distribution of $g(\gamma)$ as a function of γ for the 2-year nominal forward in the one-day window case. Values of γ for which the 2.5% quantile lies below zero and the 97.5% quantile lies above zero are included in the 95% confidence interval. This method for constructing confidence intervals is referred to as the Fieller method by Staiger, Stock, and Watson (1997) as it is an extension of an

³⁴Recall that the Rigobon estimator—equation (16)—is a ratio with the difference in the variance of Δi_t between the treatment sample and the control sample in the denominator. If the distribution of this difference has significant mass in the vicinity of zero, the sampling distribution of the estimator will have significant mass at large positive and negative values.

approach proposed by Fieller (1954). We use a bootstrap to estimate the joint distribution of Δcov and Δvar . Our approach is therefore similar to the grid bootstrap proposed by Hansen (1999) for a different application.

This more sophisticated procedure for constructing confidence intervals is not important for our baseline estimator based on changes in the policy news shock over a 30-minute window. In this case, the weak-IV robust confidence intervals coincide closely with the standard non-parametric bootstrap confidence interval reported in Table A.2. However, this weak-IV robust procedure is very important for the Rigobon estimator when the policy news shock is measured over a 1-day window.

D A Simple New Keynesian Model

This section lays out micro-foundations for the simple New Keynesian business cycle model discussed in section 4 in the main text. See Woodford (2003) and Gali (2008) for thorough expositions of New Keynesian models.

D.1 Households

The economy is populated by a continuum of household types indexed by x . A household's type indicates the type of labor supplied by that household. Households of type x seek to maximize their utility given by

$$E_0 \sum_{t=0}^{\infty} \beta^t [u(C_t, \xi_t) - v(L_t(x), \xi_t)], \quad (18)$$

where β denotes the household's subjective discount factor, C_t denotes household consumption of a composite consumption good, $L_t(x)$ denotes household supply of differentiated labor input x , and ξ_t denotes a vector of preference shocks. There are an equal (large) number of households of each type. The composite consumption good in expression (18) is an index given by

$$C_t = \left[\int_0^1 c_t(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}}, \quad (19)$$

where $c_t(z)$ denotes consumption of products of variety z . The parameter $\theta > 1$ denotes the elasticity of substitution between different varieties.

Households have access to complete financial markets. Households of type x face a flow budget constraint given by

$$P_t C_t + E_t[M_{t,t+1} B_{t+1}(x)] \leq B_t(x) + W_t(x) L_t(x) + \int_0^1 \Xi_t(z) dz - T_t, \quad (20)$$

where P_t is a price index that gives the minimum price of a unit of the consumption good C_t , $B_{t+1}(x)$ is a random variable that denotes the state contingent payoff of the portfolio of financial securities held by households of type x at the beginning of period $t+1$, $M_{t,t+1}$ is the stochastic discount factor that prices these payoffs in period t ,³⁵ $W_t(x)$ denotes the wage rate received by households of type x in period t , $\Xi_t(z)$ denotes the profits of firm z in period t , and T_t is a lump-sum tax levied by the government. To rule out Ponzi schemes, household debt cannot exceed the present value of future income in any state of the world.

Households face a decision in each period about how much to spend on consumption, how many hours of labor to supply, how much to consume of each differentiated good produced in the economy and what portfolio of assets to purchase. Optimal choice regarding the trade-off between current consumption and consumption in different states in the future yields the following consumption Euler equation:

$$\frac{u_c(C_{t+j}, \xi_{t+j})}{u_c(C_t, \xi_t)} = \frac{M_{t,t+j} P_{t+j}}{\beta^j P_t} \quad (21)$$

as well as a standard transversality condition. Subscripts on the function u denote partial derivatives. Equation (21) holds state-by-state for all $j > 0$. Optimal choice regarding the intratemporal trade-off between current consumption and current labor supply yields a labor supply equation:

$$\frac{v_\ell(L_t(x), \xi_t)}{u_c(C_t, \xi_t)} = \frac{W_t(x)}{P_t}. \quad (22)$$

Households optimally choose to minimize the cost of attaining the level of consumption C_t . This implies the following demand curves for each of the differentiated products produced in the economy:

$$c_t(z) = C_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta}, \quad (23)$$

where $p_t(z)$ denotes the price of product z and

$$P_t = \left[\int_0^1 p_t(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}. \quad (24)$$

D.2 Firms

There are a continuum of firms indexed by z in the economy. Firm z specializes in the production of differentiated good z , the output of which we denote $y_t(z)$. For simplicity, labor is the only variable

³⁵The stochastic discount factor $M_{t,t+1}$ is a random variable over states in period $t+1$. For each such state it equals the price of the Arrow-Debreu asset that pays off in that state divided by the conditional probability of that state. See Cochrane (2005) for a detailed discussion.

factor of production used by firms. Each firm is endowed with a fixed, non-depreciating stock of capital. The production function of firm z is

$$y_t(z) = A_t f(L_t(z)), \quad (25)$$

where A_t denotes aggregate productivity. The function f is increasing and concave. It is concave because there are diminishing marginal return to labor given the fixed amount of other inputs employed at the firm. We follow Woodford (2003) in introducing heterogeneous labor markets. Firm belongs to an industry x . There are many firms in each industry. The goods in industry x are produced using labor of type x and all firms in industry x change prices at the same time. This heterogeneous labor market structure is a strong source of real rigidities in price setting.

