How Should Monetary Policy Respond to Changes in the Relative Price of Oil? Considering Supply and Demand Shocks

Michael Plante

Research Department
Working Paper 1202

Federal Reserve Bank of Dallas

HOW SHOULD MONETARY POLICY RESPOND TO CHANGES IN THE RELATIVE PRICE OF OIL? CONSIDERING SUPPLY AND DEMAND SHOCKS.

MICHAEL PLANTE

ABSTRACT. This paper examines optimal monetary policy in a New Keynesian model where the relative price of oil is affected by exogenous supply shocks and a productivity-driven demand shock. When wages are flexible, stabilizing core inflation is optimal and the nominal rate rises (falls) in response to a demand (supply) shock. When both prices and wages are sticky, core inflation falls (rises) in response to the demand (supply) shock. Stabilizing CPI inflation generates small welfare losses only if the demand shock is the main driver of oil prices. Based on a VAR estimated using post-1986 data for the U.S., both shocks have had minimal impacts on core inflation. The federal funds rate rises in response to the demand shock but falls in response to the supply shock, consistent with the predictions of the theoretical model for a policy that stabilizes core inflation.

JEL Classifications: E31, E52, Q43
Keywords: oil prices, optimal monetary policy, inflation

1. Introduction

A long tradition in the literature on monetary policy and oil prices has been to assume that the relative price of oil is exogenous to the model being considered.[1] One of the motivations for this setup was the view that exogenous supply shocks, particularly due to OPEC, were the fundamental driver of oil prices. Kilian (2009), however, has provided evidence that demand shocks are also important in determining the price of oil, and that macroeconomic aggregates may respond differently to these shocks even though they also cause the price of oil to increase.

This raises an interesting question about whether or not monetary policy should respond differently to a rise in the price of oil driven by demand rather than supply. One way to answer this question is to solve for an optimal monetary policy and examine the impulse response functions of interest rates and inflation variables to see if they differ in important ways in response to supply and demand shocks. The approach taken in this paper is to consider the optimal responses that come from a welfare maximizing policy done from the timeless perspective with a non-distorted steady state.

For this type of optimal policy, Bodenstein, Erceg, and Guerrieri (2008) showed that when wages and prices are sticky an exogenous supply shock brings about a rise in core inflation and a decrease in nominal wage inflation. The supply shock reduces the marginal product of labor so the real wage should fall, and higher core inflation helps bring about this adjustment. Echoing the results provided in Aoki (2001), there is no explicit attempt to stabilize the price of oil because it is a flexible price. As such, there is no distortion associated with changes in that price and, therefore, no reason to stabilize the inflation rate of that price. As a consequence, there is also no explicit attempt to stabilize CPI inflation.

In this paper, I show that there are fundamental differences in the optimal responses when the increase in the price of oil is due to a productivity-driven demand shock instead of a supply shock. While core inflation initially rises in response to an exogenous oil supply shock, it falls in response to the demand shock. Even though the productivity shock drives up the price of oil, on net it increases the marginal product of labor, the opposite of what occurs with the supply shock. This calls for core inflation to decrease so as to help push the real wage up.

---

[1]Examples of this approach include Leduc and Sill (2004), Dhawan and Jeske (2007), Blanchard and Gali (2010), and the previous version of this paper, Plante (2009). Further examples of the exogenous price assumption can be found in literature that explores how oil prices affect the macroeconomy, including, but not limited to, Finn (2000), Rotemberg and Woodford (1996), and Kim and Loungani (1992). Some recent work has begun to model endogenous oil prices, including Bodenstein, Erceg, and Guerrieri (2008), Nakov and Pescatori (2010).

Some important differences remain even in the simpler case when wages are flexible. While core inflation is perfectly stabilized regardless of the shock, the nominal interest rate adjusts quite differently. In response to the demand shock the rate rises, but it falls in response to the exogenous supply shock.

The welfare implications of some alternative policy rules that stabilize core inflation, CPI inflation, or nominal wage inflation are also examined. Policy rules that stabilize core or nominal wage inflation produce relatively minor welfare losses in all of the cases considered, so long as the response to inflation is not too weak. The costs of stabilizing CPI inflation, however, depend upon the relative importance of supply and demand shocks. When productivity shocks are the sole driver of oil prices, stabilizing CPI inflation performs relatively well. When exogenous oil supply shocks drive the relative price of oil, however, this policy produces high losses compared to stabilizing core inflation or nominal wage inflation.

The final contribution of this paper is to take the model's predictions about monetary policy to the data using a modified version of the VAR introduced in Kilian (2009). The modified model can identify exogenous oil supply shocks, demand shocks driven by global economic activity, and a demand shock driven by unexpectedly strong economic activity in the United States. Data from the post-1986 era is used to estimate the model.

Impulse response functions from the VAR show minimal movements in core inflation in response to both the exogenous supply shock and the shock to U.S. real GDP. An oil supply shock causes a rise in core inflation of about 10 basis points in the first month. In response to the demand shock, core inflation initially falls by a trivial amount. But, the initial responses and the ones following that are not statistically different from 0 for either shock.

The federal funds rate adjusts differently in response to the two shocks. The month to month changes are small, but the cumulative impacts show the federal funds rate falling in response to the exogenous supply shock but increasing in response to the demand shock. These results are similar to the findings in Kilian and Lewis (2011), which showed that the federal funds rate had a tendency to fall in response to a supply shock and increase in response to a demand shock driven by global economic activity.

Interestingly, the movement in the funds rate is qualitatively similar to the response the theoretical model predicts should occur when the central bank stabilizes core inflation. This finding suggests that monetary policy has effectively distinguished between different shocks that affect the price of oil and successfully stabilized core inflation. While this may not be the fully optimal policy in response to both shocks, the losses predicted by the theoretical model under this policy are relatively small in nature.

The rest of the paper is organized as follows. Section two outlines the model. Section three presents the results about the optimal policy and the policy rules. Empirical results are presented in section four. Section five concludes.

2. THE MODEL

The model used is a New Keynesian model modified to incorporate firm and household demand for oil, capital accumulation, and an exogenous oil supply. The following exposition introduces those equations necessary for understanding the special features of the model or the results presented in later sections. All other equations and derivations are relegated to the technical appendix.

2.1. **Production.** There is a continuum of firms of measure 1 in the intermediate goods sector, with firms being indexed by i. Each firm produces a specific good, denoted by y_{it}, which is sold at price p_{it}. The final good, Y_t, is produced using the intermediate goods under the standard assumptions. The elasticity of substitution between the various intermediate goods is given by θ, with $\theta > 1$.

In this model the price of the final good, P_t, is also the price of the final, non-oil consumption good and is therefore referred to as the core CPI. The equation for core inflation is

$$\Pi_t = \frac{P_t}{P_{t-1}}. \tag{1}$$

The final good is the numeraire so all nominal variables are deflated by P_t and these are referred to as real variables.

In the intermediate goods sector, each firm i uses capital, k_{it}, an aggregated labor input, h_{it}, and oil, o_{it}^f to produce y_{it}. The real prices of capital, labor, and oil are denoted as R_t, W_t, and P_t^o, respectively. Technology is CES and the functional form is

$$y_{it} = A \left[a_1 \left(z_t h_{it} \right)^{\frac{\eta-1}{\eta}} + a_2 o_{it}^{f \, \frac{\eta-1}{\eta}} + a_3 k_{it}^{\frac{\eta-1}{\eta}} \right]^{\frac{\eta}{\eta-1}},$$

where η is the elasticity of substitution between the inputs, A a scaling parameter, and a_1, a_2, and a_3 are distribution parameters. The variable z_t is a temporary productivity shock. The first order conditions for the inputs come from solving the usual cost-minimization problem.

The real unit cost function for this technology is

$$\phi_t = \frac{1}{A} \left[a_1^\eta \left(\frac{W_t}{z_t} \right)^{1-\eta} + a_2^\eta P_t^{o\,1-\eta} + a_3^\eta R_t^{1-\eta} \right]^{\frac{1}{1-\eta}}. \tag{2}$$

The assumption that there is no factor specificity at the firm level implies that the unit cost function is taken as given by firms in the intermediate goods sector.

Price setting follows Calvo (1983) in that only a set fraction of firms change their prices each period, with the probability of a price remaining fixed given by ω. Firms that are able to change their price do so by solving the standard profit maximization problem, given by

$$\max_{p_{it}} \Gamma_{it} = E_t \sum_{j=0}^{\infty} \omega^j \frac{\beta^j \lambda_{t+j}}{\lambda_t} \left[(1+\tau^p) \left(\frac{p_{it}}{P_{t+j}} \right)^{1-\theta} Y_{t+j} - \left(\frac{p_{it}}{P_{t+j}} \right)^{-\theta} \phi_{t+j} Y_{t+j} \right], \qquad (3)$$

where β is the discount factor of the representative household and λ_t is equal to the marginal utility of aggregate consumption. In line with Bodenstein, Erceg, and Guerrieri (2008), I assume the existence of a subsidy, τ^p, which removes the distortion due to monopolistic competition in the steady state.

2.2. **The Labor Input.** Following Schmitt-Grohe and Uribe (2006) the labor input h_{it} is an aggregate of a continuum of labor types indexed by j. The labor types are imperfect substitutes for each other with the elasticity of substitution given by $\theta_n > 1$. Demand by firm i for labor type j is denoted as h_{it}^j. The setup for producing h_{it} from h_{it}^j is analogous to the one used when producing Y_t from y_{it}. That is,

$$h_{it} = \left[\int_0^1 h_{it}^{j \frac{\theta_n-1}{\theta_n}} dj \right]^{\frac{\theta_n}{\theta_n-1}}.$$

The demand for labor type j by firm i is given by

$$h_{it}^j = \left(\frac{w_t^j}{W_t} \right)^{-\theta_n} h_{it},$$

where $w_t^j = \frac{\tilde{w}_t^j}{P_t}$ is just the real wage of labor type j. The variable W_t is an aggregate real wage index given by

$$W_t = \left[\int_0^1 w_t^{j \ 1-\theta_n} dj \right]^{\frac{1}{1-\theta_n}}.$$

Since W_t is the real wage, nominal wage inflation is given by

$$\Pi_t^w = \frac{W_t}{W_{t-1}} \Pi_t. \qquad (4)$$

As with prices, only a set fraction of nominal wages change each period, with ω_n being the probability of a wage being fixed in the current period. A first order condition for the reset wages comes directly from the agent's optimization problem. As the job types are ex-ante identical, all wages reset in the same period are equal in equilibrium. The optimal real wage

5

choice is denoted as $\frac{\tilde{w}_t^*}{P_t} = w_t^*$. Any real wage which is changed in period t is given by

$$w_t^j = w_t^*,$$

while wages that remain fixed are given by

$$w_t^j = \frac{w_{t-1}^j}{\Pi_t}.$$

Note that the unchanged wages are deflated by core inflation because they are in real terms, not nominal, and over time the real wage may rise or fall depending on what happens with core inflation.

