

High Frequency Identification of Monetary Non-Neutrality

Emi Nakamura and Jón Steinsson*

Columbia University

December 20, 2013

Abstract

We provide new evidence on the responsiveness of real interest rates and inflation to monetary shocks. Our identifying assumption is that the increase in the volatility of interest rate news in a 30-minute window surrounding scheduled Federal Reserve announcements arises from news about monetary policy. Nominal and real interest rates respond roughly one-for-one several years out into the term structure at these times, implying that changes in expected inflation are small. At longer horizons, the response of expected inflation grows. Accounting for “background noise” in interest rates on FOMC days is crucial in identifying the effects of monetary policy on interest rates, particularly at longer horizons. We show that in conventional business cycle models with nominal rigidities our estimates imply that monetary non-neutrality is large. We also find evidence that FOMC announcements provide the public with information not only about monetary policy but also about the evolution of exogenous economic fundamentals.

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

JEL Classification: E30, E40, E50

*We thank Matthieu Bellon, Vlad Bouchouev, Nicolas Crouzet, Jesse Garret and Shaowen Luo, for excellent research assistance. We thank Michael Abrahams, Tobias Adrian, Richard K. Crump, Matthias Fleckenstein, Michael Fleming, Refet Gurkaynak, Hanno Lustig, Emanuel Moench, and Eric Swanson for generously sharing data and programs with us. We thank Marco Bassetto, 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. 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.

Early work using this approach assumed that no other shocks affect interest rates on FOMC announcement days. We weaken this identification assumption in two ways. First, we follow Gurkaynak, Sack, and Swanson (2005) in considering changes in interest rates over a 30-minute window around FOMC announcements (see also, Fleming and Piazzesi, 2005). Second, we allow for the possibility that other shocks affect interest rates even within this 30-minute window. To separate the effects of monetary shocks from other shocks at the time of FOMC announcements, we employ a heteroskedasticity-based estimator developed by Rigobon (2003) and Rigobon and Sack (2004). Our identifying assumption is that the increase in volatility of interest rates at the time of FOMC announcements is due to monetary news. In other words, we assume that what is special about the 30-minutes around FOMC announcements is that the volatility of monetary shocks rises, while the volatility of other shocks is the same as at other times.¹

¹Wright (2012) uses Rigobon's heteroskedasticity-based estimation approach to identify the effects of unconventional monetary policy on interest rates during the recent period over which short-term nominal interest rates have been at their zero lower bound.

We estimate the effect of monetary shocks on nominal interest rates, real interest rates and expected inflation. For this purpose, we use data on nominal Treasuries and Treasury Inflation Protected Securities (TIPS). 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 105 basis points, raises the 2-year real TIPS yield by 100 basis points. The effect of this shock on the 2-year instantaneous real forward rate is 86 basis points. The impact of the shock then falls monotonically at longer horizons to 72 basis points at 3 years, 39 basis points at 5 years, and 9 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.²

We can infer the response of market expectations about inflation by taking the difference between the response of nominal and real rates. At horizons of 2 and 3 years, the response of this “break-even” measure of inflation to our monetary shock is essentially zero. At longer horizons, the response of break-even inflation grows modestly and becomes significantly negative. Overall, our results thus indicate that monetary shocks that have large and persistent effects on real interest rates yield relatively small and very delayed effects on expected inflation.

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.³ To address this possibility, we study the effect of our monetary shocks on expected real rates using direct measures of expectations from *Blue Chip Economic Indicators*. While our estimates based on this approach are less precise than those based on asset prices, they support a similar time-pattern of effects on real interest rates and a small inflation response. We also consider the response of risk-adjusted, expected future nominal and real rates implied by the affine term structure model of Abrahams et al. (2013) to our monetary shock. The response of these risk-adjusted, interest rates are very similar to the response of the raw interest rates: large and persistent movements in risk-adjusted real interest rates and small movements in inflation. Furthermore, we

²Hanson and Stein (2012) employ a similar high-frequency approach to study the impact of monetary shocks on long-term real interest rates. Our results differ significantly from theirs in that their measure of monetary shocks has a substantial effect on real forwards even at the 10-year horizon. A key difference is that their monetary shock measure is the 2-day change in the 2-year nominal yield around FOMC days and they do not employ the heteroskedasticity-based estimation approach we employ to account for “background noise” in interest rates.

³A key point is that constant, or slowly moving, risk premia do not affect our results, since our identification is based on changes in bond yields at the time of FOMC announcements.

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.⁴

In the second half of the paper, we interpret this empirical evidence through the lens of New Keynesian business cycle models. In this analysis, we seek to answer two questions: 1) What structural parameters does our evidence provide information about? and 2) How much monetary non-neutrality does our evidence imply?

We start by using the textbook, three-equation, New Keynesian model to build intuition. We show that 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, there are two forces at play here. First, the Euler equation implies that an increase in real interest rates leads to a decrease in output. The strength of this force is governed by the intertemporal elasticity of substitution (IES). Second, the resulting decrease in output leads firms to reduce their prices, generating a fall in inflation. The strength of this force is governed by the extent of nominal and real rigidities. If the response of inflation to a monetary shock is small relative to the response of real interest rates, this implies that output does not respond much to real interest rates (small IES), or prices do not respond much to output (large nominal and real rigidities), or both.

The textbook New Keynesian model implies that inflation is purely forward looking. This means that the largest response of inflation should be immediately following the monetary shock, when all the high real interest rates are in the future. The response of inflation should then dissipate as the response of real interest rates dies out. In contrast to this, the response of inflation that we estimate in the data is initially close to zero and builds over time. This suggests a model with a substantial degree of inflation inertia—i.e., a lagged inflation term in the Phillips curve.

Building on this intuition, we quantitatively assess the degree of monetary non-neutrality implied by our evidence using the workhorse business cycle model proposed by Christiano, Eichenbaum and Evans (2005, CEE) and further developed by Altig et al. (2011, ACEL). We estimate key parameters of this model using a simulated method-of-moments approach. The moments we use in the estimation are the responses of nominal and real interest rates to our monetary shocks. This empirical approach is analogous to the impulse response matching approach used by Rotemberg and Woodford (1997), CEE, and ACEL. A key difference is that our empirical impulse responses are estimated using the

⁴Hanson and Stein (2012) present a behavioral model in which “search for yield” generates significant risk premium effects of monetary shocks that dissipate over time.

high-frequency identification approach described above, as opposed to a structural VAR.

Our estimates imply that monetary non-neutrality is large. Output responds about three times as much as inflation to a standard monetary shock for our estimates. This ratio is 2.6 for the parameters obtained by ACEL and 1.7 for the parameters obtained by CEE. On this metric, our estimates, thus, imply somewhat more monetary non-neutrality than does the VAR based evidence used by CEE and ACEL.

In the above discussion, we have made the conventional assumption that FOMC announcements change the private sector's beliefs about current and future monetary policy but do not provide the private sector with new information about current or future exogenous shocks such as productivity growth. In other words, we assume that the Fed does not have an informational advantage vis-à-vis the private sector. In section 5, we present evidence indicating that, in fact, FOMC announcements may convey such information. We show that private forecasts of output growth—from *Blue Chip Economic Indicators*—increase in response to FOMC announcements that raise nominal and real interest rates.⁵

If FOMC announcements only convey information about monetary policy, output growth should fall when the FOMC surprises the market by raising interest rates. However, if surprise increases of interest rates by the FOMC are interpreted by the private sector as indications that the FOMC is more optimistic about future economic developments than the private sector had previously thought, this may lead the private sector to update its own beliefs about future economic developments. We show how this notion can be captured parsimoniously within the New Keynesian model by assuming that FOMC announcements convey information to the private sector about the future evolution of the natural rate of interest. We then recalibrate our model to match the response of expected output growth in addition to the responses of interest rates and expected inflation. This alternative calibration implies somewhat less monetary non-neutrality than the parameter estimates obtained under the conventional assumption that FOMC announcements only contain information about current and future monetary policy. Nevertheless, even in this calibration, the degree of monetary non-neutrality is large.

In recent related work, Gertler and Karadi (2013) combine high frequency identification and VAR methods to identify the effects of monetary shocks on macro variables and credit spreads. They find

⁵See Romer and Romer (2000) for additional evidence that monetary policy shocks convey information about current and future exogenous shocks.

that monetary policy shocks have large effects on credit spreads and argue that it is important to incorporate financial frictions in macro models to understand the effects of monetary policy on the economy. 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.

The paper proceeds as follows. Section 2 describes the data we use in our analysis. Section 3 describes the construction of our policy news shock and presents our main empirical results regarding the response nominal and real interest rates and inflation to the policy news shock. 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. Section 5 presents our evidence on the response of output growth to our monetary shocks and shows that the New Keynesian model can match this additional piece of evidence if Fed announcement are interpreted as conveying information not only about monetary policy but also about the evolution of exogenous economic fundamentals. Section 6 concludes.