Firm z acts to maximize its value,

$$E_t \sum_{j=0}^{\infty} M_{t,t+j} [p_{t+j}(z) y_{t+j}(z) - W_{t+j}(x) L_{t+j}(z)]. \quad (26)$$

Firm z must satisfy demand for its product given by equation (23). Firm z is therefore subject to the following constraint:

$$C_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta} \leq A_t f(L_t(z)). \quad (27)$$

Firm z takes its industry wage $W_t(x)$ as given. Optimal choice of labor demand by the firm is given by

$$W_t(x) = A_t f_\ell(L_t(z)) S_t(z), \quad (28)$$

where $S_t(z)$ denotes the firm's nominal marginal cost (the Lagrange multiplier on equation (27) in the firm's constrained optimization problem).

Firm z can reoptimize its price with probability $1 - \alpha$ as in Calvo (1983). With probability α it must keep its price unchanged. Optimal price setting by firm z in periods when it can change its price implies

$$p_t(z) = \frac{\theta}{\theta - 1} E_t \sum_{j=0}^{\infty} \frac{\alpha^j M_{t,t+j} y_{t+j}(z)}{\sum_{k=0}^{\infty} \alpha^k M_{t,t+k} y_{t+k}(z)} S_{t+j}(z). \quad (29)$$

Intuitively, the firm sets its price equal to a constant markup over a weighted average of current and expected future marginal cost.

D.3 A Linear Approximation of Private Sector Behavior

We seek a linear approximation of the equation describing private sector behavior around a zero-growth, zero-inflation steady state. We start by deriving a log-linear approximation for the consumption Euler equation that related consumption growth and a one-period, riskless, nominal bond. This

equation takes the form $E_t[M_{t,t+1}(1 + i_t)] = 1$, where i_t denotes the yield on a one-period, riskless, nominal bond. Using equation (21) to plug in for $M_{t,t+1}$ and rearranging terms yields

$$E_t \left[\beta U_c(C_{t+1}, \xi_{t+1}) \frac{P_t}{P_{t+1}} \right] = \frac{U_c(C_t, \xi_t)}{1 + i_t}. \quad (30)$$

The zero-growth, zero-inflation steady state of this equation is $\beta(1 + \bar{i})$. A first order Taylor series approximation of equation (30) is

$$\hat{c}_t = E_t \hat{c}_{t+1} - \sigma(\hat{i}_t - E_t \hat{\pi}_{t+1}) - \sigma E_t \Delta \hat{\xi}_{ct+1}, \quad (31)$$

where $\hat{c}_t = (C_t - C)/C$, $\hat{\pi}_t = \pi_t - 1$, $\hat{i}_t = (1 + i_t - 1 - \bar{i})/(1 + \bar{i})$, and $\hat{\xi}_{ct} = (U_{cc}/U_c)(\xi_t - 1)$. The parameter $\sigma = -U_c/(U_{cc}C)$ denotes the intertemporal elasticity of substitution of households.

We next linearize labor demand, labor supply, and the production function and combine these equations to get an expression for the marginal costs in period $t + j$ of a firm that last changed its price in period t . Let $\ell_{t,t+j}(x)$ denote the percent deviation from steady state in period $t + j$ of hours worked for workers in industry x that last was able to change prices in period t . Let other industry level variables be defined analogously. We assume that $f(L_t(x)) = L_t^a(x)$.

A linear approximation of labor demand—equation (28)—in period $t + j$ for industry x that was last able to change its prices in period t is then

$$\hat{w}_{t,t+j}(x) = \hat{a}_{t+j} - (1 - a)\hat{\ell}_{t,t+j}(x) + \hat{s}_{t,t+j}(x), \quad (32)$$

where $\hat{w}_{t,t+j}(x)$ and $\hat{s}_{t,t+j}(x)$ denote the percentage deviation of real wages and real marginal costs, respectively, from their steady state values.

A linear approximation of labor supply—equation (22)—in period $t + j$ for industry x that was last able to change its prices in period t is

$$\hat{w}_{t,t+j}(x) = \eta^{-1} \hat{\ell}_{t,t+j}(x) + \sigma^{-1} \hat{c}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j}, \quad (33)$$

where $\hat{\xi}_{\ell,t+j} = (V_{\ell\xi}/V_\ell)(\xi_t - 1)$. The parameter $\eta = V_\ell/(V_{\ell\ell}L)$ is the Frisch elasticity of labor supply.

A linear approximation of the production function—equation (25)—in period $t + j$ for industry x that was last able to change its prices in period t is

$$\hat{y}_{t,t+j}(x) = \hat{a}_{t+j} + a\hat{\ell}_{t,t+j}(x). \quad (34)$$

Combining labor demand and labor supply—equations (32) and (33)—to eliminate $\hat{w}_{t,t+j}(x)$ yields

$$\hat{s}_{t,t+j}(x) = (\eta^{-1} + 1 - a)\hat{\ell}_{t,t+j}(x) + \sigma^{-1} \hat{c}_{t+j} - \hat{a}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j}.$$

Using the production function—equation (34)—to eliminate $\hat{\ell}_{t,t+j}(x)$ yields

$$\hat{s}_{t,t+j}(x) = \omega \hat{y}_{t,t+j}(x) + \sigma^{-1} \hat{c}_{t+j} - (\omega + 1) \hat{a}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j}, \quad (35)$$

where $\omega = (\eta^{-1} + 1 - a)/a$.