2.3. Aggregation.

Aggregation is done linearly across firms and labor types. This introduces two state variables, Δ_t and Δ_t^w, which measure price and real wage dispersion, respectively. The equation for the price dispersion variable is

$$\Delta_t = \int_0^1 \left(\frac{p_{it}}{P_t}\right)^{-\theta} di \tag{5}$$

while the equation for real wage dispersion is

$$\Delta_t^w = \int_0^1 \left(\frac{w_t^j}{W_t}\right)^{-\theta_n} dj. \tag{6}$$

These variables are important because price and wage dispersion bring about deadweight losses through the inefficient use of the various goods and labor types. Optimal policy will, therefore, focus on minimizing these distortions in order to reduce the welfare losses they generate. This implies that there will be an emphasis on stabilizing core and/or nominal wage inflation in response to shocks.

2.4. The Household.

Utility from the final, non-oil consumption good, C_t, and from oil products, O_t^h, is aggregated with a CES function

$$\left(C_t^{\frac{\nu-1}{\nu}} + \kappa_2 O_t^{h\,\frac{\nu-1}{\nu}}\right)^{\frac{\nu}{\nu-1}},$$

where ν is the elasticity of substitution between non-oil and oil consumption and κ_2 is a distribution parameter. Per period utility is given by

$$\left[\frac{\left(C_t^{\frac{\nu-1}{\nu}} + \kappa_2 O_t^{h\,\frac{\nu-1}{\nu}}\right)^{\left(\frac{\nu}{\nu-1}\right)\left(1-\frac{1}{\tau}\right)}}{1-\frac{1}{\tau}} - \kappa_1 \frac{N_t^{1+\frac{1}{\mu}}}{1+\frac{1}{\mu}}\right],$$

where τ is the intertemporal elasticity of substitution, μ is the wage elasticity of labor supply, and $N_t = \int_0^1 n_t^j dj$ is aggregate labor supplied by the agent.

6

For convenience, I re-write per period utility in terms of real aggregate consumption expenditure, $X_t = C_t + P_t^o O_t^h$. This can be done by setting up the indirect utility version of the aggregator function.[2] In this form per period utility is equal to

$$\frac{\left(\frac{X_t}{P_t^{CPI}}\right)^{1-\frac{1}{\tau}}}{1-\frac{1}{\tau}} - \kappa_1 \frac{N_t^{1+\frac{1}{\mu}}}{1+\frac{1}{\mu}}. \tag{7}$$

The P_t^{CPI} term is simply the CPI, $\tilde{P}_t^{CPI}$, deflated by P_t. As the aggregator is CES the exact equation for this term is

$$P_t^{CPI} = \left(1 + \kappa_2^\nu P_t^{o1-\nu}\right)^{\frac{1}{1-\nu}}. \tag{8}$$

CPI inflation is therefore

$$\Pi_t^{CPI} = \Pi_t \left[\frac{1 + \kappa_2^\nu P_t^{o\,1-\nu}}{1 + \kappa_2^\nu P_{t-1}^{o\,1-\nu}}\right]^{\frac{1}{1-\nu}}. \tag{9}$$

As expected, CPI inflation consists of a core inflation component as well as a term that represents, in a slightly complicated manner, changes in the relative price of oil. This is clearer in the special case when the aggregator is Cobb-Douglas, in which case CPI inflation is equal to

$$\Pi_t^{CPI} = \Pi_t \left(\frac{P_t^o}{P_{t-1}^o}\right)^{\gamma_o},$$

where γ_o is the weight of oil products in the Cobb-Douglas aggregator.

In addition to a short-term nominal bond, I assume the existence of two other bonds: a real bond indexed to the core CPI and a real bond indexed to the CPI. The first-order conditions for these bonds provide the equations that link up the nominal interest rate with the two real interest rates. These bonds are in net-zero supply and do not affect the results, but their first-order conditions will be necessary to discuss the results in sections 3 and 4.

The agent's budget constraint is

$$\begin{aligned}
X_t + B_t + B_t^{CORE} + P_t^{CPI} B_t^{CPI} + I_t^k &= (1+\tau^w)\int_0^1 w_t^j n_t^j dj + R_t K_t + T_t \\
&\quad + \Gamma_t + \frac{I_{t-1}}{\Pi_t} B_{t-1} + R_{t-1}^{CORE} B_{t-1}^{CORE} \\
&\quad + P_t^{CPI} R_{t-1}^{CPI} B_{t-1}^{CPI} + P_t^o O_t^s.
\end{aligned} \tag{10}$$

As in Bodenstein, Erceg, and Guerrieri (2008), I assume the existence of a subsidy, $1 + \tau^w$, which removes the distortion in the steady state due to monopolistic competition in labor markets.

---

[2]This is done for expositional purposes only and has no impact on the results. Nor does the fact that the non-oil consumption good is the numeraire.

In the budget constraint B_t is real holdings of the nominal bond, I_{t-1} is the nominal return on that bond, I_t^k is aggregate investment spending on capital goods, Γ_t is aggregated real profits, T is lump sum transfers, and I_{t-1} is the nominal return on the bond. Holdings of the bond indexed to core inflation are denoted as B_t^{CORE} and its real return denoted as R_t^{CORE}. The bond indexed to the CPI is denoted as B_t^{CPI} and its real return is denoted as R_t^{CPI}. Note that since the numeraire is the non-oil consumption good, there is a P_t^{CPI} term multiplying holdings of the CPI-indexed bond.[3]

The law of motion for capital is given by

$$K_{t+1} - K_t = I_t^k - \delta_k K_t. \tag{11}$$

2.5. **Interest Rates on Bonds in the Model.** The fact that the relative price of the two consumption goods is not necessarily equal to 1 implies there are two ways to index bonds to inflation in the model, and hence two different real interest rates. One is linked to core inflation and the other linked to CPI inflation. This is different from the basic New Keynesian model, where there is simply one real interest rate linked to the price of the final good.

In the standard New Keynesian model with log preferences ($\tau = 1$), the log-linearized first order conditions for consumption, C_t, the nominal bond, and a real bond indexed to P_t are given by

$$
\begin{aligned}
-\hat{C}_t &= \hat{\lambda}_t, \\
\hat{I}_t &= \hat{\lambda}_t - E_t\hat{\lambda}_{t+1} + E_t\hat{\Pi}_{t+1}, \\
\hat{R}_t^{CORE} &= \hat{\lambda}_t - E_t\hat{\lambda}_{t+1}
\end{aligned}
$$

where λ_t is the multiplier on the household's budget constraint, equal to the marginal utility of consumption. Variables in hats are log-deviations of variables from their steady state values. The first order condition for the nominal bond and the real bond can be combined to provide the usual (log-linearized) Fisher equation which says that deviations in the nominal interest rate are equal to deviations in the real interest rate and expected inflation. Deviations in the real interest rate are driven by the growth rate of consumption.

For the model used here, assuming log preferences for simplicity, the first order conditions for aggregate consumption expenditure, the nominal bond, the bond indexed to the core

---

[3]The technical appendix describes how to go from the nominal budget constraint to the real budget constraint, denominated in terms of the final good / non-oil consumption good.

CPI, and the bond indexed to the CPI are, respectively,

$$X_t^{-1} = \lambda_t, \tag{12}$$

$$\lambda_t = \beta E_t \frac{I_t \lambda_{t+1}}{\Pi_{t+1}}, \tag{13}$$

$$\lambda_t = \beta E_t R_t^{CORE} \lambda_{t+1}, \tag{14}$$

$$\lambda_t = \beta E_t \frac{\Pi_{t+1}^{CPI}}{\Pi_{t+1}} R_t^{CPI} \lambda_{t+1}. \tag{15}$$

The log-linearized versions of these equations are

$$-\hat{X}_t = \hat{\lambda}_t,$$

$$\hat{I}_t = \hat{\lambda}_t - E_t \hat{\lambda}_{t+1} + E_t \hat{\Pi}_{t+1},$$

$$\hat{R}_t^{CORE} = \hat{\lambda}_t - E_t \hat{\lambda}_{t+1},$$

$$\hat{R}_t^{CPI} = \hat{\lambda}_t - E_t \hat{\lambda}_{t+1} + E_t \left(\hat{\Pi}_{t+1} - \hat{\Pi}_{t+1}^{CPI} \right).$$

As in the standard model, the first order condition for the nominal bond and the bond linked to the core CPI can be combined again to give the usual log-linearized Fisher equation,

$$\hat{I}_t = \hat{R}_t^{CORE} + E_t \hat{\Pi}_{t+1}. \tag{16}$$

There is also a Fisher equation linking up the nominal interest rate and the real return on the bond indexed to the CPI,

$$\hat{I}_t = \hat{R}_t^{CPI} + E_t \hat{\Pi}_{t+1}^{CPI}. \tag{17}$$

Equations (16) and (17) imply that the model is still Wicksellian in that stabilizing inflation requires the nominal interest rate to track the real rate. But, the particular real interest rate that needs to be tracked depends upon the inflation rate being stabilized. If core inflation is being stabilized then the nominal interest rate should track R_t^{CORE}, whereas if CPI inflation is being stabilized then the nominal interest rate should track R_t^{CPI}.