2 Data

We use data on interest rates from several sources. First, 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 Mercantile Exchange). Fed Funds futures have been traded since 1988, while Eurodollar futures began trading in the early 1980's. 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. The effective Federal Funds Rate is the rate that is quoted by the Federal Reserve Bank of New York on every business day. The Fed Funds future can, thus, be used to construct market based expectations of the average Fed Funds rate over the month in question.⁶

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

⁶See 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.

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.⁷

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.⁸ 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 2000 to avoid relying on data from the period when TIPS liquidity was limited.

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 use data on the level of the S&P500 stock price index obtained from *Yahoo Finance*. We

⁷See 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.

⁸This holds unless cumulative inflation is negative, in which case no adjustment is made for the principle payment.

use data on inflation swaps from *Bloomberg*. We use data on expectations of future nominal interest rates and inflation 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. Finally, 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 (2013).

3 Empirical Analysis

Our goal in this section is to identify the effect of the monetary news contained in scheduled FOMC announcements on nominal and real interest rates and expected inflation. Our identification approach makes use of the discontinuous increase in the volatility of monetary shocks at the time of FOMC announcements. We therefore consider changes in interest rates in a narrow window around FOMC announcements. We consider two time intervals. The first is a 30-minute window from 10 minutes before the FOMC announcement to 20 minutes after it. The second is a 1-day window from the close of markets the day before the FOMC meeting to the close of markets the day of the FOMC meeting.

The increased use of “forward guidance” has been an important development in the conduct of monetary policy over the past 15 years (Gurkaynak, Sack, and Swanson, 2005; Campbell et al., 2012). Forward guidance refers to the fact that, in their post-meeting announcements, the FOMC conveys information not only about immediate changes in the Federal Funds Rate but also about likely changes in monetary policy at later dates. In fact, over the last 15 years, changes in the Federal Funds Rate have often been largely anticipated by markets once they occur, while FOMC announcements have come to focus more and more on guiding expectations about future changes in the Federal Funds Rate. Motivated by these developments, we construct a measure of monetary policy news Δi_t by taking the first principle component of changes in five interest rates of maturity less than one year which can be inferred from futures data. We use Federal Funds futures and Eurodollar futures to infer changes in the market’s expectations about the Federal Funds rate immediately following the FOMC meeting, the Federal Funds rate immediately following the next FOMC meeting, and the 3-month Eurodollar interest rate at horizons of two, three and four quarters.⁹ We refer to Δi_t

⁹More 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.

as the “policy news shock.”¹⁰ The scale of the policy news shock is arbitrary. For convenience, we rescale it such that an OLS regression of the 1-year Treasury yield on the policy news shock yields a coefficient of one. Appendix A provides details about the construction of the policy news shock.¹¹

3.1 Identification

If we were confident that movements in the policy news shock Δi_t over the windows of time we consider around FOMC announcements were due to monetary shocks and nothing else, then this variable would constitute a pure measure of monetary shocks. We could thus regress any other variable of interest on the policy news shock to assess the effect of monetary shocks on that variable. This is the approach taken by Cook and Hahn (1989), Kuttner (2001) and Cochrane and Piazzesi (2002) (all with a one-day window) and more recently by Hanson and Stein (2012) (with a two-day window). A potential concern with this approach is that other shocks may occur over the course of FOMC days. Interest rates fluctuate substantially on non-FOMC days. This suggests that other shocks than FOMC announcements affect interest rates on FOMC days. There is no way of knowing whether these other shocks are monetary shocks or non-monetary shocks.

We would, therefore, like to allow for “background noise” in interest rates on both FOMC and non-FOMC announcement days. To this end we adopt a heteroskedasticity-based estimator of monetary shocks developed by Rigobon (2003) and Rigobon and Sack (2004). Let ϵ_t denote a pure monetary shock and suppose that movements in the policy news shock we measure in the data is governed both by monetary and non-monetary shocks:

$$\Delta i_t = \alpha_i + \epsilon_t + \beta_i \eta_t, \tag{1}$$

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

¹⁰Our policy news shock variable is closely related to the “path factor” considered by Gurkaynak, Sack, and Swanson (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.

¹¹The construction of the policy news shock uses changes in Fed Funds futures and Eurodollar futures to measure changes in market expectations about future Federal Funds rates. Piazzesi 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 for these short term interest rates.

shocks:

$$\Delta s_t = \alpha_s + \gamma \epsilon_t + \beta_s \eta_t. \quad (2)$$

Our objective is to estimate γ , which should be interpreted as the impact of the pure monetary shock ϵ_t on Δs_t relative to its effect 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 $\sigma_{\eta,R1} = \sigma_{\eta,R2}$.

We show in Appendix B that given these assumptions γ is given by

$$\gamma = \frac{\text{cov}_{R1}(\Delta i_t, \Delta s_t) - \text{cov}_{R2}(\Delta i_t, \Delta s_t)}{\text{var}_{R1}(\Delta i_t) - \text{var}_{R2}(\Delta i_t)}. \quad (3)$$

Notice that if we set the variance of the “background noise” η_t to zero, then this estimator 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 .

As we discuss above, we present results where the policy news shock is constructed using 30-minute and 1-day time intervals surrounding FOMC announcements. Our control samples are then 30-minute or 1-day intervals that are chosen to be as comparable as possible except that they do not include FOMC announcements. Specifically, in the case of 30-minute windows, we choose the same 30-minute window (from 2:05pm to 2:35pm) on all non-FOMC Tuesdays and Wednesdays as our control sample (since scheduled FOMC meetings tend to occur on Tuesdays and Wednesdays), and in the case of 1-day windows, we choose all non-FOMC Tuesdays and Wednesdays as our control sample.¹² For our treatment sample, we focus on only scheduled FOMC meetings, since unscheduled meetings may occur in reaction to other shocks and thus be endogenous. In all cases, the outcome variables are measured over a 1-day window. Our sample period starts on January 1st 2000 and extends to January 25th 2012. We drop data before 2000 because of concerns about liquidity of TIPS and because very few TIPS securities were trading at the time. In our baseline analysis, we drop

¹²For the case of 30-minute windows, we have also tried using a 30-minute window one hour before FOMC announcements on FOMC days as our control sample. This yields very similar results.

the second half of 2008 and the first half of 2009 to avoid the period when disruption of financial markets in the Great Recession was most severe.

3.2 Main Estimates

Table 1 presents our baseline estimates of monetary shocks on nominal and real interest rates and inflation. The first column presents the effects of the policy news shock on nominal Treasury interest rates. By construction, the policy news shock has large effects on nominal yields. The effect of our policy news shock on the zero-coupon 2-year Treasury yield is 105 basis points, and declines monotonically to 29 basis points at 10 years. 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. A 2-year instantaneous forward rate (say) is the 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 maturities. For maturities of 2, 3, 5, and 10 years, its effects on forward rates are 100, 60, 13 and -13 basis points, respectively. We show below that the negative effect on long-horizon nominal interest rates reflects a decline in long-horizon inflation expectations.

The second column of Table 1 presents the effects of the policy news shock on real interest rates measured using TIPS. While the effects on nominal rates are by construction, the impact of monetary shocks on real interest rates is not. In neoclassical models of the economy, the Fed controls the nominal interest rate but has no impact on real interest rates. Our estimate of the impact of our policy news shock on the 2-year real yield is 100 basis points, and the impact on the 3-year real yield is 94 basis points. Once again, the time-path of effects is easier to interpret using evidence on instantaneous forward rates. The effect of the shock on the 2-year real forward rate is 86 basis points. It falls monotonically at longer horizons to 72 basis points at 3 years, 39 basis points at 5 years, and 9 basis point at 10 years (which is not statistically significantly different from zero). 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 expected inflation as measured by the break-even difference between Treasury rates and TIPS rates. The first several rows provide estimates based on bond yields, which indicate that the response of expected 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 inflation rates to get estimates of expected inflation at points in time in the future. The response of expected 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 -12, -27 and -22 basis points. It is only the responses at 5 and 10 years that are statistically significantly different from zero. Our evidence thus points to expected inflation responding quite gradually to monetary shocks that have a substantial effect on real interest rates. In section 4 below, we discuss what we can infer about the structure of the economy from these estimates.

Much of the earlier literature that uses high frequency identification to estimate the effect of monetary shocks, focuses on the impact of FOMC announcements on market expectations about the level of the Federal Funds Rate immediately following the announcement (e.g., Kuttner, 2001). The disadvantage of this approach is that it captures less of the variation in interest rates in response to monetary shocks than the policy news shock we construct. The remaining columns of Table 1, nevertheless, present estimates based on this approach. The conclusions are very similar. Nominal and real rates respond by roughly the same amount at horizons out to about 3 years. At longer horizons, the response of nominal rates is smaller than real rates, implying that inflation falls.¹³

3.3 Alternative Estimates

Table 2 compares our baseline methodology to alternative methods of identifying the monetary policy shock. The top panel presents results based on the Rigobon estimator, while the bottom panel reports results based on OLS. The policy news shock is measured over a 30-minute window in the first two columns, but a one-day window in the middle two columns. The last two columns present results where changes in a longer-term interest rate—the two-year nominal yield—are used as the monetary policy shock. In Table 2, we assess statistical significance based on confidence intervals that are constructed using a more sophisticated procedure than we use in our baseline results. We describe the details and motivation for this procedure later in this section.

