Australian National University

Crawford School of Public Policy

CAMA

Centre for Applied Macroeconomic Analysis

Estimating Monetary Policy Rules When Nominal Interest Rates Are Stuck at Zero

CAMA Working Paper 53/2013
August 2013

Jinill Kim
Korea University and
Centre for Applied Macroeconomic Analysis (CAMA), ANU

Seth Pruitt
Federal Reserve Board

Abstract

Did the Federal Reserve's response to economic fundamentals change with the onset of the Global Financial Crisis? Estimation of a monetary policy rule to answer this question faces a censoring problem since the interest rate target has been set at the zero lower bound since late 2008. Surveys by forecasters allow us to sidestep the problem and to use a conventional regression. We find that the Fed's inflation response has decreased and that the unemployment response has remained as strong; this suggests that the Federal Reserve's commitment to stable inflation has become weaker in the eyes of the professional forecasters.

Keywords

monetary policy, policy rule, survey data, market perceptions, censoring, zero lower bound, Blue Chip survey.

JEL Classification

E53, E58

Address for correspondence:

(E) cama.admin@anu.edu.au

Estimating Monetary Policy Rules When Nominal Interest Rates Are Stuck at Zero*

Jinill Kim†
Korea University, CAMA

Seth Pruitt‡
Federal Reserve Board

July 2013

Abstract

Did the Federal Reserve's response to economic fundamentals change with the onset of the Global Financial Crisis? Estimation of a monetary policy rule to answer this question faces a censoring problem since the interest rate target has been set at the zero lower bound since late 2008. Surveys by forecasters allow us to sidestep the problem and to use a conventional regression. We find that the Fed's inflation response has decreased and that the unemployment response has remained as strong; this suggests that the Federal Reserve's commitment to stable inflation has become weaker in the eyes of the professional forecasters.

Keywords: monetary policy, policy rule, survey data, market perceptions, censoring, zero lower bound, Blue Chip survey.
JEL codes:E53, E58.

*First draft June 2010. Previous drafts of this paper were circulated under the title "Forecaster-Perceived Monetary Policy Rules under the Zero Lower Bound." We thank Stefano Eusepi, Ryo Kato, Bae-Keun Kim, Andy Levin, Ricardo Nunes, Nick Turner, Yuchiro Waki, and Jonathan Wright for helpful discussions, Roger Koenker for code, as well as presentation participants at the 18th Australian Macroeconomics Workshop 2013, the University of Tokyo, the Bank of Japan, the 2012 Fall Meeting of the Japan Society of Monetary Economics, the 2011 Society for Nonlinear Dynamics and Econometrics Annual Symposium and the Federal Reserve Board. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. The first author acknowledges the financial support from the National Research Foundation of Korea Grant funded by the Korean Government (NRF-2012S1A5A8024344).

†jinillkim@korea.ac.kr
‡seth.j.pruitt@frb.gov

1 Introduction

Did the (Global Financial) Crisis somehow alter the Federal Reserve's behavior? Since the 1990s, simple monetary policy rules for the nominal short rate have been widely used to analyze the behavior of central banks—most notably the Federal Reserve's behavior following Taylor (1993). The short rate is critical for the dynamics of the whole economy, not only by representing the reaction of a central bank to the current state of the economy but also by providing information of the economy's future to financial markets. Previous empirical literature such as Clarida, Gali and Gertler (2000) finds that the Federal Reserve's policy rule significantly changed around the singular event of high U.S. inflation at the end of the 1970s—we ask whether or not the singular event of the Global Financial Crisis has caused another such break in monetary policy. To empirically answer this question, an estimate of the policy rule depends on meaningful covariation between nominal short rates and economic fundamentals such as the inflation and unemployment rates. But since December 2008—shortly after the collapse of Lehman Brothers in September—Federal Reserve has been targeting its policy rate between zero and 25 basis points: U.S. nominal short rates have been stuck at their zero lower bound (ZLB).[1] That is, policy rates are censored at zero and realized nominal short rates cannot provide the meaningful variation required to answer our research question.[2]

[1] In this paper, we assume that a nominal interest rate is effectively at its ZLB as far as it is below 25 basis points—see Bernanke and Reinhart (2004) for reasons why the short-term rate had better not be pushed down all the way to zero. In model simulations, the ZLB creates difficulty since it implies nonlinearity in the setting of monetary policy.

[2] This is our censoring problem: Even when both the rate recommended by a simple rule and so the desired target rate are negative, the ZLB-constrained target rate announced by the Fed stays weakly positive. See Rudebusch (2009) and Curdia and Woodford (2011) among others for the argument that desired target rates are negative. The censorship is with regards to the targeted policy rate, not the actual observed nominal short rate. We use the terms "policy rate" and "short rate" interchangeably in the remainder of the paper.

A textbook way of tackling a censoring problem is to apply limited dependent variable methods to the historical data. However, these methods cannot define a powerful hypothesis test of whether or not policy response parameters changed because after the Crisis began the short rate has *always* been censored.[3] The main idea of this paper is to avoid the censoring problem altogether by using data on agent expectations obtained from surveys of economic forecasts, and this paper attempts to answer whether or not Fed policy has changed since the Global Financial Crisis. Even while realized short rates were at the ZLB, one-year-ahead forecasts of the short rate stayed well above zero until August 2011, when a calendar-based forward guidance was introduced by the Federal Open Market Committee (FOMC). By using the one-year-ahead forecasts available in various issues of the Blue Chip Economic Indicators, we avoid the ZLB-induced censoring problem and estimate the Federal Reserve's post-Crisis monetary policy rule using conventional regression methods. No other paper, to our knowledge, has so straightforwardly evaluated the Federal Reserve's Taylor-type policy response function since the Crisis began. Therefore our results shed light on an important economic question that cannot be satisfactorily answered using conventional historical data.

In summary, the forecast data provide economically-sensible and statistically-significant answers to our research question. Historical data, on the other hand, do not: This is because the historical data either include the ZLB-censored observations—in which case there are

[3] To make this concrete, suppose we placed our break *at* December 2008. Likelihood ratio tests of the break require a maximum likelihood post-break estimate, but the estimate would be based *only* on censored observations—this cannot be consistent (for example, premises of Newey and McFadden 1994 Theorem 2.1 cannot be satisfied). The same problem exists for any later breakpoint, and breakpoints prior to December 2008 would be based on post-break estimates from just a handful of observations, thus seriously diminishing their power. Therefore, historical data do not satisfactorily answer whether or not a policy shift took place in response to the Global Financial Crisis. One might be able to alter our question to a partially identified context and conduct modified statistical tests. However, that would be a different question than we answer here, which is manifest as a simple parameter break test. Statistical analysis using partial identification is beyond the scope of this paper, but it is of interest to see if partially-identified models of the censored historical data confirm this paper's results.

statistical and economic problems with the results—or end when the ZLB begins—in which case statistical power is obviously compromised given the small number of observations available between the onset of the Global Financial Crisis and the start of the ZLB period. In addition to providing one-year-ahead forecasts in each issue, some issues of the Blue Chip Economic Indicators include long-horizon forecasts of inflation, unemployment and short rate. We take these forecasts as *data* for respective long-run values of these variables, and argue that having data to discipline the inflation target and market-perceived NAIRU is a further advantage of using these forecast data. In contrast, conventional policy rule estimation using historical data does not uniquely identify these values—absent additional assumptions—which can lead to economically unreasonable results that we explain below.

An intellectual price we pay for our simple solution to the ZLB-censorship problem is that the object estimated in this paper is market participants' *perception* of how the Federal Reserve sets policy. However, this is often assumed to coincide with the actual policy rule (by rational expectations) or else is interesting on its own as an object, as argued by Hamilton, Pruitt and Borger (2011). In either case, it is important for economists to understand financial markets' perception of the Fed's behavior since the Global Financial Crisis.

Our results using forecast data broadly show that after the Crisis the Federal Reserve's perceived inflation response significantly decreased, while its perceived unemployment response remained strong or may even have strengthened a bit. This conclusion successfully surmounts a difficulty inherent to our research question: The Crisis does not have an official start date. There are several months in 2007 and 2008 during which important events occurred that may have led to a shift in perceptions of the Federal Reserve's short-rate policy.

Therefore, we consider results for several potential breakpoints.[4] Nevertheless, there are quantitative differences between different breakpoints' results which we can only modestly alleviate by employing more sophisticated statistical econometric methods. We investigate in detail the features of the data, particularly the chaotic period between August 2007 and December 2008, which contribute to the variation in our estimates. We then argue for a simple solution to the results' sensitivity: Exclude the chaotic months. We do so on grounds that parts of the 2007:08–2008:12 period largely reflected forecasters' chaotic adjustment to unprecedented events. This leads to our preferred set of results.

In terms of central bank communication, our findings provide evidence that the Federal Reserve's commitment to stable inflation has become much weaker in the eyes of the professional forecasters—and probably the financial markets as well. This result could be interpreted as either bad (the loss of inflation-fighting credibility) or good (successful forward guidance), and we do not advocate one interpretation over the other. At the same time, professional forecasters view the Federal Reserve as firmly as ever, if not more, committed to closing the unemployment gap.

This paper touches on several strands of previous work. Taylor (1993) and Clarida, Gali and Gertler (2000) are the seminal references on constructing and detecting breaks in the Fed's policy rule using actual data. Hamilton, Pruitt and Borger (2012) instead estimate the market-perceived rule, as we focus on, using futures prices and lagged actual variables to instrument for market expectations. In the same vein, Ang, Boivin, Dong and Loo-Kung

[4] Our results show that a break occurred somewhere between August 2007 and December 2008, and it is not our aim to definitively date the Crisis using Federal Reserve policy. Hence we do not *search* for the break date by statistical methods. Instead, we try our best to make sure—as much as possible—that our main conclusions are robust to uncertainty regarding the exact breakpoint, using important events to clue us into the potential breakpoints that should be considered.

(2011) estimate a time-varying policy response using yield curve data which could also be viewed as capturing the market-perceived rule. Mitchell and Pearce (2010) use Wall Street economists' forecasts, and Carvalho and Nechio (2012) use consumers' forecasts to estimate the Fed's perceived policy rule. Fendel, Frenkel and Rulke (2011) use professionals' forecasts to estimate perceived policy rules in emerging markets. Survey data have also been used for structural estimation by Coibion and Gorodnichenko (2012) to assess informational rigidities, by Devereux, Smith and Yetman (2012) to assess cross-country risk-sharing on a panel of professionals' expectations, and by Hirose and Kurozumi (2012) to identify news shocks. None of the above papers analyze whether or not the Global Financial Crisis changed the Fed's monetary policy. Gerlach and Lewis (2011) assess if the Crisis affected the policy rule of the European Central Bank, and Swanson and Williams (2012) assess how the Crisis affected Treasury prices' response to data releases.

The plan of the paper is as follows. Sections 2 and 3 describe our data and estimation methods. In Section 4, we present our main results, analyze their robustness, and use our estimates to understand the forces which market forecasters believe shape Federal Reserve policy, as well as the effects of Federal Reserve communication on markets' beliefs. Section 5 concludes.

2 Data

The main idea of this paper is to use data on agents' *expectations* to answer our research question: Has the Fed reaction function changed since the Global Financial Crisis? A key advantage of such data is that they may forecast future economic conditions that warrant uncensored policy rates, even when actual short rates at the time of the forecast are stuck at the ZLB. This means that all the necessary ingredients—dependent and explanatory variables—retain meaningful variability that can be used to estimate the key policy response coefficients. In turn, we can in the usual manner form statisical hypothesis tests addressing our research question.

A brief digression on terminology may be helpful. This paper deals with two types of data. The first type of data we refer to as *actual*, *historical* or *realized* data are prices or measurements that actually prevailed in the U.S. experience. Examples are the 3-month T-Bill rate during December 2005 or the unemployment rate measured by the Bureau of Labor Statistics (BLS) for December 2005—these are actual, historical or realized data. The second type of data we refer to as *expectations* or *forecast* data are forecasts of future economic conditions. An example is the 3-month T-Bill rate expected to prevail in 2006Q4, as forecasted in December 2005. Studies such as Ang, Bekaert and Wei (2007) suggest that surveys contain remarkably accurate forecasts, and we make the assumption that the survey data we use accurately represent market expectations.

