Measuring the equilibrium real interest rate

Alejandro Justiniano and Giorgio E. Primiceri

Introduction and summary

In conducting monetary policy, policymakers find it useful to monitor the performance of the economy relative to some benchmark. For instance, the policy decision whether to raise or lower the short-term nominal interest rate might be affected by the deviations of current inflation from policymakers' comfort zone, of output from potential output, and of the real interest rate (current nominal rate minus expected future inflation) from its equilibrium value (the rate that would be consistent with output at its potential level). Unfortunately, these benchmark concepts are not directly observed in the data, but can only be defined in the context of a specific theoretical framework.

Over the past decade, the *new Keynesian model* has become the workhorse for the analysis of monetary policy. This model departs from the neoclassical framework of the 1980s by assuming imperfect competition in goods and labor markets and "sticky" (meaning rigid or inflexible) prices and wages—neoclassical models assume prices and wages are flexible and adjust quickly. These ingredients in the new Keynesian model alter the transmission of fundamental shocks perturbing the economy and allow monetary policy to have temporary real effects.

The *equilibrium real interest rate* is a crucial concept in the new Keynesian class of models. This rate represents the real rate of return required to keep the economy's output equal to *potential output*, which, in turn, is the level of output consistent with flexible prices and wages and constant markups in goods and labor markets (Woodford, 2003; and Galí, 2008).[1] Meanwhile, the difference between the *ex ante real interest rate*—the nominal interest rate minus expected inflation—and the equilibrium real interest rate is defined as the *real interest rate gap*.

In the new Keynesian model, the real interest rate (RIR hereafter) gap is central to the determination of output and inflation. Loosely speaking, if this RIR gap is positive, output will decline relative to potential. This is because people will be inclined to postpone spending decisions today to take advantage of higher returns to savings. All else being equal, a negative output gap will then put downward pressures on prices and wages because of weaker aggregate demand. Conversely, a negative RIR gap will typically be associated with a positive output gap, setting in motion inflationary forces—higher demand leads to higher prices.

The main policy implication of this observation is that policymakers concerned with maintaining output close to its potential level should set short-term nominal interest rates—the policy instrument of most central banks—in order to minimize the RIR gap. In the absence of a trade-off between stabilizing inflation and output, this simple policy prescription would also completely stabilize inflation. In practice, however, there may well be a trade-off between the two objectives of output and inflation stabilization.[2] Nonetheless, the equilibrium RIR constitutes a natural benchmark for the conduct of monetary policy, and the RIR gap can be viewed as providing some indication of the stance of monetary policy (Neiss and Nelson, 2003).

While the equilibrium RIR is theoretically appealing, its use in guiding monetary policy decisions faces at least two major hurdles. First and foremost, the equilibrium RIR is not directly observable in the data, limiting its usefulness as a target for monetary policy in practice.[3] Moreover, rather than being constant, the

Alejandro Justiniano is a senior economist in the Economic Research Department at the Federal Reserve Bank of Chicago. Giorgio E. Primiceri is an assistant professor in the Department of Economics at Northwestern University.
The authors are grateful to Anna Paulson, Richard Porter, Spencer Krane, and seminar participants at the Federal Reserve Bank of Chicago for helpful comments.

equilibrium RIR fluctuates over time in response to a variety of shocks to preferences and technology that perturb the economy.

Second, setting nominal interest rates to track the equilibrium RIR may not be feasible at times because of the existence of the *zero bound*; that is, nominal interest rates cannot be set lower than zero. Indeed, the equilibrium RIR may fall enough to induce a positive RIR gap, even with the nominal interest rate at zero. Output would then decline below potential, engendering deflation. In this way, the gap helps us to gauge the constraint imposed by the zero bound on monetary policy. With short-term nominal interest rates now at historically low levels in the United States and a number of other industrialized countries, this scenario is receiving a lot of attention from both the academic community and policymakers.

Given the importance that the equilibrium RIR plays for the design of monetary policy in modern macroeconomic models, our purpose in this article is to provide an estimate of this unobservable variable. We do so by inferring it from an empirical new Keynesian model fitted to U.S. quarterly data on a few key macroeconomic variables from 1962:Q1 through 2008:Q4.[4]

Specifically, our analysis accomplishes three objectives. First, we describe the historical evolution of the equilibrium RIR. We find that this rate has been negative at times, particularly in the late 1970s and, most interestingly, during the latest recession.

Second, we estimate the short-term RIR gap as the difference between the current (as opposed to future) ex ante RIR and the equilibrium RIR. This provides some indication of the stance of monetary policy. Consistent with the anecdotal view, the estimated short-term RIR gap suggests that policy was loose during most of the 1970s. In contrast, policy would seem to have been tight at the end of our sample. However, this mostly reflects the zero bound problem—policymakers' inability to lower short-term nominal interest rates below zero—and provides a rationale for the nonconventional policy measures undertaken by the Federal Reserve during the most recent recession, such as direct purchases of longer-term securities and the creation of special facilities and programs (for example, the Term Asset-Backed Securities Loan Facility, or TALF) intended to increase access to credit.

Finally, we compare the evolution of the short-term and long-term RIR gaps, where the latter is defined as the sum of the current and expected future short-term RIR gaps or, alternatively, the difference between the ex ante long-term RIR and the equilibrium long-term RIR. Long-term rates reflect the path of current and expected future short-term rates. Therefore, long-term gaps summarize private expectations about future macroeconomic outcomes and monetary policy, providing a more forward-looking measure of the policy stance. For instance, according to this measure, policy was not loose in the 2002–06 period, which preceded the recent economic downturn. This characterization of the policy stance contrasts with what is suggested by the short-term RIR gap and, in particular, with the view of several commentators (see, for instance, Taylor, 2007).

Several papers have tackled the estimation of the equilibrium RIR before, most notably Laubach and Williams (2003) and Kozicki and Clark (2005). In contrast to these earlier studies, our estimate of the equilibrium RIR is based on a *micro-founded* model, which builds on the optimizing behavior of households and firms seeking to maximize their utility and profits. In this respect, this article is related to the approach of Neiss and Nelson (2003), Amisano and Tristani (2008), and, in particular, Edge, Kiley, and Laforte (2008). However, in contrast to these earlier studies, we stress the importance of both current and expected future RIR gaps for the determination of macroeconomic outcomes.

As with all empirical work based on structural models, our results may be sensitive to some aspects of the model specification. To illustrate this point, we compare our results across two models that differ in scale, shocks, and transmission mechanisms of these disturbances.