One important difference between the model in this paper and the basic model is that it is not the growth rate of consumption itself which drives R_t^{CORE} but instead the growth rate in aggregate consumption denominated in terms of the non-oil consumption good. The growth rate of the aggregate consumption basket, denominated in terms of the CPI, is linked to R_t^{CPI}. This suggests that intuition about consumption smoothing and the behavior it implies for the real interest rate should apply to R_t^{CPI} but may not apply to R_t^{CORE}. The importance of this will become clearer in section 3.

9

2.6. **Oil Supply.** The supply of oil is given by an exogenous process,

$$\ln O_t^s = (1 - \rho_o)O^s + \rho_o O_{t-1}^s + \epsilon_t^o, \tag{18}$$

where ϵ_t^o is a mean zero, i.i.d. shock with a standard deviation of σ_o. The assumption of an exogenous supply of oil follows Bodenstein, Erceg, and Guerrieri (2008). Unlike that model, there is just a temporary shock to the supply instead of a temporary and a (near) permanent component.

The relative price of oil is determined endogenously by the market clearing condition,

$$O_t^f + O_t^h = O_t^s. \tag{19}$$

2.7. **Monetary Policy.** Monetary policy is given by the Ramsey optimal solution or by a simple Taylor-type rule. The Ramsey problem maximizes the expected sum of discounted utility subject to all of the equilibrium equations in the model, such as the agent's first order conditions, the resource constraint, and the law of motions for the exogenous processes. The optimal policy is done under full commitment from the timeless perspective. This problem produces its own set of first order conditions which, along with the original equilibrium conditions, provide a set of non-linear equations that can then be solved using standard techniques.

For cases where the Ramsey solution is not used, monetary policy follows a simple Taylor type rule of the form

$$\ln(\frac{I_t}{I}) = \alpha_\pi \ln(\frac{\Pi_t}{\Pi}) \tag{20}$$

where the parameters I, and Π are the steady state values of the nominal interest rate and core inflation, respectively. Rules that replace Π_t with Π_t^w or Π_t^{CPI} are also considered.

As shown in the previous version of this paper, there was an inverse relationship between the volatility of the variable in the rule and the calibration of α_π. This continues to hold in the model used here. Consequently, these policy rules can be considered as alternative policies that focus on stabilizing the particular inflation variable in the rule.

2.8. **Calculation of the Welfare Losses and Solution Method.** Welfare losses are calculated as the amount of aggregate real consumption required to produce the same welfare in a model with sub-optimal policy as the Ramsey optimal policy. To make this clearer, define per period utility under the Ramsey optimal solution as $U(\frac{X_t^r}{P_t^{CPI}}, N_t^r)$ and in a non-optimal solution as $U(\frac{X_t^n}{P_t^{CPI}}, N_t^n)$. Then the conditional welfare losses, λ^c, are implicitly given by

$$E_t \sum_{t=0}^{\infty} \beta^t \left\{ U\left[\frac{X_t^n}{P_t^{CPI}}, N_t^n\right] \right\} = E_t \sum_{t=0}^{\infty} \beta^t \left\{ U\left[(1 - \lambda^c)\frac{X_t^r}{P_t^{CPI}}, N_t^r\right] \right\}.$$

10

As shown in Schmitt-Grohe and Uribe (2006) it is possible to derive an equation for λ^c. In order to be able to distinguish the welfare implications of different policies it is necessary to solve a second order approximation of the model. I use the method of Schmitt-Grohe and Uribe (2004) and the code written by the authors in conjunction with that paper to solve the model. Further details can be found in that paper.

2.9. **Calibration.** The model is calibrated to an initial steady state using a calibration that follows Bodenstein, Erceg, and Guerrieri (2008) in most regards. Table 1 shows the values for the model's parameters. The elasticity of substitution between intermediate goods, θ, and the elasticity of substitution between various labor types, θ_n, are both set to 6. Both the probability of a price and a wage being fixed in any given period, ω and ω_n, are set to .75. Log preferences are used so τ is set to 1. The steady state inflation rate is the Ramsey optimal steady state inflation rate. As this paper abstracts from money demand the the optimal steady state gross inflation rate is 1. Real GDP is calibrated to unity and steady state real aggregate consumption expenditure, investment spending, household demand for oil, and firm demand for oil are set as percentage shares of GDP. Household and firm demand for oil products are set to 5 percent of GDP and 2 percent of GDP, respectively. These roughly match the averages found in NIPA data for energy usage by firms and households from 1987 to 2008.

The one significant departure in the calibration is the elasticity of substitution parameters, ν and η. In Bodenstein, Erceg, and Guerrieri (2008) these are set to 1. Here, the elasticity of substitution between C_t and O_t^h, ν, is set to .25. This makes C_t and O_t^h Edgeworth compliments and ensures that the price elasticity of demand for oil products is fairly low. The elasticity of substitution in production, η, is calibrated to set the compensated price elasticity of demand for oil. Empirical studies tend to find fairly small sizes, in absolute terms, for both this elasticity and other closely related elasticities. I calibrate the value of η to .25 which sets the own price elasticity of oil to about $-.25$. The motivation for these choices comes from the calibration of the shock processes.

The law of motion for the technology shock is

$$\ln z_t \;=\; \rho_z \ln z_{t-1} + \epsilon_t^z, \tag{21}$$

where ϵ_t^z is an i.i.d white noise shocks with standard deviation σ_z. The parameters for the oil supply process also need to be calibrated. The parameters ρ_z and ρ_o are set to .80. The standard deviations of the shocks are set so that volatility of real GDP and the relative

price of oil in the model match that found in U.S. data from 1987 to the present.[4] For the model with sticky wages and prices, this led to a calibration of σ_o to .003545 and σ_z to .0175. Without low price elasticities of demand these volatilities would have to unrealistically large to generate sufficient volatility in the price of oil.

3. Theoretical Results

This section presents results for the optimal policy and the policy rules for the cases where only prices are sticky and where both wages and prices are sticky. Comparing the results from the two provides strong intuition about what drives the optimal policy in a very general manner. It also gives a broader set of theoretical results to take to the data.

3.1. **Optimal Policy with Flexible Wages.** It is well known that when sticky prices are the only distortion that there is no tradeoff between stabilizing inflation and the output gap. It has also been shown that a policy that fully stabilizes inflation makes the equilibrium nominal interest rate change one-for-one with changes in the real interest rate.[5]

These results also hold in the model used here. As discussed in Bodenstein, Erceg, and Guerrieri (2008), and presented in a more general form in Aoki (2001), the inclusion of the oil sector does not add any distortions to this model since the price of that good is flexible. Since this price is flexible there is no reason to stabilize its inflation rate. This holds true regardless if supply or demand shocks are driving the price of oil. The analysis, however, is complicated by the fact that there are multiple inflation rates that could be stabilized and multiple real interest rates in the model.

Since wages are flexible, the relevant inflation variable to be stabilized is core inflation. This is because deviations in core inflation from its target level generate price dispersion, which will lead to welfare losses. As policy is done from the timeless perspective, core inflation is stabilized perfectly in response to all shocks.

Given that the optimal policy stabilizes core inflation, from equation (16) we can see that the relevant real interest rate that needs to be tracked is R_t^{CORE}, not R_t^{CPI}. In equilibrium, therefore, the log-deviations of the nominal interest rate will be given by

$$\hat{I}_t = \hat{R}_t^{CORE}.$$

[4]The model is simulated and the simulated data is then logged and HP filtered. A different method would lead to different calibrations for the volatilities but changing the volatilities does not impact the qualitative features of the impulse response functions.

[5]See Woodford (2003) and the numerous references therein for more discussion on this feature of the model.

The interesting question is whether or not R_t^{CORE} responds differently to supply and demand shocks. Figure 1, which plots the responses of the nominal interest rate and the two real interest rates to a negative oil supply shock and a positive productivity shock, shows that this is indeed the case. The top panel shows the deviations of I_t, the middle panel R_t^{CORE}, and the bottom panel R_t^{CPI}. The shocks are one-standard deviation in size. The responses have been annualized.

In response to a negative supply shock, the nominal interest rate is initially lowered and remains below average for around two years. But, in response to the productivity shock the rate is increased and remains above average for almost two years.

The movement in the nominal rate is being driven by the response of R_t^{CORE}. To see what drives R_t^{CORE} more explicitly, note that the first order conditions for aggregate consumption is given by

$$\lambda_t = \frac{1}{X_t},$$

where $X_t = C_t + P^o O_t^h$. This variable measures aggregate consumption expenditure, in terms of the non-oil good. In response to a negative oil supply shock, X_t initially rises and then falls overtime. This occurs because oil is demanded inelastically by the household, so that spending by households, denominated in terms of the non-oil consumption good, rises.[6] As such, the multiplier falls initially and then rises over time, bringing about the behavior of R^{CORE} seen in figure 1.

This may seem at odds with the usual intuition that the agent would want to smooth consumption over time in response to a shock like the supply shock. Such smoothing should cause the real interest rate to increase. But, the key point to remember is that the interest rate linked to the overall consumption basket is given by R_t^{CPI}, not R_t^{CORE}. As shown in figure 1, R_t^{CPI} rises in response to both shocks, as expected.

3.2. **Policy Rules with Flexible Wages.** Given that stabilizing core inflation is the optimal policy, how costly would it be to pursue a different policy, say of stabilizing CPI inflation, and would the results differ significantly if one shock or the other was more important in driving oil prices?

To answer these questions, the welfare losses for the three simple policy rules introduced earlier are calculated. This is done for three cases: one where both shocks hit the economy, one where only oil supply shocks hit the economy (σ_z set to 0), and one where only productivity shocks hit the economy (σ_o set to 0).

---

[6]This does not imply that utility is higher for the household, since utility is given by $\frac{X_t}{P_t^{CPI}}$.

Figure 2 plots the losses as a function of α_π. The top panel shows the case with both shocks, the middle panel the case with only supply shocks, and the bottom panel the case with only the productivity shock.

Under all three cases, the rules that stabilize core inflation or nominal wage inflation produce essentially zero welfare losses, so long as α_π is away from the lower bound of 1.[7]

The rule that stabilizes CPI inflation, however, only performs well when productivity shocks are the sole shock to hit the economy. Why is this so? Theoretically speaking, stabilizing CPI inflation always produces welfare losses in this model since this policy forces the core CPI to adjust in response to movements in the relative price of oil. How large these losses are, though, depends upon how big the movements in the relative price of oil are.

