Monetary Policy Responses to Oil Price Fluctuations

MARTIN BODENSTEIN, LUCA GUERRIERI, and LUTZ KILIAN*

The paper provides the first quantitative analysis of how U.S. monetary policy responses should differ depending on the source of the observed oil price fluctuations. It presents three main sets of results. First, the paper proposes a novel decomposition of the marginal cost of production that highlights the role of each factor input for the evolution of inflation. Second, conditional on an estimated interest rate policy reaction function, the paper demonstrates that no two structural shocks induce the same monetary policy response, even after controlling for the impact response of the real price of oil, and quantifies these differences. Third, the paper shows that the policy responses implied by a policy rule, whose coefficients were chosen to maximize U.S. welfare, differ substantially from the policy response implied by the same rule estimated on historical data. Among a wide range of rules, a rule that is easily implementable and that nearly maximizes U.S. welfare involves the Federal Reserve putting zero weight on the price of oil and responding to wage inflation without interest rate smoothing. [JEL E31, E43, F41, Q43]

IMF Economic Review (2012) **60**, 470–504. doi:10.1057/imfer.2012.19

*Martin Bodenstein is an economist in the Economic Research and Development Department at the Asian Development Bank in Manila. Luca Guerrieri is a senior economist at the Federal Reserve Board in Washington D.C. Lutz Kilian, Professor of Economics, received his Ph.D. in Economics from the University of Pennsylvania in 1996 and his M.A. in Development Banking from The American University in 1988. The authors thank the editors, two anonymous referees, and Giovanni Lombardo, our discussant, for helpful comments. They benefited from conversations with Matt Canzoneri, Chris Erceg, Dale Henderson, and Sylvain Leduc. 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.

The recent volatility in global commodity prices and in the price of oil, in particular, has created renewed interest in the question of how monetary policymakers should respond to oil price fluctuations. We provide the first quantitative analysis of how U.S. monetary policy responses should differ depending on the source of the observed oil price fluctuations. We examine in detail how the structural shocks underlying fluctuations in the real price of oil are transmitted to U.S. real activity and inflation. We also explore how the monetary policy response to these shocks may be optimized within the context of commonly studied instrument rules. An important feature of our analysis is that we rely on a global dynamic stochastic general equilibrium (DSGE) model with endogenous real oil prices and real exchange rates.

We present three main sets of results. First, we provide a novel decomposition of the marginal cost of production that highlights the role of each factor input for the evolution of inflation. We show that, in a New-Keynesian setting with price and wage rigidities, the labor market plays a central role in the adjustments following changes in the price of oil. Second, conditional on the estimated interest rate policy reaction function, we demonstrate that no two structural shocks induce the same monetary policy response. Even after controlling for the magnitude of the impact effect on the real price of oil, the magnitude, shape, and sometimes even the sign of the policy responses will differ. Our results contradict the popular notion in the literature that an increase in the real price of oil driven by foreign demand is just like an exogenous oil supply shock from the point of view of other oil importers (see, for example, Blanchard and Galí, 2010, p. 384). Third, we construct welfare optimal policy rules within the class of interest rate policy reaction functions and show that the optimal policy responses to a given structural shock differ substantially from the responses implied by the policy rule estimated using historical data. The estimated rule responds to an inflation term and to a lagged interest rate term, but assigns essentially zero weight to the output gap. In contrast, the optimized rule responds to the output gap, but attaches no weight to inflation, or the price of oil, and involves no interest rate smoothing. While a rule that responds to the output gap, as defined in the model, is difficult to implement in practice, we show that an alternative rule that puts high weight on wage inflation, but places virtually no weight on the price of oil, does almost as well in terms of welfare.

Much of the existing analysis on the conduct of monetary policy in the face of oil price fluctuations relies on the counterfactual premise that the real price of crude oil is exogenous with respect to the U.S. economy (see, for example, Leduc and Sill, 2004; Carlstrom and Fuerst, 2006; Dhawan and Jeske, 2007; Plante, 2009a, b; Winkler, 2009; Montoro, 2010; Kormilitsina, 2011; Natal, 2012). Even those DSGE studies that have endogenized the real price of oil have made strong and unrealistic simplifying assumptions about the determination of the price of oil in global markets (see, for example, Backus and Crucini, 1998), have ignored monetary policy (see, for example, Backus and Crucini, 1998; Balke, Brown, and Yücel, 2010; Bodenstein, Erceg, and Guerrieri, 2011; Nakov and Nuño, 2011), or have ignored the open economy

aspect of the transmission of oil price shocks (see, for example, Bodenstein, Erceg, and Guerrieri, 2008; Nakov and Pescatori, 2010a, b). We show that some of the familiar results from closed-economy models carry through to our open economy with endogenous oil prices, while others do not. For instance, the “divine coincidence” between stabilizing the output gap and inflation, applicable to some closed-economy models, does not hold in our model.¹

Kilian, Rebucci, and Spatafora (2009) and Bodenstein, Erceg, and Guerrieri (2011) recently have demonstrated empirically and theoretically that one cannot understand the global effects of oil demand and oil supply shocks without considering their effects on exchange rates, asset prices, oil and non-oil trade balances, and capital accounts. Understanding the relationship between the economy and oil markets requires a multicountry DSGE model. Our analysis builds on the global model recently proposed by Bodenstein and Guerrieri (2011), which includes a net oil-importer (the United States) and a net oil exporter (the rest of the world).

The transmission channels and sources of oil price fluctuations are potentially quite different in closed and in open economy settings. One reason is that the price of oil is determined endogenously in global oil markets. By contrast closed economy models assume exogenous oil price fluctuations. Hence, the optimal response of monetary policy will be different depending on the source of the oil price shock in our setting, even conditioning on the same magnitude of the initial oil price increase. This point was first discussed in Kilian (2009), but little is known about how policymakers’ response should change depending on the nature of the oil price shock. Our model accounts for 15 distinct structural shocks that shift oil demand or oil supply including, for example, shocks to the intensity of oil use at home and abroad, shocks to aggregate productivity at home and abroad, shocks to the global production of crude oil at home and abroad as well as shocks to exogenous spending, markups, monetary policy, consumption preferences, and trade, among others.²

The second reason why the open economy analysis is different is that the effects on real output in the home country are affected by trade channels. Furthermore, under incomplete markets, headline and core inflation are influenced by the different responses of the non-oil terms of trade for oil-importing and oil-exporting countries.

When it comes to domestic goods price inflation, the standard New-Keynesian Phillips curve underlying our model provides a familiar framework for understanding the propagation channels of movements in the real price of oil. We show that domestic goods inflation is related to appropriately weighted gaps between the rental rate and the marginal product of each factor input.

¹In a partial equilibrium setting, Monacelli (2012) shows that open economy considerations break the “divine coincidence.”

²We do not consider speculative oil demand shocks, as discussed in Alquist and Kilian (2010), Fattouh and others, 2012 and Kilian and Murphy (2010) because there is no empirical evidence that speculation mattered for the fluctuations in the real price of oil between 2003 and 2010 and because including speculation would considerably complicate the model.

When nominal rigidities are absent, these gaps never open up. But with nominal price rigidities, even abstracting from sticky wages, these gaps can be sizable. We show that the gap between the real wage and the marginal product of labor is a key contributor to the rise in marginal cost and domestic price inflation in the wake of changes in the price of oil.

Bodenstein and Guerrieri (2011) used the same model as in this paper to quantify the effects of oil demand and oil supply shocks on U.S. macroeconomic aggregates and to explain the evolution of the real price of oil. In contrast, our focus is on the response of monetary policy to these shocks. We quantify the interest rate policy responses under two scenarios. In the first scenario, we employ a policy rule with coefficients fixed at their estimated values. In the second scenario, we choose the coefficients of the U.S. policy rule to maximize domestic welfare, taking the policy rule in the rest of the world as given. By terminating the estimation sample in the third quarter of 2008, our analysis deliberately abstracts from the presence of a zero lower bound following the financial crisis of 2008 and from quantitative easing policies. Rather our focus is on characterizing the appropriate policy responses in the presence of shocks that shift the demand for oil or the supply of oil during normal times. The results are intended to provide a first benchmark for policy discussions. Further refinements of the global DSGE model and a discussion of the role of international policy coordination are left for future research.³

The remainder of this paper is organized as follows. In Section I, we review the literature on the relationship between oil prices and monetary policy. Section II outlines the DSGE model on which our analysis is based. Further details on this model can be found in the not-for-publication appendix. In Section III, we illustrate how optimal monetary policy responses to oil price fluctuations depend on the source of the oil price fluctuations. The concluding remarks are in Section IV.

I. Oil Prices and Monetary Policy

The literature on the relationship between the real price of oil and monetary policy dates back to the 1980s. There is a consensus that causality in this relationship may run from events in oil markets to monetary policy as well as from shifts in monetary policy to the supply of oil and the demand for oil in global markets. Barsky and Kilian (2002), for example, discuss in the context of the experience of the 1970s and early 1980s how an exogenous shift in the global monetary policy regime may cause a shift in the demand for crude oil and hence in the real price of crude oil. Kilian (2010) and Erceg, Guerrieri, and Kamin (2011) explain why this explanation does not fit the more recent data, and indeed much of the literature since the 1990s has focused on the

³Bodenstein, Guerrieri, and Gust (2010) discuss the implications of the zero lower bound on nominal interest rates for the effects of exogenous oil price shocks.

reverse direction of causality from oil prices to monetary policy. Notably, Bernanke, Gertler, and Watson (1997) in an influential empirical paper attributed the severity of the 1974 and 1982 recessions to the Federal Reserve's direct response to the preceding oil price shocks. Recent research has cast doubt on their empirical analysis and on the theoretical premise of their analysis (see, for example, Kilian and Lewis (2011) and the references therein). There is no compelling evidence that the Federal Reserve was responding mechanically to oil price shocks beyond the response to the inflation and real output fluctuations associated with such shocks.