The article is organized as follows. First, we provide a brief description of our baseline model economy. Then, we describe the data and the estimation approach. Next, we present the main results—that is, we present our estimates of the equilibrium RIR and RIR gaps. We also discuss the robustness of these estimates when inferred from a larger-scale model. We conclude with a few comments and caveats to our analysis, particularly with regard to the current economic situation. More specifically, we note how the larger-scale model also suggests the presence of positive short-term and long-term RIR gaps for the fourth quarter of 2008. This provides a further rationale for the Federal Reserve's response to the current crisis with nonconventional measures to ease monetary policy. We do, however, emphasize the need to enhance these models' ability to capture the interplay between the financial sector and the real economy, particularly in light of recent events.

The model

In this section, we sketch our baseline new Keynesian model and analyze two of its key equilibrium relations—the aggregate demand and supply equations. The presentation is mostly narrative, with most of the technical details relegated to the appendix. Interested readers can refer to Justiniano and Primiceri (2008) for greater details on the model, or they can see the comprehensive treatment of new Keynesian models in Woodford (2003) and Galí (2008), as well as the excellent primer by Galí and Gertler (2007). For simplicity, relative to Justiniano and Primiceri (2008), the model here abstracts from the roles of habit formation, indexation, and endogenous capital accumulation. We present the results based on a larger-scale model with these additional features as a robustness check in a later section.

There are five types of agents in our model economy: 1) households, 2) employment agencies, 3) firms producing intermediate goods, 4) firms producing final goods, and 5) the monetary authority. We now briefly describe the behavior of each of them.

Households

We assume that we have a large number of households seeking to maximize their stream of current and expected future utility, which depends positively on their consumption of a single final good and negatively on the number of hours they work for the production of intermediate goods. Each household is the sole supplier of a specialized type of labor that it sells to the employment agencies in exchange for wages. Rather than taking wages as given—as under the neoclassical assumption of perfect competition—each household has some market power and can post its wage. This, in turn, determines the amount of their specialized labor demanded by the employment agencies.

We introduce sticky wages in the labor market by assuming that at each point in time only a random fraction of households can change their posted wage. Hence, when setting its wage, each household takes into consideration not only current but also future labor demand and costs of working. For example, if future labor demand is expected to rise, households will preemptively post higher wages, since they might not be able to do so in the near future.

Finally, all households have access to savings through two types of assets: one-period government bonds and state-contingent securities, which pay only if a certain future state is realized. The former are used to smooth consumption over time. State-contingent securities serve instead to insure against the idiosyncratic risk arising from the uncertainty about the length of time before households will be able to reset their wages.

Employment agencies

Employment agencies mediate the demand and supply of labor between households and firms producing intermediate goods. Their role is to purchase all types of specialized labor supplied by households and bundle them into a single homogenous labor input sold to intermediate goods firms. Employment agencies operate in a perfectly competitive market, taking the wage received for the labor bundle as given and making zero profits.

Intermediate goods producers

A large number of intermediate goods producers combine technology with labor inputs purchased from employment agencies to produce differentiated intermediate goods, which are then sold to final goods producers. Each of the intermediate goods producers has some market power and can therefore post the price of its good. This, in turn, determines the amount of its output demanded by the final goods producers.

We introduce sticky prices in the goods market by assuming that at each point in time only a random fraction of firms can change their posted price. Hence, when setting its price, each firm takes into consideration not only current but also future demand and marginal costs, where the latter depend on wages. For example, if future demand is expected to rise, producers will preemptively increase prices, since they might not be able to adjust them in the near future.

Final goods producers

Final goods producers mediate between intermediate goods producers and households. They produce the final good by bundling all intermediate goods into a single final homogenous commodity purchased by households. Final goods firms maximize profits as well, but in contrast to the intermediate goods producers, they operate under perfect competition, taking the price for the final good as given and making zero profits.

Monetary authority

The central bank determines monetary policy by setting the short-term nominal interest rate in response to price inflation and output growth. This interest rate rule is a variant of the instrument rule proposed by Taylor (1993), the Taylor rule, which approximates the historical behavior of the U.S. federal funds rate. According to this rule, nominal interest rates rise more than one-to-one with inflation and fall in response to output contractions.

Demand, supply, and the equilibrium RIR

Before presenting our estimation results, we highlight the main insights of the two crucial equilibrium relations in the model. This helps explain the roles of the equilibrium RIR and RIR gaps in the determination of output and inflation.

Aggregate demand

In the model, aggregate spending is determined by the behavior of the representative household, which seeks to smooth consumption over time by investing its savings in one-period government bonds. This optimizing behavior results in the following (log-linearized) *aggregate demand* equation, which is also known as the *IS* equation:

$$1)\quad \hat{y}_t = E_t\hat{y}_{t+1} - \hat{r}_t,$$

where y_t and r_t are output and the RIR, respectively, and the hat symbol (^) denotes deviations from the equilibrium level. Hence, $\hat{y}_t$ denotes the output gap, and $\hat{r}_t$ stands for the short-term RIR gap. Intuitively, according to the aggregate demand equation, fluctuations in the short-term RIR gap induce deviations of the output gap from its expected future value, $E_t\hat{y}_{t+1}$, where the operator E_t denotes households' expectation of future values conditional on the information available today.

Equation 1 can be iterated forward to express the output gap today only as a function of the current and expected future short-term RIR gaps. This procedure yields the expression

$$2)\quad \hat{y}_t = -\sum_{j=0}^{+\infty} E_t\hat{r}_{t+j},$$

by which the output gap is negatively associated with the long-term RIR gap. The latter corresponds to the sum of current and expected future short-term RIR gaps.[5] Notice, therefore, that if the long-run RIR gap is negative, the output gap will be positive, and vice versa.

Aggregate supply

In terms of the supply side, intermediate goods firms set prices according to the current and expected future evolution of marginal costs and demand conditions. Profit-maximizing behavior results in the following (log-linearized) *aggregate supply* or *Phillips curve* equation:

$$3)\quad \pi_t = \beta E_t\pi_{t+1} + \kappa s_t + \lambda_{\pi,t},$$

where π_t and s_t stand for price inflation and real marginal costs, respectively, and $\lambda_{\pi,t}$ is a markup shock that represents exogenous variation to the level of markup desired by intermediate goods producers. Finally, β is a constant very close to one that represents the temporal discount factor, and κ is a positive constant that is inversely related to the degree of price stickiness. Intuitively, inflation exceeds its expected future level either if real marginal costs increase or if intermediate goods firms change their desired markup of prices over marginal costs for other reasons exogenous to the model.

