

Oil Shocks and External Adjustment*

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

May 17, 2007

Abstract

This paper investigates how oil price shocks affect the trade balance and terms of trade in a two country DSGE model. Given that oil shocks may exert very different wealth effects on oil importers and exporters, the response of the external sector depends critically on the structure of financial risk-sharing. With incomplete markets, higher oil prices can induce a sizeable deterioration in an oil-importer's terms of trade, and associated improvement in its non-oil trade balance. The magnitude of the terms of trade response hinges on several structural parameters that affect the divergence in wealth effects across oil importers and exporters, including the elasticity of substitution between oil and other inputs, and the discount factor. By contrast, cross-country wealth differences effectively disappear under complete markets, with the implication that oil shocks have essentially no effect on the non-oil terms of trade or on the non-oil trade balance.

Keywords: oil-price shocks, trade, DSGE models

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

* The authors thank Roc Armenter, David Backus, David Bowman, Brett Berger, Chris Carroll, Don Coletti (our discussant), Jon Faust, Joe Gagnon, Enrique Mendoza, Trevor Reeve, and seminar participants at the Bank of Canada, the IMF, the Federal Reserve Board, the Federal Reserve Bank of New York, and Johns Hopkins University. We are particularly grateful to Christopher Gust for many helpful conversations, and to Hilary Croke for excellent research assistance. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

1 Introduction

The tripling of world oil prices over the past few years has stimulated renewed interest in the macroeconomic consequences of oil shocks. From an open economy perspective, questions about how oil price shocks affect trade and the exchange rate seem especially relevant against the backdrop of a large and persistent U.S. trade deficit. In particular, it is of interest to identify the channels that bring about trade adjustment in response to an oil shock, and whether the adjustment process is likely to entail sizeable downward pressure on the dollar.¹

A second reason for analyzing oil shocks is that they may generate large shifts in purchasing power between oil-exporting and oil-importing countries. Because the magnitude of such wealth effects depends on the availability of risk-sharing arrangements across countries, studying oil shocks may prove insightful in assessing how alternative financial market structures may influence macroeconomic outcomes. While this has been an important issue in the academic literature since the influential work of Cole and Obstfeld (1991), the literature is somewhat ambiguous about whether differences in financial market structure have pronounced effects. This ambiguity may reflect a nearly exclusive focus on technology shocks, which may generate smaller wealth differences across countries than plausibly-calibrated oil shocks.²

¹A large body of research has attempted to isolate empirically the effects of energy supply disruptions on real output and prices; examples include Hamilton (1983), Hamilton (2003), Hamilton and Herrera (2004). Bernanke, Gertler, and Watson (1997) pointed out that the economy's reaction to oil shocks could be critically influenced by the monetary policy reaction. In a recent working paper Kilian, Rebucci, and Spatafora (2007) used a structural VAR to analyze the reaction of the oil and non-oil trade balance to oil demand and supply shocks.

²Baxter and Crucini (1995) found substantial differences in macroeconomic responses to a persistent country-specific technology shock when comparing complete and incomplete markets, though differences were small in the case of temporary shocks. Schmitt-Grohe and Uribe (2003) concluded that responses under incomplete markets looked nearly the same as under complete markets, at least for preferences that precluded significant wealth effects on labor supply. Finally, Corsetti, Dedola, and Leduc (2004) found considerable divergence across

In this paper, we investigate how a persistent rise in oil prices affects trade and the real exchange rate in the context of a two country DSGE model. The general equilibrium framework is useful in helping to identify structural features of the economy that influence the response of trade and the exchange rate to the oil shock. Moreover, it provides a natural framework for assessing the implications of the two alternative financial market structures of complete and incomplete markets that are often contrasted in the literature.

Our modeling framework builds on the open economy DSGE model of Backus and Crucini (1998).³ In our model, each country produces a distinct tradable output good which may be used as an intermediate input in the production of consumption or investment goods. Oil serves as both a factor of production, and is consumed directly in the household consumption bundle (with each production function having a constant elasticity of substitution form). One country is an oil-importer, reflecting its relatively low endowment of energy, while the other is an oil-exporter.