Taking logs of consumer demand—equation (23)—in period $t + j$ for industry x what was last able to change its prices in period t yields

$$\hat{y}_{t,t+j}(z) = -\theta \hat{p}_t(x) + \theta \sum_{k=1}^j \hat{\pi}_{t+k} + \hat{y}_{t+j}, \quad (36)$$

where we use the fact that $Y_t = C_t$ and $y_t(x) = c_t(x)$. Plugging this equation into equation (35) and again using the fact that $Y_t = C_t$ yields

$$\hat{s}_{t,t+j}(x) = -\omega \theta \hat{p}_t(x) + \omega \theta \sum_{k=1}^j \hat{\pi}_{t+k} + (\omega + \sigma^{-1}) \hat{y}_{t+j} - (\omega + 1) \hat{a}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j} \quad (37)$$

It is useful to derive the level of output that would prevail if all prices were flexible. Since our model does not have any industry specific shocks (other than the opportunity to change prices), marginal costs of all firms are the same when prices are flexible. Firm price setting in this case yields $p_t(x) = \mu S_t$, where $\mu = \theta/(\theta - 1)$. This implies that all prices are equal and that $S_t/P_t = 1/\mu$. Since real marginal cost is a constant, we have $\hat{s}_t = 0$. The flexible price version of equation (37) is then

$$(\omega + \sigma^{-1}) \hat{y}_t^n = (\omega + 1) \hat{a}_t - \hat{\xi}_{\ell,t} + \hat{\xi}_{c,t}, \quad (38)$$

where we use the fact that output in all industries is the same under flexible prices and $\hat{y}_t = \hat{c}_t$ and denote the rate of output under flexible prices as y_t^n . We will refer to y_t^n as the natural rate of output.

Combining equations (37) and (38) yields

$$\hat{s}_{t,t+j}(x) = -\omega \theta \hat{p}_t(x) + \omega \theta \sum_{k=1}^j \hat{\pi}_{t+k} + (\omega + \sigma^{-1}) (\hat{y}_{t+j} - \hat{y}_{t+j}^n) \quad (39)$$

We next linearize the price setting equation—equation (29). This yields:

$$\sum_{j=0}^{\infty} (\alpha\beta)^j \hat{p}_t(x) - \sum_{j=0}^{\infty} (\alpha\beta)^j E_t \hat{s}_{t,t+j}(x) - \sum_{j=1}^{\infty} (\alpha\beta)^j \sum_{k=1}^j E_t \hat{\pi}_{t+k} = 0.$$

Manipulation of this equation yields

$$\hat{p}_t(x) = (1 - \alpha\beta) \sum_{j=0}^{\infty} (\alpha\beta)^j E_t \hat{s}_{t,t+j}(x) + \alpha\beta \sum_{j=1}^{\infty} (\alpha\beta)^j E_t \hat{\pi}_{t+j}. \quad (40)$$

Using equation (39) to eliminate $\hat{s}_{t,t+j}(x)$ in equation (40) and manipulating the resulting equation yields

$$\hat{p}_t(x) = (1 - \alpha\beta)\zeta \sum_{j=0}^{\infty} (\alpha\beta)^j E_t(\hat{y}_{t+j} - \hat{y}_{t+j}^n) + \alpha\beta \sum_{j=1}^{\infty} (\alpha\beta)^j E_t \hat{\pi}_{t+j}, \quad (41)$$

where $\zeta = (\omega + \sigma^{-1})/(1 + \omega\theta)$. A linear approximation of the expression for the price index—equation (24)—yields

$$\hat{\pi}_t = \frac{1 - \alpha}{\alpha} \hat{p}_t(x). \quad (42)$$

Using this last equation to replace $\hat{p}_t(x)$ in equation (41) yields

$$\hat{\pi}_t = \kappa\zeta \sum_{j=0}^{\infty} (\alpha\beta)^j E_t(\hat{y}_{t+j} - \hat{y}_{t+j}^n) + (1 - \alpha)\beta \sum_{j=1}^{\infty} (\alpha\beta)^j E_t \hat{\pi}_{t+j},$$

where $\kappa = (1 - \alpha)(1 - \alpha\beta)/\alpha$. Quasi-differencing the resulting equation yields

$$\hat{\pi}_t - \alpha\beta E_t \hat{\pi}_{t+1} = \kappa\zeta(\hat{y}_t - \hat{y}_t^n) + (1 - \alpha)\beta E_t \hat{\pi}_{t+1},$$

which implies

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa\zeta(\hat{y}_t - \hat{y}_t^n). \quad (43)$$

Finally, we rewrite the household's Euler equation—equation (31) in terms of the output gap:

$$y_t - y_t^n = E_t(y_{t+1} - y_{t+1}^n) - \sigma(\hat{i}_t - E_t \hat{\pi}_{t+1} - r_t^n), \quad (44)$$

where r_t^n denotes the “natural rate of interest” as is given by

$$r_t^n = E_t \Delta \xi_{c,t+1} + \frac{1}{\sigma} E_t \Delta y_{t+1}^n. \quad (45)$$