To highlight the importance of the RIR gap for inflation determination, we briefly analyze a special case of our model obtained by assuming perfectly flexible wages. Under this assumption, real marginal costs are proportional to the output gap. Hence, all else being equal, a positive output gap will cause inflation to rise relative to its expected future level. Moreover, if the output gap is projected to remain positive in the future, expected future inflation will also increase, further fueling the rise in current inflation. That is, current and expected future RIR gaps engender pressures on prices through their effects on aggregate demand. This crucial insight also holds in our general model with wage rigidities, although with sticky wages the link between the output gap and real marginal costs is more complex.

RIR gaps and monetary policy

Equations 1 and 3 highlight the importance of RIR gaps for output and inflation determination. Current and future expected deviations of ex ante RIRs from their corresponding equilibrium values affect the output gap, which, in turn, influences the inflation rate. Since the ex ante RIRs depend on the nominal interest rates set by the monetary authority, the conduct of monetary policy is central to the behavior of the RIR gaps and, hence, output and inflation.

Consider, for instance, a central bank that seeks to stabilize price inflation and the output gap. Absent any markup shocks ($\lambda_{\pi,t}$), the central bank can achieve full stabilization of both output and inflation by committing to set nominal interest rates according to an appropriate instrument rule that delivers a zero RIR gap at every point in time.

However, as we mentioned in the introduction, tracking the equilibrium RIR may not be feasible when the zero bound on nominal interest rates becomes binding. Put another way, sometimes the equilibrium RIR may fall enough that, even with the short-term nominal interest rate at zero, positive RIR gaps would emerge. In this case, according to the model, output would decline relative to potential and inflation would fall.

Even abstracting from the zero bound, in practice optimal monetary policy is more involved than the simple prescription of tracking the equilibrium RIR. This is due to the fact that markup shocks bring about a trade-off between stabilizing the output gap and inflation.[6] Nonetheless, despite these considerations, the equilibrium RIR remains an important reference point for the conduct of monetary policy, assuming that it can be accurately estimated and forecasted. This is the task we undertake next.

Model solution and estimation

In this section, we provide a brief overview of the approach that we adopt to estimate the model's parameters and to infer the evolution of the latent (unobservable) variables. The discussion is somewhat technical, although we do not aim to provide a comprehensive overview of the techniques we used. For more details on these techniques, interested readers should refer to An and Schorfheide (2007).

Model solution and state-space representation

The model we described in the preceding section has a solution of the form

$$4) \quad \tilde{\xi}_t = G(\theta)\tilde{\xi}_{t-1} + M(\theta)\tilde{\varepsilon}_t,$$

where the state vector $\tilde{\xi}_t$ collects all variables except for the shocks. The elements of $\tilde{\xi}_t$ are expressed in (log) deviations from the model's nonstochastic steady state, which corresponds to the constant values of all variables that the economy would converge to in the absence of shocks. The shocks inducing temporary deviations from the steady state are stacked in the vector $\tilde{\varepsilon}_t$. Meanwhile, $G(\theta)$ and $M(\theta)$ are matrices whose elements are functions of the vector of model structural parameters, denoted by θ. Our goal is to estimate these parameters and to uncover the historical behavior of the unobserved variables in the state vector.

In fact, while some elements of the state vector are directly observed in the data (for instance, inflation and output), others are not (such as the equilibrium RIR and expected inflation). Therefore, in order to estimate the model, equation 4 must be combined with an additional set of equations specifying which elements of the state vector are observed in the data.

The general form of this additional set of equations is

$$5) \quad x_t = Z\left(\tilde{\xi}_t + C(\theta)\right),$$

where Z is a matrix mapping the elements of $\tilde{\xi}$ into x_t (the vector of observable data) and where C is a vector of constant terms (which may depend on θ) representing the steady state of the observable elements of $\tilde{\xi}$. Equations 4 and 5 constitute the transition and measurement equations of a linear *state-space* model.

Data

We estimate the model, using five series of U.S. quarterly data: 1) real per capita gross domestic product (GDP), 2) per capita hours worked, 3) real per capita wages, 4) quarterly inflation, and 5) the short-term nominal interest rate. We construct real GDP by dividing nominal GDP by the population aged 22–65 and the GDP deflator.[7] For hours, we use a measure of hours in all sectors of the economy following Francis and Ramey (2008). This is also our source for the population series. Real wages correspond to nominal compensation of employees from the U.S. Bureau of Economic Analysis's national income and product accounts (NIPAs), divided by hours and the GDP deflator; for the nominal interest rates, we use the effective federal funds rate. The sample period spans 1962:Q1 through 2008:Q4.[8] We do not de-mean or de-trend any series.

Bayesian inference

The state-space representation of the model allows us to use a very powerful algorithm known as the Kalman filter to estimate the parameters θ and retrieve the most likely path of the unobservable elements of $\tilde{\xi}_t$. We discuss each in turn.

A natural way to estimate the model is to find the value of the parameters θ that maximizes the *likelihood* function. The likelihood function summarizes all information about θ contained in a sample of data and plays a pivotal role in econometrics and statistics. The likelihood function of our state-space model can be evaluated using the Kalman filter.

In practice, however, the likelihood function associated with most modern macroeconomic models is typically a complicated nonlinear function of the model parameters. This makes finding a unique value that maximizes the likelihood a rather arduous task. For this reason, most of the recent literature estimating macro models has turned to Bayesian methods, which discipline the set of plausible values for θ through the use of prior information.

Bayesian inference then seeks to characterize the distribution of θ that results from combining the likelihood function with the prior information. This is known as the *posterior* distribution, from which we can compute the location of a parameter (mean or median) and a measure of uncertainty. For instance, the uncertainty surrounding θ can be conveyed by reporting posterior probability bands that contain the range of values that parameters are likely to take with, say, 99 percent probability.

Prior beliefs about the elements of θ may be informed by theory or simply reflect and summarize

FIGURE 1

Equilibrium real interest rate, 1962–2008

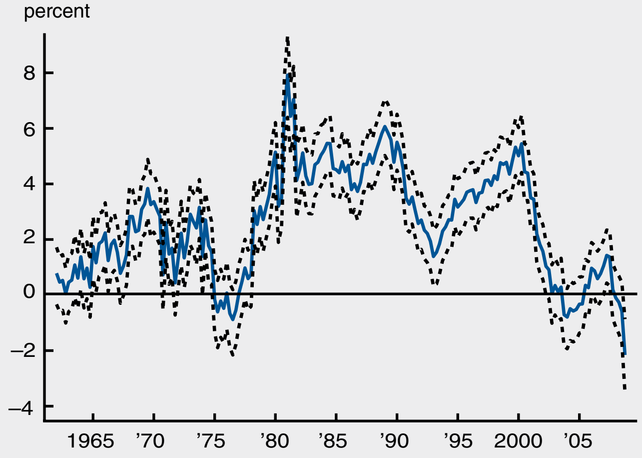

Note: The dashed lines are the 99 percent posterior probability bands.
Sources: Authors' calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

information not contained in the estimation sample. In practice, this prior information is formulated by specifying a certain distribution for each element of the parameter vector, centered at a particular value (mean) and with an associated measure of uncertainty (standard deviation).