We generalize the Backus and Crucini model by allowing for convex costs of adjusting the share of oil used in production and consumption. Adjustment costs provide a tractable way of capturing the putty-clay nature of oil demand, and allow our model to account for the substantial wedge between the short and long-run elasticity of oil demand that appears to be a robust finding of the empirical literature. Finally, we depart from the complete markets framework of Backus and Crucini insofar as we assume that financial markets are incomplete across national borders

the alternative financial structures when the price elasticity of import and export demand was sufficiently low, i.e., well below unity, but much less divergence for higher elasticities (in the range of unity or higher). Taken together, these results suggest that there are at least plausible conditions (fairly low Frisch elasticities, and trade price elasticities near unity) under which a technology shock has similar effects under either financial market structure, even if much more divergence can occur under alternative calibrations.

³An extensive literature has analyzed the effects of oil shocks in a closed economy DSGE setting, including Kim and Loungani (1992), Rotemberg and Woodford (1996), and Finn (2000). A survey of this literature is provided by Jones, Leiby, and Paik (2004).

in our benchmark specification.⁴

Under incomplete financial markets, we find that the effects of an oil shock on the nonoil terms of trade and the overall trade balance depend on the magnitude of the divergence in wealth effects between the oil-importing and oil-exporting country blocks. Specifically, our simulations focus on the implications of a highly persistent increase in the oil price (generated by an adverse supply shock in the oil-exporting country). Under our benchmark calibration, the oil importing country experiences a highly persistent deterioration in the oil component of its trade balance. The only way for the oil-importer to satisfy its intertemporal trade balance condition is for its non-oil component to improve by a sufficient amount, which requires some initial worsening of its nonoil terms of trade (or real exchange rate, as the latter adjusts proportionately).

Several structural factors play a major role in determining the relative wealth effects across countries by influencing the magnitude of transfers from oil-importers to exporters, their persistence, and how they are discounted. Notably, a low time-path of the oil price elasticity of demand implies that the oil importer runs more persistent deficits on the oil component of its trade balance, while the oil-exporter experiences a more persistent revenue windfall. This magnifies the difference in wealth effects across countries, and hence require a much larger depreciation of the oil-importer's real exchange rate in order to satisfy the intertemporal current account balance condition. Moreover, we show that the depreciation is larger when the steady state real interest rate is higher, and the oil importer has a smaller oil endowment.

We then proceed to show that the responses of the trade balance and real exchange rate are dramatically different under complete markets. In response to an oil

⁴Backus and Crucini used their model to explain instability in the time series correlation between output and the terms of trade that they attributed to changes in the relative importance of oil and technology shocks. Given this focus, their results would seem likely to hold under either financial market structure.

price hike, the oil importer receives an insurance transfer that enables it to satisfy its intertemporal current account balance constraint without having to run an eventual non-oil trade surplus. Thus, oil shocks have essentially no effect on the non-oil terms of trade or non-oil trade balance. In addition, the structural parameters identified above as markedly affecting cross-country wealth effects under incomplete markets have essentially no impact on the terms of trade or the non-oil balance.

Finally, we conduct sensitivity analysis to the particular type of price shock assumed, and to our benchmark assumption that wages and prices are fully flexible. While our benchmark model assumes that the adverse supply shock pushes up the oil price permanently – under the assumption that foreign exporter reduces supply enough to support the higher price – our results are quite similar under the assumption that the price rises endogenously in response to an exogenous contraction in foreign supply. We also show that although nominal rigidities in prices and wages affect the path of domestic absorption in the short-run, their inclusion has little consequence for the trade and real exchange rate responses that are our principal focus.

The remainder of the paper is organized as follows. Section 2 describes our benchmark model, and Section 3 the calibration. Section 4 provides simulation results for our benchmark model under incomplete markets, while Section 5 considers complete markets. Finally, Section 6 assesses the sensitivity of our results to an alternative type of oil price shock, and Section 7 to the inclusion of nominal rigidities. Section 8 concludes.