Our main historical data include the civilian unemployment rate from BLS, the ex-food/energy consumer price index (core CPI) from BLS, and 90-day U.S. Treasury bill rate from the Federal Reserve Board. In subsequent analysis, we also use two-year and ten-year U.S. Treasury bond yields from the Federal Reserve Board. We construct actual core-CPI inflation as the 12-month log growth rate in the core-CPI index.[5]

Our main forecast data source is the Blue Chip Economic Indicators. Collected monthly

[5]CPI excluding food and energy prices is used to measure inflationary pressures to which policy responds since it strips out price fluctuations (food and energy prices) that are typically considered temporary. Congruently, core-CPI inflation is a better forecaster of future CPI inflation than headline inflation itself. Hence we find it reasonable to associate realized core inflation (in the rule estimated on historical data) with expected headline inflation (in the rule estimated on forecast data).

since the late 1970s, these surveys provide short-to-medium horizon (one- to six-quarter ahead) forecasts of the annualized percentage change in the CPI, the 3-month T-Bill rate, and the unemployment rate. The survey is released on the 10th of the eponymous month but collected through the end of the previous month—therefore the "April 2000" survey actually represents information known to the market until the end of March 2000. Having forecasts for particular future quarters is important because we want to maintain a constant forecast horizon as much as possible. Strictly speaking, the forecast horizons vary by a couple of months since the surveys are monthly and forecast horizons are quarterly.[6] For example, if we assign the quarterly value to the middle month, then the April 2000 survey contains a 2001:Q2 "four-quarter-ahead" forecast that is 13 months ahead, and the June 2000 survey contains a 2001:Q2 "four-quarter-ahead" forecast that is 11 months ahead. Since our paper's main idea is to use professional forecasts to estimate professionals' perception of Federal Reserve policy response, we would like to maintain, as much as possible, a constant forecast horizon so as to minimize superfluous additional dimensions to the analysis. The fore mentioned data limitations prohibit us from strictly doing so, but we have found that these small differences in forecast horizon do not alter our results (in the pre-Crisis period during which we have a sufficient number of observations). Hereafter, we refer to the four-quarter-ahead forecasts of a month's survey as "one-year-ahead" forecasts.

The forecasts we use are the median forecast from a panel of between 30 and 60 professional forecasters in the financial services, consultancy, and academic industries, with emphasis on forecasters in the financial services industry.[7] Biannually the survey also publishes long-

[6]Another popular survey data is the Survey of Professional Forecasters which is published quarterly. Romer and Romer (2000) examined Federal Reserve forecasts against these private forecasts.

[7]The Blue Chip Economic Indicators show the full panel of annual forecasts for the current and following year. Using this full panel would advantageously allow us to control for panel composition issues discussed

horizon forecasts (the five-year average, five years ahead) for these variables, which we convert to a monthly measure by linear interpolation. We use the 3-month T-Bill rate as our primary policy measure so that we have both its medium- and long-horizon forecasts.[8] We use CPI inflation and unemployment to represent inflation and activity pressures, respectively. We use data starting in 1986 when all variables are observed, and these series are plotted in Figures 1–3.

Figure 1 demonstrates the main idea of our paper. The solid line represents the one-year-ahead forecast for the short rate, and this forecast remains above the ZLB even after December 2008 when the realized short rate (the dashed line) hits the ZLB. It is inconsequential to our exercise that this forecast indicates that market expectations exhibited serially correlated forecast errors after the onset of the Global Financial Crisis. What is important for our analysis is that the forecasts manifest a systematic relationship between short rate on one side (the solid line in Figure 1) and inflation and unemployment on the other side (the solid lines in Figures 2 and 3).

The long-horizon forecast of the short rate, inflation rate and unemployment rate are plotted by the dash-dotted lines in Figures 1–3, respectively. Taking the difference between

by Engelberg, Manski and Williams (2011), but disadvantageously confound our control of the forecast horizon. To be concrete about the disadvantage, consider the January 2000 and December 2000 surveys: Both contain individual forecasts for variable growth rates realized in 2000 and 2001. This means that the January survey contains individual 12-month- and 24-month-ahead forecasts, while the December survey contains individual 1-month- and 13-month-ahead forecasts. It is unclear how to adjust the forecast horizon for these two surveys and the surveys for months in-between. Blue Chip Financial Forecasts, a separate but similar survey, gives a full panel of individual forecasts for quarterly horizons that is available electronically since 2001—see footnote 8 for further comparison between the two surveys.

[8]It is conventional to measure policy by the federal funds rate, but this variable is not in our data. Instead, we note that over our sample period the correlation between the funds rate and the 3-month T-Bill rate is greater than 0.99 and so this data limitation is negligible. Blue Chip Financial Forecasts *do* contain forecasts for the fed funds rate and the GDP growth rate, but do not include an unemployment rate forecast and do not provide long-horizon forecasts. Therefore, we opt for using Blue Chip Economic Indicators since it forecasts unemployment and 3-month T-Bill rates, the latter of which are extremely similar to the fed funds rates.

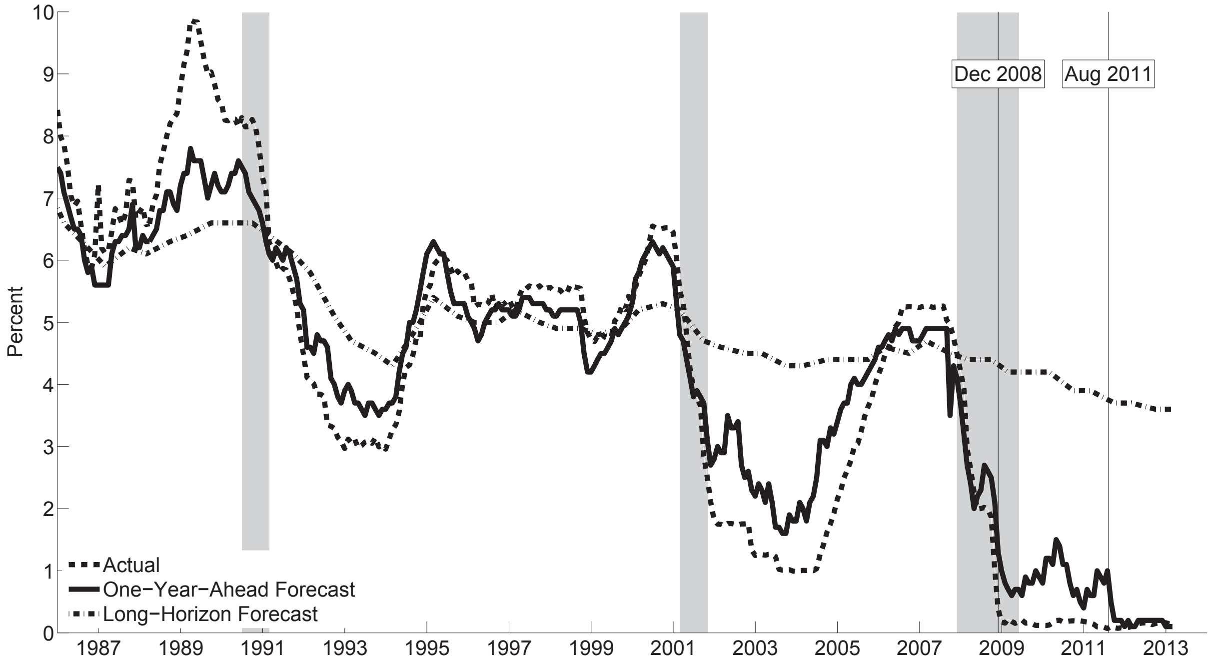

Figure 1: Short-Term Interest Rates

Notes: Data for January 1986 to February 2013. Actual values from the Federal Reserve Board, dashed line; median one-year-ahead forecast from Blue Chip Economic Indicators, solid line; long-horizon forecasts from Blue Chip Economic Indicators, dash-dotted line. Vertical lines are at December 2008 when the ZLB period began and at August 2011 when the FOMC introduced calendar-based forward guidance. Shaded are the NBER recession periods.

the long-horizon nominal-rate forecast and the long-horizon inflation-rate forecast yields the long-horizon real-rate forecast, by the Fisher identity. Figure 1 shows the long-horizon forecasts of the short rate have trended down over the past three decades; much of this follows the fall in the long-horizon forecasts of inflation plotted in Figure 2. Meanwhile, the long-horizon forecast of the unemployment rate hovers around familiar values of five to six percent, in Figure 3.

Going back to Figure 1, its right tail demonstrates why we end our estimation sample in August 2011. At the August 2011 FOMC meeting, the Committee announced that "excep-

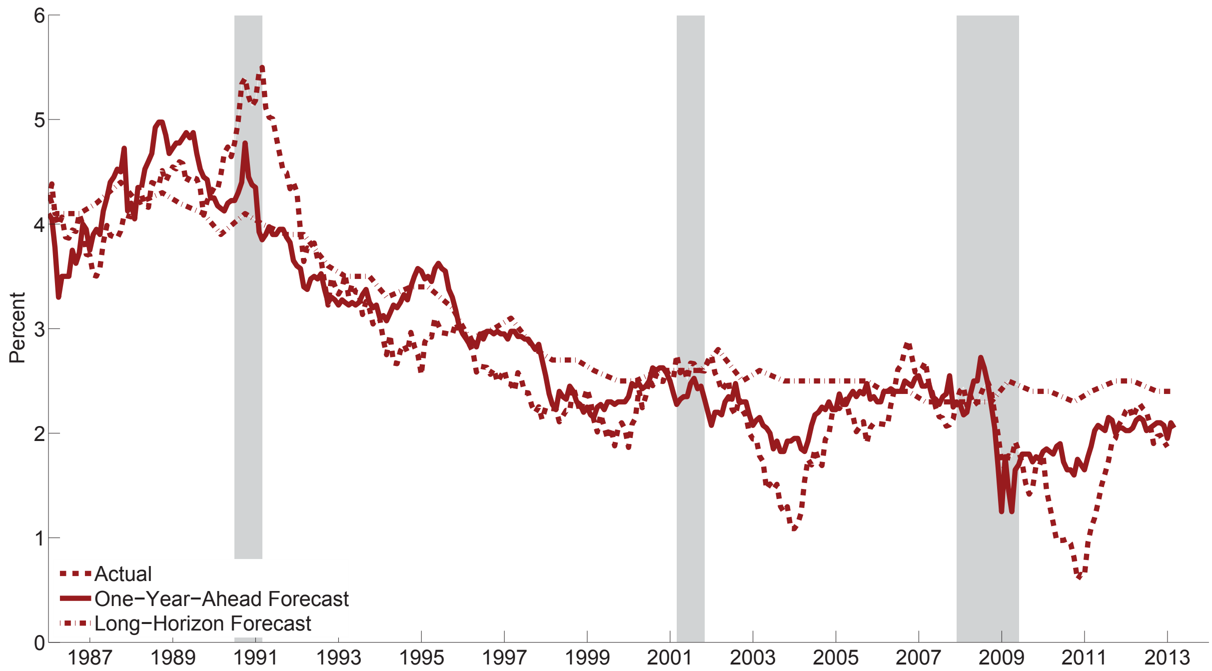

Figure 2: Inflation Rates

Notes: Data for January 1986 to February 2013. Core inflation from BLS, dashed line; median one-year-ahead inflation forecast from Blue Chip Economic Indicators, solid line; long-horizon inflation forecasts from Blue Chip Economic Indicators, dash-dotted line. Shaded are the NBER recession periods.

tionally low" short rates would likely be held "at least through mid-2013."[9] Our surveys—as well as price-based measures of market expectations such as fed funds futures or eurodollar contracts—respond immediately to this policy statement. Surveys published between September 2011 and the current draft (March 2013) exhibit medium-horizon short rate forecasts at the ZLB. Therefore this paper investigates the evidence of a discrete change in the Federal Reserve's (perceived) policy response to inflation and unemployment using forecast data somewhere between the onset of the Global Financial Crisis and August 2011.

[9]Earlier, FOMC introduced the "extended period" language in March 2009, instead of "for some time". Subsequently, the Committee extended the calendar-based forward guidance to late 2014 in January 2012 and to mid-2015 in September 2012. In December 2012, FOMC switched to its new threshold-based forward guidance (regarding 6.5% unemployment and 2.5% inflation) instead of the calendar-based forward guidance.