Once we have estimated the model's parameters, we can employ the Kalman filter to sequentially and systematically update our guess for the unobserved elements of the state vector. More precisely, at each point in time, our guess for $\tilde{\xi}_t$, based on data available in the previous quarter, is updated after we observe the data for the current period. This *filtered* (or one-sided) estimate for the state vector forms the basis for our guess on the value of the state vector next period, which we also update once we have data for the next quarter, and so on.

Having followed this procedure for all periods, we can go back and revise the filtered estimate of $\tilde{\xi}_t$, conditional not only on information up to time t but also on the whole sample of data. We call the state vector emerging from this procedure the *smoothed* (or two-sided) estimate. We analyze these estimates in the next section.

Equilibrium RIR and RIR gaps in the estimated model

We do not report the estimated parameters in this article. They are similar to those of Justiniano and Primiceri (2008), who use a longer sample. Here, we focus on our estimates of the equilibrium RIR and the RIR gaps.

The equilibrium RIR

Figure 1 plots the smoothed estimate of the equilibrium RIR (solid blue line). It is also important to characterize the uncertainty surrounding the estimated equilibrium RIR, particularly since this is cited as a possible concern regarding its usefulness for monetary policy analysis. Therefore, we also report uncertainty bands (dashed black lines), which represent the values this variable is likely to have taken with 99 percent probability. We first highlight a few properties of the smoothed estimate and later discuss these probability bands.

The first thing to notice is that the inferred equilibrium RIR has fluctuated substantially over our sample, with a standard deviation of 1.94 percent around a mean of 2.6 percent (annualized).[9]

A second interesting feature of figure 1 is that the equilibrium RIR has turned negative in a few instances. This occurred around 1975 and the end of 2008—two recession dates, as determined by the National Bureau of Economic Research—and during the 2003–04 period. These episodes were characterized by a substantial decline in the federal funds rate in response to weak economic conditions. However, the 2008 episode is the only one in our sample for which the uncertainty bands are completely below zero.

Indeed, the third interesting observation is that the equilibrium RIR has plummeted in the latest part of the sample. In particular, during the latest recession, the equilibrium RIR seems to have recorded by far its largest decline, with an estimate for 2008:Q4 of roughly –2.15 percent.

The tightness of the posterior probability bands deserves some comment. In particular, the precision with which the equilibrium RIR is estimated perhaps seems implausible, especially considering that these bands account for the uncertainty surrounding both the unobserved states and the model parameters. It is important to keep in mind, however, that these probability bands abstract from model uncertainty. That is, alternative specifications of the model (for example, a different historical characterization of U.S. monetary policy or a model with additional propagation mechanisms and/or shocks) might deliver different estimates of the equilibrium RIR. We return to this issue in the section explaining the larger-scale model.

This being said, the cross-sectional dispersion at different points in time is larger than perhaps suggested visually by figure 1. For example, figure 2 plots the posterior distribution of the equilibrium RIR for the last point in the sample, 2008:Q4. Values of the equilibrium RIR are on the horizontal axis, with the vertical line drawn at the median of –2.15 percent, which coincides with the estimate reported in the previous figure. Notice that this distribution has a range from roughly –4 percent to –0.5 percent, with hardly any weight assigned to values close to zero. Therefore, our model-based estimates suggest that it is quite likely that the equilibrium RIR became negative in 2008. To what extent did this induce positive RIR gaps? We address this key issue next.

The short-term RIR gap

The ex ante RIR is given by the difference between the nominal interest rate and the inflation rate expected for next quarter. While the former is directly observable in our data, the latter is part of the unobservable state vector and must be backed out using the Kalman filter.

Figure 3 shows the smoothed estimate of the ex ante RIR (blue line) together with the equilibrium RIR (black line). The mean of the ex ante RIR is 2.37 percent (annualized) with a standard deviation of 2.45 percent. These statistics are similar to those corresponding to the equilibrium RIR. The overall contours of these two series coincide, although they have differed at times.

In order to highlight the discrepancies between the ex ante RIR and the equilibrium RIR, figure 4 plots their difference together with its 99 percent probability bands. We refer to this difference as the short-term RIR gap, in order to distinguish it from the long-term gap that we analyze next. Note that the short-term gap has also fluctuated considerably over time, with an average of –0.33 percent and a standard deviation of 1.28 percent.

As we noted earlier, the short-term RIR gap is commonly taken as a measure of the monetary policy stance. And indeed,