References

- ABRAHAM, M., T. ADRIAN, R. K. CRUMP, AND E. MOENCH (2015): “Decomposing Real and Nominal Yield Curves,” Federal Reserve Bank of New York Staff Report No. 570.
- ALTIG, D., L. J. CHRISTIANO, M. EICHENBAUM, AND J. LINDE (2011): “Firm-Specific Capital, Nominal Rigidities and the Business Cycle,” *Review of Economic Dynamics*, 14, 225–247.
- BARAKCHIAN, S. M., AND C. CROWE (2010): “Monetary Policy Matters: New Evidence Based on a New Shock Measure,” IMP Working Paper 10/230.
- BEECHEY, M. J., AND J. H. WRIGHT (2009): “The High-Frequency Impact of News on Long-Term Yields and Forward Rates: Is It Real?,” *Journal of Monetary Economics*, 56, 535–544.
- BERNANKE, B. S., AND K. N. KUTTNER (2005): “What Explains the Stock Market’s Reaction to Federal Reserve Policy,” *Journal of Finance*, 60(3), 1221–1257.
- CALVO, G. A. (1983): “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, 12, 383–398.
- CAMPBELL, J. R., C. L. EVANS, J. D. FISHER, AND A. JUSTINIANO (2012): “Macroeconomic Effects of Federal Reserve Forward Guidance,” *Brookings Papers on Economic Activity*, 2012(1), 1–80.
- CHRISTIANO, L. J., M. EICHENBAUM, AND C. L. EVANS (1999): “Monetary Policy Shocks: What Have We Learned and to What End?,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, pp. 65–148, Amsterdam, Holland. Elsevier.
- (2005): “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, 115, 1–45.
- COCHRANE, J. H. (2005): *Asset Pricing*. Princeton University Press, Princeton, NJ, second edn.
- COCHRANE, J. H., AND M. PIAZZESI (2002): “The Fed and Interest Rates: A High-Frequency Identification,” *American Economic Review*, 92(2), 90–95.
- COOK, T., AND T. HAHN (1989): “The Effect of Changes in the Federal Funds Rate Target on Market Interest Rates in the 1970s,” *Journal of Monetary Economics*, 24(3), 331–351.
- CUMMINS, J., K. HASSETT, AND R. G. HUBBARD (1994): “A Reconsideration of Investment Behavior Using Tax Reforms as Natural Experiments,” *Brookings Papers on Economic Activity*, 1994(2), 1–59.
- ELLINGSEN, T., AND U. SODERSTROM (2001): “Monetary Policy and Market Interest Rates,” *American Economic Review*, 91(5), 1594–1607.
- FAUST, J., E. T. SWANSON, AND J. H. WRIGHT (2004): “Do Federal Reserve Policy Surprises Reveal Superior Information about the Economy?,” *Contributions to Macroeconomics*, 4(1), 1–29.
- FIELLER, E. (1954): “Some Problems in Interval Estimation,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 16(2), 175–185.

- FLECKENSTEIN, M., F. A. LONGSTAFF, AND H. LUSTIG (2014): “The TIPS-Treasury Bond Puzzle,” *Journal of Finance*, 69, 2151–2197.
- GAGNON, J., M. RASKIN, J. REMACHE, AND B. SACK (2011): “Large-Scale Asset Purchases by the Federal Reserve: Did They Work?,” *Federal Reserve Bank of New York Economic Policy Review*, May, 41–59.
- GALI, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ.
- GERTLER, M., AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7(1), 44–76.
- GILCHRIST, S., D. LOPEZ-SALIDO, AND E. ZAKRAJSEK (2015): “Monetary Policy and Real Borrowing Costs at the Zero Lower Bound,” *American Economic Journal: Macroeconomics*, 7(1), 77–109.
- GILCHRIST, S., AND E. ZAKRAJSEK (2012): “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review*, 102(4), 1692–1720.
- GONCALVES, S., AND L. KILIAN (2004): “Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form,” *Journal of Econometrics*, 123(2004), 89–120.
- GURKAYNAK, R. S., B. SACK, AND E. T. SWANSON (2005): “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking*, 1, 55–93.
- (2007): “Market-Based Measures of Monetary Policy Expectations,” *Journal of Business and Economic Statistics*, 25(2), 201–212.
- GURKAYNAK, R. S., B. SACK, AND J. H. WRIGHT (2010): “The TIPS Yield Curve and Inflation Compensation,” *American Economic Journal: Macroeconomics*, 2(1), 70–92.
- HANSEN, B. E. (1999): “The Grid Bootstrap and the Autoregressive Model,” *Review of Economics and Statistics*, 81(4), 594–607.
- HANSON, S. G., AND J. C. STEIN (2015): “Monetary Policy and Long-Term Real Rates,” *Journal of Financial Economics*, 115.
- KIM, D. H., AND J. H. WRIGHT (2005): “An Arbitrage-Free Three-Factor Term Structure Model and the Recent Behavior of Long-Term Yields and Distant-Horizon Forward Rates,” Working Paper.
- KRISHNAMURTHY, A., AND A. VISSING-JORGENSEN (2011): “The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy,” *Brookings Papers on Economic Activity*, 2011, 215–265.
- KUTTNER, K. N. (2001): “Monetary policy surprises and interest rates: Evidence from the Fed funds futures market,” *Journal of Monetary Economics*, 47, 523–544.
- MAVROEIDIS, S., M. PLAGBORG-MOLLER, AND J. H. STOCK (2014): “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” *Journal of Economic Literature*, 52.

- MERTENS, K., AND M. O. RAVN (2013): “The Dynamic Effects of Personal and Corporate Income Tax Change in the United States,” *American Economic Review*, 103(4), 1212–1237.
- PIAZZESI, M., AND E. T. SWANSON (2008): “Future Prices as Risk-Adjusted Forecasts of Monetary Policy,” *Journal of Monetary Economics*, 55, 677–691.
- RIGOBON, R. (2003): “Identification through Heteroskedasticity,” *The Review of Economics and Statistics*, 85(4), 777–792.
- RIGOBON, R., AND B. SACK (2004): “The impact of monetary policy on asset prices,” *Journal of Monetary Economics*, 51, 1553–1575.
- ROMER, C. D., AND D. H. ROMER (2000): “Federal Reserve Information and the Behavior of Interest Rates,” *American Economic Review*, 90(3), 429–457.
- (2004): “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 94(4), 1055–1084.
- ROSA, C. (2012): “How “Unconventional” Are Large-Scale Asset Purchases? The Impact of Monetary Policy on Asset Prices,” Federal Reserve Bank of New York Staff Report No. 560.
- ROTEMBERG, J. J., AND M. WOODFORD (1997): “An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy,” in *NBER Macroeconomics Annual 1997*, ed. by B. S. Bernanke, and J. J. Rotemberg, pp. 297–346, Cambridge, MA. MIT Press.
- RUDEBUSCH, G. D. (2006): “Monetary Policy Inertia: Fact or Fiction,” *International Journal of Central Banking*, 2(4), 85–135.
- STAIGER, D. O., J. H. STOCK, AND M. W. WATSON (1997): “How Precise Are Estimates of the Natural Rate of Unemployment,” in *Reducing Inflation: Motivation and Strategy*, ed. by C. D. Romer, and D. H. Romer, pp. 195–246, Chicago, IL. University of Chicago Press.
- STOCK, J. H., AND M. W. WATSON (2012): “Disentangling the Channels of the 2007-2009 Recession,” *Brookings Papers on Economic Activity*, 42(1), 81–135.
- WOODFORD, M. (2003): *Interest and Prices*. Princeton University Press, Princeton, NJ.
- WRIGHT, J. H. (2012): “What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound,” *Economic Journal*, 122, F447–F466.