Bernanke, Gertler, and Watson's (1997) empirical work has stimulated a large DSGE model literature on the normative question of how monetary policymakers should respond to oil price shocks. Much of this optimal monetary policy literature has focused on models in which the policymaker follows a conventional interest rate policy rule within a closed economy (see, for example, Leduc and Sill, 2004; Carlstrom and Fuerst, 2006). One of the key questions in this literature has been what inflation measure the central bank should focus on. Another key question has been the existence, or not, of a trade-off between stabilizing inflation and the welfare relevant output gap. For example, Bodenstein, Erceg, and Guerrieri (2008) and Natal (2012) largely agree that dual mandate instrument rules based on core inflation measures come close to replicating welfare maximizing policies, as long as they are not overly aggressive in stabilizing core inflation. In related research, Plante (2009a) finds that optimal monetary policy should stabilize a weighted average of core and nominal wage inflation. Winkler (2009) considers anticipated and unanticipated (deterministic) oil price shocks and also finds that optimal policy cannot stabilize at the same time prices, wages, and the welfare-relevant output gap; indeed, following an oil price shock, optimal policy requires a larger output drop than under a traditional Taylor rule.

II. Model Description

The model used in this paper is borrowed from Bodenstein and Guerrieri (2011). As in Backus and Crucini (1998), the model encompasses international trade in oil and non-oil goods. In addition, it incorporates the nominal and real rigidities that Christiano, Eichenbaum, and Evans (2005) and Smets and Wouters (2007) have found to be empirically relevant in closed economy models.

Here, we sketch the key features of the model, and the determinants of oil demand and supply in particular.⁴ There are two countries: the home country (country 1) and the foreign country (country 2). We estimate the model using U.S. data for the home country and aggregate data for the principal trading partners of the United States for the foreign bloc. Because the structure of the country blocs is symmetric, we focus on the home country in describing the

⁴An online appendix gives a more detailed description of the model.

model. Country-specific values for the parameters allow for differences in population size, oil shares in production and consumption, oil endowments, expenditure shares, and in non-oil and oil trade flows. Although asset markets are complete at the country level, asset markets are incomplete internationally. The assumption of incomplete asset markets across countries is a key ingredient in generating country-specific wealth effects in response to shocks that affect the real price of oil.

In each country, a continuum of firms produces differentiated varieties of an intermediate good under monopolistic competition. Each firm utilizes capital, labor, and oil and acts in perfectly competitive factor markets. The production technology is characterized by a nested constant-elasticity of substitution specification. Since capital is owned by households and rented out to firms, the cost minimization problem of firm i that intends to produce overall output $Y_{1,t}(i)$ can be written as:

$$\min_{K_{1,t}(i), L_{1,t}(i), O_{1,t}^y(i), V_{1,t}(i)} R_{1,t}^k K_{1,t}(i) + W_{1,t} L_{1,t}(i) + P_{1,t}^o O_{1,t}^y(i) \quad (1)$$

s.t.

$$Y_{1,t}(i) = \left((\omega_1^{yy})^{\frac{\rho_1^o}{1+\rho_1^o}} (V_{1,t}(i))^{\frac{1}{1+\rho_1^o}} + (\omega_1^{oy})^{\frac{\rho_1^o}{1+\rho_1^o}} (\mu_{zo}^t Z_{1,t}^o O_{1,t}^y(i))^{\frac{1}{1+\rho_1^o}} \right)^{1+\rho_1^o} \quad (2)$$

$$V_{1,t}(i) = \left((\omega_1^k)^{\frac{\rho_1^y}{1+\rho_1^y}} (K_{1,t}(i))^{\frac{1}{1+\rho_1^y}} + (\omega_1^l)^{\frac{\rho_1^y}{1+\rho_1^y}} (\mu_z^t Z_{1,t} L_{1,t}(i))^{\frac{1}{1+\rho_1^y}} \right)^{1+\rho_1^y}. \quad (3)$$

Utilizing capital $K_{1,t}(i)$ and labor services $L_{1,t}(i)$, the firm produces a “value-added” input $V_{1,t}(i)$, which is then combined with oil $O_{1,t}^y(i)$ to produce variety i of the domestic non-oil good, $Y_{1,t}(i)$. The rental rates of capital, labor, and oil are, respectively: $R_{1,t}^k$, $W_{1,t}$, and $P_{1,t}^o$.

The quasi-share parameter ω_1^{oy} determines the importance of oil purchases in the output of firms, and the parameter ρ_1^o determines the price elasticity of demand for oil. The term $Z_{1,t}$ represents a stochastic process for the evolution of productivity while μ_z denotes constant labor augmenting technological progress. The term $Z_{1,t}^o$ represents a stochastic process that influences the oil intensity of production, while the term μ_{zo}^t can capture a secular decline in oil intensity.

Goods prices are determined by Calvo-Yun staggered contracts. Trade occurs at the level of intermediate goods. Within each country the varieties are aggregated into a (non-oil) consumption and an investment good. Households consume oil, the non-oil consumption good, save and invest, and supply differentiated labor services under monopolistic competition. Wages are determined by Calvo-Yun staggered contracts.

The consumption basket $C_{1,t}$ that enters the households’ utility is produced by perfectly competitive consumption distributors whose production

function mirrors the preferences of households over home and foreign non-oil goods and oil.⁵ The cost minimization problem of a representative distributor that produces the consumption good $C_{1,t}$ can be written as:

$$\min_{\substack{C_{1,t}^d, M_{1,t}^c \\ C_{1,t}^{ne}, O_{1,t}^c}} P_{1,t}^d C_{1,t}^d + P_{1,t}^m M_{1,t}^c + P_{1,t}^o O_{1,t}^c \quad \text{subject to} \\ C_{1,t} = \left((\omega_1^{cc})^{\frac{\rho_1^o}{1+\rho_1^o}} (C_{1,t}^{ne})^{\frac{1}{1+\rho_1^o}} + (\omega_1^{oc})^{\frac{\rho_1^o}{1+\rho_1^o}} (\mu_{zo}^t Z_{1,t}^o O_{1,t}^c)^{\frac{1}{1+\rho_1^o}} \right)^{1+\rho_1^o} \quad (4)$$

$$C_{1,t}^{ne} = \left((\omega_1^{mc})^{\frac{\rho_1^c}{1+\rho_1^c}} (C_{1,t}^d)^{\frac{1}{1+\rho_1^c}} + (\omega_1^{cc})^{\frac{\rho_1^c}{1+\rho_1^c}} (Z_{1,t}^m M_{1,t}^c)^{\frac{1}{1+\rho_1^c}} \right)^{1+\rho_1^c}. \quad (5)$$

The representative distribution firm produces a non-oil aggregate $C_{1,t}^{ne}$ from the home and foreign intermediate consumption aggregates $C_{1,t}^d$ and $M_{1,t}^c$, which is then combined with oil $O_{1,t}^c$ to produce the final consumption good in the home country $C_{1,t}$.

The parameter ω_1^{oc} determines the ratio of oil purchases to the output of the firm. The price elasticity of oil demand ρ_1^o in the consumption aggregate (4) coincides with the one in the production function (2). The same shock $Z_{1,t}^o$ that affects oil intensity in production also affects the oil intensity of consumption. μ_{zo}^t denotes a constant rate of oil efficiency gains. The quasi-share ω_1^{mc} determines the importance of non-oil imports in the non-oil aggregate. The elasticity of substitution between the home and foreign intermediate good is denoted by ρ_1^c . The term $Z_{1,t}^m$ captures a trade shock. In our estimation, this shock accounts for the volatility of non-oil goods trade that is not explained by the remaining shocks.

The distributors sell the consumption aggregate at the price $P_{1,t}^c$ under perfect competition. Thus, $P_{1,t}^c$ coincides with the Lagrange multiplier on Equation (4) in the cost minimization problem of a distributor. The price of the non-oil consumption good $C_{1,t}^{ne}$ is referred to as the “core” price level $P_{1,t}^{ne}$.

Each period the home and foreign countries are endowed with exogenous supplies of oil $Y_{1,t}^o$ and $Y_{2,t}^o$, respectively. The two endowments are governed by distinct stochastic processes. With both domestic and foreign oil supply determined exogenously, the oil price $P_{1,t}^o$ adjusts endogenously to clear the world oil market:

$$Y_{1,t}^o + \frac{1}{\zeta_1} Y_{2,t}^o = O_{1,t} + \frac{1}{\zeta_1} O_{2,t}, \quad (6)$$

⁵For convenience, we suppress firm-specific indices as all the distributors behave identically in equilibrium.

where $O_{i,t} = O_{i,t}^y + O_{i,t}^c$. For the oil market to clear, the sum of home and foreign oil production must equal the sum of home and foreign oil consumption by firms and households.⁶

Our model of the oil market focuses on the demand side of the market, while keeping the supply side deliberately simple, similar to Backus and Crucini (1998). This approach is in line with overwhelming empirical evidence in recent years that the large fluctuations in the real price of oil have been driven by demand shocks (see, for example, Bodenstein and Guerrieri, 2011; Kilian, 2009; Kilian and Hicks, 2011; Kilian and Murphy, 2010, 2012). A number of recent DSGE studies have imposed more structure on the supply side of the crude oil market, often focusing on models of imperfect competition (see, for example, Nakov and Pescatori, 2010a, b; Balke, Brown, and Yücel, 2010; Nakov and Nuño, 2011). Finding direct empirical evidence in favor of such models is difficult, given the paucity of relevant data (see, for example, Smith, 2005; Almoguera, Douglas, and Herrera, 2011). Although it is not difficult to design elaborate models of endogenous oil production decisions, without reliable data on reserves, exploration, drilling, and other investment activities that could be used to pin down the parameters of this process, it is difficult to estimate the parameters of such models reliably. Given the lack of a consensus on how to model the supply side of the global crude oil market, we treat oil production as exogenous.