2 Model Description

2.1 Model Overview

Our model is comprised of two countries, a home country (calibrated based on U.S. data) and a foreign country (rest-of-the-world). The structure of each country

block is symmetric, although our calibration allows for differences in population size and in the per capita oil endowment. Each country specializes in the production of a final good that is an imperfect substitute for the final good produced in the other country. Production requires capital, labor, and oil. The consumption bundle entering the household utility function depends on consumption of the domestically-produced good, on imports of the foreign good, and oil. However, it is convenient for expositional purposes to assume that this composite consumption bundle is simply produced by a competitive distribution sector (with a productive structure that mirrors household preferences over the three goods). While asset markets are complete at the country level, we assume that asset markets are incomplete internationally. Finally, both the United States and foreign country are endowed with a non-storable flow supply of oil each period.

2.2 Households

The utility functional of a typical member of household h is

$$\mathbb{E}_t \sum_{j=0}^{\infty} \beta^j \left\{ \frac{1}{1-\sigma} \left(C_{t+j}(h) - \varkappa \frac{C_{t+j-1}}{\zeta} \right)^{1-\sigma} + \frac{\chi_0}{1-\chi} (1 - N_{t+j}(h))^{1-\chi} \right\}, \quad (1)$$

where the discount factor β satisfies $0 < \beta < 1$. As in Smets and Wouters (2003), we allow for the possibility of external habits in consumption, so that the utility that each member of household h derives from its consumption flow $C_t(h)$ depends on the level of aggregate consumption in the previous period. The term ζ is a scale factor determining population size. The period utility function also depends on an individual's current leisure $1 - N_t(h)$.

Each member of household h faces a flow budget constraint in period t which states that his combined expenditure on goods and on the net accumulation of

financial assets must equal his disposable income:

$$\begin{aligned} & P_{Ct}C_t(h) + P_{It}I_t(h) + \frac{e_t P_{Bt}^* B_{Ft+1}(h)}{\phi_{Bt}} - e_t B_{Ft}(h) \\ &= W_t(h) N_t(h) + R_{Kt}K_t(h) + \Gamma_t(h) + TR_t(h) - T_t(h) - P_{Dt}\phi_{It}(h). \end{aligned} \quad (2)$$

Final consumption goods are purchased at the price P_{Ct} , and final investment goods at the price P_{It} . Investment in physical capital augments the per capita capital stock $K_{t+1}(h)$ according to a linear transition law of the form:

$$K_{t+1}(h) = (1 - \delta)K_t(h) + I_t(h), \quad (3)$$

where δ is the depreciation rate of capital.

Households accumulate financial assets by purchasing state-contingent domestic bonds, and a non state-contingent foreign bond (but given the representative agent structure, we omit terms involving the former from the household's budget constraint). Thus, the term $B_{Ft+1}(h)$ in the budget constraint represents the quantity of the non-state contingent bond purchased by a typical member of household h at time t that pays one unit of foreign currency in the subsequent period, P_{Bt}^* is the foreign currency price of the bond, and e_t is the exchange rate expressed in units of home currency per unit of foreign currency. To ensure that net foreign assets are stationary, we follow Turnovsky (1985) and assume there is an intermediation cost ϕ_{Bt} paid by households in the home country for purchases of foreign bonds. Specifically, the intermediation costs depend on the ratio of economy-wide holdings of net foreign assets to nominal output and are given by:

$$\phi_{Bt} = \exp\left(-\phi_b \left(\frac{e_t B_{Ft+1}}{P_{Dt} Y_t}\right)\right). \quad (4)$$

If the home economy has an overall net lender position internationally, then a household will earn a lower return on any holdings of foreign bonds. By contrast, if the economy has a net debtor position, a household will pay a higher return on any foreign debt.

