PRELIMINARY

Optimal Monetary Policy in a Model with Distinct Core and Headline Inflation Rates*

Martin Bodenstein, Christopher J. Erceg, and Luca Guerrieri**
Federal Reserve Board

September 2007

Abstract

This paper investigates the relative merits of responding to forecasts of headline or core inflation within the context of a stylized DSGE model with an energy sector. We derive the optimal policy in our model using a quadratic approximation to the welfare function of households. The optimal policy response to an adverse energy supply shock implies a rise in core inflation, a larger rise in headline inflation, and a decline in wage inflation. The optimal policy is well-approximated by policies that stabilize the output gap, but also by a wide array of "dual mandate" policies that aren't overly aggressive in stabilizing core inflation. Finally, we show that policies that react to a forecast of headline inflation following energy shocks perceived as temporary imply markedly different effects than policies that react to a forecast of core. Notably, the headline forecast policy implies greater volatility in core inflation and the output gap, and may even induce greater variation in headline inflation if agents are uncertain about the persistence of the shock.

Keywords: energy-price shocks, monetary policy tradeoffs, DSGE models

JEL Classification: E32

** Corresponding author. Telephone (202) 452 2550. E-mail Luca.Guerrieri@frb.gov.

* The authors thank Brian Doyle, Marvin Goodfriend, Chris Gust, Steve Kamin, Paul Wood, for insightful comments and discussions, and Hilary Croke for excellent research assistance. The views expressed 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.

1 Introduction

The pronounced divergence between headline and core inflation rates in response to the substantial energy price hikes of the past few years has intensified debate over which inflation measure is the more appropriate focus of policy. Some central banks, such as the Bank of England, have an explicit goal of adjusting policy so that their forecast of headline inflation reverts to target over a two to three year horizon. Other central banks appear more concerned with the behavior of core inflation.[1]

In this paper, we use a stylized optimization-based DSGE model to assess the implications of alternative monetary policies in response to energy shocks, as well as to compare such policies to the "optimal" policy that maximizes the utility of households in our economy. Our model economy is formulated to allow a distinction between core and headline inflation, as in recent work by Blanchard and Galí (2007). In particular, goods comprising the core basket are produced by monopolistically-competitive retailers that set prices in staggered Calvo-style contracts, so that core prices are sticky. By contrast, energy prices are flexible, and determined to equate the sum of household and firm demand to the available supply (a stochastic endowment process). Because wages are sticky, shocks that depress energy supply raise marginal cost even for a policy that keeps output at potential, pushing up both core and headline price inflation; thus, monetary policy faces a stabilization tradeoff.

For a special case of our model in which preferences are logarithmic in consumption and production is Cobb-Douglas, we derive a quadratic approximation to welfare following the the seminal analysis of Rotemberg and Woodford (1997). As in the model of Erceg, Henderson, and Levin (2000), the deviation of welfare from its Pareto-optimal level depends on the variance of price inflation, wage inflation, and the output gap: but importantly, it is core price inflation that matters for wel-

[1]Aucremanne and Wouters (1999) provide an overview of alternative statistical measures of core inflation that may be useful for monetary policy. Notable early contributions include Bryan and Pike (1991), Bryan and Cecchetti (1994), and Quah and Vehey (1995).

fare, since the price of energy is assumed to be completely flexible.[2] Welfare also turns out to depend on the price markup, which affects the intratemporal allocation of energy across household and firms, although this term turns out to have small quantitative implications for optimal policy.

Given that the supply block of our model also parallels Erceg, Henderson, and Levin (2000) quite closely, some key prescriptions about optimal policy derived in that setting for technology shocks carry over to energy supply shocks. In particular, in the case of a contractionary supply shock that pushes up the energy price, it is optimal for the required fall in the real wage to occur through a combination of a temporary rise in core inflation, and fall in the wage inflation (with headline inflation increasing by more than core). Policies that keep output near potential turn out to be nearly optimal even though the output gap receives a small weight in the welfare function, reflecting such policies perform well in achieving the wage and price adjustment that occurs under the optimal policy. Finally, because the variance tradeoff frontier, or "Taylor" Curve (Taylor (1979)) between inflation and output gap volatility is very flat, even simple targeting rules derived from an ad hoc objective function over realized (core) inflation and output gap volatility come close to replicating macroeconomic responses under the optimal policy: accordingly, "dual mandate" policies perform very well using our social welfare metric for a wide array of policymaker preferences, except if the policymaker puts essentially no weight on output gap variation.

We then turn attention to the performance of simple forecast-based rules that respond either to forecasts of headline or core inflation.[3] We show that such rules

[2]This implication corroborates the insight of Goodfriend and King (1997) and King and Wolman (1998). These authors argued that monetary policy should be concerned with stabilizing the components of the price index that are sticky, rather than the overall index. Aoki (2001) formalized this insight by constructing a model with both a flexible-price sector, and a sticky-price sector. Because the welfare losses depend exclusively on the variance of the sticky price components, the optimal policy consists of stabilizing those components.

[3]In related work, Dhawan and Jeske (2007) compared the performance of alternative rules that respond to

may have markedly different macroeconomic implications in the case of a temporary but persistent change in energy prices. In particular, we illustrate how a temporary rise in the energy price requires a much more accommodative policy under the rule that responds to a headline inflation forecast, as a consequence of expected mean-reversion in the energy price. As a result, such a policy contributes to a much larger rise in core inflation and the output gap than a policy of stabilizing a forecast of core (which more closely resembles the optimal policy). In the presence of uncertainty about the persistence of the shock, we illustrate how a focus on headline inflation may even contribute to greater volatility in realized headline inflation than the alternative of focusing on core.

This paper is organized as follows. Section 2 provides some empirical motivation. Section 3 describes our workhorse model, and Section 4 the calibration. Section 5 describes our derivation of the welfare function, while Section 6 analyzes optimal policy and various simple rules. Section 7 concludes.

2 Headline and Core Inflation Data and Forecasts

Figure 1 reports the evolution of headline CPI inflation and core CPI inflation between 1990 and the second quarter of 2007 (both series are reported as four quarter changes). The core inflation measure excludes food and energy prices. Headline inflation is clearly much more volatile than the core CPI inflation rate. Despite the higher weight on food compared with energy in the headline CPI basket (the weights are 15% and 9%, respectively), the energy component accounts for most of the higher volatility of headline inflation. It is also evident from the figure that headline and core inflation exhibit fairly persistent divergence during certain periods, including in both the early 1990s, and in the wake of the large escalation in energy prices of

either realized core or headline inflation. Our analysis differs insofar as we focus on forecasts of headline and core inflation, and compare suboptimal policies to the full-commitment optimum.

the past few years.

Forecasts of headline and core inflation have frequently shown large and fairly persistent gaps, especially following energy price changes that were expected to be transitory. Figure 2 provides some examples of episodes in which the Federal Reserve Greenbook forecasts for headline and core inflation exhibited a persistent gap. In each panel, the solid lines show data or Greenbook projections for periods prior to final editing of the Greenbook, while the dotted line denotes the forecast.[4] For each of the episodes, it is clear that energy price movements dominated the gap between the forecast of core and headline CPI inflation. For example, the Greenbook narrative dated 05/09/2000 is especially insightful in providing a rationale for the staff's forecast. Several developments that had pushed up energy prices sharply during the previous year – including higher oil prices, regulatory changes in California, capacity problems in electricity production, and the transition to MTBE as additive in gasoline – were expected to continue to push up headline inflation (measured as a four quarter change) over the remainder of that year. However, as these pressures were expected to be partly reversed, energy prices were expected to drop sharply in the subsequent two years, with the headline inflation rate expected to average roughly 3/4 percentage point less than core.

3 Model Description

3.1 Households and Wage Setting

There is a continuum of monopolistically competitive households indexed by $h \in [0, 1]$, each of which supplies a differentiated labor service to an intermediate goods-producing sector (the only producers demanding labor services in our framework). It is convenient to assume that a representative labor aggregator (or "employment

[4]The data underlying each panel was obtained from the Greenbooks dated 02/04/1987, 09/27/1989, 03/17/1993, and 05/09/2001. Greenbooks for 2001 are the latest publicly available.

agency") combines households' labor hours in the same proportions as firms would choose. Thus, the aggregator's demand for each household's labor is equal to the sum of firms' demands. The aggregate labor index L_t has the Dixit-Stiglitz form:

$$L_t = \left[\int_0^1 \left(N_t \left(h \right) \right)^{\frac{1}{1+\theta_w}} dh \right]^{1+\theta_w}, \tag{1}$$

where $\theta_w > 0$ and $N_t(h)$ is hours worked by each member of household h. The aggregator minimizes the cost of producing a given amount of the aggregate labor index, taking each household's wage rate $W_t(h)$ as given, and then sells units of the labor index to the production sector at their unit cost W_t:

$$W_t = \left[\int_0^1 W_t \left(h \right)^{\frac{-1}{\theta_w}} dh \right]^{-\theta_w}. \tag{2}$$

It is natural to interpret W_t as the aggregate wage index. The aggregator's demand for the labor services of household h is given by

$$N_t \left(h \right) = \left[\frac{W_t \left(h \right)}{W_t} \right]^{-\frac{1+\theta_w}{\theta_w}} L_t. \tag{3}$$

The utility functional of household h is

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j U_{t+j} \left(h \right), \tag{4}$$

where the period utility function $U_t(h)$ has the time separable form:

$$U_t \left(h \right) = \frac{\sigma}{\sigma - 1} C_{t+j} \left(h \right)^{\frac{\sigma-1}{\sigma}} - \frac{\chi_0}{1+\chi} N_{t+j} \left(h \right)^{1+\chi}, \tag{5}$$

Here $C_t(h)$ and $N_t(h)$ denote each household's total consumption and hours of labor in period t, respectively. The intertemporal elasticity of consumption, σ, satisfies $\sigma > 0$, and we assume that $0 < \beta < 1$, $\chi > 0$, and $\chi_0 > 0$.