Monetary policy follows a modified version of the interest rate reaction function suggested by Taylor (1993):

$$\begin{aligned} i_{1,t} = & \bar{i}_1 + \gamma_1^i (i_{1,t-1} - \bar{i}_1) + (1 - \gamma_1^i) \left[(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) \right. \\ & \left. + \gamma_1^\pi (\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^y y_{1,t}^{gap} \right] + \varepsilon_{1,t}^i. \end{aligned} \quad (7)$$

The terms $\bar{i}_1$ and $\bar{\pi}_1^{core}$ are the steady-state values for the nominal interest rate and inflation, respectively. The inflation rate $\pi_{1,t}^{core}$ is expressed as the logarithmic percentage change of the core price level, that is, inflation in non-oil consumer prices $\pi_{1,t}^{core} = \log(P_{1,t}^{ne}/P_{1,t-1}^{ne})$. The term $y_{1,t}^{gap}$ denotes the log deviation of gross output from the value of gross output in a model that excludes nominal rigidities, but is otherwise identical to the one described.⁷ The parameter γ_1^i allows for interest rate smoothing. The term $\varepsilon_{1,t}^i$ may reflect a time varying inflation target or any other stochastic innovation to the monetary policy rule.

⁶Because all variables are expressed in per capita terms, foreign variables are scaled by the relative population size of the home country $1/\zeta_1$.

⁷We effectively setup two parallel models, with and without the nominal rigidities in both countries. With the exclusion of the shock processes, we allow all of the state variables in the two models to differ, including the capital stocks.

The preferences of the representative household are given by:

$$E_t \sum_{j=0}^{\infty} \beta_1^j \times \left\{ \frac{Z_{1,t}^c}{1 - \sigma_1} \left(Z_{1,t}^c C_{1,t+j} - \kappa_1 C_{1,t+j-1}^A \right)^{1-\sigma_1} + \frac{\chi_{0,1}}{1 - \chi_1} (1 - L_{1,t+j})^{1-\chi_1} \right\}. \quad (8)$$

E_t denotes the expectation conditional on information available at time t . The variables $C_{1,t}$ and $L_{1,t}$ represent consumption and hours worked, respectively. The parameter σ_1 is used to determine the intertemporal elasticity of substitution, χ_1 the Frisch elasticity of labor supply, $\chi_{0,1}$ the steady-state number of hours worked. The term $Z_{1,t+j}^c$ represents a preference shock to consumption. In addition, a household's utility from consumption is affected by the presence of external consumption habits, parameterized by κ_1 . $C_{1,t-1}^A$ is the per capita aggregate consumption level. In every period t , household h maximizes the utility functional (8) with respect to consumption, labor supply, investment, end-of-period capital stock, and holdings of domestic and foreign bonds, subject the budget constraint, and the law of motion for capital. In doing so, prices, wages, and net transfers are taken as given.

The model encompasses an unusually rich stochastic structure. Table 1 summarizes the 15 separate sources of shocks in our model. In addition to the shocks described above, the model allows for shocks to U.S. investment, wage and price markups, and government spending. The modeling of these shocks follows Smets and Wouters (2007). The investment-specific technology shock governs the relationship between current investment and its impact on the capital stock of the economy. Price markup shocks are modeled as raising or lowering the elasticity of substitution between product varieties. Wage markup shocks follow the same structure and affect the elasticity between differentiated labor inputs. Shocks to government spending are expressed in terms of shocks to the government spending to GDP ratio.

III. Model Results

The model is estimated by the method of maximum likelihood based on quarterly data for 1984:Q1 through 2008:Q4 using 15 observed series: the log of U.S. and foreign GDP, U.S. and foreign oil production, the U.S. dollar price of oil (deflated by the U.S. GDP deflator), U.S. hours worked per capita, and the real dollar trade-weighted exchange rate; the GDP share of U.S. private consumption expenditures, the GDP share of U.S. oil imports, the GDP share of U.S. non-oil goods imports, the GDP share of U.S. goods exports, the GDP share of U.S. fixed investment; the level of U.S.

Table 1. Shocks Processes

Shock	Stochastic process
Home shocks	
Neutral technology	$\ln(Z_{1,t}) = (1 + \rho_1^z - \rho_2^z) \ln(Z_{1,t-1}) - \rho_1^z \ln(Z_{1,t-2}) + \sigma_1^z \varepsilon_{1,t}^z$
Investment	$\ln(Z_{1,t}^i) = \rho_1^i \ln(Z_{1,t-1}^i) + \sigma_1^i \varepsilon_{1,t}^i$
Consumption	$\ln(Z_{1,t}^c) = \rho_1^c \ln(Z_{1,t-1}^c) + \sigma_1^c \varepsilon_{1,t}^c$
Government spending	$\ln(Z_{1,t}^g) = \rho_1^g \ln(Z_{1,t-1}^g) + \sigma_1^g \varepsilon_{1,t}^g$
Price markup	$\hat{\theta}_{1,t}^p = \rho_1^p \hat{\theta}_{1,t-1}^p + \sigma_1^p \varepsilon_{1,t}^p$
Wage markup	$\hat{\theta}_{1,t}^w = \rho_1^w \hat{\theta}_{1,t-1}^w + \sigma_1^w \varepsilon_{1,t}^w$
Monetary policy	$\bar{\pi}_{1,t}^{core} = \rho_1^{\pi} \bar{\pi}_{1,t-1}^{core} + \sigma_1^{\pi} \varepsilon_{1,t}^{\pi}$
Oil-specific shocks	
Home oil supply	$\ln(Y_{1,t}^o) = (1 + \rho_{11}^{yo} - \rho_{21}^{yo}) \ln(Y_{1,t-1}^o) - \rho_{11}^{yo} \ln(Y_{1,t-2}^o) + \sigma_1^{yo} \varepsilon_{1,t}^{yo}$
Foreign oil supply	$\ln(Y_{2,t}^o) = (1 + \rho_{12}^{yo} - \rho_{22}^{yo}) \ln(Y_{2,t-1}^o) - \rho_{12}^{yo} \ln(Y_{2,t-2}^o) + \sigma_2^{yo} \varepsilon_{2,t}^{yo}$
Home oil intensity	$\ln(Z_{1,t}^o) = (1 + \rho_1^{zo} - \rho_2^{zo}) \ln(Z_{1,t-1}^o) - \rho_1^{zo} \ln(Z_{1,t-2}^o) + \sigma_1^{zo} \varepsilon_{1,t}^{zo}$
Foreign oil intensity	$\ln(Z_{2,t}^o) = (1 + \rho_1^{zo} - \rho_2^{zo}) \ln(Z_{2,t-1}^o) - \rho_1^{zo} \ln(Z_{2,t-2}^o) + \sigma_2^{zo} \varepsilon_{2,t}^{zo}$
Other open-economy shocks	
Foreign neutral technology	$\ln(Z_{2,t}) = (1 + \rho_1^z - \rho_2^z) \ln(Z_{2,t-1}) - \rho_1^z \ln(Z_{2,t-2}) + \sigma_2^z \varepsilon_{2,t}^z$
Home import	$\ln(Z_{1,t}^m) = (1 + \rho_1^{zm} - \rho_2^{zm}) \ln(Z_{1,t-1}^m) - \rho_1^{zm} \ln(Z_{1,t-2}^m) + \sigma_1^{zm} \varepsilon_{1,t}^{zm}$
Foreign import	$\ln(Z_{2,t}^m) = (1 + \rho_1^{zm} - \rho_2^{zm}) \ln(Z_{2,t-1}^m) - \rho_2^{zm} \ln(Z_{2,t-2}^m) + \sigma_2^{zm} \varepsilon_{2,t}^{zm}$
Foreign consumption	$\ln(Z_{2,t}^c) = \rho_2^c \ln(Z_{2,t-1}^c) + \sigma_2^c \varepsilon_{2,t}^c$

For shocks that occur in both countries, we impose that the autoregressive coefficients are identical except in the case of oil supply shocks.

core PCE inflation, U.S. wage inflation, and the U.S. federal funds rate. Tables 2 and 3 provide information on the calibrated and estimated parameters, respectively.

The estimate of the price elasticity of oil demand, at a value of 0.42, is close to recent empirical estimates.⁸ The estimate of 1.8 for the elasticity of

⁸It is still widely believed that the short-run price elasticity of oil demand is close to zero. This consensus is based on reduced-form regression estimates that are known to be biased toward zero. These traditional estimates are invalid. Recently, a number of studies have provided properly identified estimates of the short-run price elasticity of oil demand from

Table 2. Steady-State Ratios and Calibrated Parameters

Parameter	Used to Determine	Parameter	Used to Determine
Parameters common across countries			
$\beta = 0.99$	Discount factor	$\sigma = 1$	Intertemporal consumption elasticity
$\delta = 0.025$	Depreciation rate of capital	$\rho_v = -2$	K-L sub. elasticity (0.5)
$g = 0.18$	Steady-state government consumption share of GDP	$N_{ss} = 0.33$	Steady-state labor share to fix χ_0
$\mu_o = 1.0026$	Trend growth in oil supply		
Parameters not common across countries			
$\omega_k = 1.54$	Parameter on K in value added (home)	$\omega_k^* = 1.60$	Parameter on K in value added (foreign)
$\omega_{oy} = 0.026$	Weight on oil in production (home)	$\omega_{oy}^* = 0.057$	Weight on oil in production (foreign)
$\omega_{oc} = 0.021$	Weight on oil in consumption (home)	$\omega_{oc}^* = 0.041$	Weight on oil in consumption (foreign)
$\omega_{mc} = 0.068$	Weight on imports in consumption (home)	$\omega_{mc}^* = 0.039$	Weight on imports in consumption (foreign)
$\omega_{mi} = 0.40$	Weight on imports in investment (home)	$\omega_{mi}^* = 0.25$	Weight on imports in investment (foreign)
Parameters specific to home country			
$\zeta = 1/2$	Relative size of home country	$\frac{y_{oss}^l}{o_{yss}^l + o_{css}^l} = 0.3$	Steady-state ratio oil production to consumption (home)
$\phi_b = 0.0001$	Curvature of bond intermed. cost		

substitution between domestic and foreign goods is also close to typical estimates based on aggregate data. The estimates of the parameters in the monetary policy rule are in line with other estimates of new Keynesian models. The estimate of 0.7 for the interest rate smoothing parameter represents the most glaring departure from the rule in Taylor (1993), but is in line with typical estimates based on new-Keynesian DSGE models. Moving to the wage- and price-setting equations, the Calvo parameter for wages is estimated at 0.89. The Calvo parameter for prices is estimated at 0.88. Bodenstein and Guerrieri (2011) find little evidence in favor of lagged indexation for either prices or wages. An online appendix and Bodenstein and Guerrieri (2011) provide further details and discussion of the parameter estimates.

structural econometric models. The latter studies, regardless of methodology, yield much higher elasticity estimates that are similar in magnitude to our estimate in this paper (see, for example, Kilian and Murphy (2010) and the references therein).