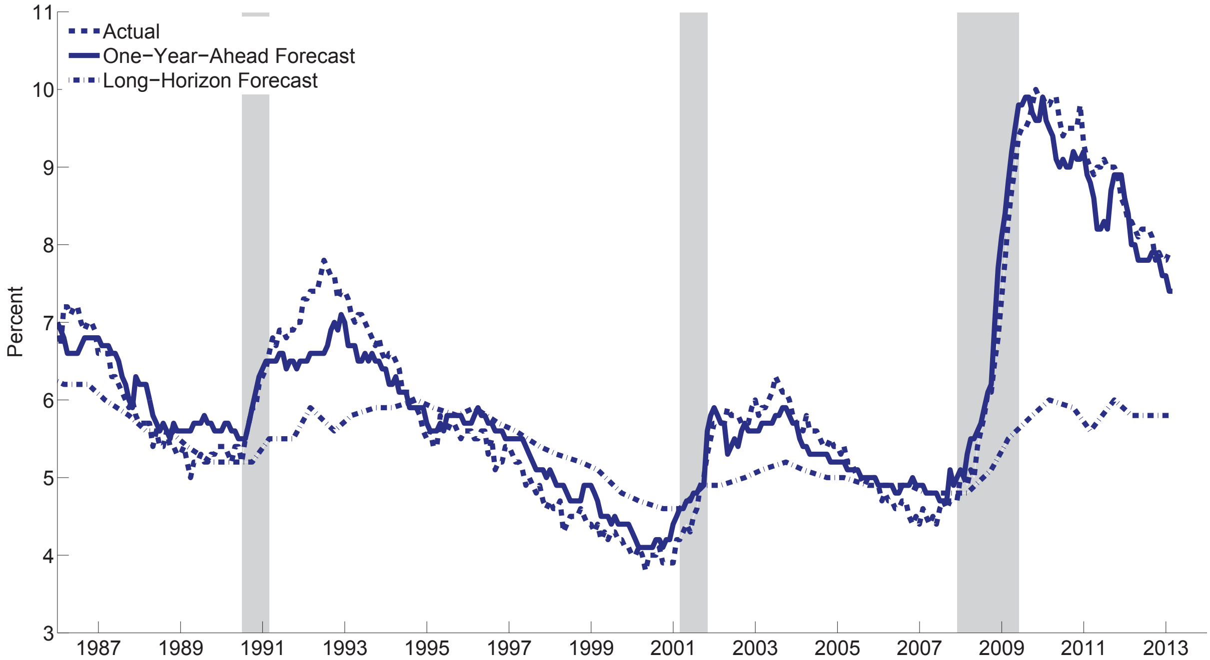

Figure 3: Unemployment Rates

Notes: Data for January 1986 to February 2013. Actual values from BLS, dashed line; median one-year-ahead forecast from Blue Chip Economic Indicators, solid line; long-horizon forecasts from Blue Chip Economic Indicators, dash-dotted line. Shaded are the NBER recession periods.

3 Estimating Policy Rules

To represent the the Federal Reserve's reaction to economic fundamentals, the popular Taylor-type policy rule specification is

$$i_t = (1-\rho)\left[(r_t^* + \pi_t^*) + \beta(\pi_t - \pi_t^*) + \delta(u_t - u_t^*)\right] + \rho i_{t-1}, \tag{1}$$

where the subscript t corresponds to the monthly frequency of the data. This specification follows, among others, Boivin (2006) in using unemployment to measure real activity. It is similar in spirit to rules using the output gap—as in Taylor (1993)—and reflects the limitations of our survey data. The policy rate i_t responds to the inflation response coefficient β times the deviation of inflation π_t from its long-run target π_t^*; the unemployment response coefficient δ times the deviation of unemployment u_t from the non-accelerating inflation rate of unemployment (NAIRU) u_t^*; and the lagged short rate via "gradual adjustment" dictated by ρ. When both inflation and unemployment are at their long-run values of π_t^* and u_t^*, the nominal short rate converges to the economy's equilibrium real rate r_t^* plus inflation π_t^* as implied by the Fisher identity. The *-variables are written with t subscripts to allow for variation in the inflation target or natural rate of unemployment—we only allow for such time variation when we impose that these long-run variables are observable as long-horizon forecasts, and otherwise assume that they are constants.

In existing literature, implementations of equation (1) use historical data. An obvious problem is that the *-variables—the inflation target, the NAIRU, and the equilibrium real rate—are typically under-identified by the historical data.[10] What is typically estimated is the reduced form

$$i_t = c_0 + c_1\pi_t + c_2 u_t + c_3 i_{t-1}.$$

Note that $c_1 = \beta$, $c_2 = \delta$ and $c_3 = \rho$, but $c_0 = (1-\rho)(r_t^* + \pi_t^*) - \beta\pi_t^* - \delta u_t^*$ which highlights the identification problem: The intercept term is a function of the unobserved *-variables *and* the response parameters β, δ.

We now discuss using *forecast data* or *historical data* in suitable modifications of (1). We then formally state the maximum likelihood problem used for estimation, which is simply OLS regressions run with dummy variables to allow for parameter breaks.

[10]This need not be the case in larger-model specifications—see for instance Cogley and Sbordone (2008).

3.1 Rules Using Forecasts Data

Using forecast data instead of historical data provides us at least two distinctive features. First, the forecast data include one-year-ahead forecasts for unemployment, inflation, and the short rate. As mentioned in Section 2, these forecasts are not censored even after the ZLB started to bind the historical short rate. Second, these data include long-horizon forecasts of these three variables. In macroeconomic models with well-defined steady-state growth paths, long-horizon forecasts are synonymous with the values for the equilibrium nominal interest rate, monetary policy's inflation target, and the NAIRU. Therefore, the survey of forecasts gives us a unique opportunity to pin down the value of the *-variables using *data*. This is an advantage of these data, and they can be used to avoid the identification problem intrinsic to conventional Taylor-type policy rule estimation using historical data.

The forecast-data-based rule we estimate is

$$i_{E,t} = (r_t^* + \pi_t^*) + \beta(\pi_{E,t} - \pi_t^*) + \delta(u_{E,t} - u_t^*). \tag{2}$$

Equation (2) takes *expectations* as data, and we assume that our median-forecast data reveal the representative forecaster's conditional expectation. For example: $i_{E,t}$ is the month t expectation of the short rate that will prevail during the quarter one year after month t; π_t^* is the month t belief of the Federal Reserve's inflation target. Therefore, variables with the E, t subscript do vary at a monthly frequency (from survey to survey), but pertain to a quarterly time period about one year in the future, as discussed in Section 2. We opt *not* to denote the *-variables using the "E" subscript because they have no counterparts in historical data and hence need no distinguishing. At least three features of our main specification are worth mentioning.

First, we have made a particular horizon choice: Our main results use the the one-year-ahead forecast.[11] We do this to prevent, as much as possible, the effect of the ZLB that is operational on shorter-horizon forecasts of the short rate.

Second, (2) includes no parametric constant to be estimated and instead pins down all movements of the policy rate forecast to the inflation response $\beta\pi_{E,t}$, unemployment response $\delta u_{E,t}$, and shifts in the *-variables $(r_t^* + \pi_t^*) - \beta\pi_t^* - \delta u_t^*$. This more tightly adheres the model to the data, and is possible because the Blue Chip surveys provide long-horizon forecasts that naturally represent forecasters' beliefs about the long-run values of the respective variables.

Third, we have zeroed out the gradual adjustment parameter ρ. Regarding this choice, we point out that Rudebusch (2006) argues that serially-correlated residuals account for the persistence of short rates, reflecting variables used by policymakers but excluded by our simple analysis. Including a gradual adjustment term serves to soak up residual variation, but without providing any additional economic rationale for *what forces* move the expected policy rate. Hence, we prefer to explicitly evaluate the residual left after inflation gaps and unemployment gaps have been controlled for, which we do in Section 4.2 in order to discuss which omitted variables may have been involved in the Taylor-type specification. In particular, we connect the estimated residuals to business-cycle-frequency financial-market risk following Atkeson and Kehoe (2009).[12]

[11]Recall our loose usage of the phrase "one-year-ahead" discussed in Section 2.

[12]In addition, Hamilton, Pruitt and Borger (2011) find that responses to past inflation and real activity differ such that gradual adjustment may not be the most useful mechanism for smoothing the short rate's dynamics. Our main results demonstrate that (2) fits the survey data extremely well without gradual adjustment; furthermore, we wish to avoid the effects of the ZLB on our forecast data, and using a shorter horizon's short-rate forecast (in the adjustment term) works against this aim.

3.2 Rules Using Historical Data

Conventional policy rule estimates based on historical data are obtained via the equation

$$i_{A,t} = C + \beta\pi_{A,t} + \delta u_{A,t}, \tag{3}$$

where $i_{A,t}, \pi_{A,t}, u_{A,t}$ are the *actual* values of the policy, inflation and unemployment rates in month t. Absent data to discipline their evolution, we assume that r^*, π^*, u^* are time-invariant constants and do not constrain them to take on the values given by the long-horizon forecast data.

3.3 Estimation Methodology

To answer our main question we test for parameter breaks around the onset of the Global Financial Crisis in a simple least squares specification. We estimate the regression

$$i_t = \sum_{b=1}^{B} \mathbb{1}(t \in \tau_b) Rule\,(\beta_b, \delta_b, \pi_t, u_t, \pi_t^*, u_t^*) + \epsilon_t, \tag{4}$$

where B is one more than the number of breaks allowed and τ_b is the subset of time periods between break periods $b-1$ and b with the convention that break period 0 is the month prior to the start of our data. $Rule(\bullet)$ denotes the right-hand-side of either (2) or (3) depending on whether we are using forecast or historical data, respectively.

When we estimate Equation 4 using forecast data we use (2) and so

$$Rule\,(\beta_b, \delta_b, \pi_t, u_t, \pi_t^*, u_t^*) = (r_t^* + \pi_t^*) + \beta_b(\pi_{E,t} - \pi_t^*) + \delta_b(u_{E,t} - u_t^*)$$

and the dependent variable is $i_{E,t}$. This can be estimated by ordinary least squares, run without a constant and using dummy variables to control for breaks, of the expected cyclical policy rate $i_{E,t} - r_t^* - \pi_t^*$ on the expected inflation gap $\pi_{E,t} - \pi_t^*$ and expected unemployment gap $u_{E,t} - u^*_t$. The residual from (2) is not a "shock" to monetary policy; instead, the estimated ϵ_t represents an *expected* future policy response that is uncorrelated with future inflation or real activity, a point to which we return in Section 4.2 below.

When we estimate Equation 4 using historical data we use (3) and so

$$Rule\left(\beta_b, \delta_b, \pi_t, u_t, \pi^*_t, u^*_t\right) = C_b + \beta_b \pi_{A,t} + \delta_b u_{A,t}$$

and the dependent variable is $i_{A,t}$. This can be estimated by ordinary least squares, including a constant and using dummy variables to control for breaks, of the actual policy rate $i_{A,t}$ on inflation $\pi_{A,t}$ and unemployment $u_{A,t}$. Recall that the constant will break if we allow for $B > 1$. The residuals from (3) resemble what previous literature has interpreted as monetary policy "shocks."

We use Chow tests to provide a statistical answer to our main research question, so we must be careful in their construction. In particular, it is well-known—as in Davidson and MacKinnon (1993)—that Chow tests are quite sensitive to heteroskedasticity. MacKinnon (1989) shows convincing Monte Carlo evidence to this effect: In an experiment similar to our own setting, MacKinnion finds that the standard Chow test has extremely low power when heteroskedasticity is present.[13] Therefore we follow Davidson and MacKinnon's (1993, Chapter 11.6) construction of heteroskedasticity-robust F-tests based on heteroskedasticity-robust Gauss-Newton regression results—when we report robust Chow test results, these are the statistics we report. In practice, using heteroskedasticity-robust tests lead to different inference than unadjusted F-tests, and we note a few places where this robustness makes a qualitative difference. The fact that robust tests give different inference is unsurprising given

[13] For example, a 5%-size test rejects the false null hypothesis in only 6% of the simulations. (Table 2, $\theta = 0.2$, $n = 200$, $\sigma_1/\sigma_2 = 1/4$).

the heteroskedastic nature of our estimated residuals, which we further explore in Section 4.2. We report heteroskedasticity and autocorrelation robust standard errors calculated as in Newey and West (1987) using three lags in the Bartlett kernel.

4 Estimation Results

We start with the results that postulate the Global Financial Crisis to begin in August 2007 when the TED spread (the difference between interest rates on interbank loans and U.S. Treasury bills) spiked, the large multinational bank BNP Paribas unprecedentedly froze withdrawals from large funds, and average repo haircuts on structured debt began to rise (Gorton and Metrick, forthcoming).[14] This follows Hamilton, Pruitt and Borger (2011) in dating the beginning of the Crisis. It is possible that this month does not definitively mark the start of the Crisis and so we explore the robustness of our results to this choice. We investigate results with a break at March 2008 (when Bear Stearns failed), September 2008 (when Lehman Brothers failed) and December 2008 (when the Fed funds rate hit the effective ZLB for the first time ever).

In this section, we compare results obtained using historical data with our main results using forecast data from the Blue Chip Economic Indicators. We find qualitative agreement between the two approaches on the pre-Crisis sample, supporting the idea that forecast and historical data reveal similar policy response functions when both types of data are uncensored. Post-Crisis, uncensored forecast data yields significantly different results than using censored historical data. Using the forecast data post-Crisis, we find evidence that

[14]In addition, Gorton (2010) dates August 2007 as the beginning of the most recent U.S. banking panic. Hamilton, Pruitt and Borger's (2011) analysis of market-perceived policy rules ends in August 2007 owing to unprecedented dislocations in the Fed funds market.

the inflation response has vanished and but the unemployment response remained strong, if not became stronger. We verify that these main results hold, qualitatively, at alternative starting-points for the Crisis.