A household's total consumption in each period depends in turn on its purchases both of a composite nonenergy consumption good C_{Nt}, and of energy O_{Ct}, according to the CES aggregator:

$$C_t = \left(\left(1 - \omega_{OC} \right)^{\frac{\rho_o}{1+\rho_o}} C_{Nt}^{\frac{1}{1+\rho_o}} + \omega_{OC}^{\frac{\rho_o}{1+\rho_o}} O_{Ct}^{\frac{1}{1+\rho_o}} \right)^{1+\rho_o}, \tag{6}$$

The parameter ω_{OC} in equation (9) determines the importance of energy in the household's consumption bundle, and equals the share of energy in total consumption in the model's steady state. The parameter ρ_o determines the elasticity of substitution between energy and the nonenergy consumption good. We denote this elasticity by $\eta_o = \frac{\rho_o}{1+\rho_o}$.

Household h faces a flow budget constraint in period t which requires that combined expenditure on goods and on the net accumulation of financial assets must equal disposable income:

$$
\begin{aligned}
&P_{Nt}C_{Nt}(h) + P_{Ot}O_{Ct}(h) + \int_s \xi_{t,t+1}B_{t+1}(h) - B_t(h) = \\
&(1+\tau_w)W_t(h)\,N_t(h) + P_{0t}Y_{0t} + \Gamma_t(h) - T_t(h)\,. \qquad (7)
\end{aligned}
$$

Household h's expenditures consist of purchases of the nonenergy consumption good at a unit price of P_{Nt}, of energy at a unit price of P_{0t}, and of net purchases of state-contingent claims. Each household earns labor income of $(1+\tau_W)W_t(h)\,N_t(h)$, where τ_W is an employment subsidy designed to allow the flexible wage economy to be efficient. Each household also derives income from receiving an aliquot share Y_{Ot} of the economy's random flow endowment of energy. Households sell their energy endowment at a price of P_{0t} either to other households, or to those firms in the economy (called "wholesalers" below) that use energy as a productive input. Finally, each household receives an aliquot share $\Gamma_t(h)$ of the profits of all firms as well as lump sum transfers, $T_t(h)$ from the government. In every period t, household h maximizes the utility functional (5) with respect to its consumption of energy, the nonenergy good, and its holdings of contingent claims, taking the prices of consumption goods (energy and nonenergy) and of bonds as given.

We assume that household wages are determined by Calvo-style staggered contracts subject to static wage indexation. In particular, with probability $1-\xi_w$, each member of a household is allowed to reoptimize its wage contract. If a household is not allowed to optimize its wage rate, we assume each household member resets its wage according to $W_t(h) = \pi W_{t-1}(h)$, where π is the steady state inflation rate. Household h chooses the value of $W_t(h)$ to maximize its utility functional (5), yielding the following first-order condition:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \xi_w^j \{ \frac{(1+\tau_w)}{(1+\theta_w)} \frac{\Lambda_{t+j}}{P_{Ct+j}} \pi^{t+j} W_t(h) - \chi_{0t+j} N_{t+j}(h)^{\chi} \} N_{t+j}(h) = 0, \tag{8}$$

where Λ_t is the marginal value of a unit of consumption. Roughly speaking, equation (8) says that the household chooses its contract wage to equate the present discounted value of working an additional unit of time to the discounted marginal cost. The employment subsidy τ_W is chosen to exactly offset the monopolistic distortion θ_W, so that the household's marginal rate of substitution would equal the consumption real wage in the absence of nominal wage rigidities.

3.2 Firms and Price Setting

Production of the Nonenergy Consumption Good The nonenergy consumption good C_{Nt} is a composite of a continuum of different "retail" goods. The composite good is produced by a representative firm (or "bundler") according to the technology:

$$C_{Nt} = Y_{Nt} = \left[\int_0^1 \left(Y_{Nt}\left(f\right) \right)^{\frac{1}{1+\theta_p}} df \right]^{1+\theta_p}, \tag{9}$$

where $\theta_p > 0$ and $Y_{Nt}(f)$ is the input of retail good f. The representative firm purchases the underlying retail goods in competitive markets at prices $P_{Nt}(f)$, and sells the composite good to households at a price of P_{Nt} (which it takes as given). The representative firm's demand for retail good f is:

$$Y_{Nt}\left(f\right) = \left[\frac{P_{Nt}\left(f\right)}{P_{Nt}} \right]^{-\frac{1+\theta_p}{\theta_p}} Y_{Nt}. \tag{10}$$

Retail goods are produced by monopolistically competitive firms. Each retailer f purchases a homogenous "wholesale" good Y_{Wt}, and transforms it into a particular type of retail good according to a simple linear production function:

$$Y_{Nt}\left(f\right) = Y_{Wt}\left(f\right) \tag{11}$$

where $Y_{Wt}(f)$ denotes purchases of the wholesale good by producer f. Similarly to households, retailers set the price of their respective output goods in Calvo-style staggered contracts. Thus, each retailer f faces a constant probability, $1 - \xi_p$, of receiving a signal to re-optimize its contract price, $P_{Nt}(f)$. Those firms not receiving a signal to re-optimize simply adjust their price by the steady state inflation rate. A firm f that receives a signal to adjust sets its contract price to maximize the profit functional:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \xi_p^j \psi_{t,t+j} \left[(1+\tau_p) \pi^{t+j} P_{Nt}(f) Y_{Nt+j}(f) - P_{Wt+j} Y_{Nt+j}(f) \right], \tag{12}$$

taking its demand schedule (10) and the price of wholesale goods P_{Wt} as given. The constant τ_p is a production subsidy that is calibrated so that the flexible price equilibrium is efficient, and $\psi_{t,t+j}$ is the stochastic discount factor.

The wholesale good Y_{Wt} is produced by a representative firm according to the technology:

$$Y_{Wt} = \left((1 - \omega_{op}^{\frac{\rho_o}{1+\rho_o}}) L_t^{\frac{1}{1+\rho_o}} + \omega_{op}^{\frac{\rho_o}{1+\rho_o}} O_{pt}^{\frac{1}{1+\rho_o}} \right)^{1+\rho_o}, \tag{13}$$

where O_{pt} is energy used in production, ω_{op} determines the importance of energy in the production of wholesale goods, and ρ_o determines the elasticity of substitution between the labor index and energy. Note that the implied elasticity of $\eta_o = \frac{\rho_o}{1+\rho_o}$ is constrained to be the same as for households. The representative firm purchases energy and hires labor in competitive factor markets (at prices P_{Ot} and W_t, respectively), and sells its output to retail firms at the wholesale price of P_{Wt}.

In addition to the constraint in equation (9) that all output of the nonenergy good is consumed by households, clearing of the wholesale market requires that the cumulative demand (of retailers) equal the available supply:

$$\int_0^1 Y_{Wt}(f) df = Y_{Wt}. \tag{14}$$

Moreover, energy market clearing implies that the energy demand of households and (wholesale) firms equal the exogenous flow endowment:

$$\int_0^1 O_{Ct}(h)dh + O_{pt} = Y_t^O. \tag{15}$$

The exogenous flow endowment, Y_{Ot}, is itself the sum of temporary and (nearly) permanent components:

$$Y_t^O = Y_{Pt}^O + Y_{Tt}^O. \tag{16}$$

The permanent and temporary components, represented as log-deviations from their steady-state values, evolve according to:

$$y_{Pt}^o = \rho_P y_{Pt-1}^o + \epsilon_{Pt}, \tag{17}$$

$$y_{Tt}^o = \rho_T y_{Tt-1}^o + \epsilon_{Tt}, . \tag{18}$$

3.3 Monetary and Fiscal Policy

We assume that the central bank follows an interest rate reaction function of the form:

$$i_t = \gamma_1 \pi_{Nt} + \gamma_2 \pi_{Ct} + \gamma_{1f} E_t \pi_{t+1} + \gamma_{2f} E_t \pi_{Ct+1} + \gamma_y (y_t - y_t^*). \tag{19}$$

In this equation, all variables are measured as percentage deviations from their steady state values. The variable i_t denotes the short-term nominal interest rate. The variable π_{Nt} is the logarithmic percentage change in the price index P_{Nt} for the nonenergy composite good, which we interpret as the core inflation rate. The variable π_{Ct} is the logarithmic percentage change in the household's price deflator for consumption P_{Ct} implied by (9), which we interpret as the headline inflation rate. The instrument rule is sufficiently general to also consider policies that react to forecasts of core inflation ($E_t \pi_{t+1}$) or headline inflation ($E_t \pi_{Ct+1}$), as well as to the true output gap ($y_t - y_t^*$).

Finally, we assume that the subsidies to firms and households are financed by lump-sum taxes. Given the Ricardian structure of the model, it is convenient to assume that the government budget is balanced each period.