FIGURE 2

Posterior distribution of equilibrium real interest rate in 2008:Q4

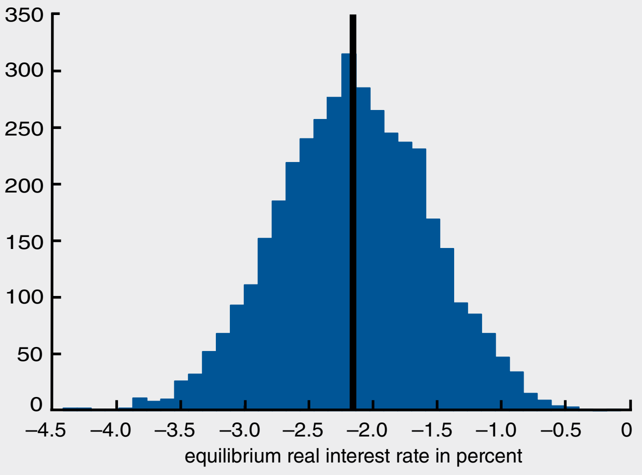

Sources: Authors' calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

FIGURE 3

Ex ante and equilibrium real interest rates, 1962–2008

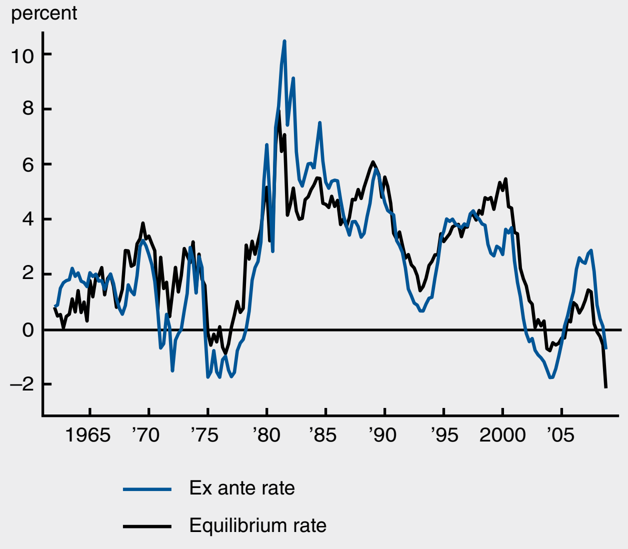

Sources: Authors' calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

FIGURE 4

Short-term real interest rate gap, 1962–2008

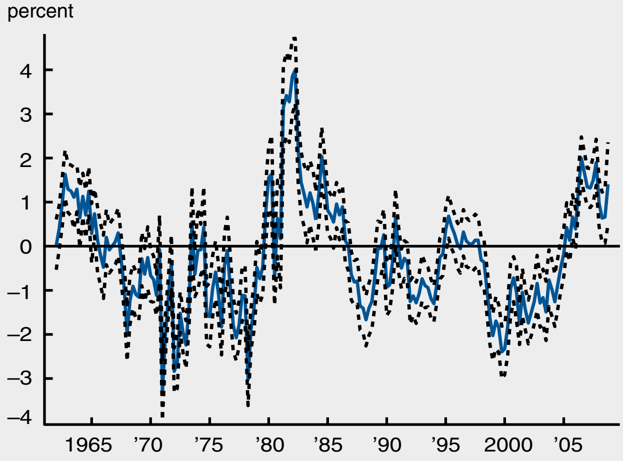

Note: The dashed lines are the 99 percent posterior probability bands.
Sources: Authors' calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

at least for some episodes, the evolution of the RIR gap aligns well with the anecdotal characterization of monetary policy that we see in the literature. For instance, according to our estimates, the equilibrium RIR exceeded the ex ante real interest rate during most of the 1970s, exactly when U.S. inflation was at historically high levels. This is consistent with the view that monetary policy during this period was characterized by an insufficient response to the rise in inflation (Clarida, Galí, and Gertler, 2000). Similarly, the significant increase in the short-term RIR gap in the early 1980s accords well with the conventional view that the disinflation in the U.S. economy was engineered by a substantial policy tightening under then-Federal Reserve Chairman Paul Volcker.

The long-term RIR gap

While the behavior of the short-term RIR gap presented in figure 4 squares quite well with the conventional view, there are a few caveats that call for caution in interpreting this gap as a good proxy for the stance of monetary policy. In particular, as we explained earlier, it is important to recognize that the whole path of expected future short-term RIR gaps—rather than just its contemporaneous value—matters for the determination of output and inflation in the new Keynesian model (see equation 2, p. 17). From this perspective, we might judge the monetary policy stance better by looking at the long-term RIR gap, which summarizes the information contained in the current and expected future values of the federal funds rate, inflation, and the equilibrium RIR.

To this end, figure 5 compares the short-term RIR gap (blue line) with the evolution of the long-term one (black line). Although the two series often move together—the correlation coefficient is equal to 0.56—the message about the stance of monetary policy implied by the two lines differs during some historical episodes.

The 2002–06 period provides an interesting example. In 2002:Q3 the federal funds rate stood at 1.75 percent, but it had declined to 1 percent by 2003:Q3, and remained there for the next three quarters. The federal funds rate then rose gradually, reaching 5.25 percent in 2006:Q3. Some have argued that monetary policy was too accommodative during this period (for example, Taylor, 2007). Although the negative value of the short-term RIR gap seems to accord with this claim (blue line), the positive value of the long-term RIR gap (black line) does not support the view that policy was too expansionary. In particular, it suggests that the private sector expected a decline of the equilibrium RIR or a monetary tightening.

The difference between short-term and long-term gaps toward the end of the sample is also informative. For instance, our estimate of the short-term RIR gap in 2008:Q4 is roughly 1.5 percent. This suggests that, according to the model, the federal funds rate of 0.5 percent was probably above the equilibrium RIR. Furthermore, it suggests that the zero bound on nominal interest rates would have been binding before additional interest rate cuts could have closed the short-term RIR gap. In addition, the estimated long-term RIR gap exceeds 3 percent. Taken at face value, this would suggest that at the end of 2008, positive short-term gaps were expected to persist and the zero bound was expected to bind beyond a single quarter.

Before we interpret this result as indicative of contractionary monetary policy, we must acknowledge that these gaps can only reflect the stance of *conventional* monetary policy. By this we mean the Federal Reserve’s management of the short-term nominal interest rate. During the current economic crisis, the Federal Reserve has also employed a variety of nonconventional policy measures; and these measures have been reflected in the changing size and composition of the Federal Reserve’s balance sheet. Our simple analysis suggests that these measures have been appropriate, insofar as both the short-term RIR and long-term RIR exceeded the equilibrium

RIR. However, these extraordinary measures are not reflected in our analysis of the short-term and long-term RIR gaps.

A larger-scale model

The baseline model can be summarized in a few simple equations that, as discussed, clearly highlight the role of the equilibrium RIR for the dynamics of output and inflation. This simplicity, however, comes at the expense of abstracting from other features that impart more realism to the model. In particular, additional shocks can be included and other mechanisms added (such as endogenous capital accumulation) through which disturbances influence the evolution of the economy. For this reason, we test the robustness of our main conclusions by using a larger-scale model estimated on a richer data set. This extended model is discussed in Justiniano and Primiceri (2008) and is based on the well-known studies of Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007).