When productivity shocks are the sole shock in the model, these price movements are minute because firms use a fairly small amount of oil. Given this, the costs of stabilizing CPI inflation are small. The cost of stabilizing CPI inflation, however, would rise if the productivity shocks were more volatile, or if the firms used more oil to produce goods and services.

3.3. **Optimal Policy with Sticky Wages.** As shown in Bodenstein, Erceg, and Guerrieri (2008), sticky wages introduce a tradeoff between stabilizing inflation and the output gap in this model. This tradeoff exists because a shock which affects the marginal product of labor requires the real wage to adjust. When wages are flexible, it is costless to have nominal wages adjust to ensure the real wage is at its optimal level. With sticky wages, however, nominal wage movements generate their own welfare losses. These losses can be mitigated, partially, through movements in core inflation. For example, any shock which requires the real wage to fall optimally brings about a rise in core inflation.

Figure 3 shows the importance of this for the case of an oil supply shock. The black line is for the case with flexible wages while the dashed line is for the case with sticky wages. When wages are flexible, there is a sharp drop in wage inflation in the initial period of the shock and no movement in core inflation. As predicted, with sticky wages core inflation rises and there is a more muted response in nominal wage inflation.

An important finding in this paper is that the response to the productivity-driven demand shock is not the same as the response to the supply shock. This can be seen in figure 4, where the top panel plots the response of core inflation to the two shocks while the bottom panel

---

[7]In a previous version of this paper, Plante (2009), I showed that the volatility of the inflation variable (around its target value) in the rule is inversely related to the value of α_π, with the volatility becoming essentially zero for larger values of α_π. This holds in this model, as well, and explains why values of α_π closer to 1 cause larger welfare losses.

plots nominal wage inflation. Core inflation immediately rises in response to the supply shock, while in response to the demand shock it falls. Similarly, the immediate impact of wage inflation to the supply shock is down, but for the demand shock it is up.

The difference is because of the opposite impact the two shocks have on the marginal product of labor, which determines the optimal response of the real wage and core inflation. An exogenous supply shock reduces the marginal product of labor. This calls for a fall in the real wage and higher core inflation. While the productivity shock raises the price of oil, its net effect on the marginal product of labor is positive. This makes it optimal for the real wage to increase in response to the demand shock and then eventually decline over time. In order to accomplish this, the optimal policy calls for core inflation to decrease at first, and then eventually rise for some length of time.

3.4. **Policy Rules with Sticky Wages.** Given that the optimal policy allows for non-zero core inflation when wages are sticky, does this significantly change which policy rules perform well? To answer this, I re-calculate the welfare losses generated by the three policy rules and plot the losses as a function of α_π in figure 5.

A visual inspection shows that the qualitative results are exactly the same as in figure 2. When supply shocks drive the price of oil, a rule that stabilizes CPI inflation performs poorly. When these shocks are shut down and the demand shock drives oil prices, this rule performs just as well as the other rules.

The reason for this finding, however, is slightly more nuanced than the previous case. Under flexible wages, stabilizing CPI inflation performs poorly because it generates price dispersion. Under sticky wages, the optimal policy itself generates some price dispersion, so this explanation may not hold.

Indeed, it turns out that what drives the results is the impact that stabilizing CPI inflation has on the real wage, not on price dispersion. More specifically, when oil supply shocks affect the relative price of oil, stabilizing CPI inflation pushes the real wage in the wrong direction. One way to see this is to re-examine the equation for CPI inflation under the Cobb-Douglas assumption,

$$\Pi_t^{CPI} = \Pi_t \left(\frac{P_t^o}{P_{t-1}^o} \right)^{\gamma_o}.$$

Under the extreme case where CPI inflation is perfectly stabilized at its steady state level, any rise in the relative price of oil forces the core CPI to adjust down. The catch is that pushing the core CPI downwards pushes the real wage up, which is the wrong direction in response to an exogenous oil supply shock. Nominal wages compensate for this by adjusting more than they would under the optimal policy, leading to unnecessary welfare losses. The

15

stickier wages are relative to prices, the worse this policy will fare when supply shocks are the main driver of oil prices.

Pushing the real wage up, however, is the correct choice when productivity shocks are the sole driver of oil prices. In that case, rising oil prices require rising real wages. Consequently, if productivity shocks drive the price of oil then a policy rule which stabilizes CPI inflation performs just as well as the other two rules.

3.5. **Sensitivity Analysis.** In this section I conduct a sensitivity analysis by considering two special cases of the model. In the first, household demand for oil products is shut down so as to highlight the role that firm demand has on the optimal policy. This is followed by a scenario where firm demand for oil is shut down so as to shed light on the role of household demand on the optimal policy. All results are for the model with sticky wages and sticky prices, and for brevity's sake the impulse response functions are omitted.[8]

3.5.1. *No Household Demand.* Compared to the baseline model, the optimal policy results are qualitatively similar when household demand for oil is abstracted from. As before, the effect of both shocks on the marginal product of labor is in the opposite direction. The oil supply shock drives up the relative price of oil, essentially reducing the demand for labor by firms. The productivity shock, on the other hand, increases the demand for labor. Consequently, in response to the supply shock core inflation rises and wage inflation falls while a demand shock causes the opposite to occur.

3.5.2. *No Firm Demand.* The results for the productivity shock are the same with or without firm demand for oil, as it increases the marginal product of labor. Therefore core inflation falls and wage inflation rises.

An oil supply shock now impacts the marginal product of labor entirely through the decisions of the household, and these are driven by whether or not oil and the non-oil consumption good are complements or substitutes. It can be shown that the two goods are Edgeworth complements when $\tau > \sigma_c$ and substitutes when $\tau < \sigma_c$. The baseline calibration therefore leads them to be complements, as τ is equal to 1 and σ_c equal to 0.25.

For the baseline calibration, an oil supply shock leads to the opposite responses in core inflation and wage inflation compared to the model with both household and firm demand. The supply shock forces consumption of the oil good to fall. Since the two goods are complements, non-oil consumption must also fall. In equilibrium, this implies a fall in the production of the final good and, consequently, a fall in the equilibrium quantity of labor.

---

[8]The impulse response functions are available upon request. Results for the model with flexible wages or for other experiments are also available upon request.

The optimal way to do this is to push up the real wage, which can be accomplished through a drop in core inflation. In an alternative calibration where τ is 1 and σ_c is 1.25, so that the two consumptions goods are substitutes, the results are the opposite with core inflation rising and nominal wage inflation falling.

4. EMPIRICAL RESULTS

The theoretical model makes predictions about the responses of several policy relevant variables, including the nominal interest rate and the rate of inflation. To summarize, when wages are flexible core inflation is stabilized at its target value and the nominal interest rate falls in response to an exogenous supply shock but rises in response to a demand shock. With sticky wages, core inflation rises in response to an exogenous supply shock but falls in response to the productivity shock.

An interesting question to ask is to what extent these predictions might be found in the data. This has been considered, indirectly, in two recent and important papers. Kilian (2009) introduced an empirical model that allows one to identify an exogenous oil supply shock, a global demand shock (essentially a demand shock for all commodities driven by global economic activity), and an oil specific-demand shock. An extension of this model, which incorporated the federal funds rate as a fourth variable, was used in Kilian and Lewis (2011). There are many interesting findings in both publications but here I focus on those results specifically related to the federal funds rate and the rate of inflation.

The findings of these papers was that the federal funds rate responds differently depending upon the underlying shock driving the price of oil. The federal funds rate has a tendency to fall when a supply shock hits, while a demand shock driven by global economic activity or by some oil specific-demand shock causes the federal funds rate to rise. It was also shown that the CPI was more affected by the demand shocks, because they appeared to have more persistent affects on the price of oil than the supply shocks.

Several interesting questions, however, remain unaddressed. First, there are no results for how core inflation responds to the different shocks. This is important because, in theory, it is the behavior of core inflation that is relevant for discussing optimal policy. Second, the empirical results do not, by themselves, provide an explanation for the different responses seen in the federal funds rate. The theoretical model might provide a coherent story about the empirical results. Finally, there is a question about whether or not the Federal Reserve responds differently if the demand shock originates from unexpectedly strong economic activity in the U.S. as opposed to a globally driven demand shock.

17

To address these issues, I modify the empirical model first developed in Kilian (2009), and later used in Kilian and Lewis (2011). In what follows, the main points of the original VAR are summarized, followed by a discussion of the modifications made to the model and the reasons for the modifications. The impulse response functions from this model are then examined and discussed in light of the results from the theoretical model.

4.1. **The Original Model.** The VAR used in Kilian (2009) is written in the following form,

$$A_0 z_t = \alpha + \sum_{i=1}^{24} A_i z_{t-i} + \epsilon_t. \tag{22}$$

The data is monthly, from 1973:1 to 2007:12, and given by $z_t = (\Delta prod_t, rea_t, rpo_t)'$ where $\Delta prod_t$ is the monthly percentage change in world oil production, rea_t is a measure of global economic activity constructed by Kilian, and rpo_t is the refiner's acquisition cost of oil deflated by the CPI. The data for world oil production comes from the Department of Energy and measures production in thousands of barrels per day. In Kilian and Lewis (2011) this model was extended by adding the federal funds rate (differenced) as a fourth variable in the system.

The three structural shocks in the vector ϵ_t are defined, in order, as

$$\epsilon_t^{\text{oil supply shock}},$$
$$\epsilon_t^{\text{aggregate demand shock}},$$
$$\epsilon_t^{\text{oil specific-demand shock}}.$$

A Cholesky decomposition is used to back out the structural shocks using the residuals. The variables are ordered as listed in z_t. This implies that world oil production responds to both demand shocks with a one-month lag, and that world economic activity responds to oil specific-demand shocks with a lag.