4 The Welfare Loss Function

To provide a normative assessment of alternative monetary policy choices, we measure social welfare as the conditional expectation of average household lifetime utility:

$$SW_0 = E_0 \sum_{i=0}^{\infty} \beta^t \left[\int_0^1 U_t\left(h\right) dh \right] = E_0 \sum_{i=0}^{\infty} \beta^t U_t \tag{20}$$

Although we conduct some exercises in the sensitivity analysis section that rely on numerical computations of social welfare, most of our analysis focuses on a particular calibration of the model for which we can derive a second-order approximation to the welfare function following the seminal analysis of Rotemberg and Woodford (1997). This calibration assumes that the subutility function over consumption is logarithmic, and that production functions between energy and other inputs (for both firms and households) are Cobb Douglas (so $\sigma = 1$ and $\eta_o = 1$). Under these conditions, the period loss function – the deviation of the period social welfare U_t from the period social welfare function under flexible prices and wages U_t^* – may be expressed as:

$$\begin{aligned} \frac{U_t - U_t^*}{U_c \bar{C}} = \quad & -\frac{1}{2} \frac{(1+\chi)(1-\omega_{oc})}{1-\omega_o} \left[g_t + \omega_o \left(\frac{O_C}{Y^O} \right) \mu_{Pt} \right]^2 \\ & -\frac{1}{2}(1-\omega_{oc}) \left(\frac{1+\theta_p}{\theta_p} \right) Var_f P_t(f) \\ & -\frac{1}{2}(1-\omega_{oc})(1-\omega_o) \left(\frac{1+\theta_w}{\theta_w} \right) \left(1 + \frac{1+\theta_w}{\theta_w} \chi \right) Var_h w_t(h) \end{aligned} \tag{21}$$

The scaling factor $U_c\bar{C}$ expresses the period loss as a percentage of steady state consumption. This period loss function is similar to that derived by Erceg, Henderson, and Levin (2000) insofar as it depends on cross-sectional dispersion in prices across monopolistically-competitive firms ("retailers"), as well as cross-sectional dispersion in wages across households. As in that model, the assumption that all firms have identical marginal costs – so that there is no factor attachment on the household side – creates a noticeable disparity between the coefficient multiplying the different dispersion terms.

The salient difference relative to Erceg, Henderson, and Levin (2000) is that the term involving the output gap also includes the price markup μ_{Pt}. This additional term reflects the costs of misallocating energy across sectors relative to what occurs in the flexible equilibrium. The relative allocation of energy across sectors is inefficient if the price markup is non-zero, because the relative price of energy faced by households – which is relative to the retail price index, i.e., $\frac{P_{Ot}}{P_{Nt}}$ – differs from the relative price of energy faced by firms – which is relative to the wholesale price index, i.e., $\frac{P_{Ot}}{P_{Wt}}$. Under flexible prices, the markup is zero, so that there is no wedge between these prices, and each sector's share of the energy endowment is simply a constant. Nevertheless, given that this term is multiplied by the energy share of production, its quantitative contribution to welfare is very small (as we note below).

Under our benchmark assumption of a Calvo-Yun contract structure, the price dispersion term $Var_f P_t(f)$ evolves according to:

$$Var_f P_t(f) = (1-\xi_p) Var_f P_{t-1}(f) + \xi_p \pi_t^2 \tag{22}$$

and similarly for wage dispersion. Substituting these expressions into (21), the discounted conditional welfare loss functional may be expressed:

$$\begin{aligned} E_0 \sum_{i=0}^{\infty} \beta^t \left[\frac{U_t - U_t^*}{U_c \bar{C}}\right] &= -\frac{1}{2}\frac{(1+\chi)(1-\omega_{oc})}{1-\omega_o} \sum_{i=0}^{\infty} \beta^t E_0 \left[g_t + \omega_o \left(\frac{O_C}{Y^O}\right)\mu_{Pt}\right]^2 \\ &-\frac{1}{2}(1-\omega_{oc})\left(\frac{1+\theta_p}{\theta_p}\right)\left(\frac{\xi_p}{(1-\beta\xi_p)(1-\xi_p)}\right)\sum_{i=0}^{\infty}\beta^t E_0 \pi_t^2 \\ &-\frac{1}{2}(1-\omega_{oc})\left(\frac{1+\theta_w}{\theta_w}\right)\left(1+\frac{1+\theta_w}{\theta_w}\chi\right)\left(\frac{\xi_w(1-\omega_o)}{(1-\beta\xi_w)(1-\xi_w)}\right)\sum_{i=0}^{\infty}\beta^t E_0 \omega_t^2 \end{aligned} \tag{23}$$

We use this welfare loss functional both to evaluate the performance of various simple instrument and targeting rules, as well as to derive the optimal full commitment targeting rule following the approach of Woodford (2003). To derive the full commitment rule, we use the time invariant first order conditions obtained from maximizing the conditional loss function (24) subject to the log-linearized behavioral equations. We refer to this full commitment rule as the optimal rule below.

5 Solution Method and Calibration

We solve the model by log-linearizing the equations around the model's steady state. To obtain the reduced-form solution of the model, we use the numerical algorithm of Anderson and Moore (1985), which provides an efficient implementation of the method proposed by Blanchard and Kahn (1980) (see also Anderson (1997))

The model is calibrated at a quarterly frequency. The discount factor β is 0.993, consistent with an (annualized) real interest rate of 3 percent. Given that we derive a quadratic approximation to the loss function in the special case of a logarithmic subutility function over consumption and for Cobb-Douglas production functions between energy and other inputs (for both firms and households), we find it convenient to set $\sigma = 1$ in our benchmark calibration, and $\eta_o = 1$. We also set $\chi = 1$, implying a Frisch elasticity of labor supply of unity. Because that these unitary elasticities are on the high end of most empirical estimates, we also explore a range of values in sensitivity analysis.

Using NIPA data, we set the energy component of consumption (gasoline and fuel, plus gas and electricity) at 6%, so that $\omega_{oc} = 0.06$. Based on value-added data at the industry level (mining and utilities), as well as data for imports of oil, gas, coal and electricity, we place the energy's share in production at 2%, which implies $\omega_{op} = 0.02$.

We assume that price and wage contracts have a duration of four quarters, so that $\xi_p = 0.75$, and $\xi_w = 0.75$. We assume that the wage markup parameter $\theta_W = 1/5$. The price markup parameter (θ_P) has no effect on the behavior of the log-linearized behavioral equations, and matters only for determining the coefficients of the optimal policy rule derived from minimizing the utility-based loss function. Accordingly, we find it useful for heuristic reasons to choose this parameter to equalize the weight on price and wage inflation volatility in the welfare loss function under our benchmark calibration, which is achieved by setting $\theta_P = 0.025$.

Finally, the persistence of the (nearly) permanent shock to oil supply (ρ_p) is arbitrarily close to unity (0.999), while the persistence of the temporary shock (ρ_p) is 0.8.

6 Results under Benchmark Calibration

6.1 The optimal policy

Figure 3 examines the effects of a temporary contraction in energy supply that is scaled so that the energy price would rise initially by 20 percent in the flexible price (and wage) equilibrium. The panels compare impulse responses of key variables derived under the optimal policy rule with corresponding responses from the flexible price and wage equilibrium. The response of employment under the optimal rule is quite small (declining only 0.1 percent in response to the 20 percent energy hike), and hence comes close to replicating the flexible price response in which employment is unaffected. This suggests – correctly, as we verify below – that policies which keep either the employment or output gap nearly constant closely approximate the optimal policy. The output gap is proportional to the employment gap, and hence also moves very little (notwithstanding a drop in the level of production due to the lower supply of energy). Real interest rates rise even in the flexible price equilibrium, reflecting that agents correctly perceive that consumption will recover as the effects of the energy shock wear off. The increase in real interest rates is slightly larger under the optimal policy (consistent with a negative output gap).

Although the rise in the energy price has an immediate depressing effect on the product real wage in the flexible price equilibrium, wages adjust gradually under the optimal policy due to price and wage rigidities. The real wage adjustment occurs through a combination of a rise in core price inflation, and falling nominal wages. To understand the mechanism underlying the rise in core inflation, it is useful to note that the price-setting equation can be expressed in log-linearized form as:

$$\pi_t = \beta \pi_{t+1|t} + \kappa_p(\psi_1(\zeta_t - \zeta_t^*) + \lambda_{mpl} x_{Lt}) \tag{24}$$

Inflation depends on the determinants of marginal cost, including the real wage gap (the difference between the actual real wage ζ_t and flexible price real wage ζ_t^*), and the employment gap x_{Lt}. The parameter ψ_1 is very close to unity, while the absolute slope of the labor demand schedule λ_{mpl} is very small (about 0.02 in our calibration). Thus, with a very small employment gap, almost all of the pressure on inflation comes from the effects of the positive real wage gap on marginal cost. Similarly, with the real wage above its flexible price level, households are willing to reduce their nominal wage.

Finally, we consider how nominal rigidities affect the allocation of energy across sectors under the optimal policy compared with the flexible price equilibrium (in which there is no change in the relative proportions of the total endowment consumed by households and firms). As discussed above, the rise in marginal cost is tantamount to a rise in the relative price of wholesale goods relative to the nonenergy (or "core") goods basket purchased by households. Accordingly, households experience a larger rise in the relative price of energy than firms, which accounts for some shift in the relative amount of energy demanded towards firms. Nevertheless, this type of resource misallocation plays little quantitative role in determining the characteristics of the optimal policy, or in contributing to welfare losses.