Most of the prior literature has used OLS and a 1-day window. This procedure implicitly makes

¹³Beechey and Wright (2009) analyze the effect of unexpected movements in the 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.

the strong identifying assumption that *only* monetary shocks occur on the day of an FOMC announcement. To assess this assumption, the middle two columns of Table 2 compare OLS (bottom panel) and the Rigobon estimator (top panel) when a 1-day window is used. The differences are substantial. While the OLS confidence intervals are quite moderate, the Rigobon confidence intervals are much wider: all the nominal forwards are statistically insignificant, as is the 5-year real forward. This indicates that there is a large amount of “background” noise in interest rates over an entire day. Clearly, the approach of using OLS with a 1-day window massively overstates the true statistical precision of the estimates.

These concerns loom even larger when longer-term interest rates are used as proxies for monetary shocks. Columns 5 and 6 present results where the monetary shock measure Δi_t is constructed as the one-day change in the two-year nominal yield. The confidence intervals are much larger using the Rigobon estimator than OLS. In fact, in some cases, the 95% confidence intervals for the Rigobon 1-day window estimator are infinite. We therefore report 90% confidence intervals. These results arise because of the large amount of background noise in longer-term interest rates.¹⁴ The increase in the volatility of longer-term interest rates associated with FOMC announcements is not large enough over a one-day horizon to accurately assess the impact of monetary shocks on these variables.

In contrast, when the policy news shock is measured over a 30-minute window, the difference between OLS and the Rigobon estimator is small, both for the point estimates and the confidence intervals. This reflects the fact that there is little background noise in interest rates over a 30-minute window. In this case, the OLS identifying assumption—that there is no background noise in interest rates—yields confidence intervals that are close to correct.

As we noted above, the confidence intervals in Table 2 are constructed using a more sophisticated procedure than we use in our baseline results. The reason is that the conventional bootstrap approach to constructing standard errors yields inaccurate confidence intervals 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.¹⁵ Figure 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

¹⁴Conversely, these concerns about background noise are not as important for shocks to the current Federal Funds rate (see, e.g., Rigobon and Sack, 2004).

¹⁵Recall that the Rigobon estimator—equation (3)—is a ratio with the difference in the variance of Δi_t between the treatment sample and the control sample in the denominator. If this difference is small, the estimator yields very large values (positive or negative depending on the whether the difference in variance is positive or negative).

Sack (2004) show that the estimator in equation (3) 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 standard errors.

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

$$g(\gamma) = \Delta \text{cov}(\Delta i_t, \Delta s_t) - \gamma \Delta \text{var}(\Delta i_t), \quad (4)$$

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 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 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.¹⁶

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 1. In fact, our baseline bootstrap approach slightly understates the statistical significance of the results relative to the more sophisticated procedure. 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.

The analysis in Tables 1 and 2 is for the sample period from Jan 1st 2000 to Jan 25th 2012, except that we drop the period spanning the height of the financial crisis in the second half of 2008

¹⁶We thank Sophocles Mavroeidis for suggesting this approach to us.

and the first half of 2009. 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. We have, however, also carried out our analysis on the full sample including the crisis, as well as a more restrictive data sample ending at the beginning of 2008. In addition to this, we have carried out our analysis for our baseline sample but including unscheduled FOMC meetings. Table A.1 presents the results of our analysis for these three alternative samples. The pre-crisis sample and the sample including unscheduled FOMC meetings yield very similar results to the baseline sample. For the full sample the response of both nominal and real rates is somewhat larger at longer horizons. In all three cases, the effect of the monetary shock on inflation is initially small and positive, but becomes increasingly negative at longer horizons.

3.4 Risk Premia 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. In this regard, it is important to keep in mind that constant, or slowly moving, risk premia will not affect our results, since our identification is based on changes in bond yields at the time of FOMC announcements. However, if risk premia change at the time of FOMC announcements this could confound our results.

We consider three pieces of evidence on this point: 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 an 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.

Blue Chip surveys professional forecasters on their beliefs about macroeconomic variables over the next two years in the first few days of every month. We study the impact of our policy news shock on survey expectations about future short-term nominal interest rates and inflation. By construction, these effects reflect expected movements in rates, as opposed to risk premium effects. We measure the change in expected interest rates for a particular quarter in the future by the change in the *Blue Chip* forecast about that quarter from one month to the next. We regress this measure

on the sum of the policy news shocks that occur 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). We use *Blue Chip* forecasts of the 3-month T-Bill rate and the GDP deflator in our analysis. We construct a measure of the expected short-term real interest rate in a particular quarter by taking the difference between the expected 3-month T-bill rate and the expected GDP deflator for that quarter. Unfortunately, *Blue Chip* asks respondents only about the current and subsequent calendar year, so fewer observations are available for longer-term expectations, leading to larger standard errors.¹⁷ The 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 of 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.

Abrahams et al. (2013) employ an affine term structure model to decompose movements in the nominal and real term structures into movements in risk-neutral expected future short rates and risk premia. In Table 4, we study the effect of our policy news shock on risk-neutral expected future short rates and risk premia using their decomposition.¹⁸ The response of the model-implied risk-neutral interest rates is very similar to the response of the raw interest rates in our baseline results. As in our baseline results, the effect on real rates is large, while the effect on expected inflation is small. In fact, the effect on expected inflation is even smaller in Table 4 than in our baseline results.

Finally, Table 5 presents the effects of our policy news shock on nominal and real interest rates over event windows of 5, 10, 20, 60, 125, and 250 trading days. While the estimates are extremely noisy, there is little evidence that the effects on interest rates tend to dissipate over time, as some theories of liquidity premia might predict. Indeed, in most cases, after a dip around 10 days the point estimates appear to grow over time (though, again, the standard errors are extremely large).

¹⁷For example, in the last quarter of the year, forecasters are only asked about their beliefs 1-year in advance; while in the first quarter they are asked about their beliefs for the next full 2-years.

¹⁸What we refer to as the risk premia here is the difference between the raw interest rate response and the model-implied risk neutral interest rate. Abrahams et al. (2013) further decompose this into a term premium, a liquidity premium, and a model error term.

3.5 Inflation Swaps

We also consider an alternative market-based measure of inflation expectations based on inflation swap data.¹⁹ Table 6 compares our estimates of the effects of the policy news shock on breakeven inflation from TIPS to that on inflation from inflation swaps. The sample period for this analysis is limited by the availability of swaps data to beginning in January 1st 2005. 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 point estimates suggest a somewhat larger “price puzzle”—i.e., positive inflation response—at shorter horizons, though statistically insignificant.

4 Evidence on Monetary Non-Neutrality

To more clearly interpret our evidence on monetary non-neutrality, we follow in the tradition of work by Rotemberg and Woodford (1997), Christiano, Eichenbaum, and Evans (2005), and others who fit structural models of monetary policy to evidence on the response of real variables to monetary shocks. Unlike this earlier work, we focus on fitting the response of the real interest rate and inflation to monetary shocks. The key advantage of looking at these variables is that we are able to use high-frequency data to obtain estimates of the effects of monetary shocks.

We begin by developing intuition for what parameters of the New Keynesian model can be identified using our evidence. We do this in the context of a textbook, three-equation, New Keynesian model. We then analyze the quantitative implications of our empirical results for monetary non-neutrality in the workhorse medium-scale business cycle model proposed by Christiano, Eichenbaum, and Evans (2005) and developed further by Altig et al. (2011). In this section, we take the conventional view that FOMC announcements convey information only about monetary policy. In section 5, we turn our attention to the possibility that our monetary policy shocks may provide information about both monetary policy and current and future exogenous shocks such as productivity growth.

¹⁹An 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 (2013) for an analysis of the differences between break-even inflation from TIPS and inflation swaps.

4.1 Intuition in a Simple New Keynesian Model

4.1.1 Private Sector Behavior

Consider a setting in which private sector behavior can be described by the following Euler equation and Phillips curve:

$$\hat{x}_t = E_t \hat{x}_{t+1} - \sigma(\hat{\lambda}_t - E_t \hat{\pi}_{t+1} - \hat{r}_t^n), \quad (5)$$

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa \zeta \hat{x}_t. \quad (6)$$

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, $\hat{\lambda}_t$ 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 C 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 (5) is common to both Real Business Cycle and New Keynesian models, and describes how household’s 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

We specify monetary policy as a simplified Taylor rule of the form

$$\hat{\lambda}_t - E_t \hat{\pi}_{t+1} = \bar{r}_t + \phi_\pi \hat{\pi}_t. \quad (7)$$

The first term in the rule is a time varying intercept term. We can think of the monetary authority as using this term to track variation in the natural rate of interest r_t^n . The second term is a conventional endogenous feedback term implying that the monetary authority raises the real interest rate as

inflation increases. If the monetary authority is successful at varying \bar{r}_t so that it tracks r_t^n , inflation will be stable at zero and the endogenous feedback term will not come into play. In this section, we view FOMC announcements as conveying information about the future path of \bar{r}_t .