Table 3. Estimation Results

	Estimate
ρ_1^z , Technology, growth AR coefficient	0.2163
ρ_2^z , Technology, level error correlation coefficient	0.0001
σ_1^z , U.S. Technology, standard deviation of innovation	0.0066
σ_1^z , Foreign Technology, standard deviation of innovation	0.0108
ρ_1^{zi} , U.S. Investment Technology, AR coefficient	0.9059
σ_1^{zi} , U.S. Investment Technology standard deviation of innovation	0.0269
ρ_1^{zg} , U.S. Government Expenditure, AR coefficient	0.9980
σ_1^{zg} , U.S. Government Expenditure standard deviation of innovation	0.0246
ρ_{11}^{vo} , U.S. Oil Supply, growth AR coefficient	0.1236
ρ_{21}^{vo} , U.S. Oil Supply, level error correlation coefficient	0.0001
σ_1^{vo} , U.S. Oil Supply, standard deviation of innovation	0.0253
ρ_{12}^{vo} , Foreign Oil Supply, growth AR coefficient	0.0001
ρ_{22}^{vo} , Foreign Oil Supply, level error correlation coefficient	0.0378
σ_2^{vo} , Foreign Oil Supply, standard deviation of innovation	0.0181
ρ_{11}^{eo} , Oil Efficiency, growth AR coefficient	0.0001
ρ_{21}^{eo} , Oil Efficiency, level error correlation coefficient	0.0145
σ_1^{eo} , U.S. Oil Efficiency, standard deviation of innovation	0.0470
σ_2^{eo} , Foreign Oil Efficiency, standard deviation of innovation	0.1269
ρ_1^c , Consumption Shock, AR(1) coefficient	0.9188
σ_1^c , U.S. Consumption, standard deviation of innovation	0.6484
σ_2^c , Foreign Consumption, standard deviation of innovation	0.7174
ρ_1^m , Import, growth AR coefficient	0.0001
ρ_1^m , Import, level error correlation coefficient	0.0019
σ_1^m , U.S. Import, standard deviation of innovation	0.0263
σ_2^m , Foreign Import, standard deviation of innovation	0.0412
ρ_1^w , U.S. Wage Markup, AR(1) coefficient	0.9768
σ_1^w , U.S. Wage Markup, standard deviation of innovation	3.6988
ρ_1^p , U.S. Price Markup, AR(1) coefficient	0.7401
σ_1^p , U.S. Price Markup, standard deviation of innovation	0.4774
ρ_1^p , U.S. Monetary Policy, AR(1) coefficient	0.4026
σ_{11}^p , U.S. Monetary Policy, standard deviation of innovation	0.0217
$((1 + \rho_1^q)/\rho_1^q)$, Oil Substitution Elasticity	0.4225
$((1 + \rho_1^c)/\rho_1^c)$, Trade Substitution Elasticity	1.7570
μ_z , Growth Rate of Technology (gross)	1.0058
K_1 , Habits in Consumption	0.6512
γ_1^i , Policy Rate Smoothing	0.6553
γ_1^π , Weight on Inflation in Monetary Policy Rule	0.1907
γ_1^y , Weight on Output Gap in Monetary Policy Rule	0.0000
ξ_1^p , Calvo Price Parameter	0.8140
ξ_1^w , Calvo Wage Parameter	0.8900
ι_p , Lagged Price Indexation	0.0000
ι_w , Lagged Wage Indexation	0.0000
π_1^{core} , Steady-State Inflation	1.0114
ψ_1^i , Investment Adjustment Cost	3.5154
ξ , Determines Labor Supply Elasticity ($1/2\xi$)	59.5402

The lower bound for the coefficient on the level error correction components is 0.0001.

The Channels of Transmission to Inflation and Output in the Net Oil-Importing Economy

A key question in recent years has been how U.S. monetary policymakers should respond to an increase in the real price of oil driven by increased demand for oil from emerging markets in particular. A number of recent studies using a variety of methods have shown that positive foreign oil intensity shocks are one of the key determinants of the surge in the real price of oil between 2003 and mid-2008 (see, for example, Kilian, 2009; Kilian and Hicks, 2011; Bodenstein and Guerrieri, 2011). For example, Kilian (2009) noted that the global demand for oil depends not only on the pace of overall growth, but also on how intensely oil is used in producing domestic real output. Thus models concerned with changes in real GDP or in aggregate productivity alone will be unable to explain the extent of the surge in the real price of oil between 2003 and mid-2008.

Given the empirical importance of foreign oil intensity shocks in our estimated model, in the next sub-section we examine in detail how these shocks are transmitted to U.S. real activity and inflation, and what determines the appropriate U.S. monetary policy response. We offer a novel decomposition of the domestic marginal cost of production that highlights the role of each factor input in the evolution of domestic inflation. While it is not possible (and indeed not necessary) to analyze each of the structural shock in our model in the same detail, further below we provide a comparison of how the dynamic response of the real price of oil and of the U.S. interest rate differs for key structural shocks in the model. We show that not only the pattern and magnitude, but even the sign of the monetary policy response may differ depending on the origin of the oil price fluctuations.

The Effects of a Foreign Oil Intensity Shock

Figure 1 illustrates the effects of a one-standard deviation shock that pushes up foreign oil intensity through a change in $Z_{2,t}^o$.⁹ As foreign oil demand expands, the real price of oil in U.S. consumption units increases. Upon impact, the price rises 15 percent. The half life of the response is close to 5 years. Home oil demand contracts as both households and firms substitute away from the more expensive oil input.

Eventually, lower oil use leads to a fall in the current and future marginal product of capital, causing investment, consumption, and gross output to fall. However, in the short run, the shock does not unequivocally lead to a fall in output. These short-term output dynamics are explained by real rigidities, the behavior of net exports, and monetary policy.

First, real rigidities prevent consumption and investment from adjusting immediately, as can be inferred from the response of domestic absorption.

⁹When the oil substitution elasticity is less than 1, an increase in foreign oil intensity is brought about by a decline in $Z_{2,t}^o$.

Absent real rigidities, output would fall in the short run despite the net export channel and role of monetary policy.

Second, non-oil net exports expand for a net oil importer like the United States and hence act toward pushing output upward. Given that $(1 + \rho_1^o)/\rho_1^o$ is well below unity, an oil price increase results in a marked deterioration of the oil trade balance. With incomplete international financial markets, the deterioration in the oil trade balance implies substantially different wealth effects across countries. As the negative wealth effect is larger for the oil importer, the home non-oil terms of trade worsen and induce an expansion in non-oil net exports. As shown in Bodenstein, Erceg, and Guerrieri (2011), the effect on net exports is more pronounced for lower values of the oil price elasticity.¹⁰

Third, nominal rigidities and monetary policy also play an important role in shaping the short-term output response. In Figure 1, realized output expands, whereas potential output contracts. In the presence of pronounced real rigidities that make the economy relatively insensitive to changes in the real interest rate in the short run, large swings in the real interest rate occur in the potential economy in order to curb domestic absorption. Consequently, potential absorption drops substantially more than realized absorption. By contrast, the smoothing component of the estimated historical monetary policy rule generates a gradual increase in real rates that ends up overshooting the increase in potential rates. The relative movements of realized and potential output mirror those of the interest rate movements. Over time, as potential real rates fall more sharply than realized rates, potential output recovers more quickly and “leapfrogs” realized output. Figure 1 also reveals that the initial expansion in realized output is associated with an expansion in hours worked. By contrast, the demand for oil and capital falls uniformly.

Inflation Dynamics

The systematic response of monetary policy to inflation associated with unexpected oil price movements has been the subject of intense scrutiny. One hypothesis advanced by Bruno and Sachs (1985) points to interactions between wages and prices that could lead to persistent inflation increases as a possible mechanism for why monetary policy could deepen the effects of shocks that drive the price of oil upward. Other papers deemed those interactions implausible on account of evidence of falling real wages in response to higher oil prices (see, for example, Rotemberg and Woodford, 1996). Our model is a

¹⁰For each country, we define the non-oil terms of trade as the price of imports over the price of exports, expressed in a common currency; accordingly, an upward movement in Figure 1 denotes a worsening of the terms of trade for the home country. We define the exchange rate as the price of the home consumption basket over the price of the foreign consumption basket, expressed in a common currency; accordingly, an upward movement in Figure 1 denotes a depreciation for the home country.

Figure 1. The Effects of a One-Standard Deviation Increase in Foreign Oil Intensity: Deviations from the Balanced Growth Path

Notes: All inflation measures and interest rates are annualized.

reminder that falling real wages cannot be taken as a sufficient statistic for the absence of inflation pressures from the labor market.

The structure of the model helps us disentangle these various hypotheses. Figure 1 highlights that a quantitatively important channel for the initial increase in core inflation is the deterioration of the terms of trade.

The presence of imported intermediate goods in the final consumption good accounts for the wedge between core inflation and domestic goods inflation.