6.2 Comparison to inflation targeting rules that react to realized variables

We next compare the optimal policy derived from the utility-based welfare loss function with responses derived from an ad hoc loss function that penalizes variation in core inflation and the output gap:

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j (\pi_{t+j}^2 + \lambda_x x_{t+j}^2), \tag{25}$$

where λ_x is the relative weight on the output gap. The policymaker is assumed to minimize the loss function subject to the behavioral equations of the model.

Given that a key feature of the optimal policy consists in allowing core prices to play a sizeable role in bringing about adjustment in the real wage, it might be suspected that policies that react very aggressively to core inflation would depart markedly from the optimal rule. This is corroborated by the responses labeled "aggressive core" in Figure 4, where we again consider the same temporary energy supply shock as in the previous figure. The aggressive response to inflation pushes up real interest rates much more sharply than under the optimal policy. The aggressive policy succeeds in reducing marginal cost more quickly than under the optimal policy through generating a much larger output decline, which puts downward pressure on nominal wages, and also raises labor's marginal product a bit by causing an endogenous decline in energy prices. However, such an aggressive policy clearly forces too much adjustment to occur through wages and employment, and hence departs markedly from the optimal policy. Similarly, an aggressive response to headline inflation would precipitate an even larger downturn.

Although a highly aggressive response to realized inflation is clearly suboptimal, our analysis suggests that a targeting rule derived from the loss function (25) with approximately balanced weights between core inflation and the output gap performs very well relative to the optimal policy. This is also illustrated in Figure 4 for the case of equal weights between the inflation and output gap objectives (with the responses under this policy rule labeled "core with balanced weights"). Clearly, the response of core inflation is almost indistinguishable from that under the optimal policy, and the responses of employment and wage inflation are also quite close. Moreover, targeting rules that move further in the direction of output gap stabilization (so that allow λ_x to become very large) also perform very well relative to the optimum (these responses are extremely close to those of the balanced case, and hence are omitted from the figure for visual clarity).

Figure 5 examines a comparison of the same targeting rules in the case of a permanent contraction in energy supply. The responses from the simple targeting rule with balanced weights on inflation and the output gap tracks the responses under the optimal policy fairly well, though there is a more noticeable divergence at longer horizons. This divergence reflects that the simple targeting rule is more aggressive insofar as it attempts to bring the price level back to baseline in the long-run, so that all of the decline in the product real wage is achieved through lower nominal wages. This contrasts with the optimal policy, which implies that most of the real wage reduction is achieved by a long-run rise in the price level. Nevertheless, given that the simple rule achieves this long-run objective gradually, it does not contribute to markedly higher volatility in core price or wage inflation, and thus still performs quite well relative to the optimal policy.

It may seem puzzling that the objective function of the form (25) comes close to delivering responses under the optimal policy for a wide set of weights on λ_x, even though it omits putting any weight on wage inflation stabilization. This reflects two factors. First, even though output gap stabilization receives a small weight in the utility-based social welfare function – and wage inflation a relatively large weight[5] – policies that keep output near potential come close to delivering the smae wage and price responses that occur under the optimal policy. This reflects that output

[5] Wage and price inflation have equal weight in the welfare loss function under our benchmark calibration, with each over 20 times as large as the weight on the output gap.

gap stabilization implies that wages and prices adjust to the "real wage gap" of $(\zeta_t - \zeta_t^*)$ in a manner proportional to their relative degree of stickiness, which under our calibration implies noticeably faster price than wage adjustment (similar to the optimal policy). The second reason is that the variance tradeoff frontier derived from the simple loss function (25) is very flat when represented in the volatility of inflation and output space, as can be seen in Figure 6. The tradeoff frontier is obtained by minimizing the policymaker's loss function (25) over all possible values of λ_x subject to the log-linearized behavioral equations. Given that it is extremely costly to reduce inflation volatility in our model relative to the level associated with keeping output gap volatility at zero, only a policymaker with a very high relative weight on inflation (i.e., low λ_x) is inclined to try to noticeably reduce inflation volatility. Even for λ_x in the range of unity – the balanced case in Figures 4 and 5 – the high cost of reducing inflation volatility keep the policymaker fairly close to a policy of targeting the output gap exclusively.

These factors account for the substantial similarity in impulse responses between the optimal rule that minimizes the utility-based welfare criterion and those derived from the ad hoc loss functions. Although the differences do appear a bit more noticeable for the permanent shock at longer horizons, this reflects differences in the long-run response of the price and the wage level that have little consequence for the volatility of the policymaker's goal variables.

6.3 Comparison to forecast-based inflation targeting rules

As noted in the introduction, many inflation-targeting central banks have an explicit goal of adjusting policy so that their forecast of inflation reverts to some target level (or falls in a desired range) within a specified period of roughly two to three years. Typically, they express this objective in terms of a measure of headline inflation. It is argued that it is desirable to focus on headline inflation because its high degree of public visibility facilitates central bank communication. Moreover, while it is acknowledged that reacting to realized core and headline inflation would imply very different policy actions in the face of energy price shocks, it is often asserted that responding to either a forecast of headline or core inflation would have similar operational implications for the conduct of policy.

The conjecture that targeting a forecast of headline or core inflation has similar implications appears corroborated in the case of a permanent shock to the level of energy prices. As was seen in Figure 5, forecasts of headline and core inflation are essentially indistinguishable in this case, reflecting that a permanent level shock has no effect on the forecast of energy price inflation. It can easily be verified that policies that target or respond to either forecast imply similar macroeconomic responses.

However, as discussed in Section 2, there appears to be substantial evidence that some energy price shocks are temporary, and that reasonable forecasts of headline and core inflation have often exhibited pronounced divergence on this account. In the case of a temporary energy price shock, the wedge between the forecast of headline and core inflation due to the predicted change in the energy price has important implications for monetary policy: in particular, policies of responding to forecasts of core and headline inflation imply significantly different policy responses, with consequent effects on output and inflation.

This is illustrated in Figure 7, which shows the effects of the temporary energy supply shock (considered in Figures 3 and 4) under the optimal policy, and also under two Taylor-style policy rules that respond to alternative inflation forecasts. The latter rules are described by the general interest reaction function (19), with the "core PCE forecast" rule setting the coefficient on one period ahead core inflation to 1.5 (i.e., $\gamma_{1f} = 1.5$), and the "headline PCE inflation forecast" rule setting the coefficient on one period ahead headline inflation to 1.5 (i.e., $\gamma_{1c} = 1.5$). The coefficient on the output gap is set to 0.5 in each of the instrument rules ($\gamma_y = 0.5$) when expressed at an annualized rate.

As seen in Figure 7, the one quarter ahead forecast of headline inflation essentially reverts to baseline after 3 years under either rule. Thus, either policy appears to succeed in keeping a forecast of headline inflation near baseline over the "medium-term" horizon that is the focus of most inflation-targeting central banks. However, the macroeconomic effects are considerably different over shorter horizons. The policy of responding to a forecast of core inflation appears reasonably similar to the optimal policy insofar as policy tightens in response to higher near-term forecasts of core inflation. There are still some noticeable differences relative to the optimal rule: the calibrated policy rule isn't aggressive enough to keep the output gap from rising, wage inflation rises, and core inflation remains above baseline; but overall, the macroeconomic implications don't appear dramatically different. By contrast, the policy that responds to a forecast of headline inflation is much more accomodative, reflecting the expected persistent decline in energy price inflation below baseline. Low real interest rates fuel a large and persistent employment gap, core inflation remains significantly above baseline even after a year, and wage inflation exhibits a persistent increase.

Thus, the rule responding to a forecast of headline inflation performs poorly on the dimensions which the our model indicates are important for welfare; notably, rather than bringing about real wage adjustment with minimal movements in wage inflation and core price inflation, it induces large movements in both variables in the same direction. Nevertheless, it might still appear somewhat attractive from the perspective of a central bank with a strong desire to control variability in headline inflation. In particular, it is clear from the panel showing the headline inflation forecast in Figure 7 that realized headline inflation is in fact much more stable after the impact period than under the alternative policies given our assumption that agents correctly ascertain the nature of the underlying shock (recalling that this assumption implies that the one quarter ahead forecast at each date equals the realized value in the following period).[6]

But even the better ability of the headline forecast rule to reduce the variability of realized headline inflation hinges on our assumption that agents correctly perceive the underlying shock. In the realistic setting in which agents are unsure if the energy shock is temporary or permanent, headline inflation forecast targeting may cause even realized headline inflation to stray more persistently from forecast, in addition to inducing persistent deviations in core inflation and the employment gap. To model this uncertainty about the underlying shock, we adopt a signal extraction framework in which agents only observe the total energy supply y_t^O, but not the individual shock components, i.e., based on the log-linearization of equation (16)

$$y_t^O = y_{Pt}^O + y_{Tt}^O \tag{26}$$

Given that agents are assumed to understand the law of motion for these exogenous shocks, i.e., equations (17), and (18), they solve a Kalman filtering problem to make optimal forecasts of the energy price.