The assumption that the oil supply responds with a lag is grounded in the fact that changing production levels in the oil industry is very difficult to do within a one month window. Kilian (2009) argues that the assumption that economic activity responds with a lag to the oil-specific demand shock is justified due the sluggish response of economic activity seen in the data to changes in oil prices, in general. Note that this ordering also implies that oil prices can respond in one months time to the other shocks, a quite reasonable assumption.

4.2. **The Modified Model.** Any attempt to match up the predictions of the theoretical model with the data needs to take into account that the global demand shock in the original empirical model does not directly map into the demand shock in theoretical model. Taking this into account is important for at least three reasons. First, a demand shock originating

18

in the United States need not always be associated with a demand shock driven by global economic activity. Second, it would be a strong assumption to claim a priori that the Federal Reserve responds equivalently to those two demand shocks. Finally, any demand shock originating in the U.S. which, for whatever reason, is orthogonal to the global demand shock would, in the original empirical model, get shuffled into the oil-specific demand shock. My solution is to modify the model in several ways which should allow for a better match between the empirical model and the theoretical model.

Let $z_t = (\Delta prod_t, rea_t, \Delta lrgdp_t, \Delta rpo_t, \Delta FF_t, \pi_t^{core})'$, where $\Delta prod_t$ is the monthly percentage change in world oil production, rea_t is Kilian's measure of global economic activity, $\Delta lrgdp_t$ is the log of Stock and Watson's monthly real GDP for the U.S., first differenced, Δrpo_t is the log of the refiner acquisition cost of oil deflated by the core CPI, first differenced, ΔFF_t is the monthly change in the effective Federal Funds rate, and π_t^{core} is the inflation rate of the CPI excluding energy prices.[9]

The model is estimated with monthly data spanning from 1987:1 to 2008:6. This sample is chosen for two reasons. First, there was a change in US monetary policy in the early 1980s. Second, there were also important changes in the oil market in the mid-1980s, with one example being the collapse of OPEC. Starting in 1987 avoids having to deal with the potential problems that might occur if there are significant breaks in the data generating process due to one or both of those issues. Data from the recent crisis is excluded given the atypical monetary policy that has been in place since then.

The form of the VAR is given by

$$A_0 z_t = \alpha + \sum_{i=1}^{5} A_i z_{t-i} + \epsilon_t. \tag{23}$$

Only 5 lags are included instead of the original 24. This is done for two reasons. First, the sample is smaller while the number of variables included larger, so there are fewer degrees of freedom available. Second, the LR test and the AIC test picked an optimal lag length of 5 lags while the Schwartz criterion chose 1 lag as optimal. In the interest of capturing a richer set of dynamics, 5 lags were chosen. Data from the end of 1986 is used in the estimation process so useable observations run from 1987:1 to 2008:6, leading to a total of 258 observations.

---

[9]The model was also estimated using Stock and Watson's measure of monthly real GDI, logged and first differenced, instead of real GDP. The responses were similar in both cases. As an additional check the model was also estimated using the monthly GDP series from Macroeconomic Advisers, available starting in 1992:04. The error bands were wider in this case due to the smaller sample, but qualitatively the responses of core inflation and the federal funds rate remained similar.

With the changes in z_t, the structural shocks are now defined as

$$\epsilon_t^{\text{oil supply shock}},$$
$$\epsilon_t^{\text{global demand shock}},$$
$$\epsilon_t^{\text{U.S. specific-demand shock}},$$
$$\epsilon_t^{\text{oil specific-demand shock}},$$
$$\epsilon_t^{\text{fed funds shock}},$$
$$\epsilon_t^{\text{shock to core inflation}}.$$

A Cholesky decomposition, where the variables are ordered in the same way as they are found in z_t, is used to identify the structural shocks.

The justification for the ordering of the first two variables follows from the arguments given in Kilian (2009). Ordering U.S. real GDP after global economic activity is done for two reasons. First, this allows economic activity in the U.S. to endogenously respond to shocks that affect the world economy, which by definition should include the U.S. It also ensures that the structural shock to U.S. real GDP is due to unexpectedly higher activity originating in the U.S., and not simply a spill over from world economic activity. Second, this ordering ensures that a demand shock due to U.S. economic activity does not get shuffled into the oil specific-demand shock.

With oil prices ordered after the economic activity variables, the shock to the price of oil can again be interpreted as an oil specific-demand shock, as in Kilian (2009). The one difference is that U.S. specific-demand shocks are specifically parceled out of this. The two other shocks in the model are unimportant for the results of this paper and not discussed further. The results presented are insensitive to the ordering of the federal funds rate or core inflation.

4.3. **Results.** The impulse response functions for the oil supply shock and U.S. demand shock are shown in figures 6 - 9.[10] The responses are plotted over a 24 month period. In all of the graphs the solid line is the point estimate, the dashed lines the one-standard error bands, and the dots are the two-standard error bands. A response is defined as marginally significant if 0 is outside of the one-standard error band and significant if 0 is outside of the two-standard error band.

Figure 6 shows the responses due to a one-standard deviation shock to the supply of oil. The unexpected decrease in the supply of oil causes several months of rising oil prices. Core inflation initially rises, on an annualized basis, by about 10 basis points. This movement,

---

[10]Results for the other shocks are available upon request.

however, is just marginally significant and all other other movements after it are not statistically different from 0. While the federal funds rate declines, this movement is only significant starting about 5 months after the shock.

Some of the responses are choppy and the month to month movements can mask the longer run implications of the responses. Figure 7, therefore, plots the cumulative responses of the variables to the supply shock. From this viewpoint, there is a marginally significant reduction in economic activity in the U.S. due to higher oil prices, similar to the findings in Kilian (2009). The cumulative impact on core inflation is essentially 0. In figure 6, the federal funds rate declined over a number of months. The cumulative effect of these declines is a reduction in the rate that is marginally significant after 5 months, and very close to being significant over the course of two years.

Figure 8 plots the responses to a positive shock to U.S. real GDP growth. Unexpectedly strong growth in the U.S. brings about several rounds of increases in the real price of oil, which are marginally significant in the first two months. While the point estimate of core inflation shows some variation, none of the movements are statistically significant at any horizon. The response of the federal funds rate to this shock is very different from the exogenous oil supply shock. Instead of falling, there are a series of small increases over time, some of which are statistically significant.

As with the supply shock, looking at the cumulative responses gives a better view of what happens over time. Two results are statistically significant: the increases in the federal funds rate and the higher growth in real GDP. The cumulative impact on the price of oil is upwards, and this is marginally significant. The response of core inflation, while positive according to the point estimate, is not statistically different from 0 at any point in time.

Quantitatively, the impulse response functions for core inflation are similar but there are very different responses in the federal funds rate. Can the theoretical model provide a reasonable story for these results? In the case of the demand shock, the answer is very much yes. In response to the demand shock, the federal funds rate rises over time while core inflation shows no statistically significant movements. This is exactly what the model predicts should happen when monetary policy stabilizes core inflation in response to this type of shock.[11]

Analyzing the supply shock is slightly more difficult. There is an initial rise in core inflation, although this is barely significant. This is the response that would be optimal if one believes wages need to be adjusted downwards through higher core inflation. But,

---

[11]The responses in figure 1 are from the model with flexible wages. The responses from the model with sticky wages are qualitatively quite similar when core inflation is stabilized. The nominal rate rises in response to the demand shock but is lowered in response to the supply shock.

all movements in core inflation after the initial rise are not significantly different from 0. Furthermore, the federal funds rate is, cumulatively, lower over the first two years. The story behind these responses would be generally consistent with what the theoretical model says monetary policy should do to stabilize core inflation.

Taken together then, after 1986 the Federal Reserve seems to have placed a strong emphasis on stabilizing core inflation in response to both oil supply shocks and a demand shock that is driven by unexpectedly strong economic activity in the U.S. Movements in the federal funds rate show that the Federal Reserve has responded differently in response to the two different shocks. This is despite the fact that both drive up the price of oil, and therefore might naively be lumped together as an oil price shock.

There is one final issue regarding the empirical results worth addressing. Given the large degree of wage and price stickiness often found at an aggregate level, stabilizing core inflation in response to these two shocks is technically sub-optimal. Is this something to be concerned about? The results from figure 5 show that this policy is fairly innocuous in welfare terms. Even with sticky wages and prices, a policy rule that stabilized core inflation would produce trivial losses for values of α_π that have been found in the literature for U.S. monetary policy after 1986.

5. Conclusions

This paper has examined optimal monetary policy in a New Keynesian model where the relative price of oil is driven by both an exogenous supply shock and a productivity-driven demand shock. When wages are flexible, the optimal policy keeps core inflation on target regardless of which shock drives oil prices. The nominal interest rate falls in response to the supply shock and rises in response to the demand shock.

When there is a tradeoff between stabilizing inflation and output, there are important qualitative differences in the response of core inflation. This variable initially rises in response to the supply shock but falls in response to the demand shock. This occurs because the supply shock reduces the marginal product of labor while the productivity shock, while raising the price of oil, on net raises the marginal product of labor. The optimal policy, since it uses variations in core inflation to affect the real wage, consequently generates higher inflation in response to the supply shock but lower inflation in response to the demand shock.

These predictions are taken to the data using a VAR that that can identify exogenous oil supply shocks and a demand shock driven by unexpectedly strong economic activity in the U.S. The impulse response functions show that the federal funds rate responds quite different in response to the two different shocks. The federal funds rate is lowered in response to the

supply shock while in response to the demand shock it is increased. In response to both shocks core inflation responds minimally. There is some statistical evidence that it rises by a small amount over the first month or two when there is an oil supply shock. But there is no statistically significant movement in core inflation in response to the demand shock. The behavior of the federal funds rate and core inflation is close to what the theoretical model predicts should happen if the Federal Reserve is attempting to stabilize core inflation in response to the two shocks.

The findings of this paper provide further evidence of the importance of distinguishing between the source of the shock that affects the price of oil. With this in mind, one obvious shortcoming of the present paper is that the model is essentially a closed economy model. As such, it is useful for thinking about a certain sub-set of the shocks that could jointly affect the U.S. economy and the price of oil. Clearly a useful avenue for future research would be to construct a multi-country model where the implications of global demand shocks and demand shocks originating in other foreign countries on U.S. monetary policy could be explored. Given the size of this undertaking, though, this is left for future research.