When it comes to domestic goods inflation, the standard New-Keynesian Phillips curve implicit in our model provides a familiar framework for understanding the propagation channels of movements in the price of oil. Given that the estimate for the parameter governing lagged indexation, ι^p , is 0 and abstracting from the mark-up shock $\hat{\theta}_{1,t}$, one can express domestic inflation $\hat{\pi}_{1,t}$ as:

$$\hat{\pi}_{1,t} = \sum_{s=0}^{\infty} \beta_1^s \frac{(1 - \xi_1^p \beta_1)(1 - \xi_1^p)}{\xi_1^p} E_t \widehat{mc}_{1,t+s}, \quad (9)$$

where the term $\widehat{mc}_{1,t+s}$ is the marginal cost of production in log deviation from its level along the balanced growth path, $1 - \xi_1^p$ is the Calvo probability of renewing the price contract, and β is the discount factor. In words, to a first-order approximation, current inflation can be thought of as the discounted sum of current and expected marginal costs of production. The only departure from the familiar specification estimated in Galí and Gertler (2000) is that marginal cost depends on oil as an additional factor input and can be expressed as:

$$\begin{aligned} \widehat{mc}_{1,t} = & \omega_1^{oy} \left[\hat{p}_{1,t}^o - \widehat{mpo}_{1,t} \right] \\ & + \omega_1^{vy} \phi_1 \left[\hat{r}_{1,t}^k - \widehat{mpk}_{1,t} \right] + \omega_1^{vy} (1 - \phi_1) \left[\hat{w}_{1,t} - \widehat{mpl}_{1,t} \right], \end{aligned} \quad (10)$$

where $\widehat{mpo}_{1,t}$, $\widehat{mpk}_{1,t}$, $\widehat{mpl}_{1,t}$ are the marginal products of oil, capital, and labor inputs, respectively, all in log deviation from their values along the balanced growth path. ϕ_1 is a constant related to the share of capital in value added.¹¹ Accordingly, domestic goods inflation is related to appropriately weighted gaps between the rental rate and the marginal product of each factor input. When nominal rigidities are absent, these gaps never open up and the real marginal cost is constant. But with nominal price rigidities, even abstracting from sticky wages, these gaps can be sizable.

Figure 2 considers the reaction of each factor input to the same shock to oil intensity discussed thus far. The fall in the oil input in reaction to higher prices pushes up the marginal product of oil, but exerts downward pressure on the marginal product of the other factor inputs. As monetary policy does

¹¹The term ϕ_1 satisfies

$$\phi_1 = \omega_1^k \left(\frac{1}{\omega_1^k \mu_z} \frac{K_{1,0}^*}{V_{1,0}^*} \right)^{\frac{1}{1+p^v}},$$

with $K_{1,0}^*/V_{1,0}^*$ denoting the ratio of capital to value added along the balanced growth path.

Figure 2. The Effects of a One-Standard Deviation Increase in Foreign Oil Intensity on Factor Inputs: Deviations from the Balanced Growth Path

Notes: The right column shows the marginal product, the rental rate and the contribution to marginal cost for each factor of production.

not push up the real rate aggressively enough, aggregate demand only contracts gradually and the demand for oil remains so elevated that the rental rate overshoots the marginal product. However, when weighted by the appropriate share, this gap only makes a small contribution to the rise in marginal cost and inflation.¹²

¹²Note that oil adjustment costs would induce a larger gap between the rental rate and the marginal product of oil inputs, but would reduce the gaps for capital and labor.

No such drastic reduction in magnitude occurs for the gap associated with labor inputs, since they have the largest share in production. As firms shift away from using more expensive oil, they push up the relative demand for other factor inputs. The labor input is the only factor that can be adjusted immediately, so hours worked increase. Notice that the contraction of the marginal product of labor is partly linked to the increase in the labor input and partly to the fall in oil input. Sticky nominal wages ward off a large immediate rise in the rental rate for labor, but they also hinder subsequent downward adjustment toward the marginal product. The resulting persistent gap between the real wage and the marginal product of labor is the key contributor to the rise in marginal cost and domestic price inflation.

Finally, capital is predetermined in its first period and there are sizable adjustment costs for investment. As agents are forward-looking in planning investment and aggregate demand is predicted to fall, capital inputs fall uniformly. On impact, the higher demand for all factor inputs leads to a substantial rise in the rental rate for capital, but that gap quickly closes up. The overall contribution to marginal costs remains modest relative to that of labor inputs.

Based on the response of capital inputs, it is easy to see that without nominal wage rigidities, the real wage would also jump up on impact. Accordingly, sticky wages restrain the contribution of labor to the rise in marginal cost. Over time, however, sticky wages also impart persistence to the increase in marginal cost and inflation. In short, the simulation results act as a reminder that falling real wages are not sufficient for lower cost pressures on inflation. Rather the labor input can make a contribution to the increase in marginal cost and inflation even with a falling real wage, simply because the reduction in other factor inputs depresses the marginal product of labor persistently.

What Difference the Source of the Oil Price Fluctuations Makes

The discussion so far has focused on a foreign oil intensity shock. Although this shock is estimated to be the most prominent driver of recent oil price fluctuations in our model, it is by no means the only source of variation in oil prices. In this section, we illustrate differences in the magnitude, pattern, and sign of the policy responses to different structural shocks. It is important to bear in mind that every structural shock in the model has implications for either the demand for oil or the supply of oil and hence sets in motion adjustments in both the real price of oil and in domestic and foreign macroeconomic aggregates. We focus on one shock at a time. It is understood that policymakers in real life may face several oil demand and oil supply shocks at the same time, the response to which will be a weighted average of the responses shown.

It may seem that the same type of shock taking place in a different part of the world should have similar effects on the real price of oil and on U.S. monetary policy. This is not the case. In the next sub-section, we

illustrate this point for the example of domestic and foreign oil intensity shocks. When comparing the magnitude of the responses with domestic and foreign oil intensity shocks, we control for the magnitude of the oil price increase implied by these shocks. This approach facilitates the comparison, but may require considering shocks far greater than are likely to prevail in practice. In the subsequent sub-section we consider a broader array of shocks. We show that no two structural shocks are alike in that each shock induces a different monetary policy response in the United States. The policy responses implied by the estimated policy rule differ not only in their magnitude and shape, but even their sign may differ.

How the Same Type of Shock Has Different Effects Depending on Where in the World it Arises

Figure 3 compares the effects of positive domestic and foreign oil intensity shocks. For ease of comparison, the solid lines show again the responses to the one-standard deviation increase in foreign oil intensity discussed above. The dashed line shows responses to an increase in U.S. oil intensity. To make the initial price increase comparable across shocks, the magnitude of the U.S. oil intensity shock had to be magnified to approximately 12 standard deviations. Not surprisingly, this scaling of the domestic intensity shock greatly magnifies the effects on domestic activity. Controlling for the oil price increase, the response of domestic price inflation is substantially larger in the case of the domestic oil intensity shock. The decomposition of marginal cost in Equation (10) again provides a useful framework for investigating the differences.

Equally strikingly, the gap between the real oil price and the marginal product of the oil input remains small. Log-linearizing the marginal product of the oil input offers some useful clues:

$$\frac{\partial \hat{Y}_{1,t}}{\partial O_{1,t}^o} = \frac{\rho_1^o}{1 + \rho_1^o} \left(\hat{Y}_{1,t} - \hat{O}_{1,t}^o \right) + \left(1 - \frac{\rho_1^o}{1 + \rho_1^o} \right) Z_{1,t}^o. \quad (11)$$

Observe that the oil substitution elasticity (in absolute value) is $(1 + \rho_1^o)/\rho_1^o$. In our case that elasticity is estimated to be 0.42. Moreover, the term $1 - \rho_1^o/(1 + \rho_1^o)$ is approximately -1.4 and, thus, the decline in $Z_{1,t}^o$ that brings about the increase in domestic oil consumption also increases the marginal product of oil. The increase in the marginal product of oil, in conjunction with the relatively small share of oil in production, explains why the contribution of the oil gap to marginal cost in Equation (10) remains small, even in the case of this domestic shock.

As shown in Figure 4, once again, the gap that opens up between the marginal product of labor and the rental rate of labor makes the largest contribution to the increase in marginal cost and domestic price inflation. To understand why the marginal product of labor declines substantially more in

Figure 3. A Comparison of Foreign and Domestic Oil Intensity Shocks

the case of the domestic oil intensity shock, relative to when the same type of shock occurs abroad, it is useful to consider the log-approximation to the marginal product of labor given by:

$$\frac{\partial \hat{Y}_{1,t}}{\partial L_{1,t}} = \frac{\rho_1^o}{1 + \rho_1^o} (\hat{Y}_{1,t} - \hat{V}_{1,t}) + \hat{Z}_{1,t} + \frac{\rho_1^v}{1 + \rho_1^v} (\hat{V}_{1,t} - \hat{Z}_{1,t} - L_{1,t}). \quad (12)$$

In the case of the domestic oil intensity shock, Equation (2) shows that the term $Y_{1,t}$ in the equation for the marginal product of labor above is directly affected by the shock $Z_{1,t}^o$. This impact substantially magnifies the decline in the marginal product of labor and consequently

Figure 4. The Effects of a One-Standard Deviation Increase in U.S. Oil Intensity on Factor Inputs: Deviations from the Balanced Growth Path

Notes: The right column shows the marginal product, the rental rate and the contribution to marginal cost for each factor of production.

the increase in the marginal cost of production and domestic price inflation.

In contrast, in the economy without price and wage rigidities, the gap between the marginal product of labor and the real wage does not open up. In the absence of nominal rigidities, the real wage decline, in line with the larger decline in the marginal product of labor, results in a bigger reduction in labor supply and in production. Hence, the direct impact of the domestic oil intensity shock on the marginal product of labor is also connected to the larger initial output gap and initial greater increase in the Federal Funds rate shown in Figure 3.