4.1.3 What Our Evidence Identifies

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 (5) forward and get that the response of the output gap to a monetary shock is,

$$\hat{x}_t = -\sigma \sum_{j=0}^{\infty} E_t \hat{r}_{t+j} = -\sigma \hat{r}_t^\ell. \quad (8)$$

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

Similarly, we can solve forward the Phillips curve—equation (6)—and get that the response of inflation to a monetary shock is

$$\hat{\pi}_t = \kappa \zeta \sum_{j=0}^{\infty} \beta^j E_t \hat{x}_{t+j}. \quad (9)$$

Combining equations (8) and (9), we get a relationship between the response of inflation and the real interest rates:

$$\hat{\pi}_t = -\kappa \zeta \sigma \sum_{j=0}^{\infty} \beta^j E_t \hat{r}_{t+j}^\ell. \quad (10)$$

More generally, the monetary authority may act in such a way that the long-run inflation rate changes. In this case, equation (10) becomes

$$\hat{\pi}_t = -\kappa \zeta \sigma \sum_{j=0}^{\infty} \beta^j E_t \hat{r}_{t+j}^\ell + \hat{\pi}_\infty, \quad (11)$$

where $\hat{\pi}_\infty$ denotes the change in the long-run inflation rate.

We wish to draw two main conclusions from equation (11). 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 $\kappa \zeta \sigma$. In other words, such a pattern of responses implies 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 (11) 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 (11) as time passes).²⁰ This effect is illustrated in Figure 3 for particular values of the structural parameters. Figure 4 presents our estimated response of 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 (11) 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.²¹

4.2 Estimating the CEE/ACEL Model with High Frequency Data

The simple, three-equation, New Keynesian model is useful for providing intuition. However, it abstracts from many features that have been shown to be important in generating realistic business cycles. It is therefore not well suited for quantitative analysis. We next investigate the quantitative implications of our empirical evidence by estimating 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).

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.

²⁰ 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.

²¹ Notice, that this argument holds even if the monetary shock leads to a shift in the long-run inflation target of the central bank (i.e., a non-zero $\hat{\pi}_\infty$ in equation (11)).

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 standard errors by bootstrapping. Our bootstrap procedure is to re-sample the data with replacement, estimate the empirical moments using the Rigobon method 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 1000 times. Importantly, this procedure for constructing the standard errors 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 of output to real rates and the response of inflation to output—but not all aspects.

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

$$\hat{\lambda}_t = \hat{\lambda}_{t-1} + \frac{1}{k_I} \sum_{j=0}^{\infty} \beta^j E_{t-1} \hat{p}_{k,t+j}, \quad (12)$$

where $\hat{\lambda}_t$ denotes investment and $\hat{p}_{k,t}$ is the shadow value of a unit of installed capital. From equation (12), 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 only change we make to the CEE/ACEL model is that we replace the monetary policy rule in that model with the monetary policy rule we discuss above—equation (7). To be able to capture the dynamics of the response of interest rates to our policy news shock, we assume that the monetary policy shock \bar{r}_t follows an AR(2) process. We fix $\phi_\pi = 0.5$ but estimate the autoregressive roots of the monetary policy shock process, which we denote ρ_1 and ρ_2 .

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. One way to measure the degree of monetary non-neutrality is as the sum of deviations of output from steady state after a standard monetary shock ($\sum_{j=0}^{\infty} |\hat{y}_{t+j}|$). Here, we consider the monetary shock used in CEE. The first column of Table 7 presents the value of this statistic for our estimation of the CEE/ACEL model as well as CEE and ACEL’s original estimates. According to this metric, our estimates of the CEE/ACEL model imply a roughly similar degree of monetary non-neutrality to CEE and ACEL’s estimates. However, the second column of Table 7 shows that the cumulative inflation response to the monetary shock ($\sum_{j=0}^{\infty} |\hat{\pi}_{t+j}|$) is smaller for our estimates than for CEE and ACEL’s estimates.

An alternative estimate of monetary non-neutrality is the ratio of the cumulative response of output to the cumulative response of inflation. The last column of Table 7 shows that our estimates imply that output responds about three times as much as inflation to the monetary shock, while it responds 2.6 times as much for the parameters obtained by ACEL and 1.7 times as much for the parameters obtained by CEE. On this metric, our estimates, thus, imply somewhat more monetary non-neutrality than CEE and ACEL’s estimates. The relative response of output to inflation ob-

tained in CEE is close to the lower end of the confidence interval for our estimates, implying that models with less monetary non-neutrality are inconsistent with our evidence.

Figure 5 presents the response of nominal and real interest rates and inflation to our monetary policy shock. Comparing these responses to those in Figure 4, we see that the model fits the data quite well. The response of inflation is very small initially and then gradually increases. The response of nominal and real interest rates is close to identical out to about 3 years. At longer horizons, the response of nominal interest rates falls below the response of real interest rates.

Figure 6 presents the response of nominal and real interest rates and inflation for CEE’s estimates of the CEE/ACEL model in response to our monetary policy shock. These responses stand in stark contrast with the responses we estimate in the data. In particular, the response of inflation is much larger than in the data, and, as a consequence, the response of nominal interest rates largely track the response of inflation rather than largely tracking the response of real interest rates as they do in the data.

Table 8 presents our individual parameter estimates. We estimate $\rho_1 = 0.93$ and $\rho_2 = 0.62$. This matches 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 identified—but these differences are not large enough to yield sharp inference for each parameter separately. This is illustrated in Figure 7, which presents a scatterplot 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 . Our 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.²² 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 1.1. 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.²³

²²Our 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 (2013) 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 (the equivalent of ξ_p in our model) vary substantially across specifications and are symmetrically dispersed around a value of zero.

²³This lines up well with existing micro-evidence. Using variation in the price of capital associated with tax changes,

5 The Information Content of Fed Announcements

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.²⁴ 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.²⁵

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

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

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 C, 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—will affect the natural rate of interest.

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 (9). 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.

Cummins, Hubbard, and Hassett (1994) estimate an elasticity of investment with respect to a permanent change in the price of capital of 6.6.

²⁴The FOMC has some advance knowledge of industrial production data since the Federal Reserve produces these data.

²⁵Faust, 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.

Table 9 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.

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.²⁶ For simplicity, we do this within the context of the textbook model 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:

$$\hat{\pi}_t = \hat{\pi}_{t-1} + \kappa \sum_{j=0}^{\infty} \beta^j E_t \hat{m}c_{t+j}, \quad (14)$$

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:

$$\hat{x}_t = b\hat{x}_{t-1} - (1-b)\sigma \sum_{j=0}^{\infty} E_t (\hat{\imath}_{t+j} - \hat{\pi}_{t+j+1} - \hat{r}_{t+j}^n). \quad (15)$$

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 their expectations 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., $\Delta E_t \hat{r}_{t+j}^n = \psi \Delta E_t \bar{r}_{t+j}$.²⁷ Intuitively, rather than assuming

²⁶This 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).

²⁷Here ΔE_t denotes the change in expectations in the 30-minute window around the FOMC announcement.

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. 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 the evidence on expected output from Table 9 and the evidence on interest rates and expected inflation from Table 1. As in earlier sections, we consider a monetary shock with AR(2) dynamics and set the autoregressive roots to $\rho_1 = 0.93$ and $\rho_2 = 0.62$ (the values estimated in section 4.2). 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 elasticity of substitution 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 8. 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,

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.

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 ($\alpha = 0.75$).

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 $\theta = 350$ is needed to roughly match the response of inflation in Figure 8. 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 10 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 7.3% 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 4.2%.

Table 10 also presents the response of stock prices to our monetary policy shock in the model. In the calibration of our model where monetary policy conveys information about both monetary policy and exogenous economic fundamentals, stock prices fall by 10.9% in response to the FOMC announcement. If monetary policy only conveys information about monetary policy, stock prices fall by 21.2%. 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.

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.

6 Conclusion

We use the fact that a disproportionate amount of monetary news gets released at the time of regularly scheduled FOMC meetings to estimate the effects of monetary shocks on nominal and real interest rates and expected inflation. The response of nominal and real interest rates to these monetary shocks is close to one-for-one several years out into the term structure and the response of expected inflation is delayed and small relative to the response of real interest rates. Evidence from survey data and an affine term structure model indicate that this effect does not arise from movements in risk premia at the time of FOMC announcements.

Using a textbook New Keynesian model, we show that the small response of inflation relative to the response of real interest rates requires a small elasticity of output with respect to the real interest rate, a large amount of nominal and real rigidities, or both. Also, the delayed response of inflation implies a substantial degree of inflation inertia. We quantitatively assess the degree of monetary non-neutrality implied by our evidence using the workhorse monetary business cycle model of Christiano, Eichenbaum, and Evans (2005) and Altig et al. (2011). Despite the short time-period over which real interest rate data are available, and the resulting large standard errors on the interest rate responses, our analysis yields strong conclusions about monetary non-neutrality. Our evidence implies a somewhat larger degree of monetary non-neutrality than conventional VAR-based evidence.