Each member of household h earns labor income $W_t(h) N_t(h)$ and capital income $R_{Kt}K_t(h)$. Each member also receives an aliquot share $\Gamma_t(h)$ of firm profits, and

receives net transfers of $T_t(h)$. Finally, as in Christiano, Eichenbaum, and Evans (2005), it is costly to change the level of gross investment from the previous period, so that the acceleration in the capital stock is penalized:

$$\phi_{It}(h) = \frac{1}{2}\phi_i \frac{(I_t(h) - I_{t-1}(h))^2}{I_{t-1}(h)}. \quad (5)$$

Because this cost is measured in units of the domestically-produced good P_{Dt} (defined below), the adjustment cost term is premultiplied by this nominal price in the household's budget constraint.

In every period t , household h maximizes the utility functional (1) with respect to its consumption, labor supply, investment, (end-of-period) capital stock, and holdings of foreign bonds, subject to its budget constraint (2), and the transition equation for capital (3). In doing so, a household takes as given prices, taxes and transfers, and aggregate quantities such as lagged aggregate consumption.

2.3 Firms and Production

Each country produces a single distinct nonoil output good. Focusing on the home country, this output good is produced by perfectly competitive firms according to a constant returns to scale technology. The representative firm's technology can be characterized as a nested constant-elasticity of substitution (CES) specification of the form:

$$V_t = \left(\omega_k^{\frac{\rho_v}{1+\rho_v}} K_t^{\frac{1}{1+\rho_v}} + L_t^{\frac{1}{1+\rho_v}} \right)^{1+\rho_v}, \quad (6)$$

$$Y_t = \left((1 - \omega_{oy})^{\frac{\rho_{oy}}{1+\rho_{oy}}} V_t^{\frac{1}{1+\rho_{oy}}} + \omega_{oy}^{\frac{\rho_{oy}}{1+\rho_{oy}}} (\varphi_{OYt} O_{Yt})^{\frac{1}{1+\rho_{oy}}} \right)^{1+\rho_{oy}}. \quad (7)$$

Each producer utilizes capital and labor services, K_t and L_t , to make a "value-added" input V_t . This composite input is combined with oil O_{Yt} to produce the domestic output good Y_t . The factor φ_{OYt} reflects costs of adjusting the oil intensity

of the final output good, and is assumed to take the following quadratic form:

$$\varphi_{OYt} = \left[1 - \frac{\varphi_{oy}}{2} \left(\frac{O_{Yt}}{V_t} - 1 \right)^2 \right], \quad (8)$$

where $O_{Y_{t-1}}^A$ and V_{t-1}^A denote oil inputs and value added aggregated over all firms. From an aggregate perspective, this specification allows for oil use to respond quickly to gross domestic output, while potentially allowing for very slow adjustment to relative price changes.⁵

Given the presence of adjustment costs, the representative producer can be regarded as choosing a contingency plan for K_t , L_t , and O_{Yt} that minimizes the discounted expected cost of producing the domestic output good subject to equations (6)-(7). In solving this problem, the producer takes as given the rental price of capital R_{Kt} , the wage W_t , and the after-tax price of oil P_{Ot}^{Tax} . The relation between the after-tax oil price and pre-tax price P_{Ot} can be expressed:

$$P_{Ot}^{Tax} = (1 + \tau_{ov})P_{Ot} + \tau_{os}P_{Dt} \quad (9)$$

where τ_{ov} is an ad-valorem tax on oil (i.e., a tax that is a constant percentage of the oil price), and τ_{os} is a specific tax (i.e., specified in per-gallon terms, and assumed to be indexed to P_{Dt}). The representative firm sells its output to households and firms at a price P_{Dt} , which is the Lagrange multiplier from the cost-minimization problem.

Production of Consumption and Investment Goods

The consumption basket C_t that enters the household's budget constraint can be regarded as produced by perfectly competitive consumption distributors. These distributors purchase a non-oil consumption good C_{Nt} (described below) and oil

⁵This form of adjustment cost captures the spirit of a putty-clay model such as in Atkeson and Kehoe (1999) and Wei (2003). In their setting, a large variety of types of capital goods are combined with energy in different fixed proportions. Thus the short-run elasticity of substitution for oil is low. In the long run, the elasticity is higher, as firms can invest in capital goods with different fixed energy intensities.