TABLE 1
Response of Interest Rates and Inflation to the Policy News Shock

	Nominal	Real	Inflation
3M Treasury Yield	0.67 (0.14)		
6M Treasury Yield	0.85 (0.11)		
1Y Treasury Yield	1.00 (0.14)		
2Y Treasury Yield	1.10 (0.33)	1.06 (0.24)	0.04 (0.18)
3Y Treasury Yield	1.06 (0.36)	1.02 (0.25)	0.04 (0.17)
5Y Treasury Yield	0.73 (0.20)	0.64 (0.15)	0.09 (0.11)
10Y Treasury Yield	0.38 (0.17)	0.44 (0.13)	-0.06 (0.08)
2Y Treasury Inst. Forward Rate	1.14 (0.46)	0.99 (0.29)	0.15 (0.23)
3Y Treasury Inst. Forward Rate	0.82 (0.43)	0.88 (0.32)	-0.06 (0.15)
5Y Treasury Inst. Forward Rate	0.26 (0.19)	0.47 (0.17)	-0.21 (0.08)
10Y Treasury Inst. Forward Rate	-0.08 (0.18)	0.12 (0.12)	-0.20 (0.09)

Each estimate comes from a separate OLS regression. The dependent variable in each regression is the one day change in the variable stated in the left-most column. The independent variable is a change in the policy news shock over a 30 minute window around the time of FOMC announcements. The sample period is 1/1/2000 to 3/19/2014, except that we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. For 2Y and 3Y yields and real forwards, the sample starts in 2004. The sample size for the 2Y and 3Y yields and forwards is 74. The sample size for all other regressions is 106.

TABLE 2
Allowing For Background Noise in Interest Rates

	2-Year Forward		5-Year Forward		10-Year Forward	
	Nominal	Real	Nominal	Real	Nominal	Real
<i>Policy News Shock, 30-Minute Window:</i>						
OLS	1.14 [0.23, 2.04]	0.99 [0.41, 1.57]	0.26 [-0.12, 0.64]	0.47 [0.14, 0.80]	-0.08 [-0.43, 0.28]	0.12 [-0.12, 0.36]
Rigobon	1.10 [0.31, 2.36]	0.96 [0.45, 1.82]	0.22 [-0.14, 0.64]	0.46 [0.15, 0.84]	-0.12 [-0.46, 0.24]	0.11 [-0.13, 0.35]
<i>Policy News Shock, 1-Day Window:</i>						
OLS	1.24 [0.80, 1.69]	1.00 [0.57, 1.43]	0.44 [0.18, 0.70]	0.48 [0.20, 0.76]	0.05 [-0.20, 0.29]	0.15 [-0.10, 0.39]
Rigobon	0.93 [-0.64, 2.08]	0.82 [0.38, 3.20]	-0.11 [-1.23, 0.33]	0.33 [-0.07, 1.12]	-0.51 [-1.93, -0.08]	-0.04 [-0.51, 0.45]
<i>2-Year Nominal Yield, 1-Day Window</i>						
OLS	1.23 [1.07, 1.38]	0.94 [0.69, 1.20]	0.64 [0.43, 0.84]	0.54 [0.31, 0.76]	0.18 [0.01, 0.35]	0.20 [0.02, 0.38]
Rigobon (90% CI)	1.14 [0.82, 1.82]	0.97 [0.62, 2.98]	0.10 [-7.94, 0.60]	0.51 [-0.01, 7.48]	-0.79 [-10.00, -0.21]	-0.08 [-4.57, 0.38]

Each estimate comes from a separate "regression." The dependent variable in each regression is the one day change in the variable stated at the top of that row. The independent variable in the first panel of results is the 30-minute change in the policy news shock around FOMC meeting times, in the second panel it is the 1-day change in the policy news shock, and in the third panel it is the 1-day change in the 2-Year nominal yield. In each panel, we report results based on OLS and Rigobon's heteroskedasticity based estimation approach. We report a point estimate and 95% confidence intervals except in the last row which reports 90% confidence intervals. The sample of "treatment" days for the Rigobon method is all regularly scheduled FOMC meeting days from 1/1/2000 to 3/19/2014. The sample of "control" days for the Rigobon analysis is all Tuesdays and Wednesdays that are not FOMC meeting days from 1/1/2000 to 12/31/2012. In both the treatment and control samples, we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. For 2Y and 3Y yields and real forwards, the sample starts in 2004. Confidence intervals for the Rigobon method are calculated using the weak-IV robust approach discussed in the appendix with 5000 iterations.