Within this framework, Figure 7 shows the effects of a permanent contraction in energy supply that agents initially think is likely to be mostly transient. In particular, the relative size of innovations to the permanent and temporary component are calibrated so that the Kalman gain parameter is 0.10: accordingly, agents intially forecast that about 90 percent of the price hike will eventually be reversed, and only gradually revise their beliefs about the permanence of the shock. In this framework, energy price inflation is incorrectly forecast to be negative for a protracted period, so that the headline forecast rule implies a persistently accomodative policy. Real rates remain very low, and this fuels a persistent rise in both the employment gap and core inflation. Moreover, because the expected drop in energy prices never materializes, headline inflation tracks core inflation after the initial price spike, and

[6]The responses of headline inflation in the impact period are dominated by the jump in energy prices; however, they are somewhat higher under the rule responding to a forecast of headline inflation due to the larger rise in core inflation implied by that policy.

thus remains quite elevated until agents finally grasp the permanent nature of the shock.

7 Conclusion

This paper has investigated the performance of alternative monetary policy rules in response to an energy supply shock, using the optimal rule that maximizes household welfare as the benchmark of comparison. Our analysis lends suggests that even a policymaker whose rule is derived from a simple "dual mandate" objective over core inflation and the output gap is likely to perform well in approximating the socially-optimal rule, at least provided that the policymaker does not give an extremely large weight to inflation. We have also found that rules that respond to a forecast of core inflation perform reasonably well; by contrast, rules that respond to a forecast of headline generate unnecessarily large volatility in real activity, core inflation, and possibly even in headline inflation itself.

References

Anderson, G. (1997). A Reliable and Computationally Efficient Algorithm for Imposing the Saddle Point Property in Dynamic Models. Federal Reserve Board, Occasional Staff's Studies 4.

Anderson, G. and G. Moore (1985). A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models. *Economic Letters 17*, 247–52.

Aoki, K. (2001). Optimal Monetary Policy Responses to Relative-Price Changes. *Journal of Monetary Economics 48*, 55–80.

Aucremanne, L. and R. Wouters (1999). A Structural VAR Approach to Core Inflation and its Relevance for Monetary Policy. In *Measures of Underlying Inflation and their Role in the Conduct of Monetary Policy: Workshop Proceedings*. Basel: BIS.

Blanchard, O. J. and J. Galí (2007). The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s? Working Paper, Pompeu Fabra Universtiy.

Blanchard, O. J. and C. M. Kahn (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica 48*(5), 1305–1312.

Bryan, M. F. and S. G. Cecchetti (1994). Measuring Core Inflation. In N. G. Mankiw (Ed.), *Monetary Policy.*

Bryan, M. F. and C. J. Pike (1991). Median price changes: an alternative approach to measuring current monetary inflation. *Economic Commentary*, 1–4.

Dhawan, R. and K. Jeske (2007). Taylor Rules with Headline Inflation: A Bad Idea. Federal Reserve Bank of Atlanta, Working Paper Number 2007-14.

Erceg, C. J., D. W. Henderson, and A. T. Levin (2000). Optimal Monetary Policy with Staggered Wage and Price Contracts. *Journal of Monetary Economics 46*, 281–313.

Goodfriend, M. and R. King (1997). The New Neoclassical Synthesis and the Role of Monetary Policy. *NBER Macroeconomics Annual*, 231–281.

King, R. G. and A. L. Wolman (1998). What Should the Monetary Authority Do When Prices Are Sticky? In J. B. Taylor (Ed.), *Monetary Policy Rules*. The University of Chicago Press.

Quah, D. and S. Vehey (1995). Measuring Core Inflation. *The Economic Journal*, 1130–1144.

Rotemberg, J. and M. Woodford (1997). An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. *NBER Macroeconomics Annual 12*, 297–346.

Taylor, J. B. (1979). Estimation and Control of a Macroeconomic Model with Rational Expectations. *Econometrica 47*(5), 1267–1286.

Woodford, M. (2003). *Interest and Prices*. Princeton, NJ: Princeton University Press.

Figure 1: CPI Inflation

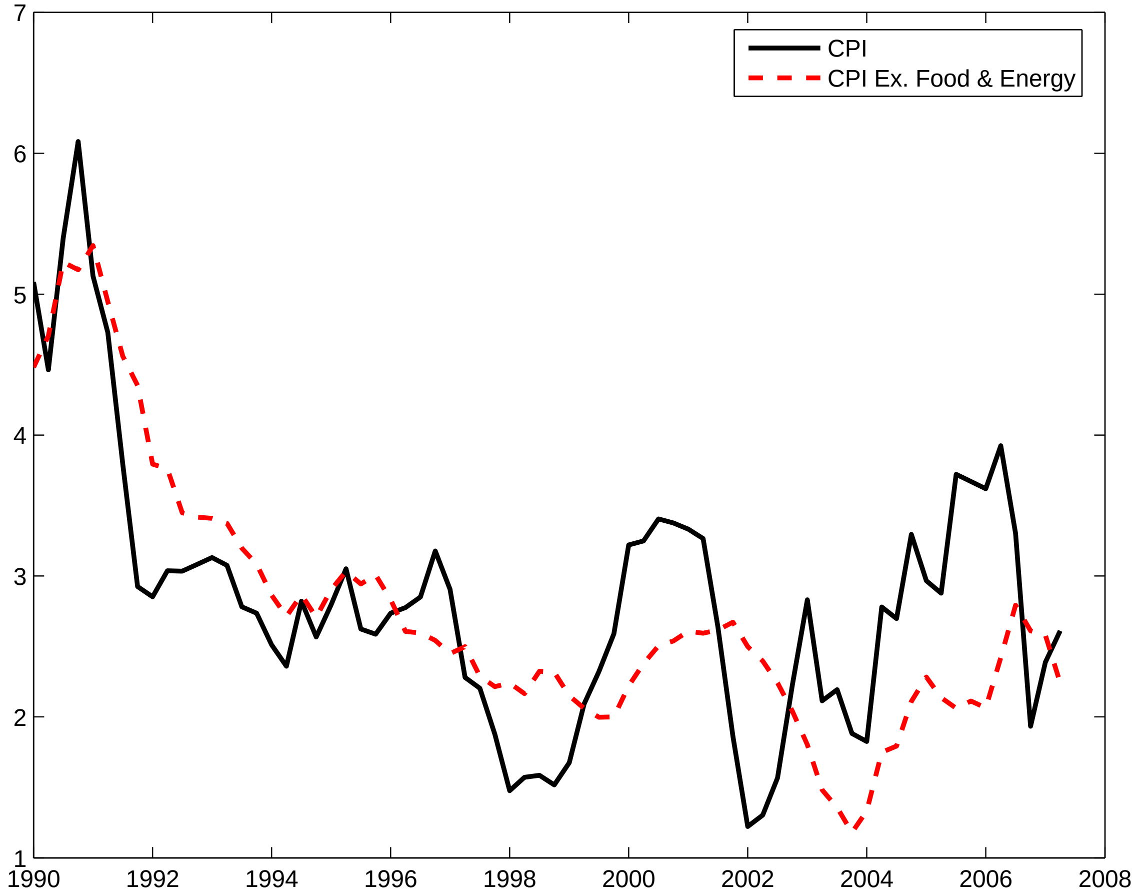

Figure 2: CPI Inflation Forecasts

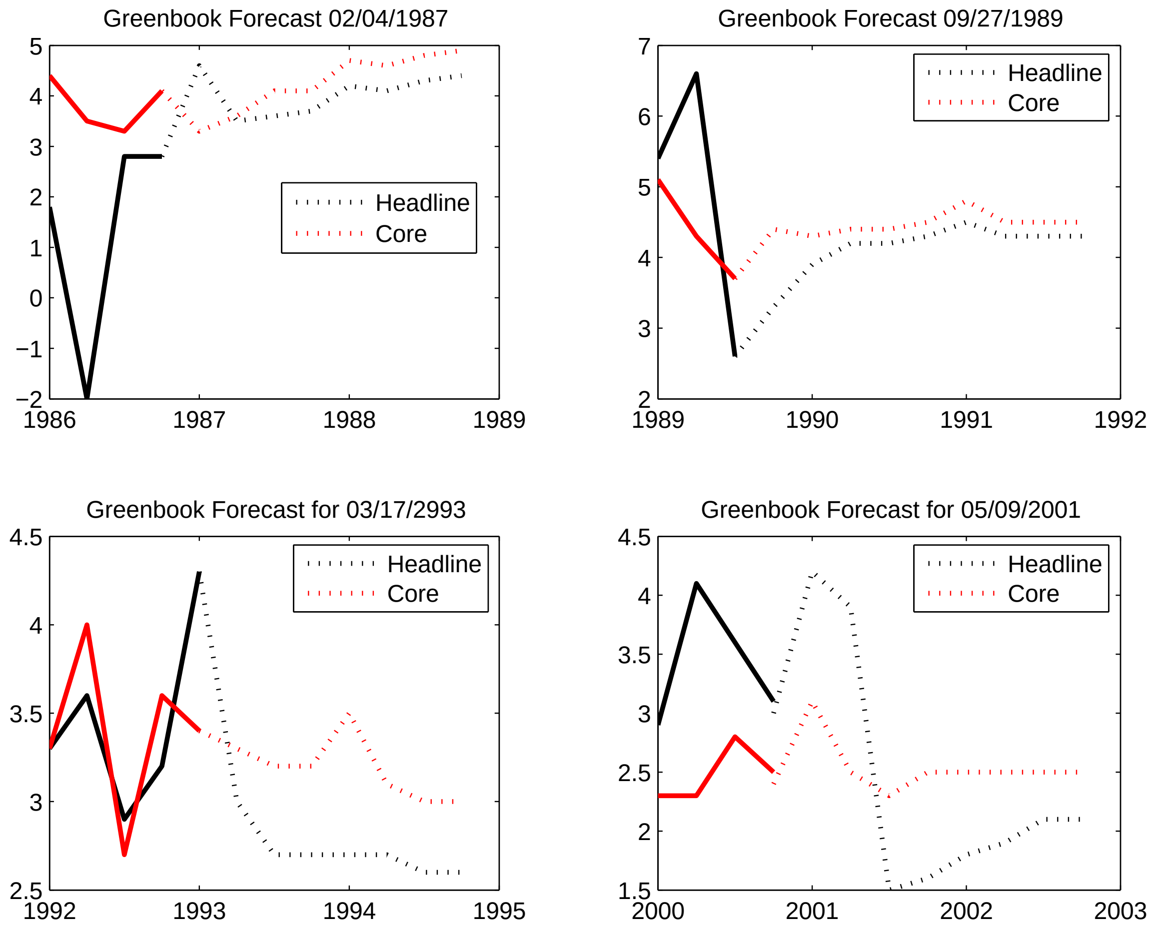

Figure 3. Contraction in energy supply: Optimal Policy vs. Flexible Price Equilibriu