We next consider the possibility that FOMC announcements provide information not only about future monetary policy but also about future exogenous shocks such as productivity. The Fed employs a legion of talented economists to interpret incoming data about the economy. So, when the Fed raises interest rates, there may be reason for optimism about the current and future state of the economy. This view is supported by the fact that survey expectations of output growth actually increase in response to a surprise increase in interest rates by the FOMC. We show how a simple New Keynesian model can match this if the FOMC announcement is assumed to affect private sector expectations about the natural rate of interest. In this case, we are able to match the small response of inflation to monetary shocks using a degree of real rigidity that is smaller than before but nevertheless substantial.

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 from the change in the price of the current month's Federal Funds Rate futures contract, 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_0 denote the average Federal Funds Rate during the month up until the point of the FOMC announcement and r_1 the average Federal Funds Rate for the remainder of the month. Then

$$\begin{aligned} f_{t-\Delta t}^1 &= \frac{d_0}{m_0} r_{-1} + \frac{m_0 - d_0}{m_0} E_{t-\Delta t} r_0, \\ f_t^1 &= \frac{d_0}{m_0} r_{-1} + \frac{m_0 - d_0}{m_0} E_t r_0. \end{aligned}$$

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 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 then 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 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 simply the change in the price of the Eurodollar futures at the time of the FOMC announcements.

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, 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).

B Derivation of Our Heteroskedasticity-Based Estimator

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,

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

C 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.

C.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

$$E_0 \sum_{t=0}^{\infty} \beta^t [u(C_t, \xi_t) - v(L_t(x), \xi_t)], \quad (16)$$

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 (16) is an index given by

$$C_t = \left[\int_0^1 c_t(z)^{\frac{\theta-1}{\theta}} dz \right]^{\frac{\theta}{\theta-1}}, \quad (17)$$

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

$$P_t C_t + E_t[M_{t,t+1} B_{t+1}(x)] \leq B_t(x) + W_t(x) L_t(x) + \int_0^1 \Xi_t(z) dz - T_t, \quad (18)$$

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 ,³¹ $W_t(x)$ 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

³¹The 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.

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:

$$\frac{u_c(C_{t+j}, \xi_{t+j})}{u_c(C_t, \xi_t)} = \frac{M_{t,t+j} P_{t+j}}{\beta^j P_t} \quad (19)$$

as well as a standard transversality condition. Subscripts on the function u denote partial derivatives. Equation (19) 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:

$$\frac{v_\ell(L_t(x), \xi_t)}{u_c(C_t, \xi_t)} = \frac{W_t(x)}{P_t}. \quad (20)$$

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:

$$c_t(z) = C_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta}, \quad (21)$$

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

$$P_t = \left[\int_0^1 p_t(z)^{1-\theta} dz \right]^{\frac{1}{1-\theta}}. \quad (22)$$

C.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 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

$$y_t(z) = A_t f(L_t(z)), \quad (23)$$

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,

$$E_t \sum_{j=0}^{\infty} M_{t,t+j} [p_{t+j}(z)y_{t+j}(z) - W_{t+j}(x)L_{t+j}(z)]. \quad (24)$$

Firm z must satisfy demand for its product given by equation (21). Firm z is therefore subject to the following constraint:

$$C_t \left(\frac{p_t(z)}{P_t} \right)^{-\theta} \leq A_t f(L_t(z)). \quad (25)$$

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

$$W_t(x) = A_t f_\ell(L_t(z)) S_t(z), \quad (26)$$

where $S_t(z)$ denotes the firm's nominal marginal cost (the Lagrange multiplier on equation (25) 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

$$p_t(z) = \frac{\theta}{\theta - 1} E_t \sum_{j=0}^{\infty} \frac{\alpha^j M_{t,t+j} y_{t+j}(z)}{\sum_{k=0}^{\infty} \alpha^k M_{t,t+k} y_{t+k}(z)} S_{t+j}(z). \quad (27)$$

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

C.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 (19) to plug in for $M_{t,t+1}$ and rearranging terms yields

$$E_t \left[\beta U_c(C_{t+1}, \xi_{t+1}) \frac{P_t}{P_{t+1}} \right] = \frac{U_c(C_t, \xi_t)}{1 + i_t}. \quad (28)$$

The zero-growth, zero-inflation steady state of this equation is $\beta(1 + \bar{i})$. A first order Taylor series approximation of equation (28) is

$$\hat{c}_t = E_t \hat{c}_{t+1} - \sigma(\hat{i}_t - E_t \hat{\pi}_{t+1}) - \sigma E_t \Delta \hat{\xi}_{ct+1}, \quad (29)$$

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 (26)—in period $t + j$ for industry x that was last able to change its prices in period t is then

$$\hat{w}_{t,t+j}(x) = \hat{a}_{t+j} - (1 - a)\hat{\ell}_{t,t+j}(x) + \hat{s}_{t,t+j}(x), \quad (30)$$

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 (20)—in period $t + j$ for industry x that was last able to change its prices in period t is

$$\hat{w}_{t,t+j}(x) = \eta^{-1}\hat{\ell}_{t,t+j}(x) + \sigma^{-1}\hat{c}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j}, \quad (31)$$

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 (23)—in period $t + j$ for industry x that was last able to change its prices in period t is

$$\hat{y}_{t,t+j}(x) = \hat{a}_{t+j} + a\hat{\ell}_{t,t+j}(x). \quad (32)$$

Combining labor demand and labor supply—equations (30) and (31)—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 (32)—to eliminate $\hat{\ell}_{t,t+j}(x)$ yields

$$\hat{s}_{t,t+j}(x) = \omega\hat{y}_{t,t+j}(x) + \sigma^{-1}\hat{c}_{t+j} - (\omega + 1)\hat{a}_{t+j} + \hat{\xi}_{\ell,t+j} - \hat{\xi}_{c,t+j}, \quad (33)$$

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

Taking logs of consumer demand—equation (21)—in period $t + j$ for industry x what was last able to change its prices in period t yields

$$\hat{y}_{t,t+j}(z) = -\theta\hat{p}_t(x) + \theta \sum_{k=1}^j \hat{\pi}_{t+k} + \hat{y}_{t+j}, \quad (34)$$

where we use the fact that $Y_t = C_t$ and $y_t(x) = c_t(x)$. Plugging this equation into equation (33) and again using the fact that $Y_t = C_t$ yields

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

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 (35) is then

$$(\omega + \sigma^{-1})\hat{y}_t^n = (\omega + 1)\hat{a}_t - \hat{\xi}_{\ell,t} + \hat{\xi}_{c,t}, \quad (36)$$

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 (35) and (36) yields

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

We next linearize the price setting equation—equation (27). 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

$$\hat{p}_t(x) = (1 - \alpha\beta)\sum_{j=0}^{\infty}(\alpha\beta)^j E_t\hat{s}_{t,t+j}(x) + \alpha\beta\sum_{j=1}^{\infty}(\alpha\beta)^j E_t\hat{\pi}_{t+j}. \quad (38)$$

Using equation (37) to eliminate $\hat{s}_{t,t+j}(x)$ in equation (38) and manipulating the resulting equation yields

$$\hat{p}_t(x) = (1 - \alpha\beta)\zeta\sum_{j=0}^{\infty}(\alpha\beta)^j E_t(\hat{y}_{t+j} - \hat{y}_{t+j}^n) + \alpha\beta\sum_{j=1}^{\infty}(\alpha\beta)^j E_t\hat{\pi}_{t+j}, \quad (39)$$

where $\zeta = (\omega + \sigma^{-1})/(1 + \omega\theta)$. A linear approximation of the expression for the price index—equation (22)—yields

$$\hat{\pi}_t = \frac{1 - \alpha}{\alpha}\hat{p}_t(x). \quad (40)$$

Using this last equation to replace $\hat{p}_t(x)$ in equation (39) 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

$$\hat{\pi}_t = \beta E_t \hat{\pi}_{t+1} + \kappa\zeta(\hat{y}_t - \hat{y}_t^n). \quad (41)$$

Finally, we rewrite the household's Euler equation—equation (29) in terms of the output gap:

$$y_t - y_t^n = E_t(y_{t+1} - y_{t+1}^n) - \sigma(\hat{i}_t - E_t \hat{\pi}_{t+1} - r_t^n), \quad (42)$$

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

$$r_t^n = E_t \Delta \xi_{c,t+1} + \frac{1}{\sigma} E_t \Delta y_{t+1}^n. \quad (43)$$

References

- ABRAHAMSON, M., T. ADRIAN, R. K. CRUMP, AND E. MOENCH (2013): “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.
- 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., R. G. HUBBARD, AND K. HASSETT (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 (2013): “The TIPS-Treasury Bond Puzzle,” *Journal of Finance*, Forthcoming.