Figure 5. The Effects of Different Shocks on U.S. Interest Rates (the shocks are scaled to induce a half percent increase in the real price of oil on impact)

The scale of the U.S. productivity shock is 1.88 standard deviations. The scale of the foreign productivity shock is 2.60 standard deviations. The scale of the U.S. consumption preference shock is 7.85 standard deviations. The scale of the foreign consumption preference shock is 0.46 standard deviations. The scale of the U.S. oil intensity shock is 0.39 standard deviations. The scale of the foreign oil intensity shock is 0.033 standard deviations. The scale of the U.S. oil supply shock is 1.26 standard deviations. The scale of the foreign oil supply shock is 0.12 standard deviations. The abbreviation “bpg” refers to “balanced growth path.”

No Two Shocks Induce the Same Policy Response

It is sometimes claimed that the origin of an oil price shock does not matter, as long as the source of the oil price shock is abroad. For example, Blanchard and Gali (2010) suggest that “if the price of oil rises as a result of, say, higher Chinese demand, this is just like an exogenous oil supply shock for the remaining countries.” Figure 5 demonstrates that this conjecture is not correct even controlling for the initial oil price increase. Not only is the subsequent response of the real price of oil different, but so is the interest rate response under the estimated policy rule. For example, if the increase in oil demand arises from a foreign productivity shock, then the interest rate is positive except on impact, whereas it is negative except for the first few quarters when the same oil price increase is driven by increased foreign oil intensity. More generally, no two structural shocks induce the same responses, regardless of the scale of the structural shocks.

Figure 6. Variation in the U.S. Federal Funds Rate Explained by Different Subsets of Shocks *Explaining the Evolution of the Interest Rate*

We conclude this section with a historical decomposition of the cumulative effects of oil supply shocks and oil intensity shocks on the U.S. federal funds rate. Although oil intensity shocks in particular explain much of the variation in the real price of oil since the mid-1980s in our model, as documented in Bodenstein and Guerrieri (2011), Figure 6 shows that these shocks explain little of the evolution of the U.S. federal funds rate. Instead, much of the historical variation in the federal funds rate is explained by the remaining shocks in the model (including spending shocks, productivity shocks, price markup shocks, trade shocks, investment shocks, and consumption preference shocks both in the United States and abroad). We conclude that oil supply and foreign oil intensity shocks have had little impact on monetary policy in the United States, to the extent that our policy rule is an adequate characterization of U.S. monetary policy.¹³ This empirical finding is also consistent with the VAR evidence in Kilian and Lewis (2011).

Optimal Monetary Policy in the Oil-Importing Economy

So far, we have focused on the responses to shocks derived under the estimated monetary policy rule. As seen above, that rule implies substantial inertia in

¹³The dotted line in Figure 6 substantiates that the monetary policy shock plays a modest role in the evolution of U.S. monetary policy.

the response of the real rate relative to the potential economy, suggesting that an optimal rule would behave quite differently. In this section we depart from the estimated model by optimizing the coefficients in the monetary policy rule with respect to a social welfare criterion. We focus on the following class of rules:

$$i_{1,t} = \bar{i}_1 + \gamma_1^i(i_{1,t-1} - \bar{i}_1) + (1 - \gamma_1^i) \left[(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^\pi(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^y y_{1,t}^{gap} \right] + \gamma_1^o(\pi_{1,t}^o - \bar{\pi}_1^o) + \gamma_1^w(\omega_{1,t} - \bar{\omega}_1) + \varepsilon_{1,t}^i \quad (13)$$

This class nests the estimated rule in Equation (7). The rule we optimize allows for two additional terms: $\pi_{1,t}^o$, oil price inflation, defined as the log difference in the nominal price of oil $P_{1,t}^o$; and $\omega_{1,t}$, wage inflation, defined as the log difference in the nominal wage rate $W_{1,t}$. In the optimization, we choose the coefficients of the monetary policy rule that govern the degree of interest rate smoothing, γ_1^i ; the strength of the responses to the deviation of inflation from target, γ_1^π ; to the output gap, γ_1^y ; and to the deviation of oil price inflation and wage inflation from their levels along the balanced growth path, respectively, γ_1^o and γ_1^w . Those coefficients are chosen so as to maximize, in expectation, the utility of the representative agent, defined in Equation (8). The monetary policy rule in the rest of the world is taken as given. A key advantage of our structural approach is that agents in the model internalize changes in the monetary policy rule.¹⁴

Optimized Rule

Table 4 presents the optimized coefficients. The first two rows are devoted to the benchmark model. Most of the coefficients are close to zero, including the coefficient on oil price inflation. This means that including the oil price in the policy rule does not improve welfare, but may lower welfare if the coefficient on the oil price is not essentially zero.

The optimized coefficient γ_1^y is equal to 1.67×10^6 . Such a large value for γ_1^y implies full stabilization of the output gap such that the standard deviation of the output gap is zero. This finding is in line with previous results in

¹⁴We compute the value of the expected discounted utility in Equation (8) by solving the entire model using second-order perturbation methods. We search for the global maximum of expected discounted utility using a sequence of numerical algorithms: (1) simulated annealing, (2) a Nelder-Mead search method, (3) a Newton-Raphson method. The stopping criteria are specified both in terms of changes in the objective function and in terms of changes in the input parameters governing the monetary policy response function. After we identify a candidate for a global maximum, we re-start the sequence of algorithms repeatedly to guard against the possibility of having identified only a local maximum. In searching for the global maximum, we put a lower bound of zero on all of the parameters in the monetary policy rule. We adapt the upper bounds on the parameters to ensure that they are at least two orders of magnitude larger than the monetary policy rule parameters at the identified maximum.

Table 4. Optimized Rules Under Alternative Model Specifications*

Rule	γ_1^i	γ_1^π	$\gamma_1^y \times 10^6$	γ_1^o	γ_1^w
Benchmark model					
Estimated	0.655	0.19	0.00	—	—
Optimized	0.000	0.02	1.67	0.01	0.003
4-quarter Calvo contracts					
Optimized	0.007	0.14	1.67	0.21	0.089
No price and wage markup shocks					
Optimized	0.075	0.08	1.67	0.21	0.007
4-quarter Calvo contracts and no price and wage markup shocks					
Optimized	0.004	0.03	1.67	0.24	0.030
No oil supply and no oil intensity shocks					
Optimized	0.074	0.05	1.67	0.10	0.002
Oil supply and oil intensity shocks only					
Optimized	0.004	0.11	17.78	0.11	0.000

*In each case, the optimized rule belongs to the following class of rules:

$$i_{1,t} = \bar{i}_1 + \gamma_1^i(i_{1,t-1} - \bar{i}_1) + (1 - \gamma_1^i) \left[\begin{array}{c} (\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^\pi(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^y y_{1,t}^{gap} \\ \gamma_1^o(\pi_{1,t}^o - \bar{\pi}_1^o) + \gamma_1^w(\omega_{1,t} - \bar{\omega}_1) \end{array} \right] + \varepsilon_{1,t}^i.$$

Bodenstein, Erceg, and Guerrieri (2008). For a stylized model nested in our setting, they showed that the optimal welfare-maximizing policy under commitment is well approximated by rules that target the output gap. Our results confirm that their previous analysis translates to simple instrument rules (as opposed to targeting rules) and applies to a large-scale empirically validated model.

The welfare losses in Table 5 relate to the changes in expected welfare relative to the optimized rule, expressed in terms of the equivalent change in permanent consumption, as a percentage of steady-state consumption.¹⁵ The table also shows the standard deviations of core inflation, wage inflation, and the output gap.¹⁶ Table 5 shows that full stabilization of the output gap does not mean that the central bank ignores the inflation objective. Relative to the estimated rule, the optimized rule also reduces the

¹⁵In the table, the welfare changes are scaled by the term $100(U_C C_1 / 1 - \beta_1)$, where U_C represents the marginal utility of consumption evaluated at the non-stochastic steady state.

¹⁶For their stylized model, Bodenstein, Erceg, and Guerrieri (2008) showed that these variables had high weights in the welfare loss function.

Table 5. A Comparison of Alternative Monetary Policy Rules: Sensitivity Analysis

Rule	U.S. Welfare Loss (change from optimized)	U.S. Core Inflation Std. Dev.	U.S. Wage Inflation Std. Dev.	U.S. Output Gap Std. Dev.
Benchmark model				
Estimated	2.99	3.41	6.24	1.15
Optimized	0	2.67	0.98	0.00
4-quarter Calvo contracts				
Estimated	1.39	3.87	11.85	0.69
Optimized	0	2.91	3.12	0.00
No price and wage markup shocks				
Estimated	0.11	3.13	4.93	0.87
Optimized	0	1.93	0.53	0.00
4-quarter Calvo contracts and no price and wage markup shocks				
Estimated	0.12	3.58	7.91	0.50
Optimized	0	1.94	0.99	0.00
No oil supply and no oil intensity shocks				
Estimated	2.99	3.36	6.16	1.13
Optimized	0	2.60	0.76	0.00
Oil supply and oil intensity shocks only				
Estimated	0.0012	0.51	1.11	0.28
Optimized	0	0.46	0.39	0.08

Notes: The losses reported are expressed as a percent of steady-state consumption. The inflation measures are annualized.

standard deviations of core inflation and of wage inflation. The standard deviation of the latter, in particular, drops drastically from 6.24 to 0.98 percent.

One of the striking results in Table 5 is the size of the welfare gains in the benchmark model. Optimizing the coefficients of the policy rule yields an increase in welfare equivalent to a permanent increase in consumption equal to 2.99 percent of steady-state consumption. Table 5 also includes some sensitivity analysis that points to the features of the benchmark model that are responsible for the large welfare gains from optimization. As before, we compare the estimated and the optimized rule. We re-optimize the rule, as we vary key elements of the model. The first change considered is a reduction in the duration of price stickiness and wage stickiness such that the Calvo coefficients imply four-quarter contracts. This change alone more than halves the welfare gained from optimization, which drops from 2.99 to 1.39 percent of steady-state consumption. However, the permanent change in consumption remains an order of magnitude larger than typically reported.¹⁷

To understand this result recall that the volatility of wage and price inflation is tightly linked to two features of the model—the average duration of the Calvo contracts and the size of the wage and price markup shocks. From the decomposition of the population variance at business cycle frequencies, price markup shocks and wage markup shocks (or labor supply shocks) account for 50 percent of the variation in output in our model. Especially when the Frisch labor supply elasticity is estimated to be close to zero, departures from the labor supply schedule implied by wage rigidities can have a large impact on welfare, as shown in Equation (20) of Bodenstein, Erceg, and Guerrieri (2008).