We then investigate the part of forecasted policy that is not associated with forecasted inflation or unemployment. We find a significant relationship between our estimated short-rate residuals and measures of investors' perceived risk, suggesting that market participants believe the Fed sets policy in response to financial market risk as well as to fundamentals at the business cycle frequency, as is argued by Atkeson and Kehoe (2009). Finally, we use our estimates, realized data, and Blue Chip forecast data to analyze the counterfactual implied short rate since the beginning of the Global Financial Crisis and discuss market perceptions of the necessary size of unconventional monetary policy.

4.1 Parameter Estimates

We first report estimates from various specifications involving the historical data, and argue that they do not satisfactorily answer our research question. When the historical data include the ZLB period, the ensuing results and statistical tests comparing parameters around the Global Financial Crisis are neither economically nor statistically sensible—owing to the fact that short rates after December 2008 are censored. Next we report estimates from historical data that exclude the ZLB period. We argue that these results lack the power necessary to make robust statements about a possible shift in Federal Reserve's short-rate policy—owing to the fact that there are so few uncensored post-Crisis observations. In large part, these historical-data results are presented in order to connect to the previous literature, motivate the use of forecast data, and provide context for the results we find on the forecast data.

Therefore, for our main results we use *forecast data* over the entire sample 1986–2011. We find robust statistical evidence across a variety of specifications that the market believes the Federal Reserve's short rate response to inflation has *weakened* significantly, but that its response to unemployment has not. In some specifications, we further find evidence for the possibility that professional forecasts believe the Federal Reserve has *strengthened* its response to unemployment in addition to weakening its inflation response.

4.1.1 Historical Data

Including the ZLB Our first set of results use the historical data over our entire sample 1986–2011, including the ZLB-censored short rates from January 2009 onwards. The model fit is quite good, and an affine function of actual inflation and unemployment explains over 85% of the variation in actual short rates.

Pre-Crisis estimates are close to values found in previous literature, as expected. For example, Clarida, Gali and Gertler (2000) estimate an inflation response of 2.01 ($se = 0.28$) while our estimate shown in Table 1 is 1.89 ($se = 0.11$) on a roughly comparable sample. On the other hand, the historical data yield a larger unemployment rate response estimate -1.26 ($se = 0.10$) than Clarida, Gali and Gertler's (2000) statistically-insignificant estimate -0.56 ($se = 0.41$).[15]

Turning to the post-Crisis estimates, we see a very different picture as both response coefficient estimates are much closer to zero. We estimate a *negative* inflation response -0.28 ($se = 0.24$) that is statistically indistinguishable from zero. The estimated unemployment response -0.57 ($se = 0.14$) remains statistically significant, but is about half the pre-Crisis

[15] These numbers are from Clarida, Gali and Gertler (2000, Table III), using the unemployment rate, during the "Volcker-Greenspan" period. Note that their specification is quarterly and includes a smoothing term.

Table 1: Estimates from Historical Data Including the ZLB Period

	pre-Crisis	post-Crisis
August 2007 Break		
Inflation response	1.893	−0.275
standard error	(0.108)	(0.237)
Chow test	< 0.001	
Unemployment response	−1.257	−0.568
standard error	(0.097)	(0.136)
Chow test	< 0.001	
Constant	6.018	1.529
standard error	(0.543)	(0.785)
R^2	0.867	

Notes: Monthly data from the Federal Reserve Board and BLS, January 1986 to August 2011, regression with a constant. Robust standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

estimate which implies that the Federal Reserve is *less accommodative* in response to unemployment after the Crisis. The Chow test indicates that both of these breaks are significant at greater than the 0.1% level. We argue that there are at least two reasons to be skeptical of these results.

First, there is a statistical reason: The historical short rate is obviously censored at the ZLB for the majority of post-Crisis observations. Running a linear regression does not account for this feature of the data. Hence, we *should expect* the post-Crisis estimates to be biased towards zero due to this limited dependent variable problem and, indeed, the evidence in Table 1 accords with this intuition. This further implies that the Chow test's statistical inference is unfounded, since it presupposes one has consistent estimates.

Second, there is an economic reason: The post-Crisis estimates imply unreasonable long-run values for the inflation target, NAIRU and the real rate. While the estimated parameters in (3) do not identify the specific values of the *-variables, they do identify a *region* in which the *-variables live. Previous literature has made use of this fact. For instance, Clarida, Gali and Gertler (2000) estimate a variant of (1) and report estimates for π^* by assuming specific values for r^* (explicitly as its average value over the period) and u_t^* (implicitly as estimated trends). We can similarly use our estimates.

In particular, combining (1) with (3) one obtains

$$r^* = C + (\beta - 1)\pi^* + \delta u^*. \tag{5}$$

We choose to get implications for r^* by calibrating values for π^*, u^* because $\beta - 1$ and δ enter multiplicatively in (5). If we instead implied either π^* or u^*, then the reciprocals of $\beta - 1$ or δ would enter their expression and could possibly give rise to extreme behavior whenever $\beta - 1$ or δ are close to zero. For our baseline, we set $u^* = 5.5$ to be consistent with the sample average of the long-horizon unemployment forecasts, and set $\pi^* = 2.5$ to balance recent Federal Reserve communication of its 2% inflation goal with higher estimates of the inflation target from previous literature.[16] We argue that a reasonable range for the real rate r^* also comes from the long-horizon forecasts, whose mean and median is 2.0 with an interquartile range of 1.8-to-2.3 and the min/max of 1.0/3.3. A real-rate-region is formed by letting β and δ vary over their respective normal-approximation 95% confidence intervals $\hat{\beta} \pm 2\widehat{se(\beta)}$ and $\hat{\delta} \pm 2 * \widehat{se(\delta)}$ for specific choices of the constant C.[17]

The top row of Figure 4 shows contour plots of the real-rate-regions implied by the pre-

[16]Recent communication of the goal is found in FOMC (2013). Clarida, Gali and Gertler (2000) estimate $\pi^* = 3.6$ on their baseline Volcker-Greenspan sample (Table II, page 157) but $\pi^* = 2.5$ for their Greenspan sample with an annual inflation horizon (Table V, page 162).

[17]Inspecting (5) we see that the r^*-region shifts as follows:

$$r^* \begin{cases} \text{increases} & \text{for} \quad C \uparrow \\ \text{increases} & \text{for} \quad \pi^* \uparrow \text{ when } \beta > 1 \quad \text{and} \quad \text{decreases} \quad \text{for} \quad \pi^* \uparrow \text{ when } \beta < 1 \\ \text{decreases} & \text{for} \quad u^* \uparrow \text{ when } \delta < 0 \quad \text{and} \quad \text{increases} \quad \text{for} \quad u^* \uparrow \text{ when } \delta > 0 \end{cases}$$

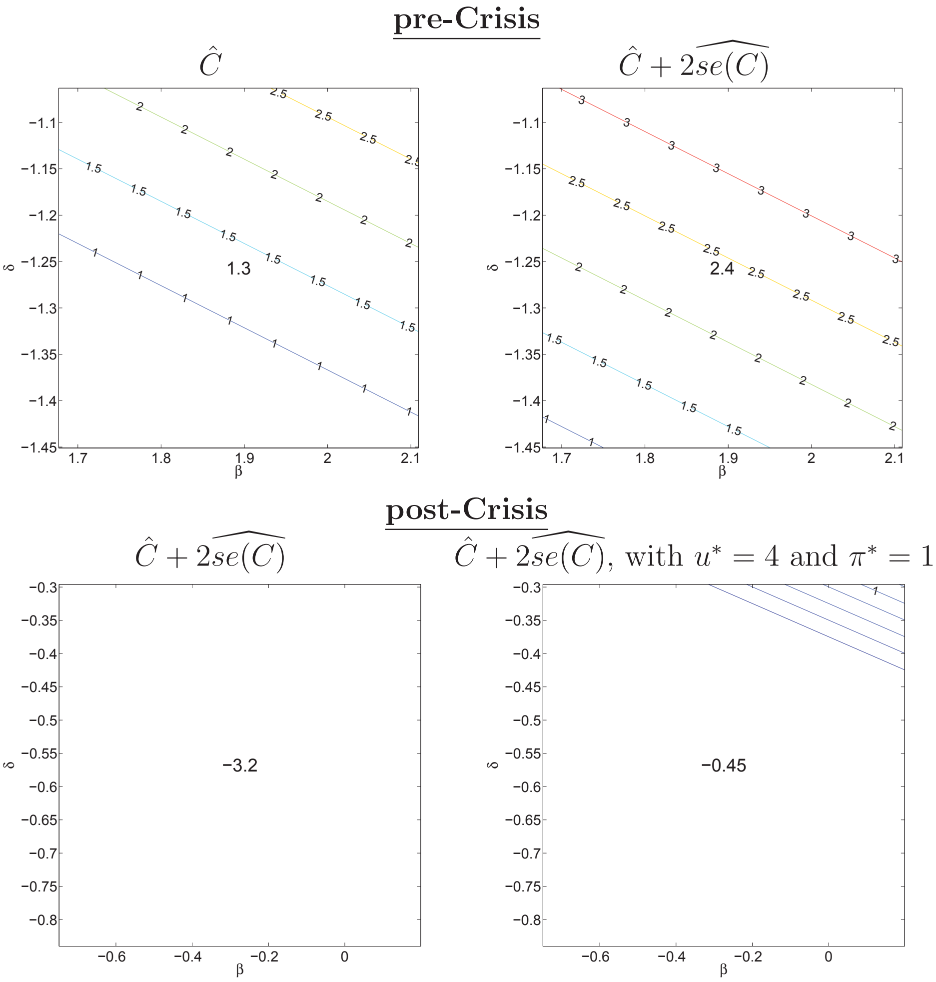

Figure 4: Implied real-rate-regions using Historical Data Including the ZLB Period

Notes: Contour plots of the real rate r^* implied by the estimates of the constant in (5) based on historical data 1986–2011 in Table 1. $u^* = 5.5$ and $\pi^* = 2.5$ unless otherwise noted. Only contours between 0.5 and 3.5 are drawn.

Crisis estimates in Table 1 for two values of the constant: The point estimate $\hat{C} \approx 6$ and the upper edge of the constant's normal-approximation 95% confidence interval $\hat{C} + 2\widehat{se(C)} \approx 7$. Note that we always obtain a higher real-rate-region for a higher value of C (see footnote 17). Either region contain reasonable values for the real rate, and the implied point estimate at $(\hat{C}, \hat{\beta}, \hat{\delta}) \approx (6, 1.9, -1.3)$ is 1.3.

Consider the real-rate-regions implied by post-Crisis estimates. The bottom-left panel of

Figure 4 shows that, in our baseline, even the upper edge of C's confidence interval implies a negative real rate. Noting the sign of the post-Crisis estimates, we can raise the implied r^* by lowering either π^* or u^* (see footnote 17). In the bottom-right panel we lower the inflation target to 1% and the NAIRU to 4%, either of which is an implausibly low value. Notwithstanding, the implied real rate is still negative and the real-rate-region is above 1% only for β, δ simultaneously at the upper edge of their confidence intervals.

In summary, by including the ZLB-censored short rates the post-Crisis historical data regression estimates in Table 1 do not imply economically reasonable values of r^*, π^* and u^* simultaneously. Moreover, there are known statistical reasons why a linear regression run on censored observations is unreliable.

Excluding the ZLB A natural possible solution to these problems is to end our data in December 2008—excluding the censored observations during the ZLB period—which leads to our second set of results, presented in Table 2. Unsurprisingly, we obtain different conclusions by excluding the censored observations—we now find that the unemployment response is quite similar before and after August 2007 (Table 2 top panel). Furthermore, the real-rate-region for the post-August-2007 estimates contains reasonable values (Figure 5 left panel).

Nevertheless, one may be concerned that there are only sixteen monthly observations in the post-Crisis regime—this paucity of post-Crisis data only worsens when we investigate the results' robustness to later Crisis start dates. Obviously, using this sample we *cannot* say anything statistically significant about a break potentially occurring in September 2008 (when Lehman collapsed) or December 2008 (when the Fed funds rate hit the ZLB) since the post-Crisis estimates are undefined. But the small number of observations is also an

Table 2: Estimates from Historical Data Excluding the ZLB Period

	pre-Crisis	post-Crisis
August 2007 Break		
Inflation response	1.893	−0.723
standard error	(0.108)	(0.957)
Chow test	0.047	
Unemployment response	−1.257	−1.352
standard error	(0.097)	(0.308)
Chow test	0.689	
Constant	6.018	11.08
standard error	(0.524)	(3.877)
R^2	0.807	
March 2008 Break		
Inflation response	1.908	1.763
standard error	(0.107)	(0.331)
Chow test	0.973	
Unemployment response	−1.243	−0.375
standard error	(0.097)	(0.112)
Chow test	0.585	
Constant	5.864	−0.517
standard error	(0.531)	(1.294)
R^2	0.549	

Notes: Monthly data from the Federal Reserve Board and BLS, January 1986 to December 2008, regression with a constant. Robust standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

issue for the earlier potential Crisis start dates. Consider taking the breakpoint to be March 2008, when Bear Stearns received a Federal Reserve bailout but subsequently collapsed. The bottom panel of Table 2 shows that the estimates are radically changed relative to the August 2007 breakpoint. We now find that the inflation response is unchanged, statistically-speaking. Meanwhile, the unemployment response point estimate is lower by two-thirds yet the Chow test cannot reject the no-break null.