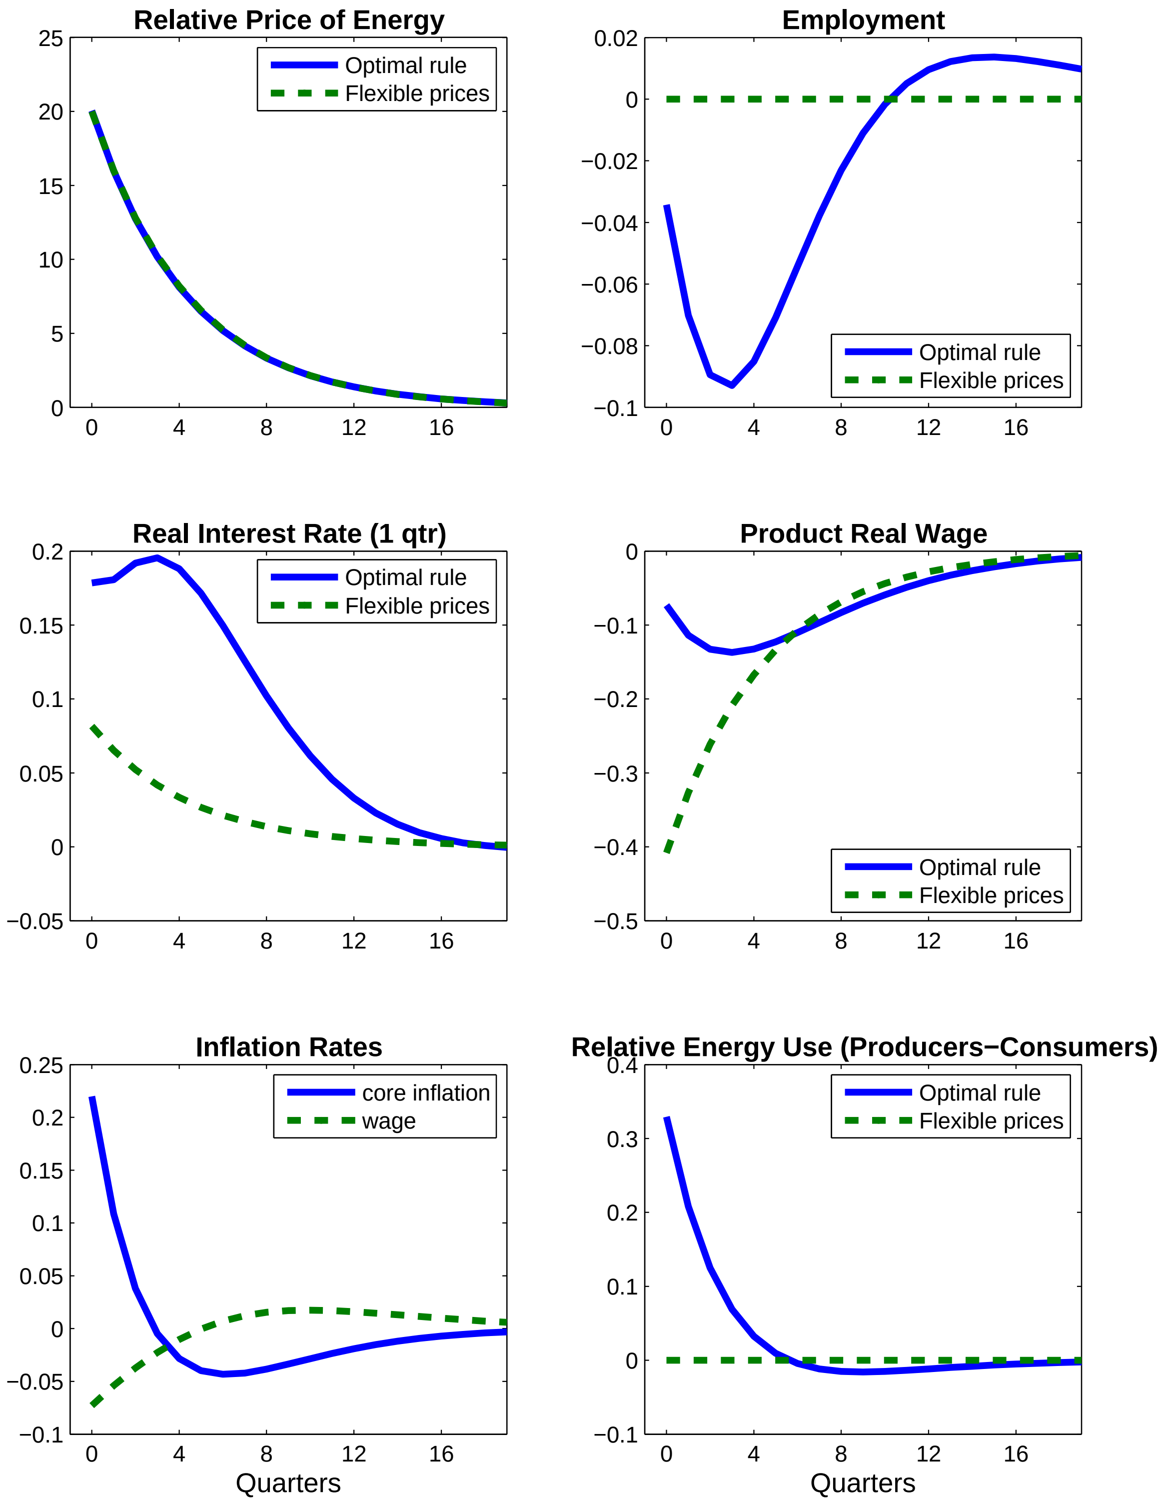

Figure 4. Temporary Fall in Energy Supply: Inflation Targeting Rules

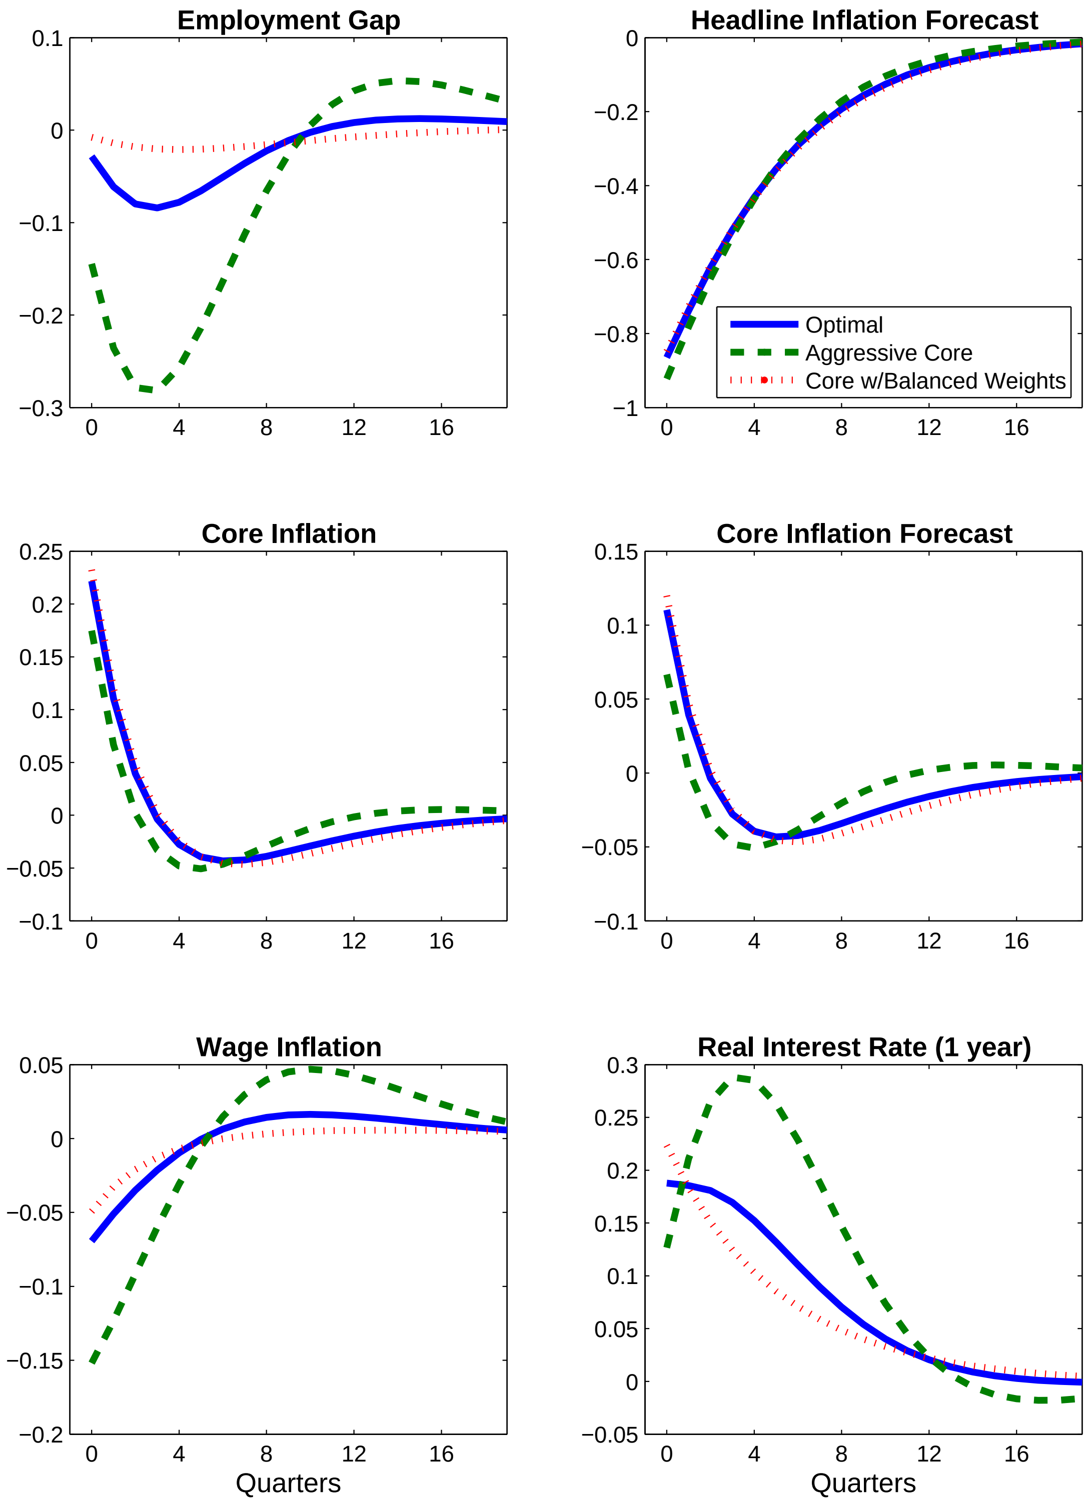

Figure 5. Permanent Fall in Energy