A second result is that when we shut off wage and price markup shocks, the welfare gain of switching to the optimized rule falls to 0.11 percent of steady-state consumption, which is the same order of magnitude as the gains typically reported in studies that characterize optimal monetary policy in a general equilibrium context. This result illustrates that the presence of wage and price markup shocks is crucial for the welfare gains.

Third, Table 5 allows us to assess the effect of explicitly including the oil market in the model. When the oil supply shocks and oil intensity shocks in the United States and abroad are excluded from the model, for example, the gains from adopting the optimized rule change little nor does the optimal policy rule change noticeably. As shown in Table 4, virtually all the weight in the policy rule remains on the output gap. When these two shocks are the only source of variation in the model, the optimized coefficients put an even greater emphasis on the stabilization of the output gap. In this case, stabilization of the output gap is still consistent with stabilization of wage inflation, but it also comes close to stabilizing core inflation—the ratio of the standard deviation of core inflation to the standard deviation of wage inflation is now close to 1, instead of 3 for the benchmark model. This result is consistent with the decomposition of inflation presented earlier. The labor market is a key channel for the transmission of shocks that affect oil prices to inflation, so stabilizing wage inflation in the face of fluctuations in oil prices also achieves stabilization of core inflation.

Responses Under the Optimized Rule

A summary statistic such as the welfare measure reported in Table 4 does not convey how policy rates change when optimizing the coefficients of the policy rule. Figures 7 and 8 address this point by showing the nominal interest rate responses for selected shocks under the estimated and the optimized policy rule. Each structural shock has been rescaled to induce a half percent increase on impact in the real price of oil under the estimated rule. Figures 7 and 8

¹⁷Another study that finds large losses from the estimated policy relative to the optimal rule is Levin and others (2006). For a closed economy model of the United States, they find that the estimated policy has a welfare cost equal to 0.56 of steady-state consumption.

Figure 7. A Comparison of the Effects of Key Shocks Affecting Oil Prices under Alternative Policy Rules (the shocks are scaled to induce a half percent increase in the real price of oil on impact)

Notes: The scale of the U.S. productivity shock is 1.88 standard deviations. The scale of the foreign productivity shock is 2.60 standard deviations. The scale of the U.S. consumption preference shock is 7.85 standard deviations. The scale of the foreign consumption preference shock is 0.46 standard deviations. The abbreviation “bgp” refers to “balanced growth path.”

show substantial differences in the path of the policy instrument variable in response to the same shock, depending on whether the estimated or the optimized rule is used. The latter difference in magnitude may easily be as high as a factor of 10 or even differ in sign.

The reason for this difference is that, in general, under the optimized policy rule, the implied response of the real interest rate is closer to the real interest rate path in the “potential economy,” defined as an economy without nominal rigidities. Figures 9 illustrates this point for the example of a foreign oil intensity shock. For ease of comparison with Figures 2–4, the shock is resized to one standard deviation. The figure shows that in line with the information in Tables 5 and 6, the output gap does not open up. The immediate drop in consumption is more pronounced under the optimized rule, while the real interest rate, after increasing more sharply initially, comes down more quickly. Furthermore, core inflation is uniformly positive and wage inflation uniformly negative. The optimized rule avoids the overshooting typical of rules with interest rate smoothing because the weight on the lagged interest rate is essentially zero.

Figure 8. A Comparison of the Effects of Key Shocks Affecting Oil Prices under Alternative Policy Rules (the shocks are scaled to induce a half percent increase in real price of oil on impact)

Notes: The scale of the U.S. oil intensity shock is 0.39 standard deviations. The scale of the foreign oil intensity shock is 0.033 standard deviations. The scale of the U.S. oil supply shock is 1.26 standard deviations. The scale of the foreign oil supply shock is 0.12 standard deviations. The abbreviation “bgp” refers to “balanced growth path.”

Other Simple Rules

To allow comparisons on the performance of simple rules across different models, we also include some results for other specifications of the interest rate rule. Table 6 allows a comparison of the welfare loss implied by the estimated rules and by other rules often considered in the literature. For example, a policy rule that responds to core inflation and excludes interest rate smoothing fares better than the estimated rule in that the welfare gain from switching to the optimized rule is smaller.¹⁸ Furthermore, rules that respond to core inflation fare better than rules that respond to headline inflation. Overall, rules that do not place a large weight on the output gap result in excess

¹⁸We assigned parameters in line with Taylor (1993). Accordingly, $\gamma_1^i = 0$, $\gamma_1^\pi = 0.5$ and $\gamma_1^y = 0.125$. The monetary policy rule in Equation (7) is expressed in terms of quarterly policy rates and inflation, while Taylor (1993) expressed them in term of annual rates. Annualizing the interest rate rule implies multiplying the coefficient on the output gap by 4 and leaving the other coefficients unchanged.

Figure 9. The Effects of a One-Standard Deviation Increase in Foreign Oil Intensity: Deviations from the Balanced Growth Path Under Alternative Policy Rules and in the Potential Economy

variation in both the output gap and wage inflation, which causes substantial welfare losses.

Table 6 also considers the optimization of other types of policy rules. The rule labeled “GDP Growth” replaces the output gap with the deviation of real GDP growth from the balanced growth path and excludes a response to wage inflation. In that case, the optimized coefficient on GDP growth is zero, but the optimized coefficient on inflation, 2.76, is

Table 6. A Comparison of Alternative Monetary Policy Rules¹

Rule	γ_1^i	γ_1^π	γ_1^y	γ_1^o	γ_1^w
Estimated	0.655	0.19	0.00	—	—
Optimized	0.000	0.02	1.67	0.01	0.00
Taylor with Core Inflation	0	0.5	0.125	0	0
Core Inflation Only	0	2	0	0	-0.07
Taylor with Headline Inflation ²	0	0.5	0.125	0	0
Headline Inflation Only ²	0	2	4.90	0	0
GDP Growth ³	0.000	2.76	0.00	—	—
No Output Gap ⁴	0.028	0.00	—	0.01	3.65×10^5

Rule	U.S. Welfare Loss (rel. to optimized)	U.S. Core Inflation Std. Dev.	U.S. Wage Inflation Std. Dev.	U.S. Output Gap Std. Dev.
Estimated	2.99	3.41	6.24	1.15
Optimized	0	2.67	0.98	0.00
Taylor with Core Inflation	2.45	2.75	3.95	0.75
Core Inflation Only	2.44	1.59	4.72	1.14
Taylor with Headline Inflation ²	2.50	2.77	3.95	0.75
Headline Inflation Only ²	2.52	1.68	4.90	1.22
GDP Growth ³	2.42	1.35	5.17	1.29
No Output Gap ⁴	0.09	2.97	0.00	0.20

¹The optimized rule belongs to the following class:

$$i_{1,t} = \bar{i}_1 + \gamma_1^i(i_{1,t-1} - \bar{i}_1) + (1 - \gamma_1^i) \left[\begin{aligned} &(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^\pi(\pi_{1,t}^{core} - \bar{\pi}_1^{core}) + \gamma_1^y \pi_{1,t}^{gap} \\ &\gamma_1^o(\pi_{1,t}^o - \bar{\pi}_1^o) + \gamma_1^w(\omega_{1,t} - \bar{\omega}_1) \end{aligned} \right] + \varepsilon_{1,t}^i.$$

The losses reported are expressed as a percent of steady-state consumption. The inflation measures are annualized.

²For these rules, headline inflation replaces core inflation.

³For this rule, GDP growth replaces the output gap and the coefficients are re-optimized.

⁴For this rule, the output gap is excluded and the coefficients are re-optimized.

large. The exclusion of wage inflation and the output gap from this class of rules causes substantial welfare losses amounting to 2.4 percent of steady-state consumption.

In the context of our model, reacting to GDP growth is unattractive. Unlike the output gap measure, GDP growth does not distinguish between efficient and inefficient fluctuations in activity. For instance, in the face of a positive productivity shock, a rule that responds to GDP growth would raise policy rates in an attempt to curb the expansion in activity. By contrast, the rule that responds to the output gap would not respond to efficient expansions, but only to the inefficient movements in output associated with the presence of nominal rigidities.

One concern with a rule that implies a response to the unobserved output gap is the difficulty of implementation. To construct the relevant output gap measure in our model, the central bank would have to estimate the complete history of shocks that hit the economy. An alternative is the last rule presented in Table 6 which replaces the output gap by wage inflation. The optimized coefficient on wage inflation is 3.65×10^5 and the optimized coefficients on the other terms are close to zero. This last rule achieves complete stabilization of wage inflation and comes close to stabilizing the output gap. It implies only modest welfare losses relative to the optimized rule and is easily implementable in practice.

IV. Conclusion

Even large and relatively closed economies like the United States and the euro area import a sizable fraction of the oil they consume. That fraction is close to half for the United States, while local production is close to nil in the euro area. Nonetheless, most of the existing analyses on the optimal design of monetary policy in the face of fluctuations in oil prices posit an autarkic environment and a simplistic stochastic structure, especially in modeling the demand side of the global crude oil market. Our analysis relaxed those assumptions.

Based on an estimated two-bloc DSGE model that encompasses trade in oil and non-oil goods, we showed how the evolution of inflation and real output and hence the conduct of monetary policy is influenced by a large variety of structural shocks that move both the real price of oil and the global economy. Our analysis highlights that the distinction between oil demand shocks and other structural shocks in the macroeconomy becomes moot once it is recognized that structural shocks simultaneously cause fluctuations in macroeconomic aggregates and in the real price of oil.