References

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

Blanchard, O., and J. Gali (2010): "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so Different from the 1970s?," in *International Dimensions of Monetary Policy*, pp. 373–428. University of Chicago Press.

Bodenstein, M., C. Erceg, and L. Guerrieri (2008): "Optimal Monetary Policy with Distinct Core and Headline Inflation Rates," *Journal of Monetary Economics*, 55, 518–533.

Calvo, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

Dhawan, R., and K. Jeske (2007): "Taylor Rules with Headline Inflation: A Bad Idea!," working paper 2008-11, Federal Reserve Bank of Atlanta.

Finn, M. G. (2000): "Perfect competition and the effects of energy price increases on economic activity," *Journal of Money, Credit, and Banking*, 32, 400–416.

Kilian, L. (2009): "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," *American Economic Review*, pp. 1053–1069.

Kilian, L., and L. T. Lewis (2011): "Does the Fed Respond to Oil Price Shocks?," *The Economic Journal*, pp. 1047–1072.

Kim, I.-M., and P. Loungani (1992): "The role of energy in real business cycle models," *Journal of Monetary Economics*, 29, 173–189.

Leduc, S., and K. Sill (2004): "A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns," *Journal of Monetary Economics*, 51, 781–808.

Nakov, A., and A. Pescatori (2010): "Monetary Policy Trade-Offs with a Dominant Oil Producer," *Journal of Money, Credit, and Banking*, 42.

Plante, M. (2009): "How Should Monetary Policy Respond to Exogenous Changes in the Relative Price of Oil?," CAEPR Working Paper.

Rotemberg, J. J., and M. Woodford (1996): "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit, and Banking*, 28, 549–577.

Schmitt-Grohe, S., and M. Uribe (2004): "Solving dynamic general equilibrium models using a second-order approximation to the policy function," *Journal of Economic Dynamics and Control*, 28, 755–775.

——— (2006): "Optimal Inflation Stabilization in a Medium-Scale Macroeconomic Model," Working Paper.