Supply: Inflation Targeting Rules

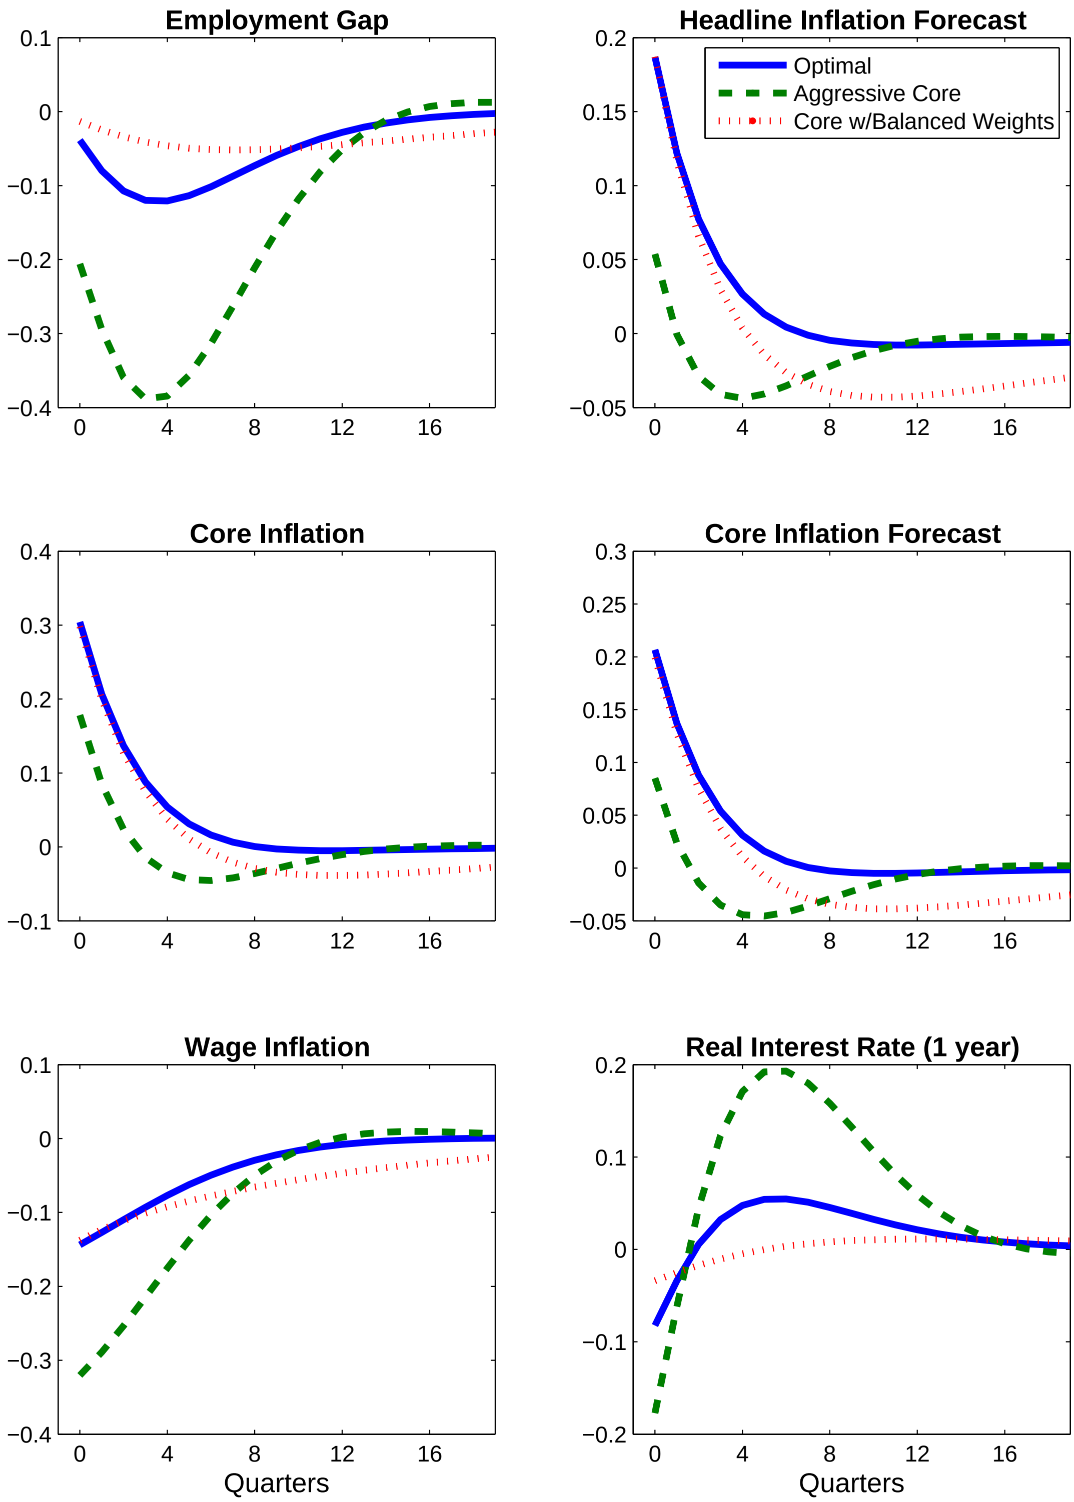

Figure 6. Temporary Fall in Energy Supply: Forecast−Based Taylor Rules

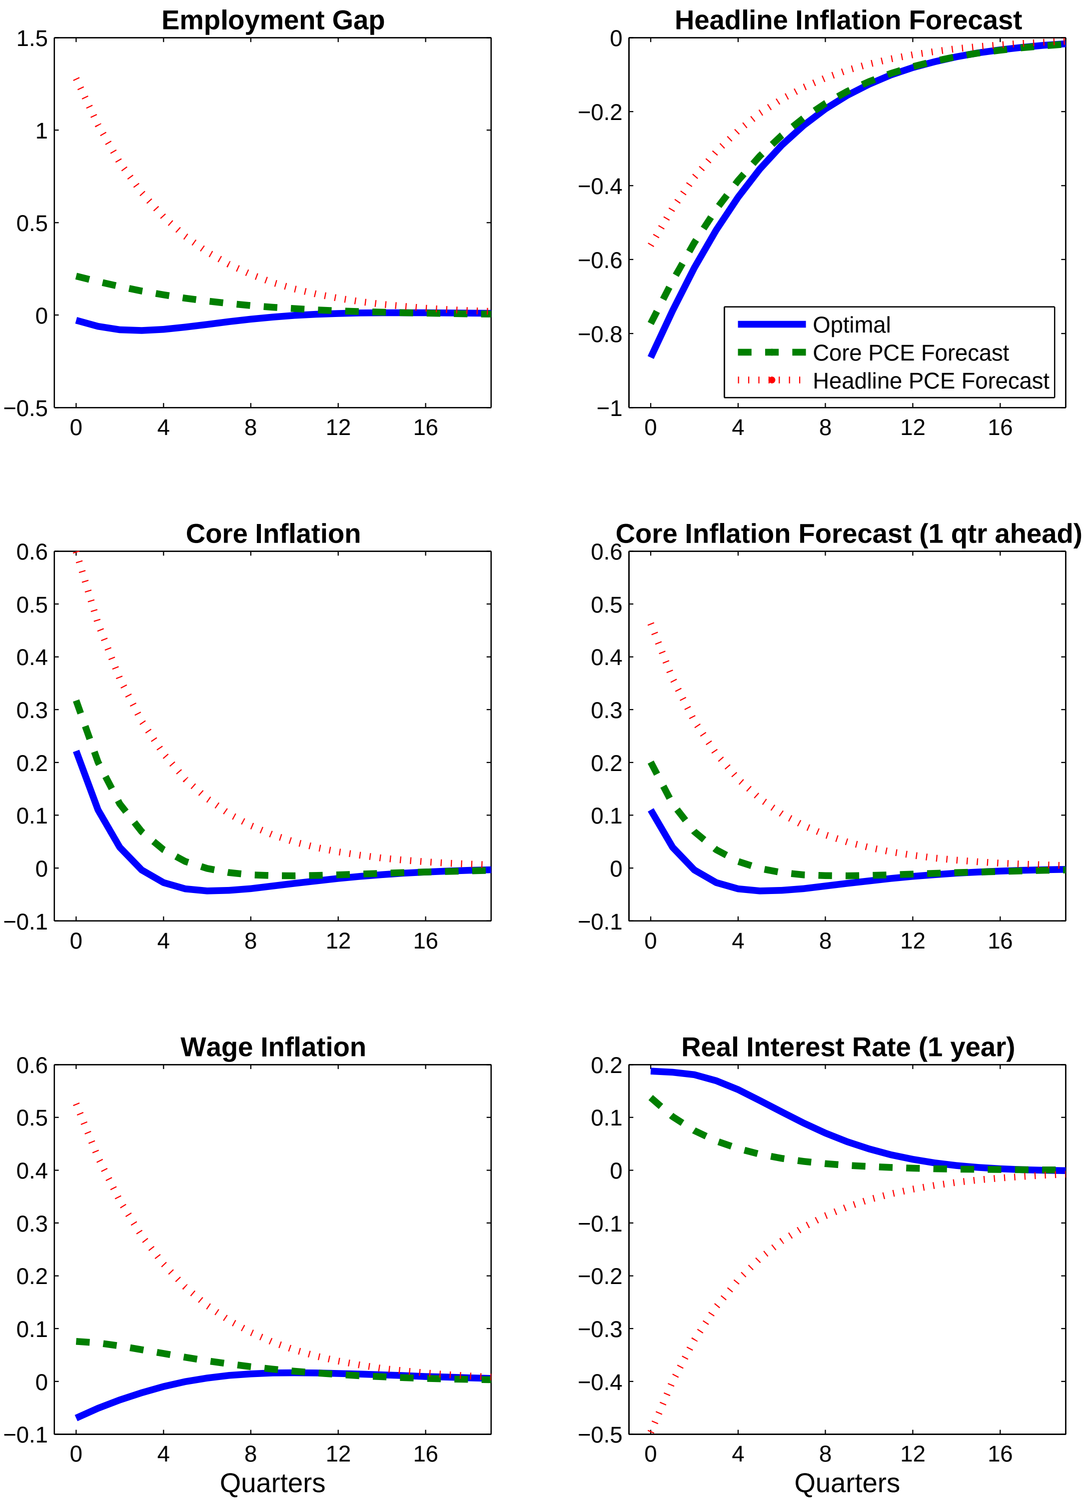