First, we showed that the labor market, in the presence of price and wage rigidities, provides a key contribution to the persistence of inflation in the face of the shocks that cause oil price fluctuations. Focusing on an estimated interest rate policy rule, we quantified how vastly different the response of policy rates can be depending on the source of the shocks, even conditioning on the same observable change in the real price of oil on impact.

Second, we constructed welfare optimal policy rules within a class of interest rate reaction functions and showed that the optimal policy responses to a given structural shock differ substantially from the responses implied by the estimated policy rule based on historical data. In particular, we showed that the monetary policy rule estimated for the United States puts a larger weight on interest rate smoothing and on inflation than the optimal rule. Under a policy rule that is optimal, in that its coefficients have been chosen to maximize social welfare, policymakers respond aggressively to the output gap. Responding aggressively to the output gap not only stabilizes output about its potential trend, but it also reduces the volatility of price and wage inflation relative to the estimated rule.

The optimized rule places a large weight on a model consistent, but unobserved output gap. When the output gap is replaced with the deviation of GDP growth from its trend, the optimized rule places virtually no weight on this output measure and a larger weight on core inflation. This rule cannot be recommended because it involves a large reduction in welfare. In contrast, a rule that excludes the output gap, but allows a response to wage inflation, implies welfare almost as high as rules based on the output gap, but has the advantage of being easily implementable. Furthermore, in the wake of fluctuations in oil prices, stabilizing wage inflation fosters the stabilization of core inflation.

The model used in this paper can be viewed as a stylized representation of the key players in global oil markets in recent years. Undoubtedly, our analysis could be refined further. Some dimensions in which the model may be lacking, include the specification of oil production decisions, the absence of valuation effects and the absence of speculative elements in the real price of oil. One also might break down the rest of the world further into distinct blocks of countries such as OPEC, OECD (other than the United States) and the emerging economies. Such extensions are not trivial given the paucity of global data. Moreover, it would be useful to extend our model to focus more directly on monetary policy decisions in oil-exporting countries. This extension is likely to require a more careful modeling of fiscal policy, however. Despite these potential limitations, our analysis constitutes the first formal study of monetary policy responses to oil price fluctuations in an open economy with endogenous oil prices. It also provides a benchmark for more refined models for policy analysis to be developed in the future.

REFERENCES

- Almoguera, P.A., C.C. Douglas, and A.M. Herrera, (2011), "Testing for the Cartel in OPEC: Noncooperative Collusion or Just Noncooperative?" *Oxford Review of Economic Policy*, Vol. 27, No. 1, pp. 144–68.
- Alquist, R., and L. Kilian, (2010), "What Do We Learn from the Price of Crude Oil Futures?" *Journal of Applied Econometrics*, Vol. 25, No. 4, pp. 539–73.
- Backus, D., and M. Crucini, (1998), "Oil Prices and the Terms of Trade," *Journal of International Economics*, Vol. 50, No. 1, pp. 185–213.
- Balke, N.S., S.P.A. Brown, and M.K. Yücel, (2010), "Oil Price Shocks and U.S. Economic Activity: An International Perspective," Discussion Paper No. 1037, Resources for the Future.
- Barsky, R.B., and L. Kilian, (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in *NBER Macroeconomics Annual 2001*, ed. by B.S. Bernanke and K.S. Rogoff (Cambridge: MIT Press), Vol. 16, pp. 137–183.
- Bernanke, B.S., M. Gertler, and M.W. Watson, (1997), "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, Vol. 1, pp. 91–142.
- Blanchard, O.J., and J. Gali, (2010), "The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s So Different from the 1970s?," in *International Dimensions of*

- Monetary Policy*, ed. by J. Gal and M. Gertler (Chicago: University of Chicago Press), pp. 373–421.
- Bodenstein, M., C.J. Erceg, and L. Guerrieri, (2008), “Optimal Monetary Policy with Distinct Core and Headline Inflation Rates,” *Journal of Monetary Economics*, Vol. 55, pp. S18–33.
- , (2011), “Oil Shocks and External Adjustment,” *Journal of International Economics*, Vol. 83, No. 2, pp. 168–84.
- Bodenstein, M., L. Guerrieri, and C.J. Gust, (2010), “Oil Shocks and the Zero Bound on Nominal Interest Rates,” *Journal of International Money and Finance*, forthcoming.
- Bodenstein, M., and L. Guerrieri, (2011), “Oil Efficiency, Demand and Prices: A Tale of Ups and Downs,” International Finance Discussion Papers No. 1031, Board of Governors of the Federal Reserve System.
- Bruno, M., and J. Sachs, (1985), *Economics of World-Wide Stagflation* (Cambridge, MA: Harvard University Press).
- Christiano, L.J., M. Eichenbaum, and C.L. Evans, (2005), “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, Vol. 113, No. 1, pp. 1–45.
- Carlstrom, C.T., and T.S. Fuerst, (2006), “Oil Prices, Monetary Policy, and Counterfactual Experiments,” *Journal of Money, Credit and Banking*, Vol. 38, No. 7, pp. 1945–58.
- Dhawan, R., and K. Jeske, (2007), *Taylor Rules with Headline Inflation: A Bad Idea*, Working Paper 2007-14, Federal Reserve Bank of Atlanta.
- Erceg, C., L. Guerrieri, and S.B. Kamin, (2011), “Did Easy Money in the Dollar Bloc Fuel the Oil Price Run-Up?” *International Journal of Central Banking*, Vol. 7, No. 1, pp. 131–60.
- Fattouh, B., L. Kilian, and L. Mahadeva, (2012), “The Role of Speculation in Oil Markets: What Have We Learned So Far?” mimeo, University of Michigan.
- Galí, J., and M. Gertler, (1999), “Inflation Dynamics: A Structural Econometric Analysis,” *Journal of Monetary Economics*, Vol. 44, No. 2, pp. 195–222.
- Kilian, L., (2009), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” *American Economic Review*, Vol. 99, No. 3, pp. 1053–69.
- , (2010), “Oil Price Shocks, Monetary Policy and Stagflation,” in *Inflation in an Era of Relative Price Shocks*, ed. by R. Fry, C. Jones, and C. Kent (Reserve Bank of Australia, Sydney, Australia), pp. 60–84.
- Kilian, L., A. Rebucci, and N. Spatafora, (2009), “Oil Shocks and External Balances,” *Journal of International Economics*, Vol. 77, No. 2, pp. 181–94.
- Kilian, L., and B. Hicks, (2011), “Did Unexpectedly Strong Economic Growth Cause the Oil Price Shocks of 2003–2008?” *Journal of Forecasting* (forthcoming).
- Kilian, L., and D.P. Murphy, (2010), “The Role of Inventories and Speculative Trading in the Global Market for Crude Oil,” mimeo, University of Michigan.
- , (2012), “Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models,” *Journal of the European Economic Association*, Vol. 10, No. 5, pp. 1166–88.
- Kilian, L., and L.T. Lewis, (2011), “Does the Fed Respond to Oil Price Shocks?” *Economic Journal*, Vol. 121, No. 555, pp. 1047–72.
- Kormilitsina, A., (2011), “Oil Price Shocks and the Optimality of Monetary Policy,” *Review of Economic Dynamics*, Vol. 14, No. 1, pp. 199–223.

- Leduc, S., and K. Sill, (2004), "A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns," *Journal of Monetary Economics*, Vol. 51, No. 4, pp. 781–808.
- Levin, A., A. Onatski, J.C. Williams, and N. Williams, (2006), "Monetary Policy Under Uncertainty in Micro-Founded Macroeconometric Models," in *NBER Macroeconomics Annual 2005*, ed. by M. Gertler and K. Rogoff (Cambridge: MIT Press), pp. 229–312.
- Monacelli, T., (2012), "Is Monetary Policy in an Open Economy Fundamentally Different?," CEPR Discussion Paper 9087.
- Montoro, C., (2010), "Oil Shocks and Optimal Monetary Policy," Working Paper No. 307, Bank for International Settlements.
- Nakov, A., and A. Pescatori, (2010a), "Monetary Policy Trade-Offs with a Dominant Oil Producer," *Journal of Money, Credit, and Banking*, Vol. 42, No. 1, pp. 1–32.
- , (2010b), "Oil and the Great Moderation," *Economic Journal*, Vol. 120, No. 543, pp. 131–156.
- Nakov, A., and G. Nuño, (2011), "Saudi Aramco and the Oil Market," Working Paper No. 1354, European Central Bank.
- Natal, J.-M., (2012), "Monetary Policy Response to Oil Price Shocks," *Journal of Money, Credit, and Banking*, Vol. 44, No. 1, pp. 53–101.
- Plante, M., (2009a), "How Should Monetary Policy Respond to Exogenous Changes in the Relative Price of Oil," CAEPR Working Papers 2009-012, Indiana University Bloomington.
- , (2009b), "Exchange Rates, Oil Price Shocks, and Monetary Policy in an Economy with Traded and Non-Traded Goods," CAEPR Working Papers 2009-016, Indiana University Bloomington.
- Rotemberg, J., and M. Woodford, (1996), "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, Vol. 28, No. 4, pp. 549–77.
- Smets, F., and R. Wouters, (2007), "Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, Vol. 97, No. 3, pp. 586–606.
- Smith, J., (2005), "Inscrutable OPEC? Behavioral Tests of the Cartel Hypothesis," *Energy Journal*, Vol. 26, No. 1, pp. 51–d-82.
- Taylor, J., (1993), "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Conference Series on Public Policy*, Vol. 39, pp. 195–214.
- Winkler, R.C., (2009), "Ramsey Monetary Policy, Oil Price Shocks and Welfare," mimeo, Christian-Albrechts-University of Kiel.

Supplementary information accompanies the paper on IMF Staff Papers website (<http://www.palgrave-journals.com/imfer>)