TABLE 3
Effects of Monetary Shocks on Survey Expectations

	Nominal	Real	Inflation
1 quarter	1.03 (0.46)	1.20 (0.48)	-0.17 (0.25)
2 quarters	1.14 (0.47)	1.58 (0.46)	-0.44 (0.22)
3 quarters	0.90 (0.50)	1.19 (0.49)	-0.29 (0.21)
4 quarters	0.83 (0.48)	1.16 (0.46)	-0.33 (0.20)
5 quarters	0.66 (0.60)	0.53 (0.61)	0.13 (0.23)
6 quarters	1.79 (0.59)	1.56 (0.62)	0.23 (0.27)
7 quarters	3.89 (1.13)	3.75 (1.16)	0.13 (0.48)

Each estimate comes from a separate OLS regression. We regress changes in survey expectations from the *Blue Chip Economic Indicators* on the policy news shock. Since the *Blue Chip* survey expectations are available at a monthly frequency, we construct a corresponding monthly measure of our policy news shock. In particular, we use any policy news shock that occurs over the month except for those that occur in the first week (because we do not know whether these occurred before or after the survey response). The dependent variable is the change in the forecasted value of a variable N quarters ahead, between this month's survey and last month's survey. We consider the effects on expected future 3-month T-Bill rates, short-term real interest rates and inflation, where the inflation rate is the GDP deflator and the short-term real interest rate is calculated as the difference between the expected 3-month T-bill rate and the expected GDP deflator for a given quarter. The sample period is January 1995 to April 2014, except that we exclude the second half of 2008 and the first half of 2009. Standard errors are in parentheses.

TABLE 4
Response of Expected Future Short Rates and Risk Premia

	Expected Future Short Rates		Risk Premia	
	Nominal	Real	Nominal	Real
2Y Treasury Yield	1.01 (0.27)	0.86 (0.17)	0.09 (0.10)	0.20 (0.18)
3Y Treasury Yield	0.93 (0.26)	0.82 (0.18)	0.13 (0.15)	0.21 (0.19)
5Y Treasury Yield	0.76 (0.16)	0.60 (0.12)	-0.04 (0.11)	0.04 (0.14)
10Y Treasury Yield	0.50 (0.11)	0.40 (0.08)	-0.12 (0.14)	0.04 (0.14)
2Y Treasury Forward Rate	0.79 (0.24)	0.73 (0.22)	0.35 (0.26)	0.26 (0.21)
3Y Treasury Forward Rate	0.61 (0.19)	0.56 (0.17)	0.21 (0.29)	0.32 (0.25)
5Y Treasury Forward Rate	0.36 (0.08)	0.33 (0.08)	-0.11 (0.17)	0.14 (0.17)
10Y Treasury Forward Rate	0.10 (0.02)	0.09 (0.02)	-0.18 (0.18)	0.04 (0.12)

Each estimate comes from a separate OLS regression. The dependent variables in the first two columns are one-day changes in risk neutral yields and forwards from Abrahams et al. (2013) -- i.e., measures of expected future short rates. The dependent variables in the later two columns are the difference between one-day changes in raw yields and forwards and one-day changes in the risk neutral yields and forwards from Abrahams et al. (2013). We refer to this difference as the risk premia. It corresponds to the term premium, liquidity premium and model error in Abrahams et al. (2013). The independent variable is a change in the policy new shock over a 30 minute window around the time of FOMC announcements. The forward rates are one-year forwards at different horizons. The sample period is 1/1/2000 to 3/19/2014, except that we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. For 2Y and 3Y yields and real forwards, the sample starts in 2004. The sample size for the 2Y and 3Y yields and forwards is 74. The sample size for all other regressions is 106.

TABLE 5
Mean Reversion

Horizon (Trading Days)	Nominal Yields			Real Yields		
	2-Year	3-Year	5-Year	2-Year	3-Year	5-Year
1	1.10 (0.33)	1.06 (0.36)	0.73 (0.20)	1.06 (0.24)	1.02 (0.25)	0.64 (0.15)
5	2.24 (1.08)	2.06 (1.06)	0.85 (0.44)	1.01 (0.66)	0.93 (0.66)	0.52 (0.37)
10	2.39 (1.11)	2.20 (1.04)	0.22 (0.72)	1.35 (0.74)	1.20 (0.66)	0.28 (0.55)
20	0.60 (1.23)	0.29 (1.15)	-0.01 (1.09)	0.88 (1.29)	0.43 (1.12)	0.04 (0.82)
60	3.41 (2.59)	2.80 (2.47)	-0.17 (1.64)	1.96 (3.09)	1.72 (2.73)	-0.10 (1.38)
125	9.42 (3.25)	8.02 (2.86)	3.20 (2.17)	6.16 (4.13)	5.22 (3.72)	2.47 (1.84)
250	13.52 (5.00)	11.56 (4.01)	4.72 (2.50)	9.58 (3.20)	8.22 (2.76)	4.13 (1.61)

This table presents the results of regressing the cumulative change in yields between the day before the FOMC announcement and 1, 5, 10, 20, 60, 125 and 250 trading days after the announcement on the policy news shock in the 30 minute interval surrounding the FOMC announcement. The first three columns present results for nominal zero coupon yields, and the next three columns present results for real zero coupon yields. Standard errors are in parentheses.