The main problem using historical data excluding the ZLB period is statistical: There are too few observations post-Crisis. While the real-rate-region (Figure 5 right panel) does

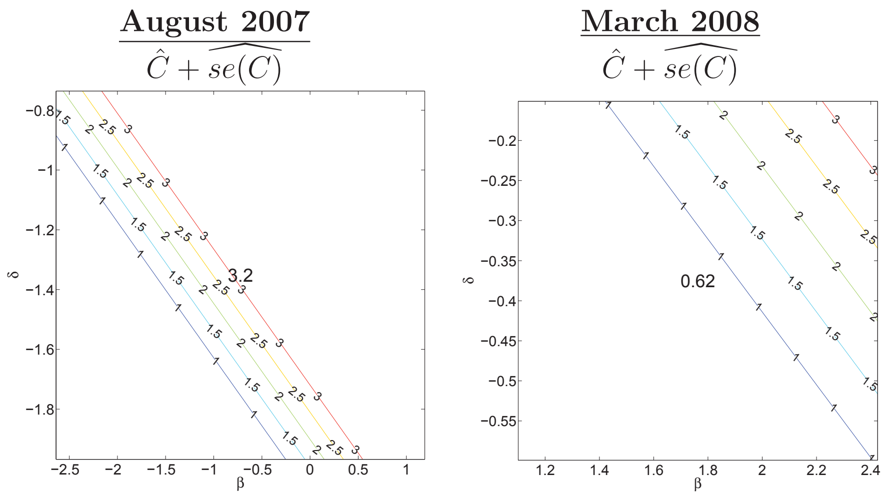

Figure 5: Implied real-rate-regions using post-Crisis estimates from Historical Data Excluding the ZLB Period

Notes: Contour plots of the real rate r^* implied by the estimates of the constant in (5) based on historical data 1986–2008 in Table 2. $u^* = 5.5$ and $\pi^* = 2.5$ unless otherwise noted. Post-Crisis estimates for two potential breakpoints: August 2007 and March 2008. Only contours between 0.5 and 3.5 are drawn.

not throw up economic warnings, as econometricians we are skeptical that eight post-break observations are sufficient to provide meaningful estimates. The small number of post-break observations contribute to the Chow test's low power, yielding an unemployment break p-value of 0.6 despite the large difference between the pre-Crisis estimate of -1.24 ($se = 0.10$) and post-Crisis estimate of -0.38 ($se = 0.11$). This simply manifests the known fact that break tests are unreliable for points near the edge of the sample.[18]

In summary, our research question is not ably answered by the historical data, even once one excludes the ZLB period. We now turn to our benchmark results using forecast data.

4.1.2 Forecast Data

Our benchmark results using forecast data are reported in Table 3. For all the breakpoints we consider, the fit is quite good and a linear function of the expected inflation gap, expected

[18]See the discussion of asymptotic behavior of test statistics in Andrews (1993, p.826).

unemployment gap, and long-run nominal rate $r^* + \pi^*$ explains over 75% of the variation in expected short rates. We view this as reasonable since simple rules are both widely used to communicate policy (described in speeches by Kohn 2007 and Bernanke 2010) and, as seen above, quite consistent with the pre-Crisis historical record. Furthermore, with the forecast data there is no need to calculate real-rate-regions and ask whether or not the parameter estimates are consistent with reasonable values for r^*, π^* and u^*—these long-run values are pinned down by the data themselves and we can see from Figures 1–3 that they are economically sensible throughout (of course, one gets the real rate by subtracting the inflation rate in Figure 2 from the nominal short rate in Figure 1).

Pre-Crisis estimates from forecast data are similar to pre-Crisis estimates from historical data. Focusing on the August 2007 breakpoint as our benchmark, the estimated inflation response of 1.97 (se = 0.26) on forecast data is quite close to the 1.89 (se = 0.11) on historical data. The unemployment response estimate of −0.61 (se = 0.11) is quite close to Clarida, Gali and Gertler's −0.56. This assures us that forecast data and historical data provide insight into the same policy rule when both are uncensored.

Table 3 provides our main finding of the perceived shift in Federal Reserve policy, which is robust to the various breakpoints considered: In the wake of the Global Financial Crisis, the central bank's response to inflation weakened but its unemployment response remained strong and perhaps got strengthened. For three of the four breakpoints considered, including the August 2007 benchmark, the downward shift in inflation response is statistically significant (p-values below 0.01, 0.01 and 0.07 for August 2007, March 2008 and September 2008, respectively) and in the remaining case the evidence is marginal (p-value just above 0.11).[19]

[19] The March 2008 and September 2008 breakpoints are the places where the robust Chow test provides

Table 3: Estimates from Forecast Data

	pre-Crisis	post-Crisis
August 2007 Break		
Inflation response	1.968	−0.894
standard error	(0.260)	(0.889)
Chow test	0.002	
Unemployment response	−0.608	−1.137
standard error	(0.112)	(0.170)
Chow test	< 0.001	
R^2	0.810	
March 2008 Break		
Inflation response	1.954	−0.770
standard error	(0.259)	(0.918)
Chow test	0.004	
Unemployment response	−0.650	−1.11
standard error	(0.117)	(0.175)
Chow test	0.002	
R^2	0.805	
September 2008 Break		
Inflation response	1.774	0.832
standard error	(0.282)	(0.364)
Chow test	0.069	
Unemployment response	−0.800	−0.796
standard error	(0.162)	(0.091)
Chow test	0.972	
R^2	0.781	
December 2008 Break		
Inflation response	1.745	0.599
standard error	(0.277)	(0.441)
Chow test	0.111	
Unemployment response	−0.775	−0.831
standard error	(0.139)	(0.108)
Chow test	0.605	
R^2	0.780	

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts, January 1986 to August 2011. One-year-ahead forecasts of inflation and unemployment minus long-horizon inflation and unemployment forecasts to construct estimated gaps, and the long-horizon real rate is used. Robust standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

better inference. The standard Chow test yields p-values of 0.32 and 0.34, respectively, for these two specifications and so does not accurately reflect the sizeable difference in point estimates, owing to its inaccurate homoskedasticity assumption.

The evidence of an upward shift in the unemployment response is mixed—for the earlier two breakpoints the shift is statistically significant at the 0.1% level while for the later two the test fails to reject.

Table 3's results are concerning for a few reasons. One, for early breakpoints we find a *negative* albeit statistically insignificant inflation response post-Crisis, while for later breakpoints this coefficient is positive and sometimes significant. Two, there is a wide variance of post-Crisis unemployment response estimates across breakpoints. These two findings may not be too surprising given the fact that the post-Crisis period has only a moderate number of monthly observations available. More surprising is the finding that *pre*-Crisis response estimates vary as much as they do, given that the various breakpoints represent a swing of only about a dozen observations combined with more than 250 observations prior to August 2007. Additionally, the Chow test evidence for the unemployment response coefficient varies substantially between significance for the early breakpoints and insignificance for the later breakpoints.

4.1.3 Pinning Down the Crisis Period

What features of the data are responsible for these estimates' differences across potential breakpoints? We investigate this graphically in Figures 6 and 7 by plotting the ingredients and outputs of the forecast data results. Figure 6 plots the forecast data entering into the regression. Recalling (2), our estimation chooses response coefficients which, multiplying the expected inflation gap $\pi_E - \pi^*$ and the expected unemployment gap $u_E - u^*$, fit as closely as possible the expected cyclical policy rate $i_E - r^* - \pi^*$.[20] We also consider two fitted values,

[20]The full history of these series is plotted in Figure A1 of the appendix.

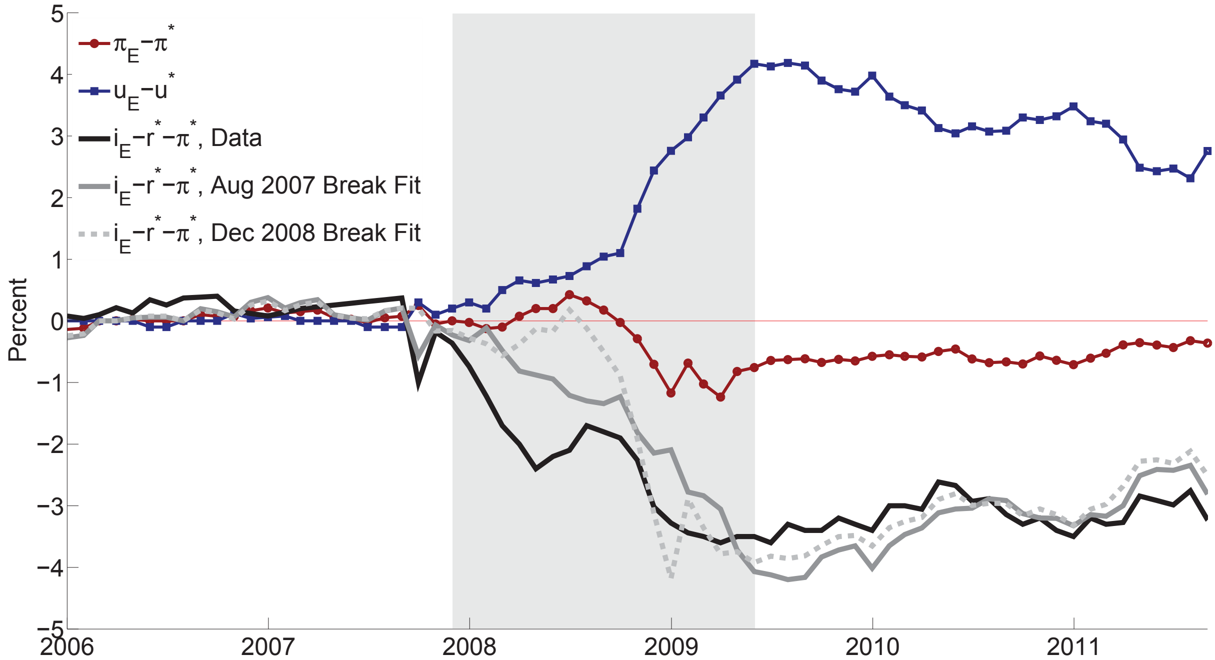

Figure 6: Expected Gaps, Expected Cyclical Policy Rate, and Fitted Values
Notes: Data from Blue Chip Economic Indicators on median one-year-ahead forecasts of the inflation gap, the unemployment gap, and the cyclical policy rate. Fitted values of expected cyclical policy rate from two break specifications, August 2007 and December 2008, in Table 3. Shaded are the NBER recession periods.

from the August 2007 break ($i_E - r^* - \pi^*$, Aug 2007 Break Fit) and alternately from the December 2008 break ($i_E - r^* - \pi^*$, Dec 2008 Break Fit) specifications in Table 2. The residuals between the actual $i_E - r^* - \pi^*$ and the fitted values are plotted in Figure 7.