FIGURE 5

Short-term and long-term real interest rate gaps, 1962–2008

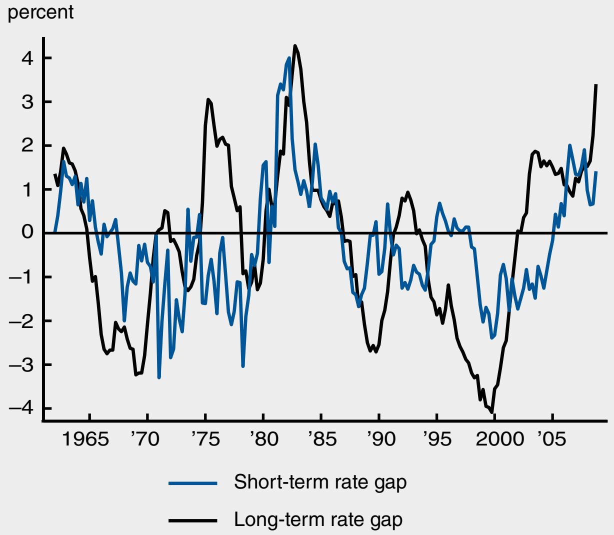

Sources: Authors’ calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

Relative to our baseline model, the larger-scale model includes the additional propagation mechanisms provided by endogenous capital accumulation, investment adjustment costs, a choice of capital utilization, habit formation in consumption, and indexation in both prices and wages. These features are essentially meant to increase the length of time for which a given shock will affect the evolution of the economy. There are three additional disturbances perturbing the model economy, specifically, shocks to the marginal efficiency of investment, to the disutility of labor, and to government spending. Finally, we estimate the model over the same sample, 1962:Q1 through 2008:Q4, but we incorporate additional data on consumption and investment.

Figure 6 reports the smoothed estimates of the equilibrium RIR and the ex ante RIR, as well as the short-term and long-term RIR gaps. In each panel, the black line reproduces the estimates from the baseline model and the blue line corresponds to estimates from the extended model.

Panel A highlights the fact that the cyclical pattern of the equilibrium RIR is very similar across models, although the equilibrium RIR is substantially more volatile in the larger-scale model.[10] One implication of this finding is that, according to the extended model, the equilibrium RIR has declined below zero more frequently than what is predicted by our baseline framework. Furthermore, the decline in the current downturn, while substantial, is not as dramatic by historical standards as suggested by the baseline model.

Since the inferred ex ante RIR (panel B) is almost identical across models, it is not surprising that the short-term RIR gap (panel C) and long-term RIR gap (panel D) are more volatile in the larger-scale model as well. Notice also that the estimates from our baseline model and larger-scale model co-move more closely in the case of the long-term gap, for which the two lines essentially overlap during the latest part of the sample.

Regarding the 2002–06 period, the discrepancy between the short-term and long-term RIR gaps is far less evident in the larger-scale model than in our baseline model. However, both measures in the larger-scale model remain positive or very close to zero. This confirms our earlier observation that policy may not have been as accommodative during this period as has been suggested (for example, Taylor, 2007).

Consistent with the baseline model, the larger-scale framework also predicts large positive short-term and long-term RIR gaps for the fourth quarter of 2008. However, the same caveats we raised earlier about interpreting these endpoint estimates as reflecting the policy stance apply to the larger-scale model as well.

FIGURE 6

Real interest rate levels and gaps in the baseline model and larger-scale model, 1962–2008

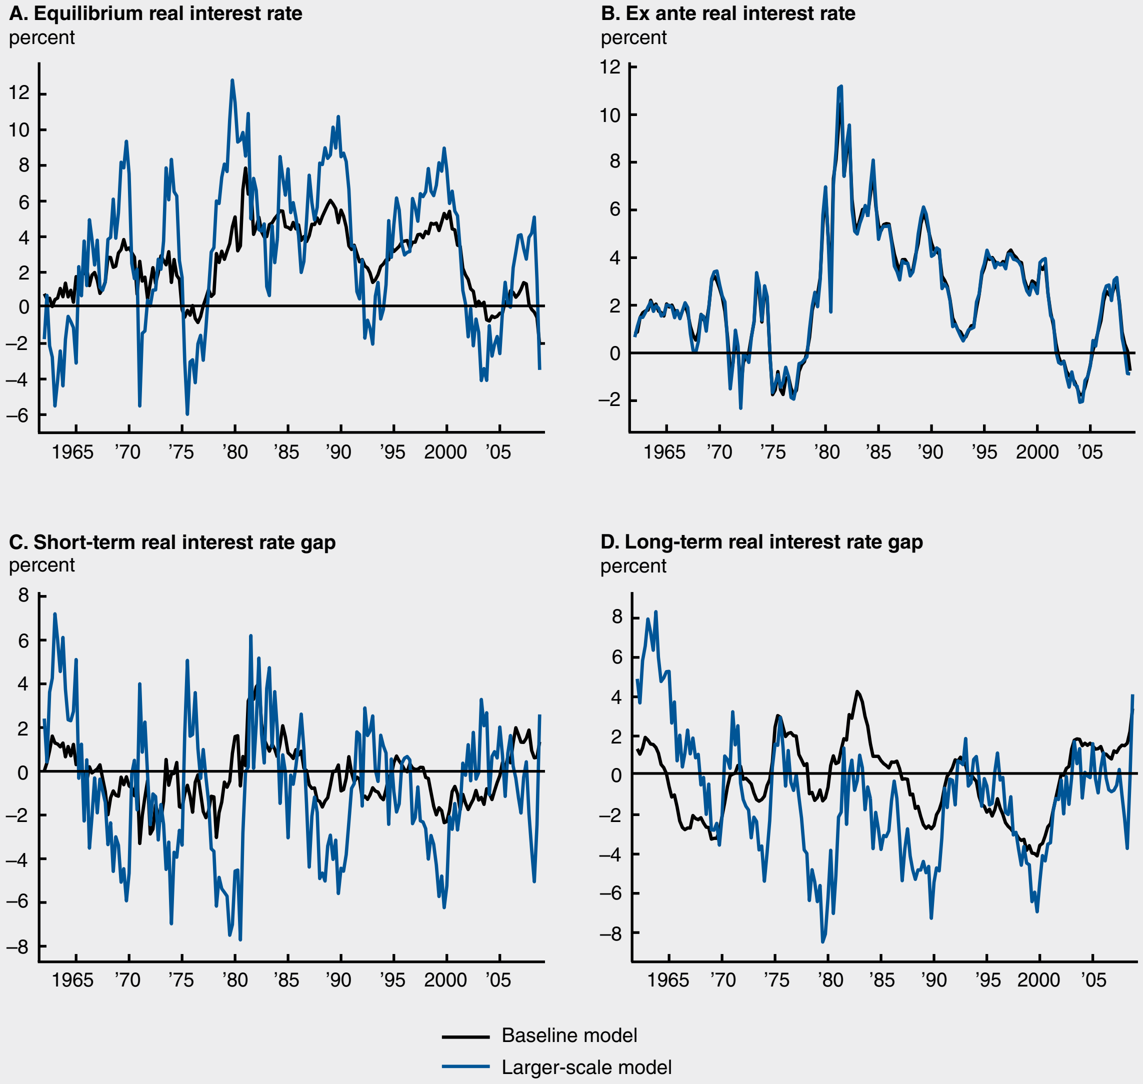

Sources: Authors' calculations based on data from Haver Analytics and the U.S. Bureau of Labor Statistics.