- FLEMING, M. J., AND M. PIAZZESI (2005): “Monetary Policy Tick-by-Tick,” Working Paper, Stanford University.
- GAGNON, J., M. RASKIN, J. REMACHE, AND B. SACK (2010): “Large-Scale Asset Purchases by the Federal Reserve: Did They Work?,” Federal Reserve Bank of New York Staff Report No. 441.
- GALI, J. (2008): *Monetary Policy, Inflation, and the Business Cycle*. Princeton University Press, Princeton, NJ.
- GERTLER, M., AND P. KARADI (2013): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” Working Paper, New York University.
- 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 (2012): “Monetary Policy and Long-Term Real Rates,” Working Paper, Harvard University.
- 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 (2013): “Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve,” Working Paper, University of Oxford.
- 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.
- 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.
- 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,” Working Paper, Johns Hopkins University.

TABLE 1
Response of Interest Rates and Inflation to Monetary Shocks

	Policy News Shock			Fed Funds Shock		
	Nominal	Real	Inflation	Nominal	Real	Inflation
3M Treasury Yield	0.68 (0.15)			0.51 (0.18)		
6M Treasury Yield	0.84 (0.12)			0.59 (0.12)		
1Y Treasury Yield	0.98 (0.15)			0.41 (0.17)		
2Y Treasury Yield	1.05 (0.37)	1.00 (0.28)	0.05 (0.20)	0.49 (0.54)	0.53 (0.37)	-0.04 (0.24)
3Y Treasury Yield	0.97 (0.41)	0.94 (0.28)	0.03 (0.20)	0.37 (0.56)	0.42 (0.35)	-0.05 (0.29)
5Y Treasury Yield	0.63 (0.21)	0.58 (0.15)	0.05 (0.12)	0.10 (0.18)	0.21 (0.14)	-0.12 (0.10)
10Y Treasury Yield	0.29 (0.18)	0.38 (0.13)	-0.09 (0.09)	-0.02 (0.14)	0.10 (0.11)	-0.12 (0.08)
2Y Treasury Inst. Forward Rate	1.00 (0.51)	0.86 (0.31)	0.13 (0.26)	0.27 (0.67)	0.31 (0.39)	-0.04 (0.37)
3Y Treasury Inst. Forward Rate	0.60 (0.46)	0.72 (0.33)	-0.12 (0.18)	0.03 (0.55)	0.12 (0.35)	-0.10 (0.34)
5Y Treasury Inst. Forward Rate	0.13 (0.19)	0.39 (0.17)	-0.27 (0.09)	-0.11 (0.15)	0.07 (0.13)	-0.18 (0.08)
10Y Treasury Inst. Forward Rate	-0.13 (0.19)	0.09 (0.13)	-0.22 (0.10)	-0.13 (0.17)	-0.03 (0.12)	-0.09 (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 (first three columns) or a change in the expected federal funds rate (last three columns) over a 30 minute window around the time of FOMC announcements. For the expected federal funds rate, this is the expected federal funds rate over the remainder of the current month unless the FOMC date in question occurs when there are 7 days or less remaining in the month, in which case it is the change in the expected federal funds rate over the next month. All results are based on Rigobon's (2003) method of identification by heteroskedasticity. The sample period is Jan 1st 2000 to Jan 25th 2012, except that we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. The "treatment" sample is a 30-minute window around all regularly scheduled FOMC announcements. The "control" sample is 2:05pm to 2:35pm on all Tuesdays and Wednesdays that are not FOMC meeting days. For 2Y and 3Y yields and real forwards, the sample starts in 2004. The sample size of the treatment sample for the 2Y and 3Y yields and forwards is 57. The sample size of the treatment sample for all other regressions is 89. Standard errors are calculated using a non-parametric bootstrap with 5000 iterations.

TABLE 2
Comparison with Alternative Methodologies

	30-Minute Window		One-Day Window		One-Day Window	
	Rigobon		Rigobon		Rigobon (90% CIs)	
	Policy News Shock		Policy News Shock		2Y Nominal Yield	
	Nominal	Real	Nominal	Real	Nominal	Real
2Y Yield	1.05	1.00	1.00	0.84	1.00	0.72
	[0.43, 1.97]	[0.58, 1.71]	[0.20, 1.77]	[0.50, 1.87]	[1.00, 1.00]	[0.46, 1.30]
2Y Forward	1.00	0.86	0.88	0.74	1.01	0.76
	[0.19, 2.24]	[0.36, 1.63]	[-0.51, 1.72]	[0.31, 2.10]	[0.48, 1.23]	[0.46, 1.55]
3Y Forward	0.60	0.72	0.55	0.61	0.80	0.72
	[-0.09, 1.73]	[0.23, 1.55]	[-0.94, 1.36]	[0.18, 2.28]	[0.02, 1.22]	[0.37, 1.93]
5Y Forward	0.13	0.39	-0.17	0.26	-0.02	0.35
	[-0.21, 0.53]	[0.10, 0.76]	[-1.25, 0.23]	[-0.12, 0.84]	[-6.31, 0.43]	[-0.25, 2.23]
10Y Forward	-0.13	0.09	-0.47	-0.05	-0.57	-0.03
	[-0.50, 0.25]	[-0.16, 0.34]	[-1.79, -0.06]	[-0.49, 0.40]	[-9.76, -0.08]	[-2.33, 0.43]
	30-Minute Window		One-Day Window		One-Day Window	
	OLS		OLS		OLS	
	Policy News Shock		Policy News Shock		2Y Nominal Yield	
	Nominal	Real	Nominal	Real	Nominal	Real
2Y Yield	1.08	1.03	1.16	0.98	1.00	0.75
	[0.42, 1.74]	[0.57, 1.50]	[0.84, 1.48]	[0.69, 1.28]	[1.00, 1.00]	[0.61, 0.88]
2Y Forward	1.03	0.89	1.18	0.93	1.15	0.83
	[0.14, 1.93]	[0.36, 1.42]	[0.74, 1.63]	[0.54, 1.32]	[1.03, 1.26]	[0.63, 1.03]
3Y Forward	0.64	0.75	0.90	0.81	0.98	0.80
	[-0.16, 1.44]	[0.18, 1.32]	[0.45, 1.34]	[0.35, 1.26]	[0.80, 1.17]	[0.53, 1.06]
5Y Forward	0.16	0.40	0.37	0.43	0.55	0.47
	[-0.21, 0.53]	[0.08, 0.72]	[0.12, 0.61]	[0.16, 0.69]	[0.36, 0.74]	[0.26, 0.67]
10Y Forward	-0.10	0.10	0.03	0.13	0.20	0.20
	[-0.47, 0.27]	[-0.15, 0.35]	[-0.22, 0.28]	[-0.12, 0.37]	[0.02, 0.38]	[0.01, 0.38]

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 or the 2-Year nominal yield over either a 30-minute window or 1-day window around FOMC announcements as described above each set of results. The estimation method is either Rigobon's (2003) method of identification by heteroskedasticity or OLS as described above each set of results. We report a point estimate and 95% confidence intervals except in the 1-day, Rigobon 2YNY case, where we report 90% confidence intervals. The sample period is Jan 1st 2000 to Jan 25th 2012, except that we drop the second half of 2008, the first half of 2009 and a 10 day period after 9/11/2001. The sample of "treatment" days for the Rigobon method is all regularly scheduled FOMC meeting day. The sample of "control" days is all Tuesdays and Wednesdays that are not FOMC meeting days. For 2Y and 3Y yields and real forwards, the sample starts in 2004. Standard errors for the Rigobon method are calculated using the weak-IV robust approach discussed in the text with 5000 iterations.

TABLE 3
Effects of Monetary Shocks on Survey Expectations

	Nominal	Real	Inflation
1 quarter	1.01 (0.49)	1.18 (0.51)	-0.16 (0.26)
2 quarters	1.12 (0.50)	1.58 (0.49)	-0.46 (0.23)
3 quarters	0.87 (0.53)	1.15 (0.51)	-0.27 (0.21)
4 quarters	0.81 (0.50)	1.13 (0.48)	-0.32 (0.21)
5 quarters	0.61 (0.64)	0.47 (0.64)	0.14 (0.24)
6 quarters	1.76 (0.63)	1.53 (0.65)	0.22 (0.28)
7 quarters	3.84 (1.19)	3.71 (1.22)	0.13 (0.51)

This table presents the results of regressing 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 calculate the sum of the policy news shocks that occur 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 January 2012, except that we exclude the second half of 2008 and the first half of 2009.

TABLE 4
Response of Expected Future Short Rates and Risk Premia

	Expected Future Short Rates		Risk Premia	
	Nominal	Real	Nominal	Real
2Y Treasury Yield	1.03 (0.28)	0.88 (0.17)	0.04 (0.10)	0.16 (0.19)
3Y Treasury Yield	0.95 (0.26)	0.83 (0.17)	0.05 (0.14)	0.14 (0.19)
5Y Treasury Yield	0.74 (0.16)	0.57 (0.11)	-0.08 (0.11)	0.03 (0.15)
10Y Treasury Yield	0.47 (0.10)	0.35 (0.07)	-0.14 (0.15)	0.04 (0.14)
2Y Treasury Forward Rate	0.78 (0.23)	0.74 (0.22)	0.25 (0.25)	0.16 (0.18)
3Y Treasury Forward Rate	0.59 (0.18)	0.53 (0.16)	0.05 (0.27)	0.22 (0.23)
5Y Treasury Forward Rate	0.32 (0.08)	0.25 (0.07)	-0.16 (0.17)	0.15 (0.17)
10Y Treasury Forward Rate	0.07 (0.02)	0.03 (0.01)	-0.17 (0.19)	0.07 (0.13)

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 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 Jan 1st 2000 to Jan 25th 2012, 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 57. The sample size for all other regressions is 89.