The most noticeable aspect is the unprecedentedly sharp fall of the actual $i_E - r^* - \pi^*$ starting in September 2007—the extraordinarily strong pace of expected policy easing following August 2007 is a key determinant of how estimates differ across breakpoints. This extraordinary behavior helps explain the shift in pre- and post-Crisis estimates for varying breakpoints. The sharp drop in $i_E - r^* - \pi^*$ leads to a large residual opening up from October 2007 through June 2008 in Figure 7, as expected short rates fell much more than estimated policy rules can explain, especially if the response coefficients are not allowed to break until

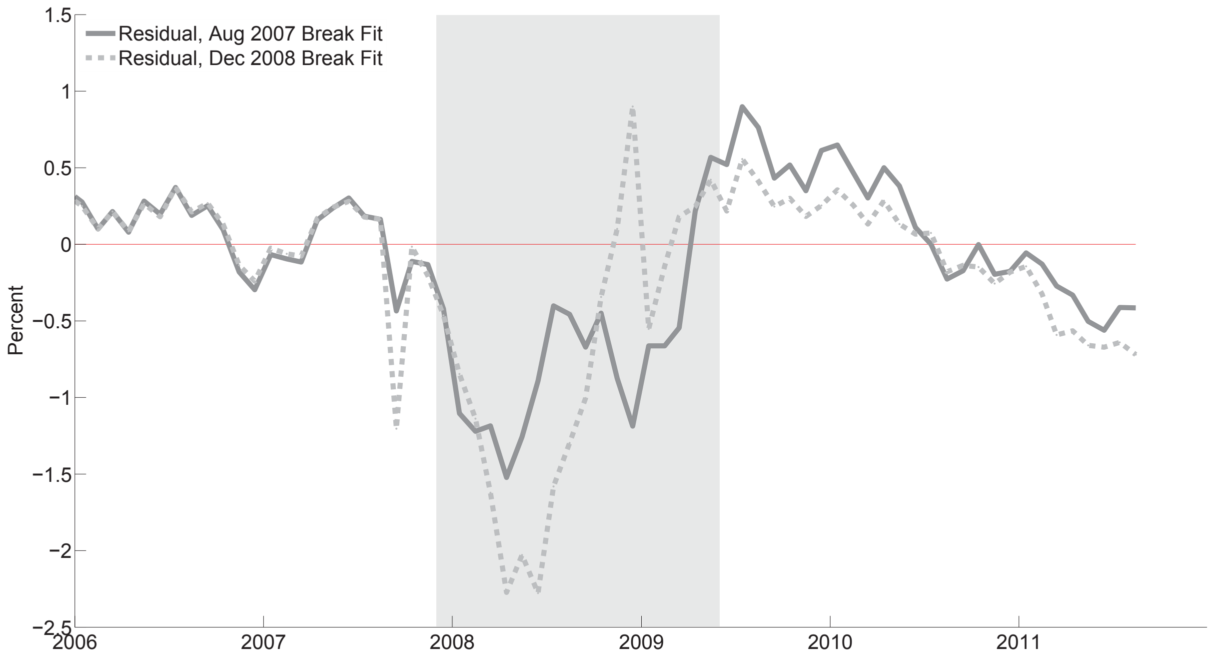

Figure 7: RESIDUALS OF EXPECTED CYCLICAL POLICY RATE

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts. Residuals of expected cyclical policy rate from two break specifications, August 2007 and December 2008, in Table 3. Shaded are the NBER recession periods.

later (the gray dashed line). Remarkably, the expected inflation gap meanwhile first hovers at zero and then *rises* through June 2008. This most likely reflects the unprecedented surge in commodity prices that occurred the beginning of 2008.

The *increase* in expected inflation pressure couples with the sharp move *down* in expected cyclical policy rates to push down on the inflation response estimate involving the period of October-2007–June-2008. For early breakpoints the October–June period pushes down on the post-Crisis β-estimate, while for later breakpoints the October–June period pushes down on the pre-Crisis β-estimate. Hence, October–June expected inflation/cyclical-policy behavior contributes to how inflation response estimates differ across breakpoints.

In addition to the October–June behavior of the data, notice the curious September 2007 datapoint immediately to the right of the vertical line labeled "Aug" in both Figures 6 and 7. This sizeable residual, which is not a data measurement error, reflects the markets' quickly-reversed jump down in one-year-ahead short rate expectations (the solid black line in Figure 1) following BNP Paribas' fund-withdrawal freezes in August 2007—the expected one-year-ahead short rate falls from 4.9 in August to 3.5 in September and bounces back to 4.3 in October. During that month liquidity started drying-up several key markets, including the fed funds market. In fact, on several days in August 2007 the effective fed funds rate deviated from target by more than 50 basis points (as much as 71 basis points on August 14), deviations last seen in the days after the September 11^{th} of 2001. Remarkably, even September 2007's sharp drop in expected one-year-ahead short rates was too optimistic a forecast, as actual short rates a year later in September 2008 would be below 2 percentage points and heading lower. Against this drastic change in the outlook for future short rates, September 2007's one-year-ahead expected inflation blips *up*, exactly opposite of what the typical, positive Taylor rule coefficient would suggest.

Taken together, the unusual October-2007–June-2008 and September-2007 data shift pre-Crisis β-estimates from 1.97 for the August 2007 breakpoint, to 1.75 for the December 2008 breakpoint. The other side of the coin is that the post-Crisis β-estimates vary from -1.14 to -0.80.

Given these features of the inflation data, the estimation seeks to fit the extraordinarily strong pace of expected policy easing by the unemployment gap. The problem it encounters is that the unemployment gap (the squared-solid blue line in Figure 6) rises at only a moderate pace from September 2007 through September 2008. This moderate pace reflects the following two facts. Although expected one-year-ahead unemployment rises steeply and mirrors expected policy's steep drop (the dashed blue line in Figure 3), during the same period the market's perceived NAIRU (the dash-dotted blue line in Figure 3) also rises substantially. Taken together, the expected unemployment gap rises only at a moderate pace. Once Lehman Brothers collapses in September 2008, expected one-year-ahead unemployment rises much more rapidly while the rise in NAIRU tapers off (Figure 3): Hence, the expected unemployment gap rises steeply (Figure 6).

All of this behavior serves to strengthen (that is, make more negative) the unemployment response estimate involving the September-2007–September-2008 period, as a modest rise in unemployment gap is relied upon (given the unusual behavior of expected inflation at the same time) to explain the sharp drop in expected cyclical policy. The September–September months shift pre-Crisis δ-estimates from -0.61 for the August 2007 breakpoint, to -0.78 for the December 2008 breakpoint.

Summing up, Table 3 allows for the conclusion that β has shifted down, but does not pin down the inflation response's post-Crisis value. Table 3 leads us to the conclusion that δ is at least no lower after the Global Financial Crisis than before, but gives mixed messages on whether or not the unemployment response is actually higher post-Crisis. This section shows why the few observations between August 2007 and December 2008 lead to our results' sensitivity. Can we or should we adjust our estimation in light of these very unusual and influential observations?

Sensitivity: Statistical Fixes There exist statistical ways of dealing with these few influential observations: Switch our estimation to a procedure that is more robust with respect to such outliers. When we do so, the pre-Crisis estimates are more stable across breakpoints, but the post-Crisis estimates still widely vary. In particular, we considered both weighted least squares (following Huber 1981) and median regression (following Koenker and Bassett 1978), either of which in different ways downweight large residuals relative to ordinary least squares.[21] The results are relegated to Tables A1 and A2 in the appendix and we summarize the results here. Using these robust estimation procedures, the pre-Crisis β-estimates vary less, from 2.00 to 1.88, and the δ-estimates vary less, from -0.59 to -0.68—as expected, the range has been compressed as the large outlier observations at the end of the pre-Crisis sample are evenly-weighted against the preceding two decades' of data.

Nevertheless, even with robust estimation techniques the post-Crisis point estimates and the Chow test evidence continue to widely vary across breakpoints. Given that the different breakpoints define post-Crisis samples varying in size from 32 to 48 observations, it is to be expected that the October-2007–June-2008 and the September-2007–September-2008 observations exert tremendous influence on the post-Crisis estimates—they contribute up to one-quarter of the post-Crisis months.

Meanwhile, Davidson and MacKinnon's (1993) robust Chow tests based on Gauss-Newton regressions carry both advantages and disadvantages which are visible here. An advantage is that they only require $\sqrt{T}$-consistent point estimates to be valid—thus we are justified in using them with any of ordinary least squares, weighted least squares, or median regression estimates, making the test quite convenient. A disadvantage, however, is that this test statistic is not specially-designed for the models one might have in mind when, say, median regression is used instead of ordinary least squares—thus the test results look quite similar

[21] For residual ϵ, weighted least squares uses the weight $\frac{1}{\max(1,|\epsilon|)}$. We use Roger Koenker's implementation of quantile regression to estimate the median regression: All residuals are weighted equally but according to their absolute value $|\epsilon|$ instead of squared value ϵ^2 as in ordinary least squares.

across breakpoints whichever estimation method we choose.

Sensitivity: Economic Fix Instead of further fine-tuning our econometric procedure to get blood from a stone, we suggest a simple theoretically-motivated fix: Exclude from our sample the months during which the Global Financial Crisis gets underway.

As we've noted, there is no definitive starting date for the crisis; we have considered several possibilities, the reason being that different important financial events occurred at different times between August 2007 and December 2008. It is perhaps unsurprising that these months' data are at once so influential to our estimates and so atypical relative to earlier and later observations. Financial markets, professional forecasters, and professional economists fell into the Crisis uncertain as to how the world was changing.

Recently, Coibion and Gorodnichenko (2012) used surveys of professional forecasts to empirically compare several economic theories of information rigidity. Their findings support the idea that economic agents face frictions in acquiring and processing information, even when that information is solely about past values of random variables. On casual inspection, it appears plausible that agents face no less friction to acquiring and processing information which causes alteration to their model of how random variables are interrelated. Coibion and Gorodnichenko's (2012) recent evidence could suggest that if the market adjusted its perception of Federal Reserve policy, it is likely to have done so slowly over time.

If this is the case, that agents gradually learn about a new policy regime, then an explicitly time-varying parameter model of the perceived policy rule, somewhat similar to Boivin (2006), would be more appropriate. However, a time-varying parameter model would present at least three unattractive features in our case. First, the model requires a characterization of the evolution of policy which would require additional assumptions for how observed outcomes are translated into agents' parameter estimates. We would like to keep our assumptions few and simple. Second, our maintained assumption is that Federal Reserve's short-rate policy is fixed except for a singular event like the Global Financial Crisis. This is a hypothesis that merely updates the now-conventional historical reading found in Clarida, Gali and Gertler (2000), that short-rate policy changed at the singular event of high U.S. inflation at the end of the 1970s (or, equivalently, with Paul Volcker's Federal Reserve Board chairmanship). Given this maintained assumption, the time-varying parameter model will overstate the amount of policy change.[22] Third, a time-varying parameter model would do nothing to counteract the unusual behavior of the data (e.g. the October-2007–June-2008 behavior mentioned above) that so forcefully affects our estimates. It is our view that those events tell us about the markets' chaotic response to incoming unprecedented events, not evolving perceptions of the Federal Reserve's short-rate policy.

Therefore, in our final and preferred specification we simply exclude the chaotic September 2007 to December 2008 months from our sample, owing to the idea that it may have taken agents time to learn about the Federal Reserve's changed policy stance. By waiting until January 2009, we hope to avoid capturing chaotic observations that only add noise to our estimates, sap power from our statistical tests, and do not capture the market's new beliefs. It is important to note that so doing leaves our results qualitatively unchanged. We report our preferred results in Table 4 that exclude data from September 2007 to December 2008. In essence, these results allow for a break anytime between August 2007 and January 2009.[23]

[22]See the estimated time-varying parameters in Boivin (2006) or Ang, Boivin, Dong and Loo-Kung (2011).

[23]Only the statistical power of the Chow tests is qualitatively changed (the p-values rise even though the point estimates differ by more) by excluding data through June 2009, the end of the NBER recession.

Table 4: ESTIMATES FROM FORECAST DATA, EXCLUDING SEPTEMBER 2007 TO DECEMBER 2008

	pre-Crisis	post-Crisis
Ordinary Least Squares		
Inflation response	1.968	0.599
standard error	(0.260)	(0.441)
Chow test	0.091	
Unemployment response	−0.608	−0.831
standard error	(0.112)	(0.108)
Chow test	0.071	
R^2	0.828	
Weighted Least Squares		
Inflation response	1.982	0.669
standard error	(0.104)	(0.584)
Chow test	0.086	
Unemployment response	−0.585	−0.814
standard error	(0.065)	(0.110)
Chow test	0.077	
Median Regression		
Inflation response	1.959	1.218
standard error	(0.105)	(0.593)
Chow test	0.068	
Unemployment response	−0.576	−0.682
standard error	(0.056)	(0.108)
Chow test	0.059	

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts, January 1986 to August 2007 and January 2009 to August 2011. One-year-ahead forecasts of inflation and unemployment minus long-horizon inflation and unemployment forecasts to construct estimated gaps, and the long-horizon real rate is used. Estimation by ordinary least squares, weighted least squares using $\frac{1}{\max(1,|\epsilon|)}$, or median regression, as indicated. Robust standard errors for ordinary least squares, non-unadjusted standard errors for weighted least squares, and bootstrapped standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

Our earlier results in Table 3 indicated that, in the wake of the Global Financial Crisis, the Federal Reserves's response to inflation weakened but its unemployment response remained strong and perhaps strengthened. Table 4 only adds specificity to the results. We find that the post-Crisis inflation response is statistically zero while the unemployment response has shifted up by a statistically-significant amount. The breaks are significant at the 10%

level for all ordinary least squares, weighted least squares, and median regression estimates. Focusing on the OLS estimates, we find that the Global Financial Crisis somehow caused the Federal Reserve's inflation response to fall from 1.97 ($se = 0.26$) to 0.60 ($se = 0.44$), and unemployment response to rise from -0.61 ($se = 0.11$) to -0.83 ($se = 0.11$).[24]

4.2 Understanding the Residuals

The solid line of Figure 8 represents the residuals from our preferred specification, the ordinary least squares results in Table 4.[25] As previously mentioned, in a conventional policy rule regression using (3) with realized data, residuals are often interpreted as capturing policy "shocks." However, in our benchmark rule using (2) with forecast data, the residuals represent a shift of expected policy that is not accounted for by a linear combination of inflation and unemployment forecasts. We find a few features noteworthy. First, Greenspan's chairmanship starting in mid-1987 was characterized by reasonably small residuals until around 2001. This means that forecasters believed that short rates largely reflected economic conditions according to a simple Taylor rule.