Overall, despite some obvious discrepancies, we view these results as an important assessment of robustness of our main findings. Furthermore, they suggest—in line with our earlier hypothesis—that model uncertainty is likely to be a crucial factor surrounding the measurement of the unobservable equilibrium RIR and related components. This source of uncertainty is sometimes ignored in studies presenting model-based estimates of the RIR, although our findings suggest that this should be a major issue for further empirical work in this area.

Conclusion

In this article, we study the evolution of the equilibrium RIR and RIR gaps, using both a prototypical new Keynesian model and a larger-scale model similar to those in Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007). Our estimates point to a substantial degree of time variation in the equilibrium RIR. Moreover, we find that this rate has sometimes become negative in the post-war period. In particular, our analysis suggests that the equilibrium RIR fell sharply below zero toward the end of 2008 (although the magnitude of this decline relative to historical standards is model dependent), resulting in positive short-term and long-term expected RIR gaps. This provides some support for the Federal Reserve's response to the current crisis with nonconventional measures to ease monetary policy.

We conclude by noting that the models we use here, even the larger-scale one, are to some extent very stylized and have some shortcomings. One of these shortcomings is the absence of an explicit theoretical framework of the financial sector and financial frictions. It would be useful to analyze how the introduction of these additional features would affect our results (as, for instance, in Christiano, Motto, and Rostagno, 2007). These features seem particularly relevant for the analysis of current economic events.

NOTES

[1]Hence, we could alternatively refer to the equilibrium real interest rate as the real interest rate at potential. We prefer the former terminology because it is more popular in the literature and policy discussions, as exemplified by the discussion in Ferguson (2004). Meanwhile, potential output is proportional, but lower than the efficient level of output. The efficient level of output is the level of output under perfect competition and, therefore, with zero markups. In the goods market, the markup is defined as the amount by which prices exceed the marginal cost of production. In the labor market, the markup is defined as the excess of wages over the marginal cost of supplying labor.

[2]Exogenous variations in desired markups, usually referred to as markup shocks, introduce such a trade-off (see, for example, Clarida, Galí, and Gertler, 1999).

[3]Potential output is not directly observable either, and the policy implications of its measurement have received substantial attention following the work of Orphanides (2001). See also Justiniano and Primiceri (2008).

[4]We also estimate the model's unknown parameters and subsequently extract all unobserved model-based variables, such as expected inflation next period.

[5]While seemingly daunting to compute, the long-run rates can be backed out from the Lagrange multiplier of the household's budget constraint.

[6]If wages are rigid, optimal monetary policy must attribute some weight to wage inflation stabilization as well.

[7]All data except for hours are from Haver Analytics. We are very grateful to Shawn Sprague, of the U.S. Bureau of Labor Statistics, for providing us the series of hours in all sectors of the economy.

[8]We use the eight years prior to the sample period to initialize the Kalman filter.

[9]This result is consistent with the large degree of time variation reported by Laubach and Williams (2003) and Edge, Kiley, and Laforte (2008), but stands in contrast to the analysis of Neiss and Nelson (2003), who argue that the equilibrium real interest rate exhibits very little volatility.

[10]The main reason the equilibrium RIR in the larger-scale model is more volatile is that this model includes habit formation.

APPENDIX: MODEL EQUATIONS

We present the main equations of the model for each of the five classes of agents described in the main text.

Households

The expected discounted stream of utility that each household j maximizes is given by

$$\text{A1)} \quad E_t \sum_{s=0}^{\infty} \beta^s b_{t+s} \left[\log C_{t+s} - \varphi \frac{L_{t+s}(j)^{1+\nu}}{1+\nu} \right],$$

where C_t denotes consumption, and the second argument of the utility function represents the marginal disutility of each household's specific labor, $L(j)$, that depends on the parameter ν, known as the inverse Frisch elasticity of labor supply. Future utility is discounted at the rate β, and b_t is a "discount factor" shock affecting both the marginal utility of consumption and the marginal disutility of labor. The logarithm of b_t is modeled as a Gaussian autoregressive process of order 1, denoted as AR(1) for short.

At every point in time t, each household's sources and uses of income must be equal, as summarized by the budget constraint

$$P_t C_t + T_t + B_t \leq R_{t-1} B_{t-1} + Q_{t-1}(j) + \prod_t + W_t(j) L_t(j),$$

where T_t is lump-sum taxes and transfers, B_t denotes holdings of government bonds, R_t is the gross nominal interest rate, $Q_t(j)$ is the net cash flow from participating in state-contingent securities that insure against idiosyncratic risk, and $\prod_t$ is the per capita profit that households get from owning the intermediate goods firms.

Following Erceg, Henderson, and Levin (2000), we permit in every period only a fraction $1 - \xi_w$ of households to reset their wages to minimize the expected discounted stream of labor disutility for the periods in which the posted wage is anticipated to remain in place,

$$E_t \sum_{s=0}^{\infty} \xi_w^s \beta^s b_{t+s} \left\{ -\varphi \frac{L_{t+s}(j)^{1+\nu}}{1+\nu} \right\}.$$

This is subject to the labor demand function of employment agencies specified next. Wages for the remaining ξ_w fraction of households are indexed to steady-state inflation and productivity.

Employment agencies

Competitive employment agencies operate in competitive markets and bundle each household's specialized labor $L_t(j)$ into a homogenous labor input according to

$$L_t = \left[\int_0^1 L_t(j)^{\frac{1}{1+\Lambda_{w,t}}} dj \right]^{1+\Lambda_{w,t}}.$$

Homogeneous labor is sold to intermediate goods firms. Profit maximization and the zero profit condition imply a specialized labor demand function,

$$L_t(j) = \left(\frac{W_t(j)}{W_t} \right)^{-\frac{1+\Lambda_{w,t}}{\Lambda_{w,t}}} L_t,$$

where $W_t(j)$ is the wage paid by the employment agencies to the household supplying labor of type j, and W_t is the hourly wage paid by intermediate goods firms for their homogenous labor input. The demand schedule for labor j is decreasing in the relative wage and depends on the elasticity of substitution among varieties of labor given by $\Lambda_{w,t}$. Notice that this elasticity is time varying, and we assume that $\log(1+\Lambda_{w,t})$ is a Gaussian independent and identically distributed (i.i.d.) process. In the literature this is referred to as the wage markup shock, and it is analyzed in detail in Justiniano and Primiceri (2008).