Woodford, M. (2003): *Interest and Prices*. Princeton University Press, New Jersey, USA.
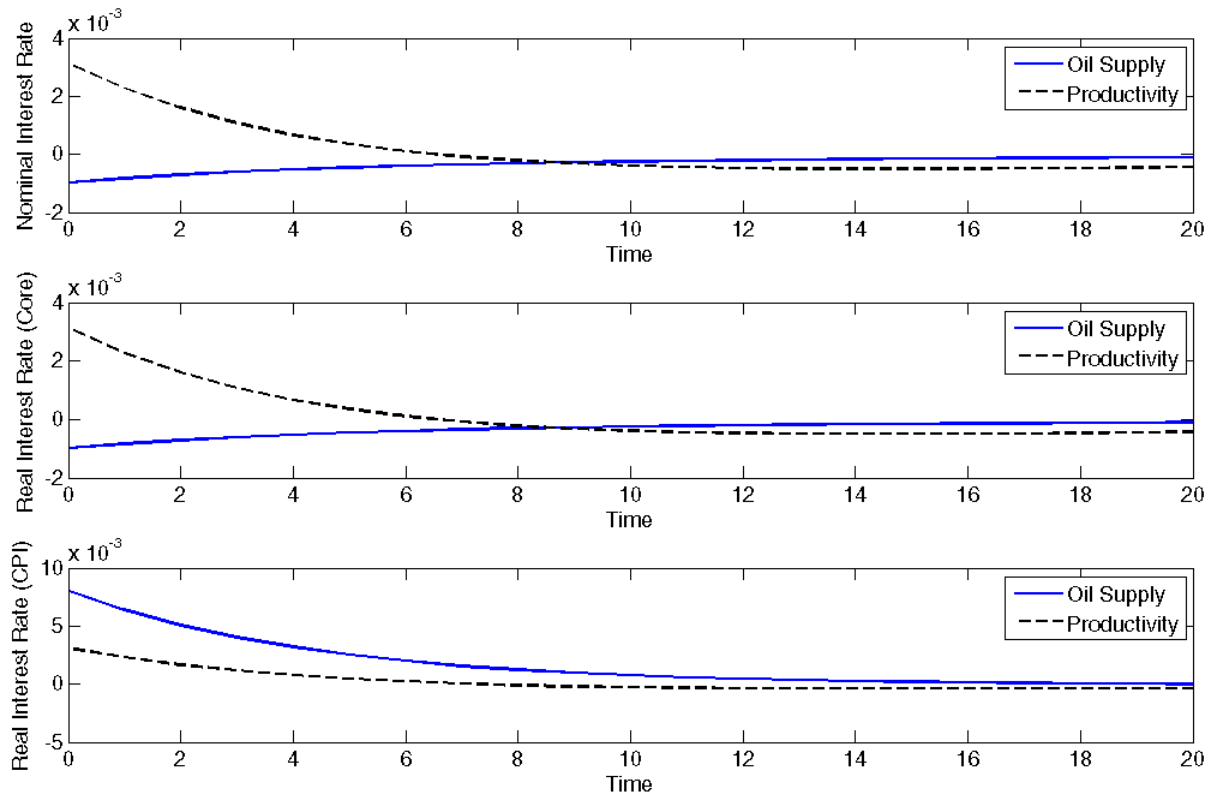
FIGURE 1. Optimal responses of nominal and real interest rates (flexible wages)
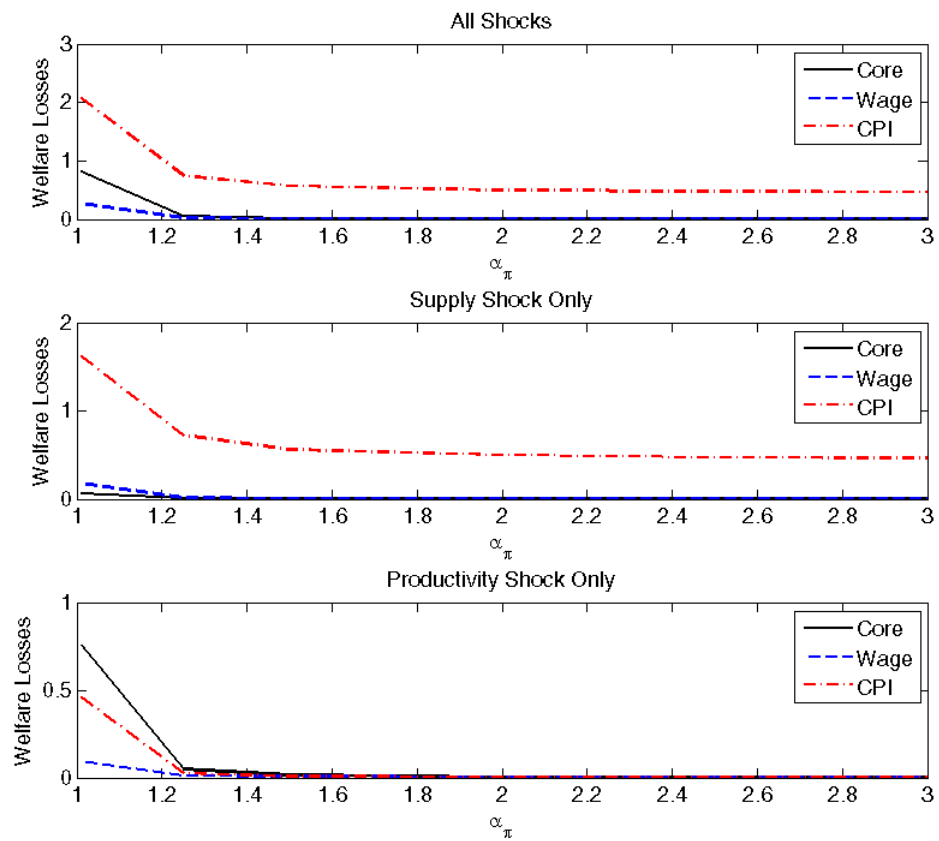
FIGURE 2. Welfare losses for different rules under flexible wages: Both shocks active (top), supply shock active (middle), productivity shock active (bottom).
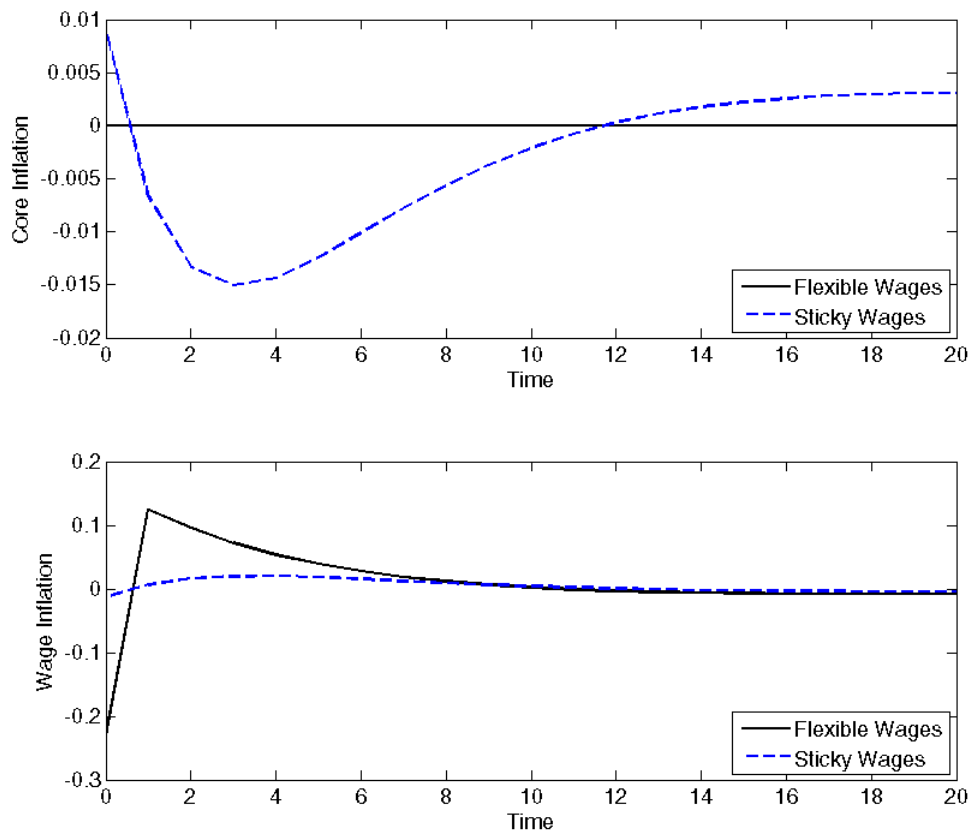
FIGURE 3. Optimal responses to a supply shock under flexible and sticky wages: core inflation (top), wage inflation (bottom)
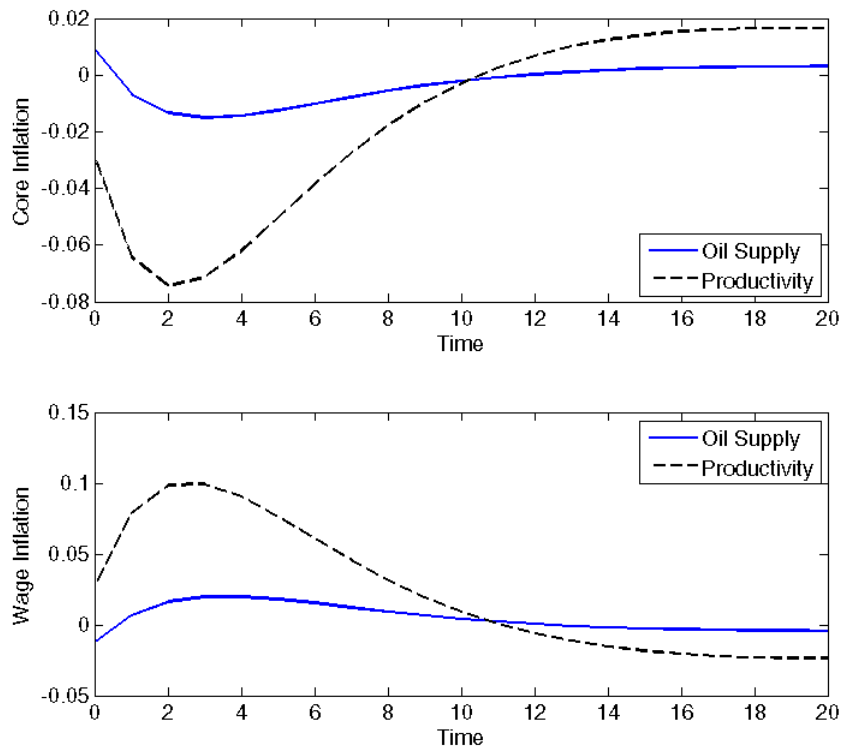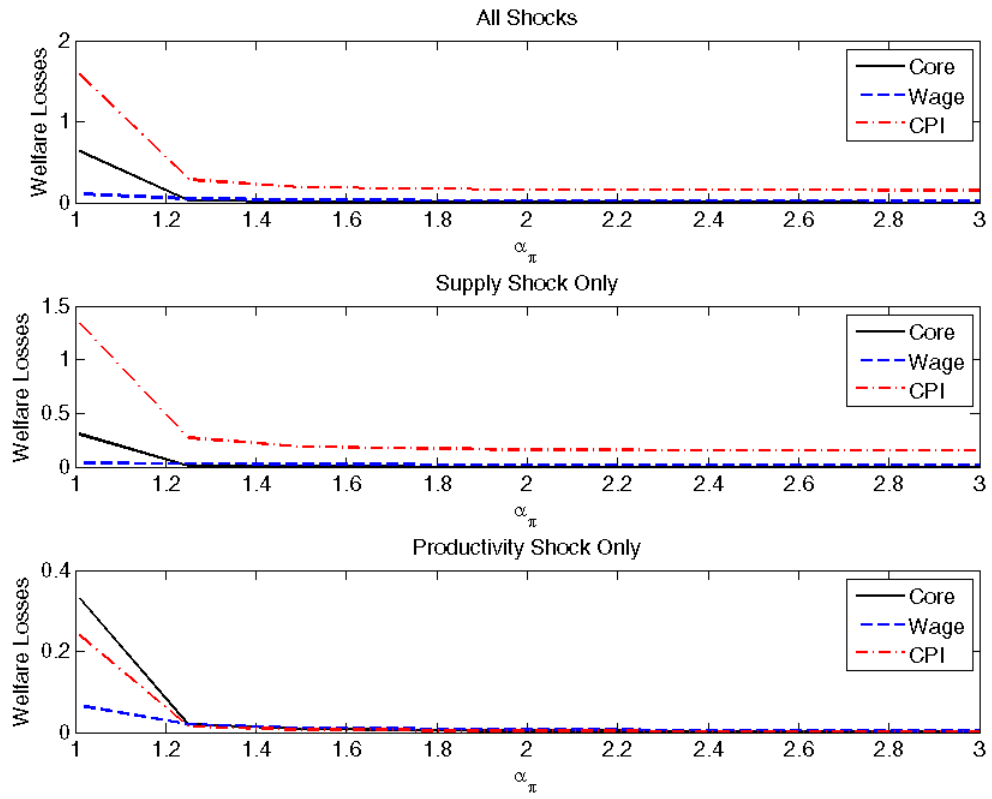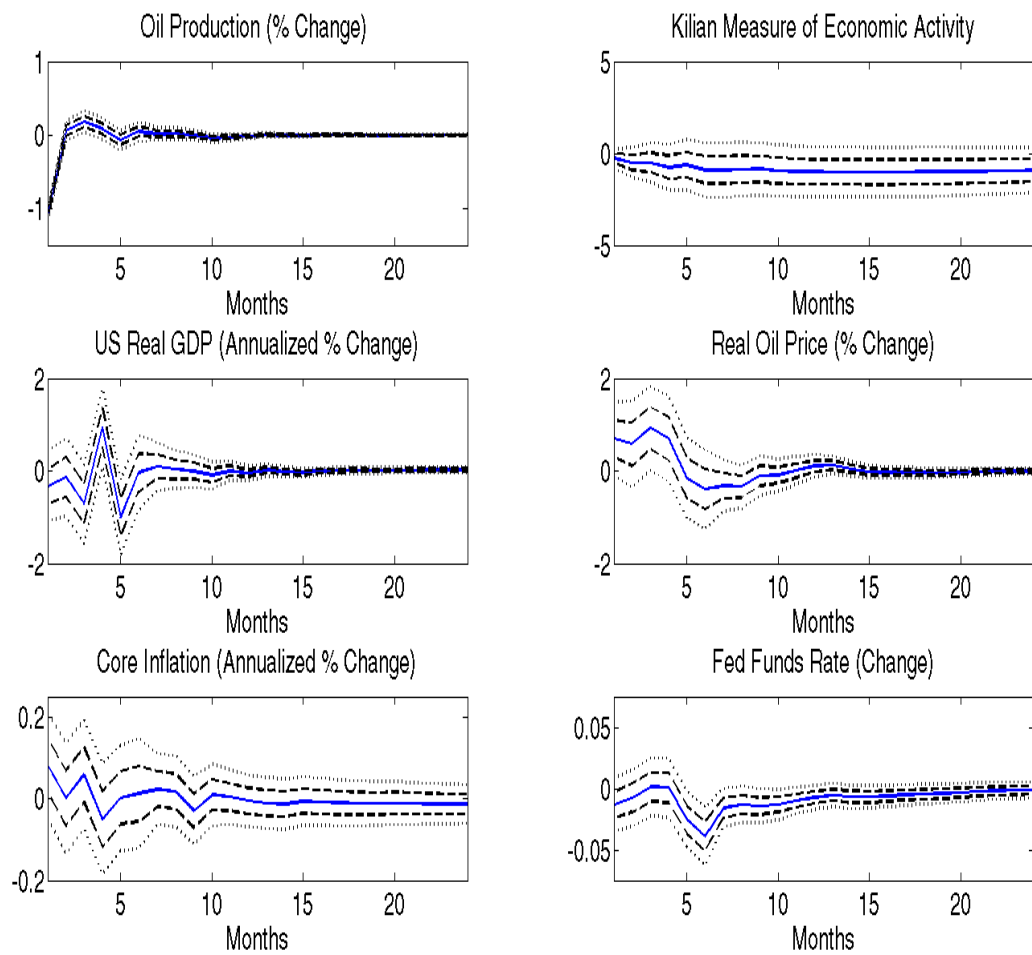
FIGURE 4. Optimal responses of core inflation and wage inflation under sticky wages: oil supply shock (top), productivity shock (bottom)

FIGURE 5. Welfare losses for different rules under sticky wages: Both shocks active (top), supply shock active (middle), productivity shock active (bottom).