TABLE 6
Monetary Non-Neutrality

	Output/Inflation
<i>Panel A: Baseline Results:</i>	
Our Estimation of CEE/ACEL Model	3.8 [2.7, 4.8]
ACEL	3.3
CEE	1.7
<i>Panel B: HF-VAR Estimation Results:</i>	
Baseline HF-VAR Estimation	3.4 [0.2, 3.8]
Output-Inflation HF-VAR Estimation	3.6 [0.1, 3.9]
Baseline + Habit HF-VAR Estimation	3.3 [0.0, 3.6]

We consider the response of the economy the monetary policy shock we estimated in the baseline case (i.e., with $\rho=0.96$ and $v=0.74$) of size 25 bp. We report the ratio of the cumulative impulse response of output to the cumulative impulse response of inflation. For both output and inflation, we sum the absolute value of the response over 500 periods after the shock. For output, we divide by four because a 1% higher level of output in all quarters of a year is equivalent to a 1% higher level of annual output for the year as a whole. For inflation, there is no need to divide by four, because a 1% higher inflation rate in all quarters of a year is equivalent to a 4% higher inflation rate on an annual basis. The first panel presents the results for our baseline estimation approach. The second panel presents the estimation results for the hybrid high-frequency-VAR estimation approach.

TABLE 7
Estimates of Structural Parameters

	Baseline	Baseline HF-VAR	Output-Inflation HF-VAR	Baseline + Habit HF-VAR
ξ_p	0.99 [0.49, 0.99]	0.95 [0.29,0.99]	0.93 [0.20,0.99]	0.99 [0.01,0.99]
ξ_w	0.90 [0.48, 0.99]	0.96 [0.08,0.99]	0.96 [0.05,0.99]	0.86 [0.07,0.99]
k_I	25.0 [0.69, 25.0]	3.89 [1.00,25.0]	14.06 [3.56,24.98]	2.66 [0.72,25.0]
ρ	0.96 [0.91, 0.99]	0.89 [0.36,0.93]	0.91 [0.27,0.94]	0.90 [0.63,0.93]
v	0.74 [0.01, 0.96]	0.01 [0.01,0.43]	0.01 [0.01,0.51]	0.02 [0.01,0.40]
b				0.94 [0.77,0.99]

The table reports our estimates of the structural parameters of the CEE/ACEL model that we estimate. We report 95% confidence intervals based on the bootstrap procedure described in the text in square brackets below the point estimate for each parameter.

TABLE 8
Response of Expected Output Growth

Exp. Output Growth in Current Qr	1.35 (1.59)
Exp. Output Growth 1 Qr Ahead	1.58 (0.61)
Exp. Output Growth 2 Qr Ahead	0.66 (0.34)
Exp. Output Growth 3 Qr Ahead	0.82 (0.26)
Exp. Output Growth 4 Qr Ahead	0.50 (0.30)
Exp. Output Growth 5 Qr Ahead	0.55 (0.27)
Exp. Output Growth 6 Qr Ahead	0.47 (0.30)
Exp. Output Growth 7 Qr Ahead	0.88 (0.66)

Each estimate comes from a separate OLS regression. We regress changes in survey expectations from the Blue Chip Economic Indicators on the policy news shock. Since the Blue Chip survey expectations are available at a monthly frequency, we construct a corresponding monthly measure of our policy news shock. In particular, we use any policy news shock that occurs over the month except for those that occur in the first week (because we do not know whether these occurred before or after the survey response). The dependent variable is the change in the forecasted value of output growth N quarters ahead, between this month's survey and last month's survey. The sample period is January 1995 to April 2014, except that we exclude the second half of 2008 and the first half of 2009. Standard errors are in parentheses.

TABLE 9
Response of Stock Prices

	Stock Prices
Response in the Data	-6.5 (3.9)
Response in the Model with News about:	
Monetary Policy Only	-23.4
Monetary Policy and Exogenous Economic Fundamentals	-12.8

TABLE A.1
Response of Interest Rates to the Policy News Shock for Different Sample Periods

	Baseline Sample		Pre-Crisis (2000-2007)		Full Sample		Baseline w/ Unsched.	
	Nominal	Real	Nominal	Real	Nominal	Real	Nominal	Real
3M Treasury Yield	0.67 (0.14)		0.76 (0.13)		0.60 (0.19)		0.76 (0.11)	
6M Treasury Yield	0.85 (0.11)		0.85 (0.12)		0.82 (0.14)		0.91 (0.10)	
1Y Treasury Yield	1.00 (0.14)		1.00 (0.14)		1.00 (0.15)		1.00 (0.13)	
2Y Treasury Yield	1.10 (0.33)	1.06 (0.24)	1.11 (0.36)	1.04 (0.24)	1.11 (0.27)	1.56 (0.31)	1.30 (0.37)	1.21 (0.26)
3Y Treasury Yield	1.06 (0.36)	1.02 (0.25)	1.03 (0.39)	0.97 (0.25)	1.14 (0.29)	1.38 (0.28)	1.26 (0.40)	1.18 (0.28)
5Y Treasury Yield	0.73 (0.20)	0.64 (0.15)	0.67 (0.20)	0.58 (0.15)	0.88 (0.20)	0.90 (0.19)	0.69 (0.16)	0.68 (0.11)
10Y Treasury Yield	0.38 (0.17)	0.44 (0.13)	0.36 (0.18)	0.44 (0.13)	0.59 (0.21)	0.67 (0.18)	0.38 (0.13)	0.43 (0.10)
2Y Tr. Inst. Forward Rate	1.14 (0.46)	0.99 (0.29)	1.07 (0.48)	0.90 (0.27)	1.25 (0.36)	0.97 (0.28)	1.34 (0.49)	1.15 (0.31)
3Y Tr. Inst. Forward Rate	0.82 (0.43)	0.88 (0.32)	0.66 (0.43)	0.76 (0.29)	1.12 (0.39)	1.07 (0.33)	1.00 (0.46)	1.03 (0.34)
5Y Tr. Inst. Forward Rate	0.26 (0.19)	0.47 (0.17)	0.20 (0.19)	0.47 (0.16)	0.55 (0.27)	0.75 (0.24)	0.27 (0.13)	0.45 (0.12)
10Y Tr. Inst. Forward Rate	-0.08 (0.18)	0.12 (0.12)	-0.01 (0.19)	0.21 (0.13)	0.11 (0.22)	0.21 (0.14)	-0.07 (0.12)	0.04 (0.11)