TABLE 5
Mean Reversion

Horizon (Trading Days)	Nominal Yields			Real Yields		
	2-Year	3-Year	5-Year	2-Year	3-Year	5-Year
1	1.14 (0.24)	1.11 (0.27)	0.97 (0.29)	1.18 (0.40)	1.11 (0.40)	0.85 (0.23)
5	0.84 (0.53)	0.76 (0.52)	0.63 (0.48)	0.89 (0.63)	0.76 (0.61)	0.25 (0.37)
10	0.12 (0.84)	-0.02 (0.87)	-0.14 (0.84)	1.29 (0.90)	1.10 (0.79)	-0.05 (0.62)
20	0.28 (1.03)	0.16 (1.15)	0.14 (1.22)	1.41 (1.47)	0.87 (1.29)	0.12 (0.94)
60	0.73 (1.82)	0.18 (1.85)	-0.39 (1.73)	2.34 (3.05)	2.03 (2.72)	-0.08 (1.43)
125	4.55 (2.92)	3.87 (2.72)	2.82 (2.36)	7.31 (4.49)	6.16 (4.13)	2.37 (2.05)
250	6.20 (3.88)	5.53 (3.36)	4.21 (2.66)	9.95 (3.39)	8.43 (2.89)	3.83 (1.66)

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
Breakeven Inflation versus Inflation Swaps

	Breakeven	Swaps
Inflation Over Next 2 Years	0.05 (0.22)	0.33 (0.40)
Inflation Over Next 3 Years	0.03 (0.21)	0.32 (0.36)
Inflation Over Next 5 Years	0.05 (0.17)	-0.05 (0.18)
Inflation Over Next 10 Years	-0.09 (0.15)	-0.23 (0.19)

Each estimate comes from a separate "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. All results are based on Rigobon's (2003) method of identification by heteroskedasticity. The sample period is Jan 1st 2005 to Jan 25th 2012, except that we drop the second half of 2008 and the first half of 2009. The sample of "treatment" days is all regularly scheduled FOMC meeting days. The sample of "control" day is all Tuesdays and Wednesdays that are not FOMC meeting days. The sample size of the treatment sample is 49. Standard errors are calculated using a non-parametric bootstrap with 5000 iterations.

TABLE 7
Monetary Non-Neutrality

	Output	Inflation	Ouput/Inflation
Our Estimation of CEE/ACEL Model	1.6 [1.5, 11.0]	0.5 [0.4, 7.1]	3.2 [1.6, 4.2]
ACEL	1.7	0.6	2.6
CEE	2.2	1.3	1.7

We consider the response of the economy to the Taylor rule shock considered in CEE. The shock is normalized such that the peak response of the nominal interest rate in our estimation of the CEE/ACEL model is 25 bp. For the output column, we sum the absolute value of the response of output over the first 500 periods after the shock and divide by four (to annualize). For the inflation column, we sum the absolute value of the response of inflation over the first 500 periods after the shock.

TABLE 8
Estimates of Structural Parameters

ξ_p	0.95 [0.79, 0.99]
ξ_w	0.92 [0.01, 0.99]
k_I	25.0 [1.1, 25.0]
ρ_1	0.93 [0.84, 0.97]
ρ_2	0.62 [0.01, 0.88]

TABLE 9
Response of Expected Output Growth

Exp. Output Growth in Current Qr	1.34 (1.69)
Exp. Output Growth 1 Qr Ahead	1.57 (0.64)
Exp. Output Growth 2 Qr Ahead	0.68 (0.35)
Exp. Output Growth 3 Qr Ahead	0.86 (0.27)
Exp. Output Growth 4 Qr Ahead	0.52 (0.32)
Exp. Output Growth 5 Qr Ahead	0.55 (0.28)
Exp. Output Growth 6 Qr Ahead	0.46 (0.31)
Exp. Output Growth 7 Qr Ahead	0.83 (0.68)

Each estimate comes from a separate "regression." The dependent variable in each regression is the one day change in the logarithm of the variable 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. All results are based on Rigobon's (2003) method of identification by heteroskedasticity. The sample is the same as in Table 1. Standard errors are calculated using a non-parametric bootstrap with 5000 iterations.


TABLE 10
Response of Stock Prices

	Stock Prices
Data	-7.3 (4.2)
News about Monetary Policy Only	-21.2
New about Monetary Policy and Exogenous Economic Fundamentals	-10.9


TABLE A1
Response of Interest Rates to Monetary Shocks 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	0.68 (0.15)		0.77 (0.14)		0.61 (0.19)		0.82 (0.15)	
6M Treasury Yield	0.84 (0.12)		0.85 (0.13)		0.81 (0.15)		0.91 (0.13)	
1Y Treasury Yield	0.98 (0.15)		0.98 (0.16)		0.99 (0.16)		0.99 (0.17)	
2Y Treasury Yield	1.05 (0.37)	1.00 (0.28)	1.08 (0.42)	1.00 (0.31)	1.09 (0.34)	1.56 (0.39)	1.30 (0.45)	1.16 (0.31)
3Y Treasury Yield	0.97 (0.41)	0.94 (0.28)	1.00 (0.46)	0.93 (0.32)	1.09 (0.37)	1.35 (0.35)	1.20 (0.47)	1.09 (0.31)
5Y Treasury Yield	0.63 (0.21)	0.58 (0.15)	0.63 (0.21)	0.56 (0.15)	0.81 (0.22)	0.86 (0.20)	0.60 (0.19)	0.62 (0.13)
10Y Treasury Yield	0.29 (0.18)	0.38 (0.13)	0.33 (0.19)	0.42 (0.14)	0.52 (0.22)	0.63 (0.19)	0.30 (0.15)	0.38 (0.12)
2Y Tr. Inst. Forward Rate	1.00 (0.51)	0.86 (0.31)	1.04 (0.55)	0.87 (0.34)	1.17 (0.44)	0.90 (0.35)	1.20 (0.57)	1.02 (0.34)
3Y Tr. Inst. Forward Rate	0.60 (0.46)	0.72 (0.33)	0.62 (0.50)	0.73 (0.35)	0.98 (0.45)	0.97 (0.39)	0.76 (0.49)	0.86 (0.34)
5Y Tr. Inst. Forward Rate	0.13 (0.19)	0.39 (0.17)	0.17 (0.20)	0.46 (0.17)	0.44 (0.28)	0.70 (0.25)	0.16 (0.15)	0.38 (0.14)
10Y Tr. Inst. Forward Rate	-0.13 (0.19)	0.09 (0.13)	-0.03 (0.20)	0.20 (0.13)	0.06 (0.24)	0.19 (0.15)	-0.11 (0.15)	0.00 (0.13)

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 the Rigobon's (2003) method of identification by heteroskedasticity. The "treatment" sample is a 30-minute window around all regularly scheduled FOMC announcements, except the last two columns where we include 30-minute windows unscheduled FOMC announcements. The "control" sample is 2:05pm to 2:35pm on all Tuesdays and Wednesdays that are not FOMC meeting days. The baseline sample period is Jan 1st 2000 to Jan 25th 2012, 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 Jan 1st 2000 to Jan 25th 2012. 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. Standard errors are calculated using a non-parametric bootstrap with 5000 iterations.



Panel A: 1-Day Window



Panel B: 30-Minute Window

Figure 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 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




Figure 3: Interest Rate and Inflation in the Simple New Keynesian Model

The parameters used in this illustrative example are: $\beta = 0.99$, $\sigma = 0.5$, $\kappa\zeta = 0.0017$, $\phi_\pi = 0.5$. We assume that \bar{r}_t has two components: 1) an AR(2) component with roots $\rho_1 = 0.93$ and $\rho_2 = 0.62$, and a permanent component (which yields a non-zero $\hat{\pi}_\infty$).