FIGURE 6. Responses to oil supply shock, error bands in dashed black (1 s.e.) and dots (2 s.e)
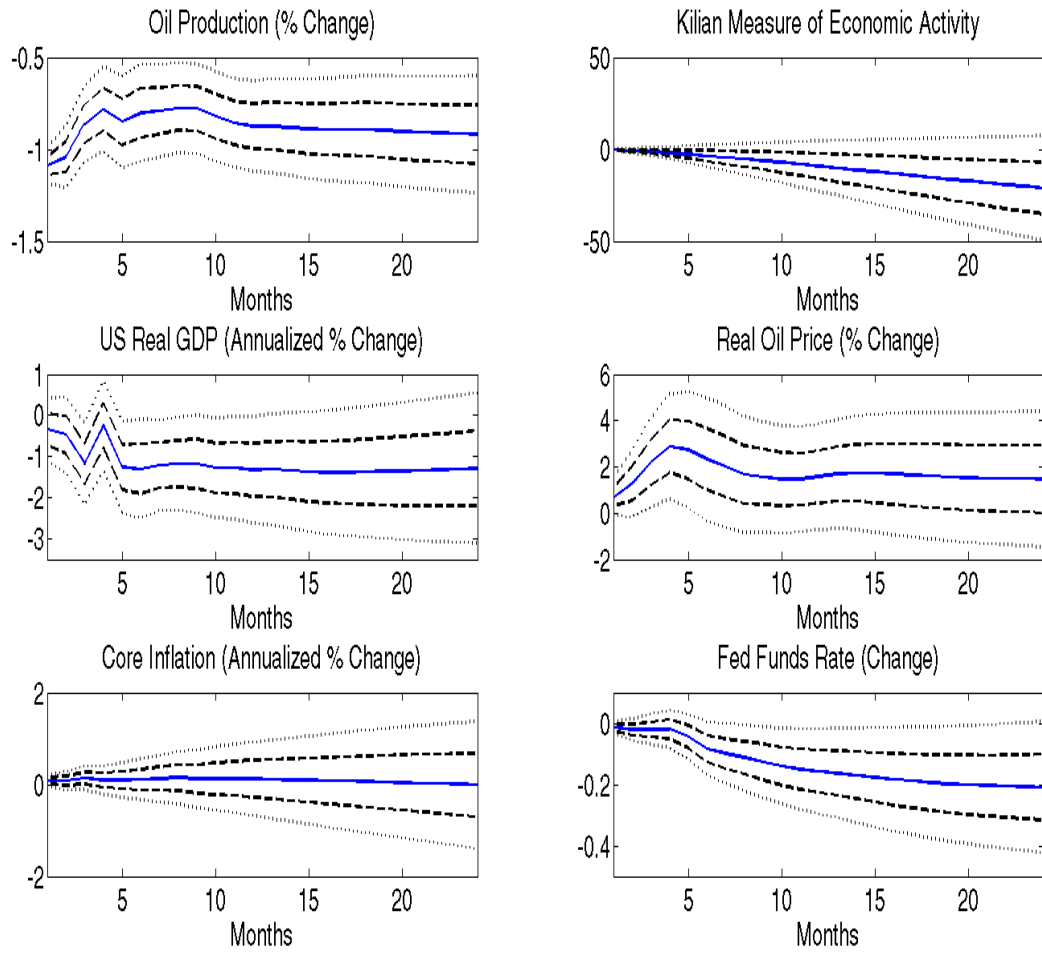
FIGURE 7. Cumulative responses to negative oil supply shock, error bands in dashed black (1 s.e.) and dots (2 s.e)
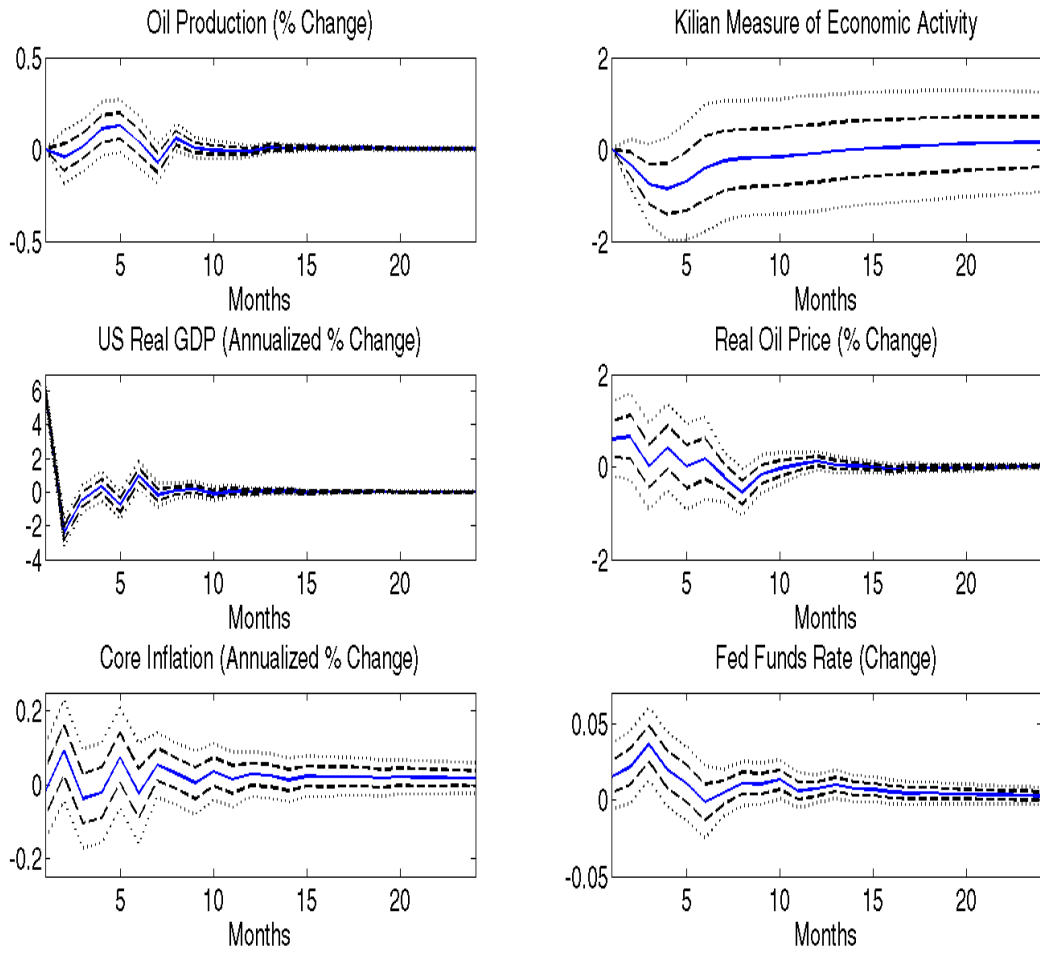
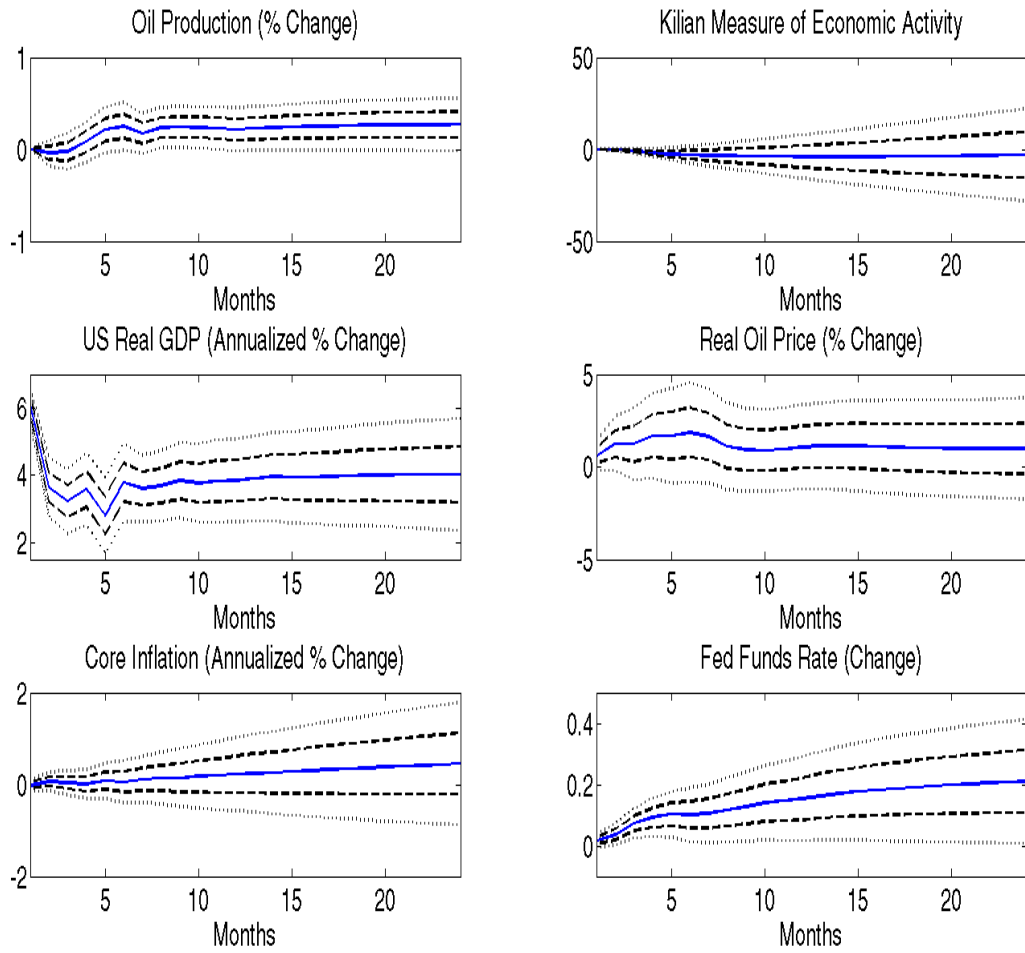FIGURE 8. Responses to US real GDP growth rate shock, error bands in dashed black (1 s.e.) and dots (2 s.e)

FIGURE 9. Cumulative responses to US real GDP growth rate shock, error bands in dashed black (1 s.e.) and dots (2 s.e)

TABLE 1. Baseline Calibration

Parameter	Description	Value
β	Discount Factor	.993
δ_k	Capital Deprecation Rate	.01785
ω	Degree of Price Stickiness	.75
ω_n	Degree of Wage Stickiness	.75
Π	Steady State Inflation	1
Y^g	Steady State Real GDP	1
X	Aggregate Consumption to GDP Ratio	.80
O^h	Household Oil to GDP Ratio	.05
O^f	Firm Oil to GDP Ratio	.02
I^k	Investment Spending to GDP Ratio	.20
θ	Elasticity of Substitution (Intermediate Goods)	6
θ_n	Elasticity of Substitution (Labor Types)	6
μ	Wage Elasticity of Labor Supply	1
ν	Elasticity of Substitution between C and O^h	.25
η	Elasticity of Substitution in Production	.25
τ	Intertemporal Elasticity of Substitution	1
ρ_o	AR(1) Coefficient of Supply Shock	.80
ρ_z	AR(1) Coefficient of Productivity Shock	.80
σ_o	Standard Deviation of Supply Shock	.003545
σ_z	Standard Deviation of Productivity Shock	.0175