Figure 7. Temporary Fall in Energy Supply: Forecast−Based Taylor Rules

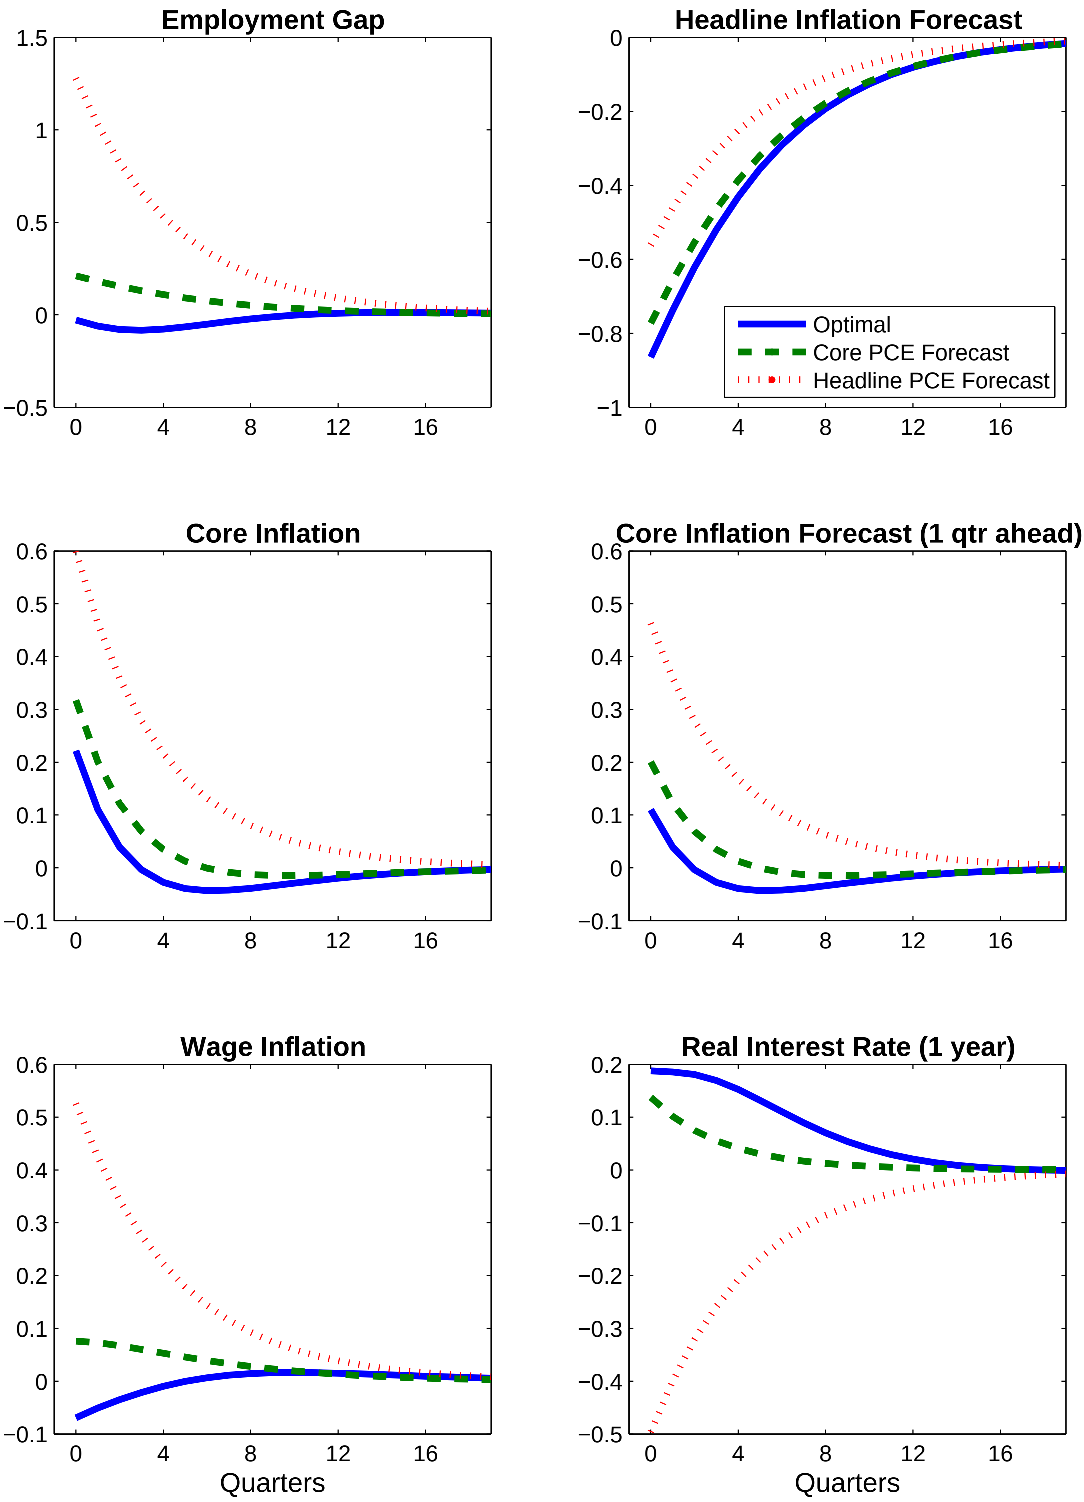

Figure 8. Fall in Energy Supply Perceived as Temporary: Imperfect Information