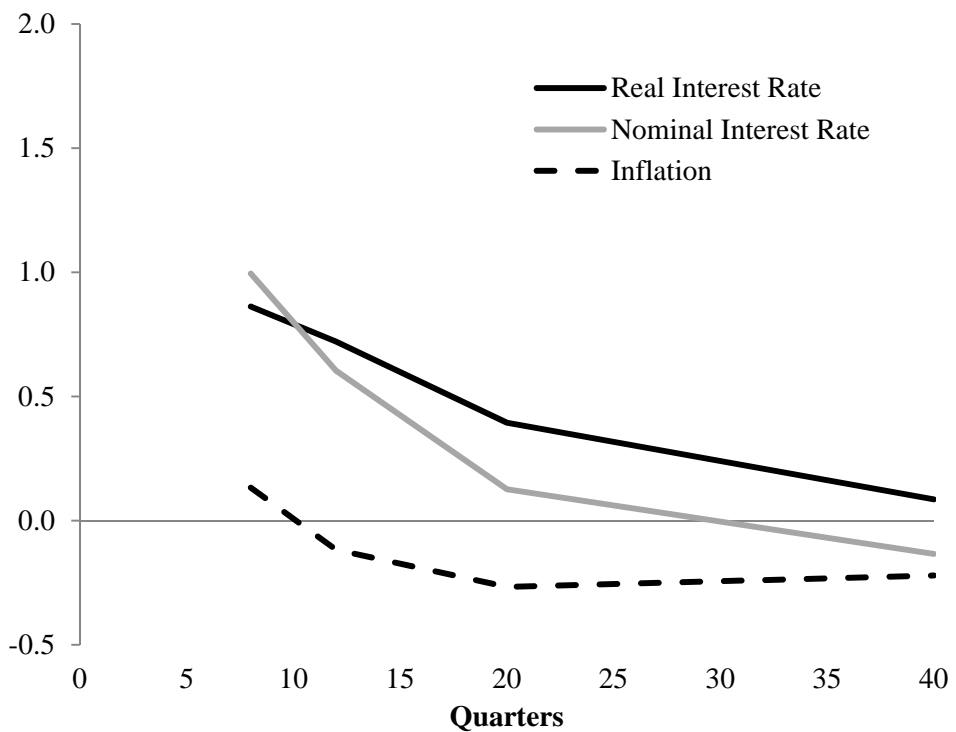


Figure 4: Interest Rates and Inflation in the Data

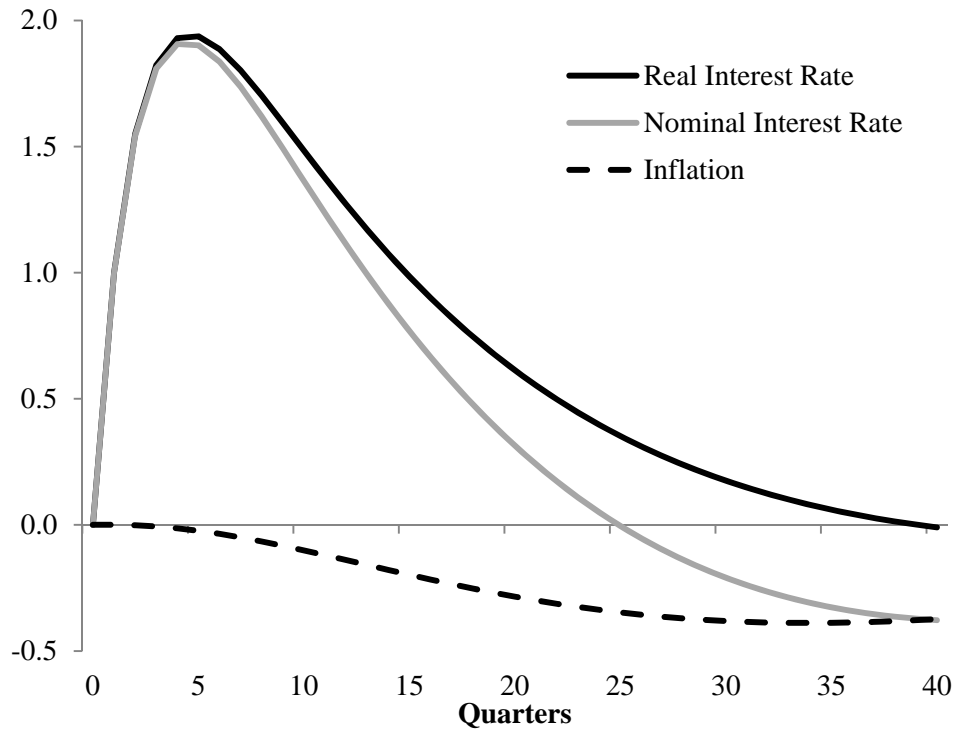


Figure 5: Response of Inflation and Interest Rates to Policy News Shock in Our Estimation of CEE/ACEL Model




Figure 6: Response of Inflation and Interest Rates to Policy News Shock in CEE/ACEL Model with CEE Parameters




Figure 7: Scatterplot of Estimated Joint Distribution of ξ_w and ξ_p

Note: The figure plots the values of ξ_w and ξ_p from the 1000 bootstrap draws we calculate.

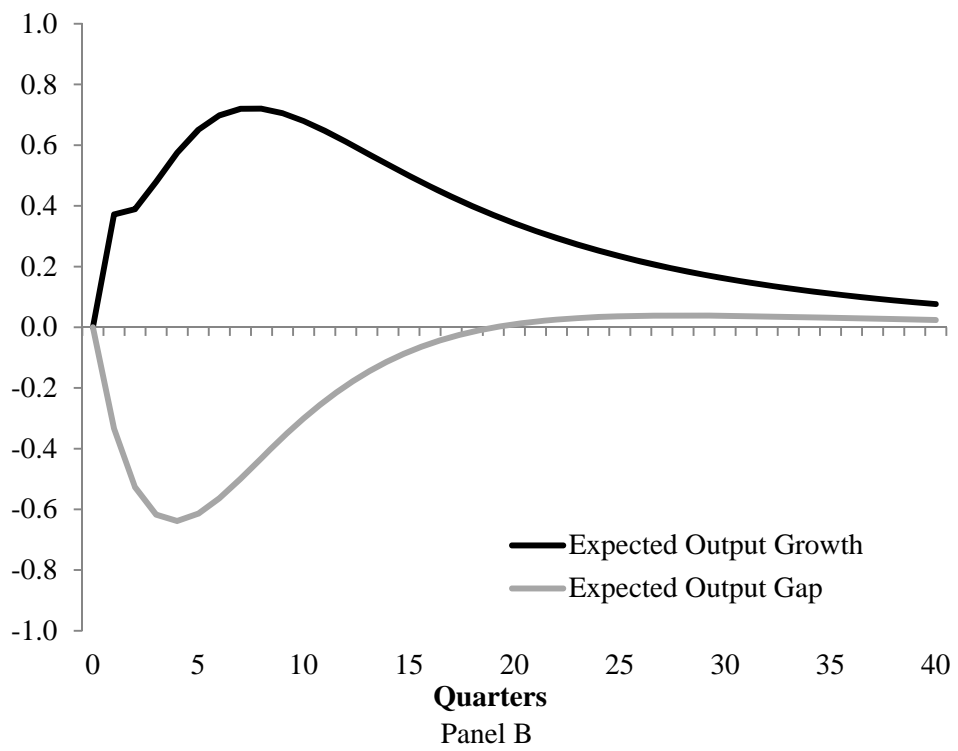
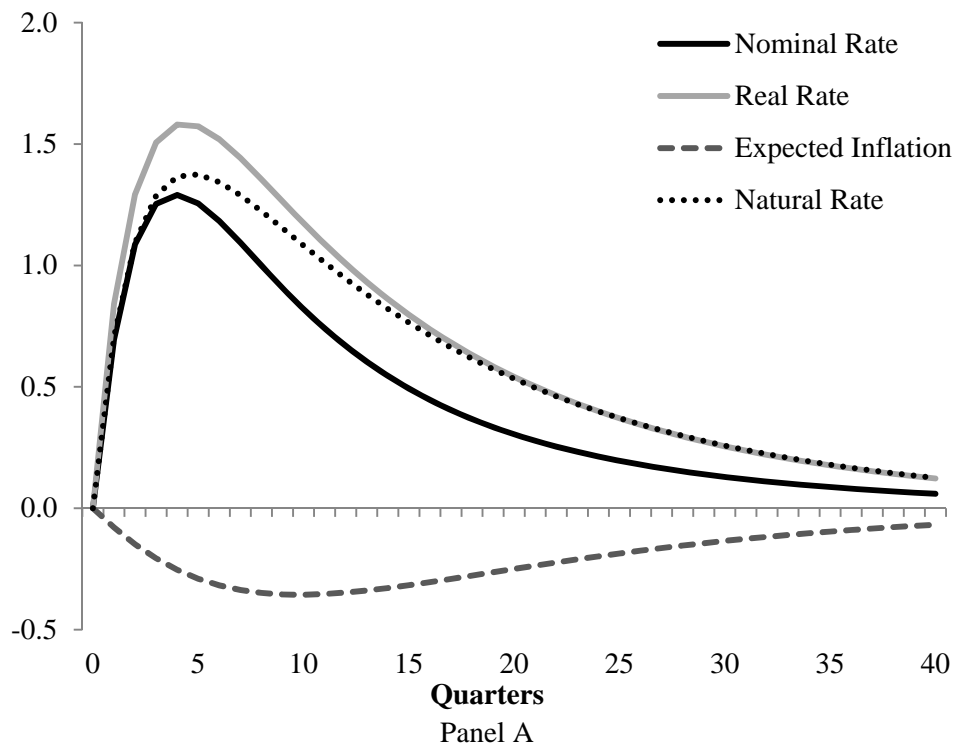


Figure 8: Responses of Interest Rates, Expected Inflation, and Expected Output when FOMC Announcements Convey Information about Both Monetary Policy and Exogenous Shocks