In the first half-decade of the 21st century, the residuals are persistently negative. This implies that market professionals during this time expected more expansionary policy than would otherwise be called for by economic conditions. Starting in 2005 the expected expansionary boost is steadily removed, and by the time of Bernanke's chairmanship starting in

[24] A possible concern is that the ZLB is affecting these estimates by holding up the average level of expected cyclical policy response, which affects these estimates since (2) is estimated without a constant. We can address this by estimating the rule in differences, where such a constant ZLB-effect would vanish. Doing so, the inflation response is still estimated to have fallen significantly (p-value $= 0.054$), and the unemployment response becomes stronger from pre-Crisis's -0.76 ($se = -0.09$) to post-Crisis's -0.94 ($se = 0.28$) albeit insignificantly.

[25] As mentioned above, our inference is robust to serial correlation in the residuals. Figure 8 makes it clear that this robustness was required.

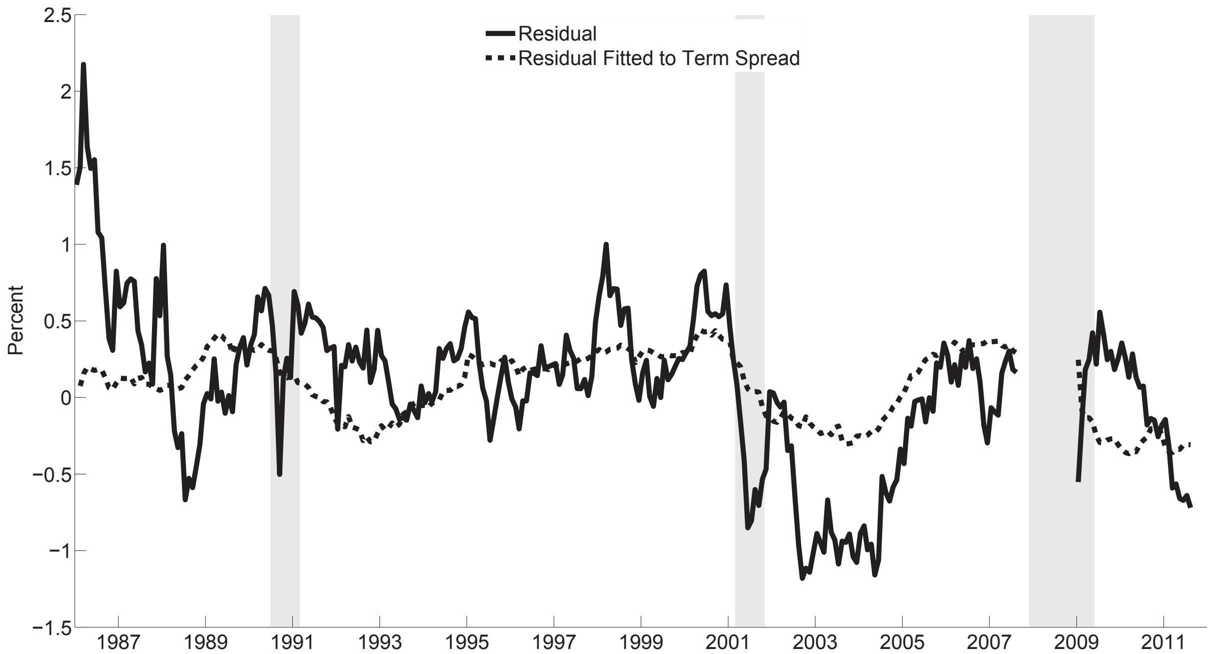

Figure 8: Residuals of the One-Year-Ahead Short-Rate Equation

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts, January 1986 to August 2007 and January 2009 to August 2011. Residual from the preferred specification, Ordinary Least Squares in Table 4, plotted as solid line. Dashed line is fitted value of regression of the residual on the actual U.S. term spread (10y-2y). Shaded are the NBER recession periods.

early 2006 the residuals are nearer to zero. After the chaotic September-2007–December-2008 period (deleted in Figure 8), the residuals trend negative as an expected expansionary boost increases throughout the ZLB period.

What accounts for the residuals' fluctuation? Atkeson and Kehoe (2009) argue that "business cycle movements of the short rate arise as a result of the Fed's endogenous policy response to exogenous business cycle fluctuations in risk." Hence, it may be that an important component of these residuals' evolution comes from forecasters' perceived risk.[26]

We proxy for risk by the term spread, defined as the difference between 10-year and 2-

[26]We also ran a regression of the residuals on nowcasts of unemployment and inflation (as well as their forecasts) to check if the shifts reflect economic conditions at the time of the forecasts or dynamic adjustment. The coefficients were statistically insignificant with minimal variation explained.

year bond yields, a common measure financial market risk at a business cycle frequency.[27] A regression of the residuals on the term spread shows a strong connection that is both economically and statistically significant: The coefficient estimate is -0.24 with robust t-statistics of -8.25 and R^2 of 0.19.[28] The estimate implies that forecasters translate a 100 basis point *increase* in the term spread with about a 25 basis point *decrease* in the future policy rate, irrespective of what inflation and the unemployment rate are forecasted to be.

The dashed line in Figure 8 plots the residuals that are fitted to the term spread. The fit is reasonably good at a business cycle frequency during the 2000s until the Crisis. Thereafter, the fitted values move in an opposite direction from the residuals. We find this interesting because one of the Fed's *unconventional* policy responses (including large-scale asset purchases) during the ZLB period has been to target medium-maturity treasuries, directly affecting the term spread. One explanation is that prior to the Crisis our policy residuals are indeed capturing the market's perceived risk. Then the term-spread fitted values might move contrarily during the Crisis because the Fed is directly lowering the measured term spread yet not lowering perceived risk.[29]

[27] We are careful here to measure the term spread in the month *before* the survey month, since the economic indicator survey is collected during the first few days of the month it is released, therefore the prior month's bond prices are what is known to survey respondents at the time they report their forecasts.

[28] As for the regressions of the residual, we again use three lags in the Bartlett kernel to calculate robust standard errors following Newey and West (1987). Similar results obtain from the residuals from any of Table 3's specifications.

[29] We also tried using the CME's VIX index to proxy for financial market conditions. A regression of the residuals on the VIX is statistically insignificant, and there is essentially no link between the VIX and the policy residuals. This is perhaps unsurprising given Atkeson and Kehoe's (2009) argument that the Fed responds to risk chiefly at a business cycle frequency, not to high-frequency financial market conditions. It may be that the VIX largely captures financial market conditions to which the Fed *does not* respond.

4.3 Unconstrained Short Rates and Asset Purchase Size

In any regression analysis, estimation results are usually used for two purposes: in-sample and out-of-sample analyses. To perform an out-of-sample analysis, it is important to obtain an in-sample estimation period that has a similar characteristic to the out-of-sample period in mind. Since it is commonly believed that many economic phenomena may have changed with the emergence of the Global Financial Crisis, it is critical to estimate policy rules with a sample after 2008—as we do in this paper.[30] We make use of estimation results based on forecast data to construct an implied unconstrained nominal short-term interest rate would be based on actual data were it not for the zero lower bound (see Rudebusch 2009).

In Figure 9, the circled-red and squared-blue lines represent the inflation gap and the unemployment gap, respectively, that are constructed using each month's long-horizon forecasts subtracted from the actual values observed. The solid black line represents the short rate implied by the post-Crisis estimates in Table 4 using those actual gaps. This line implies the unconstrained interest rate if there were not the ZLB constraint and if the estimated policy rule could set the short rate to any recommended value, including negative ones. This rate goes into the negative range by early 2009, is on the downward trajectory for most of 2009 and 2010, and reaches about -2% by late 2010 when the unemployment gap is the highest. These negative values can be related to the extra stimulus that is necessary due to the contractionary influence of the ZLB, and the FOMC's asset purchases are viewed

[30] Of course, it is necessary to assume that the in-sample relationship still holds out of sample after August 2011. We realize that there are significant changes in the monetary policy (e.g. forward guidance and large-scale asset purchases), but our working assumption is that the relationship of (2) remains stable both in sample and out of sample. As pointed out in Wright (2012), using asset purchases in identifying a policy change is also problematic since asset purchases affect asset prices at the announcement as well as at the implementation.

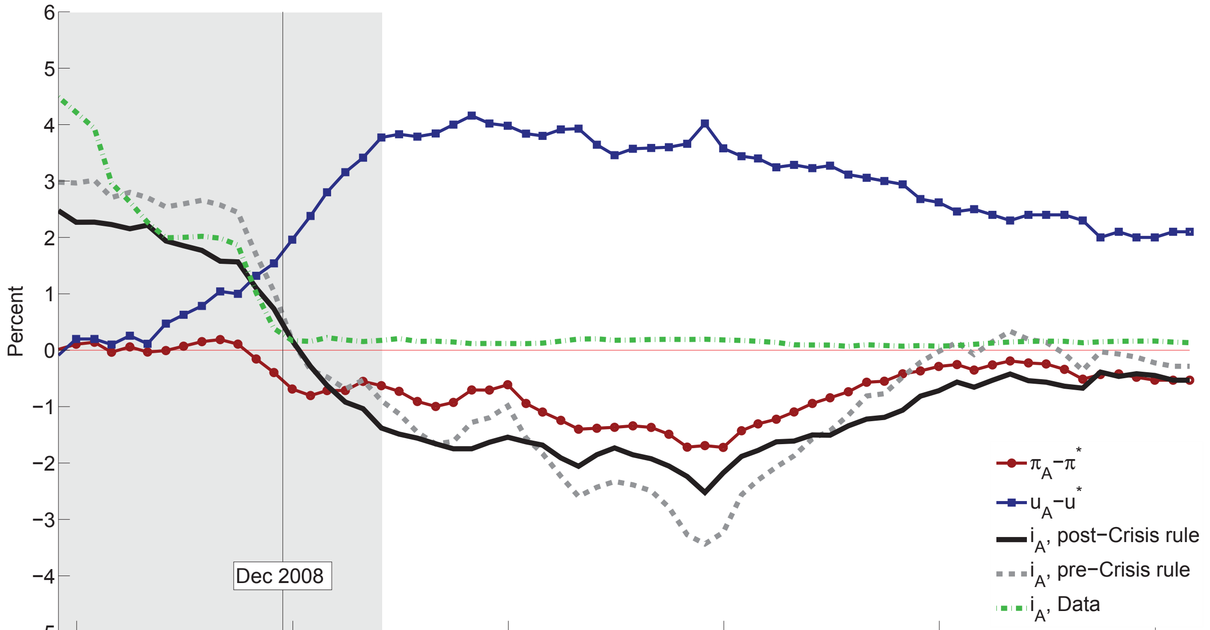

Figure 9: Actual Gaps and Implied Short Rates

Notes: Actual data combined with Blue Chip long-horizon forecasts of unemployment, inflation, and real rate to construct estimated gaps. Implied short rates use the long-horizon forecasts and gaps $\pi_A - \pi^*$ and $u_A - u^*$. Pre- and post-Crisis parameter estimates in Table 4 are used to construct implied rates "i_A, pre-Crisis rule" and "i_A, post-Crisis rule" respectively; the actual short rate is plotted as "i_A, Data." A one-time level-shift equates the implied rates to the actual short rate in December 2008, which is marked by a vertical line. Shaded are the NBER recession periods.

as providing such necessary stimulus. Since early 2011, the economy slowly recovered, and the unconstrained rate moved up slowly to zero. Our estimation sample stopped in the middle of 2011 when the FOMC introduced the calendar-based forward guidance, but our unconstrained-rate exercise could nevertheless continue.

In constructing "unconstrained" short rates, we take on board the sharp drop in rates during the Global Financial Crisis. So the implied rates we construct equal the actual short rate in December 2008, the start of the ZLB period. After this point, their evolution is determined by the response coefficient estimates applied to the actual inflation and unemployment rate gaps which have prevailed during the ZLB. We consider two such implied rates: One from the pre-Crisis response estimates and the other from the post-Crisis response estimates.

The dashed gray line represents the short rate implied by the pre-Crisis estimates in Table 4 based on the same gap values. Since the inflation response is much higher for the pre-Crisis sample than for the post-Crisis sample in Table 4, the implied rate based on the pre-Crisis rule stayed below that based on the post-Crisis rule between early-2010 and late-2011. Thereafter, the strong pre-Crisis response to inflation suggests that the desired short rate rose quickly and actually broke the ZLB in early-2012. But this did not in fact happen, consistent with the post-Crisis rule's weaker inflation response which has kept its implied short rate below the ZLB.