Each estimate comes from a separate OLS regression. The dependent variable in each regression is the one day change in the variable stated in the left-most column. The independent variable is a change in the policy news shock over a 30 minute window around regularly scheduled FOMC announcements, except the last two columns where we include unscheduled FOMC announcements. The baseline sample period is 1/1/2000 to 3/19/2014, except that we drop the second half of 2008 and the first half of 2009. The "Pre-Crisis" sample is 2000-2007. The "Full Sample" is 1/1/2000 to 3/19/2014. In all cases, we drop a 10 day period after 9/11/2001. For 2Y and 3Y yields and real forwards, the sample starts in 2004.

TABLE A.2
Response of Interest Rates and Inflation to the Policy News Shock
Rigobon's Heteroskedasticity-Based Estimator

	Nominal	Real	Inflation
3M Treasury Yield	0.69 (0.15)		
6M Treasury Yield	0.85 (0.12)		
1Y Treasury Yield	0.98 (0.15)		
2Y Treasury Yield	1.07 (0.37)	1.03 (0.29)	0.05 (0.21)
3Y Treasury Yield	1.03 (0.41)	0.99 (0.30)	0.04 (0.19)
5Y Treasury Yield	0.69 (0.21)	0.62 (0.16)	0.07 (0.12)
10Y Treasury Yield	0.34 (0.18)	0.42 (0.14)	-0.08 (0.09)
2Y Treasury Inst. Forward Rate	1.10 (0.53)	0.96 (0.34)	0.14 (0.25)
3Y Treasury Inst. Forward Rate	0.78 (0.49)	0.86 (0.38)	-0.08 (0.18)
5Y Treasury Inst. Forward Rate	0.22 (0.20)	0.46 (0.18)	-0.24 (0.09)
10Y Treasury Inst. Forward Rate	-0.12 (0.19)	0.11 (0.13)	-0.22 (0.10)

Each estimate comes from a separate "regression." The dependent variable in each regression is the one day change in the variable stated in the left-most column. The independent variable is a change in the policy news shock over a 30 minute window around the time of FOMC announcements. All results are based on Rigobon's (2003) method of identification by heteroskedasticity. The sample of "treatment" days is all regularly scheduled FOMC meeting day from 1/1/2000 to 3/19/2014. The sample of "control" days is all Tuesdays and Wednesdays that are not FOMC meeting days from 1/1/2000 to 12/31/2012. In both the treatment and control samples, we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. For 2Y and 3Y yields and real forwards, the sample starts in 2004. Standard errors are calculated using a non-parametric bootstrap with 5000 iterations.

TABLE A.3
Breakeven Inflation versus Inflation Swaps

	Breakeven	Swaps
Inflation Over Next 2 Years	-0.02 (0.18)	0.37 (0.35)
Inflation Over Next 3 Years	-0.03 (0.17)	0.41 (0.32)
Inflation Over Next 5 Years	-0.13 (0.14)	-0.02 (0.15)
Inflation Over Next 10 Years	-0.22 (0.12)	-0.17 (0.16)

Each estimate comes from a separate OLS regression. The dependent variable in each regression is the one day change in expected inflation measured either by breakeven inflation from the difference between nominal Treasuries and TIPS (first column) or from inflation swaps (second column) for the period stated in the left-most column. The independent variable is a change in the policy new shock over a 30 minute window around the time of FOMC announcements. The sample is all regularly scheduled FOMC meeting day from 1/1/2005 to 11/14/2012, except that we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001.




Figure 1: Interest Rate and Inflation in the Simple New Keynesian Model

The parameters used in this illustrative example are: $\beta = 0.99$, $\sigma = 0.5$, $\kappa\varsigma = 0.0017$, $\phi_\pi = 0.5$. We assume that $\bar{r}_t$ follows an AR(2) with roots $\rho_1 = 0.94$ and $\rho_2 = 0.70$.




Figure 2: Interest Rates and Inflation in the Data

The figure plots the response of instantaneous nominal and real forward rates and instantaneous break-even inflation to our policy news shock. These are the same estimates as are reported in Table 1.




Figure 3: The Response of Inflation and Interest Rates to the Policy News Shock in Our Estimation of CEE/ACEL Model




Figure 4: Scatterplot of Estimated Joint Distribution of ξ_w and ξ_p

The figure plots the values of ξ_w and ξ_p from the 500 bootstrap draws we calculate.







Figure 5: Responses of Interest Rates, Expected Inflation, and Expected Output when FOMC Announcements Convey Information about Both Monetary Policy and Exogenous Shocks



Panel A: 1-Day Window



Panel B: 30-Minute Window

Figure A.1: Scatterplots of Joint Sampling Distribution of $Dcov$ and $Dvar$ for 2-Year Nominal Forward Rate

Each point in the figure is a draw from our bootstrap. $Dvar$ denotes the difference in variance of our policy news shock between the treatment and control sample. $Dcov$ denotes the difference in the covariance of our policy news shocks and the 2-year nominal forward rate between the treatment and the control sample.







Figure A.2: Quantiles of the distribution of $g(\gamma)$ for different values of γ when estimating effect on the 2-year nominal forward rate using a 1-day window



A. Nominal Interest Rate



B. Real Interest Rate



C. Inflation

Figure A.3: Response of Inflation and Interest Rates to Policy News Shock in Our Estimation of CEE/ACEL Model Including 95% Confidence Intervals