Intermediate goods producers

A monopolistically competitive firm produces the intermediate good $Y_t(i)$ with the production function

$$Y_t(i) = A_t L_t(i)^{\alpha},$$

where $L_t(i)$ denotes the bundled labor input purchased from employment agencies for the production of good i, and A_t represents a productivity shock. We model A_t as nonstationary, with its growth rate following a Gaussian AR(1) process.

As in Calvo (1983), at each point in time a fraction ξ_p of firms cannot reoptimize their prices and index them to steady-state inflation. The remaining fraction $1 - \xi_p$ of firms post a new price $\tilde{P}_t(i)$ to maximize the expected discounted stream of profits for the periods in which the new price is anticipated to remain in place,

$$E_t \sum_{s=0}^{\infty} \xi_p^s \beta^s \Lambda_{t+s} \left\{ \left[\tilde{P}_t(i) \pi^s \right] Y_{t+s}(i) - W_{t+s} L_{t+s}(i) \right\},$$

where Λ_{t+s} is the marginal utility of consumption used to value future income, subject to the goods demand function specified in the next section.

Final goods producers

Perfectly competitive firms produce the final good Y_t by bundling all intermediate goods according to

$$Y_t = \left[\int_0^1 Y_t(i)^{\frac{1}{1+\Lambda_{p,t}}} di \right]^{1+\Lambda_{p,t}}.$$

Profit maximization and zero profit condition for the final goods producers imply the following demand function for the intermediate good i:

$$Y_t(i) = \left(\frac{P_t(i)}{P_t} \right)^{-\frac{1+\Lambda_{p,t}}{\Lambda_{p,t}}} Y_t,$$

where P_t corresponds to the aggregate price level. The demand schedule for intermediate good i is decreasing in its relative price, and depends on the elasticity of substitution $\Lambda_{p,t}$ among varieties of intermediate goods. This elasticity is time varying, and we assume that $\log(1+\Lambda_{p,t})$ is a Gaussian i.i.d. process. This disturbance is known as the price markup shock.

Monetary authority

The Taylor type rule for the short-term nominal interest rate, R_t, is given by

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R} \right)^{\rho_R} \left[\left(\frac{\left(\prod_{s=0}^{3} \pi_{t-s} \right)^{1/4}}{\pi_t^*} \right)^{\phi_\pi} \left(\frac{\left(Y_t / Y_{t-4} \right)^{1/4}}{e^{\gamma}} \right)^{\phi_Y} \right]^{1-\rho_R} e^{\varepsilon_{R,t}},$$

with R being the steady state for the gross nominal interest rate and $\varepsilon_{R,t}$ being a Gaussian i.i.d. monetary policy shock. The parameters ϕ_π and ϕ_Y capture how aggressively the monetary authority responds to variations in inflation and output growth over the current and previous three quarters. There is a time-varying inflation target π_t^*, which evolves exogenously according to a Gaussian AR(1) process. Finally, notice that short-term nominal interest rates are adjusted gradually, as given by ρ_R, referred to as the smoothing coefficient.

REFERENCES

Amisano, G., and O. Tristani, 2008, "Perceived productivity and the natural rate of interest," European Central Bank, mimeo.

An, S., and F. Schorfheide, 2007, "Bayesian analysis of DSGE models," *Econometric Reviews*, Vol. 26, Nos. 2–4, March, pp. 113–172.

Calvo, G., 1983, "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, Vol. 12, No. 3, September, pp. 383–398.

Christiano, L. J., M. Eichenbaum, and C. L. Evans, 2005, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of Political Economy*, Vol. 113, No. 1, February, pp. 1–45.

Christiano, L. J., R. Motto, and M. Rostagno, 2007, "Financial factors in business cycles," Northwestern University, mimeo.

Clarida, R., J. Galí, and M. Gertler, 2000, "Monetary policy rules and macroeconomic stability: Evidence and some theory," *Quarterly Journal of Economics*, Vol. 115, No. 1, February, pp. 147–180.

__________, 1999, "The science of monetary policy: A new Keynesian perspective," *Journal of Economic Literature*, Vol. 37, No. 4, December, pp. 1661–1707.

Edge, R. M., M. T. Kiley, and J.-P. Laforte, 2008, "Natural rate measures in an estimated DSGE model of the U.S. economy," *Journal of Economic Dynamics and Control*, Vol. 32, No. 8, August, pp. 2512–2535.

Erceg, C. J., D. W. Henderson, and A. T. Levin, 2000, "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, Vol. 46, No. 2, October, pp. 281–313.

Ferguson, R., 2004, "Equilibrium real interest rate: Theory and application," speech to the University of Connecticut School of Business Graduate Learning Center and the SS&C Technologies Financial Accelerator, Hartford, CT, October 29, available at www.federalreserve.gov/boarddocs/speeches/2004/20041029/default.htm.

Francis, N. R., and V. A. Ramey, 2008, "Measures of hours per capita and their implications for the technology-hours debate," University of California, San Diego, mimeo.

Galí, J., 2008, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton, NJ: Princeton University Press.

Galí, J., and M. Gertler, 2007, "Macroeconomic modeling for monetary policy evaluation," *Journal of Economic Perspectives*, Vol. 21, No. 4, Fall, pp. 25–46.

Justiniano, A., and G. E. Primiceri, 2008, "Potential and natural output," Northwestern University, mimeo.

Kozicki, S., and T. Clark, 2005, "Estimating equilibrium real interest rates in real time," *North American Journal of Economics and Finance*, Vol. 16, No. 3, December, pp. 395–413.

Laubach, T., and J. C. Williams, 2003, "Measuring the natural rate of interest," *Review of Economics and Statistics*, Vol. 85, No. 4, November, pp. 1063–1070.

Neiss, K. S., and E. Nelson, 2003, "The real interest rate gap as an inflation indicator," *Macroeconomic Dynamics*, Vol. 7, No. 2, April, pp. 239–262.

Orphanides, A., 2001, "Monetary policy rules based on real-time data," *American Economic Review*, Vol. 91, No. 4, September, pp. 964–985.

Smets, F., and R. Wouters, 2007, "Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach," *American Economic Review*, Vol. 97, No. 3, June, pp. 586–606.

Taylor, J. B., 2007, "Housing and monetary policy," in *Housing, Housing Finance, and Monetary Policy*, Jackson Hole Economic Symposium Proceedings, Federal Reserve Bank of Kansas City, pp. 463–476.

__________, 1993, "Discretion versus policy rules in practice," *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, No. 1, December, pp. 195–214.

Woodford, M., 2003, *Interest and Prices: Foundations of a Theory of Monetary Policy*, Princeton, NJ: Princeton University Press.