During the ZLB period, the Federal Reserve has conducted policy in at least two ways: through communication to the markets about the present and future course of conventional short-rate policy, and by unconventional asset purchases. These two tools most likely affect each other. Asset purchases "make up" for the inability of conventional short-rate policy to provide the desired level of stimulus, therefore it matters to know how much stimulus is lacking. A simple measure of the necessary stimulus is the unconstrained rate's area below zero. While their evolutions differ, both the pre-Crisis and post-Crisis rule suggest a similar amount of desired stimulus between January 2009 and February 2013.

In summary, the market's perception of Federal Reserve policy has changed since the Global Financial Crisis. Market participants believe that the central bank is more tolerant of inflation in the short run—consistent with popular commentators' interpretation of Federal Reserve communication during the ZLB period; furthermore, some evidence points to market participants' belief that the Federal Reserve has increased its response to the employment situation. Nonetheless, the significant change in short-rate policy does not appreciably affect the perceived stimulus provided by the Federal Reserve's large-scale asset purchases during the ZLB.

5 Conclusion

It is crucial to policymakers that they understand how their behavior is perceived by financial markets. Given the ongoing Global Financial Crisis, it is quite natural to ask: Have policymakers' behavior, or markets' perceptions thereof, changed? But the conventional method of empirically answering this question is unavailable because short rates are censored due to the zero lower bound. We get around this hurdle by using surveys of professional forecasters. The surveys allows us to use a simple regression to estimate the policy rule that has prevailed, at least in professional forecasters' perception, since the onset of the Global Financial Crisis. We find that this policy rule is significantly different from the one prevailing prior to the Global Financial Crisis.

We find that the Fed's inflation response has significantly fallen, by more than half. This could portend troublesome times ahead when the Fed does start to combat inflation, or instead this could reflect the success of the Fed's expansionary forward guidance that "a highly accommodative stance of monetary policy will remain appropriate for a considerable time after the economic recovery strengthens." Meanwhile, the Fed's response to unemployment has remained strong throughout the pre- and post-Crisis periods. Furthermore, we find evidence that the Federal Reserve responds to business-cycle frequency risk in addition to inflation and unemployment. Finally, although its short-rate policy rule has significantly changed since the Global Financial Crisis, this does not appreciably affect the amount of perceived stimulus provided by large-scale asset purchases during the ZLB.

References

Andrews, Donald W. K. 1993. "Tests for Parameter Instability and Structural Change with Unknown Change Point." *Econometrica*, 61: 821–856.

Ang, Andrew, Geert Bekaert, and Min Wei. 2007. "Do macro variables, asset markets, or surveys forecast inflation better?" *Journal of Monetary Economics*, 54(4): 1163–1212.

Ang, Andrew, Jean Boivin, Sen Dong, and Rudy Loo-Kung. 2011. "Monetary Policy Shifts and the Term Structure." *Review of Economic Studies*, 78(2): 429–457.

Atkeson, Andrew, and Patrick J. Kehoe. 2009. "On the Need for a New Approach to Analyzing Monetary Policy." In *NBER Macroeconomics Annual 2008, Volume 23. NBER Chapters*, 389–425. National Bureau of Economic Research, Inc.

Bernanke, Ben S., and Vincent R. Reinhart. 2004. "Conducting Monetary Policy at Very Low Short-Term Interest Rates." *American Economic Review*, 94(2): 85–90.

Carvalho, Carlos, and Fernanda Nechio. 2012. "Do People Understand Monetary Policy?" Federal Reserve Bank of San Francisco.

Clarida, Richard, Jordi Gali, and Mark Gertler. 2000. "Monetary Policy Rules And Macroeconomic Stability: Evidence And Some Theory." *Quarterly Journal of Economics*, 115(1): 147–180.

Cogley, Timothy, and Argia M. Sbordone. 2008. "Trend Inflation, Indexation, and Inflation Persistence in the New Keynesian Phillips Curve." *American Economic Review*, 98(5): 2101–2126.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, 120(1): 116 – 159.

Curdia, Vasco, and Michael Woodford. 2011. "The central-bank balance sheet as an instrument of monetary policy." *Journal of Monetary Economics*, 58(1): 54–79.

Davidson, Russell, and James G. MacKinnon. 1993. *Estimation and Inference in Econometrics.* Oxford University Press.

Devereux, Michael B., Gregor W. Smith, and James Yetman. 2012. "Consumption and real exchange rates in professional forecasts." *Journal of International Economics*, 86(1): 33–42.

Engelberg, Jospeh, Charles F. Manski, and Jared Williams. 2011. "Assessing the temporal variation of macroeconomic forecasts by a panel of changing composition." *Journal of Applied Econometrics*, 21: 1059–1078.

Federal Open Market Committee. 2013. "Statement on Longer-Run Goals and Monetary Policy Strategy." Federal Reserve System.

Fendel, Ralf, Michael Frenkel, and Jan-Christoph Rulke. 2011. ""Ex-ante" Taylor rules and expectation forming in emerging markets." *Journal of Comparative Economics*, 39: 230–244.

Gerlach, Stefan, and John Lewis. 2011. "ECB reaction functions and the crisis of 2008." Centre for Economic Policy Research.

Gorton, Gary. 2010. *Slapped by the Invisible Hand: The Panic of 2007.* Oxford University Press.

Gorton, Gary, and Andrew Metrick. forthcoming. "Securitized Banking and the Run on Repo." *Journal of Financial Economics.*

Hamilton, James D., Seth Pruitt, and Scott Borger. 2011. "Estimating the Market-Perceived Monetary Policy Rule." *American Economic Journal: Macroeconomics*, 3(3): 1–28.

Hirose, Yasuo, and Takushi Kurozumi. 2012. "Identifying News Shocks with Forecast Data." Australian National University.

Huber, Peter J. 1981. *Robust Statistics.* John Wiley & Sons.

Koenker, Roger, and Gilbert Bassett. 1978. "Regression Quantiles." *Econometrica*, 46(1): 33–50.

MacKinnon, James G. 1989. "Heteroskedasticity-Robust Tests for Structural Change." *Empirical Economics*, 14(2): 77–92.

Mitchell, Karlyn, and Douglas K. Pearce. 2010. "Do Wall Street economists believe in Okun's Law and the Taylor Rule?" *Journal of Economics & Finance*, 34: 196–217.

Newey, Whitney K., and Kenneth D. West. 1987. "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix." *Econometrica*, 55(3): 703–708.

Romer, Christina D, and David H Romer. 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review*, 90(3): 429–457.

Rudebusch, Glenn D. 2006. "Monetary Policy Inertia: Fact or Fiction?" *International Journal of Central Banking*, 2(4): 85–135.

Rudebusch, Glenn D. 2009. "The Fed's monetary policy response to the current crisis." *FRBSF Economic Letter*, 17.

Swanson, Eric T., and John C. Williams. 2012. "Measuring the effect of the zero lower bound on medium- and longer-term interest rates." Federal Reserve Bank of San Francisco Working Paper Series 2012-02.

Taylor, John B. 1993. "Discretion versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, XXXIX: 195–214.

Wright, Jonathan H. 2012. "What does Monetary Policy do to Long-term Interest Rates at the Zero Lower Bound?" *Economic Journal*, 122: F447–F466.

A Appendix

Table A1: Estimates from Forecast Data using Weighted Least Squares

	pre-Crisis	post-Crisis
August 2007 Break		
Inflation response	1.996	−0.750
standard error	(0.111)	(0.356)
Chow test	0.002	
Unemployment response	−0.592	−1.103
standard error	(0.069)	(0.071)
Chow test	< 0.001	
March 2008 Break		
Inflation response	1.988	−0.535
standard error	(0.112)	(0.371)
Chow test	0.005	
Unemployment response	−0.619	−1.061
standard error	(0.069)	(0.073)
Chow test	0.004	
September 2008 Break		
Inflation response	1.911	0.858
standard error	(0.110)	(0.488)
Chow test	0.089	
Unemployment response	−0.683	−0.788
standard error	(0.066)	(0.096)
Chow test	0.972	
December 2008 Break		
Inflation response	1.880	0.633
standard error	(0.107)	(0.619)
Chow test	0.132	
Unemployment response	−0.673	−0.823
standard error	(0.060)	(0.116)
Chow test	0.636	

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts, January 1986 to August 2011. One-year-ahead forecasts of inflation and unemployment minus long-horizon inflation and unemployment forecasts to construct estimated gaps, and the long-horizon real rate is used. Estimation by weighted least squares using $\frac{1}{\max(1,|\epsilon|)}$ weighting function. Standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

Table A2: Estimates from Forecast Data using Median Regression

	pre-Crisis	post-Crisis
August 2007 Break		
Inflation response	1.959	−1.407
standard error	(0.110)	(0.343)
Chow test	0.002	
Unemployment response	−0.576	−1.250
standard error	(0.069)	(0.068)
Chow test	0.001	
March 2008 Break		
Inflation response	1.943	0.383
standard error	(0.103)	(0.379
Chow test	0.006	
Unemployment response	−0.586	−0.854
standard error	(0.066)	(0.074)
Chow test	0.007	
September 2008 Break		
Inflation response	1.956	1.161
standard error	(0.085)	(0.486)
Chow test	0.076	
Unemployment response	−0.653	−0.696
standard error	(0.054)	(0.093)
Chow test	0.972	
December 2008 Break		
Inflation response	1.902	1.218
standard error	(0.094)	(0.641)
Chow test	0.115	
Unemployment response	−0.688	−0.682
standard error	(0.055)	(0.118)
Chow test	0.626	

Notes: Monthly data from Blue Chip Economic Indicators, median forecasts, January 1986 to August 2011. One-year-ahead forecasts of inflation and unemployment minus long-horizon inflation and unemployment forecasts to construct estimated gaps, and the long-horizon real rate is used. Estimation by median regression. Bootstrap standard errors are provided in parentheses. Robust Chow test p-value is for the null hypothesis that all parameters except for the indicated parameter break in the indicated month.

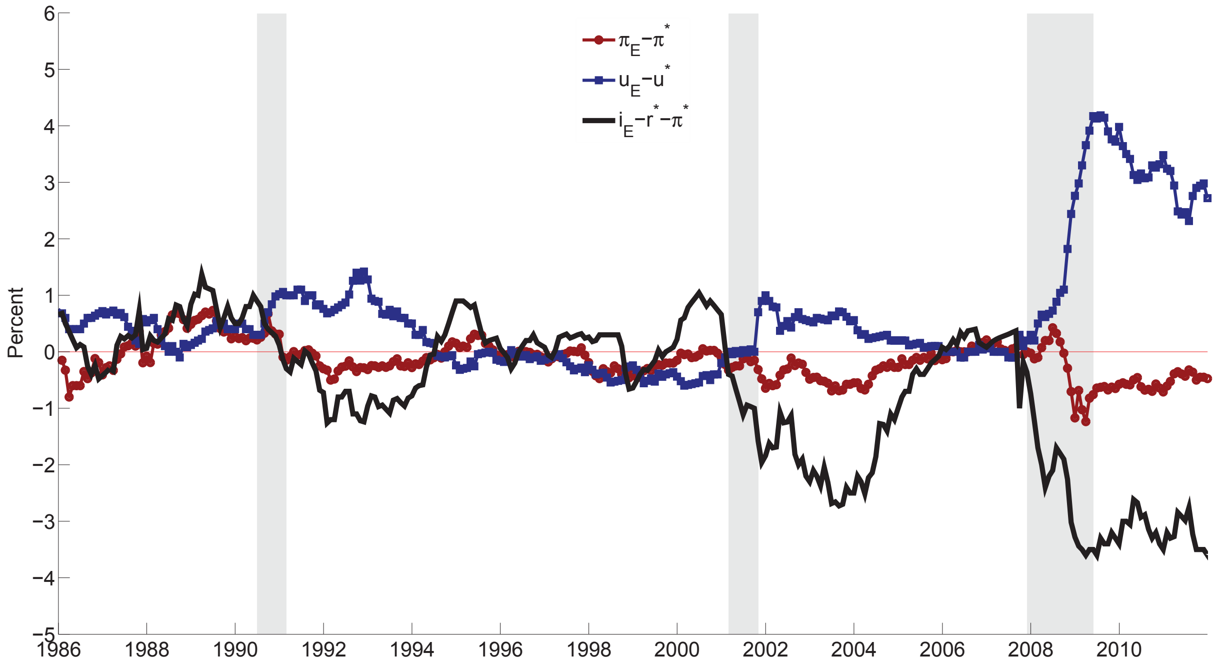

Figure A1: Expected Gaps, Expected Cyclical Policy Rate, and Fitted Values

Notes: Data from Blue Chip Economic Indicators on median one-year-ahead forecasts of the inflation gap, the unemployment gap, and the cyclical policy rate. Shaded are the NBER recession periods.