O_{Ct} as inputs in perfectly competitive input markets, and produce the composite consumption good according to a CES production function:

$$C_t = \left((1 - \omega_{oc})^{\frac{\rho_{oc}}{1+\rho_{oc}}} C_{Nt}^{\frac{1}{1+\rho_{oc}}} + \omega_{oc}^{\frac{\rho_{oc}}{1+\rho_{oc}}} (\varphi_{OCt} O_{Ct})^{\frac{1}{1+\rho_{oc}}} \right)^{1+\rho_{oc}} \quad (10)$$

where the quasi-share parameter ω_{oc} determines the importance of oil purchases in the household's composite consumption bundle, and the parameter ρ_{oc} determines the long-run price elasticity of demand for oil. The form of the production function mirrors the preferences of households over consumption of non-oil goods and oil. The term φ_{OCt} captures costs of adjusting the oil consumption (at the household level), and is assumed to have the quadratic form:

$$\varphi_{OCt} = \left[1 - \frac{\varphi_{oc}}{2} \left(\frac{\frac{O_{Ct}}{C_{Nt}}}{\frac{O_{Ct-1}^A}{C_{Nt-1}^A}} - 1 \right)^2 \right], \quad (11)$$

where O_{Ct-1}^A and C_{Nt-1}^A denote, respectively, oil consumption and non-oil consumption aggregated over all households.

Thus, households are similar to firms in that they adjust their oil demand slowly in response to changes in the relative price of oil. As in the case of firms, the presence of adjustment costs implies that the consumption distributors must solve a dynamic problem of choosing a contingency path for their inputs C_{Nt} and O_{Ct} so as to minimize their discounted expected costs of producing the consumption bundle, taking as given (after-tax) input prices P_{CNt} and P_{Ot}^{Tax} , respectively (n.b., the after-tax price of oil is assumed to be the same for households and firms). The Lagrangian from this cost-minimization problem determines the price of the consumption bundle they charge to households, i.e., P_{Ct} in the household's budget constraint given in equation (2).

Similarly, it is also convenient to regard the nonoil consumption good C_{Nt} and investment good I_t as produced by perfectly competitive distributors. Both the domestically-produced good and the foreign (nonoil)good are utilized as inputs,

though we allow for the proportion of each input to differ between nonoil consumption and investment goods. Thus, the production function for the nonoil consumption good C_{Nt} is given by:

$$C_{Nt} = \left((1 - \omega_{mc})^{\frac{\rho_c}{1+\rho_c}} C_{Dt}^{\frac{1}{1+\rho_c}} + \omega_{mc}^{\frac{\rho_c}{1+\rho_c}} (\varphi_{MCt} M_{Ct})^{\frac{1}{1+\rho_c}} \right)^{1+\rho_c}, \quad (12)$$

where C_{Dt} denotes the quantity of domestically-produced goods used as inputs by the representative nonoil consumption distributor (purchased at a price of P_{Dt}), M_{Ct} denotes imports of the foreign good (purchased at a price of P_{Mt}), and φ_{MCt} represents a cost of adjusting imports. This adjustment cost takes a quadratic form:

$$\varphi_{MCt} = \left[1 - \frac{\varphi_{mc}}{2} \left(\frac{\frac{M_{Ct}}{C_{Dt}}}{\frac{M_{Ct-1}^A}{C_{Dt-1}^A}} - 1 \right)^2 \right], \quad (13)$$

where M_{Ct-1}^A and C_{Dt-1}^A denote (nonoil) goods consumption imports and domestic consumption inputs aggregated over all households. This adjustment cost allows goods imports to adjust slowly to the terms of trade. The Lagrangian from this cost-minimization problem determines the price of the nonoil consumption good P_{CNt} .

Finally, the production function for investment goods is isomorphic to that given in equation (12), though allowing for possible differences in the import intensity of investment goods (determined by ω_{mi} , akin to ω_{mc} in equation (12)), the degree of substitutability between nonoil imports and domestically-produced goods in producing investment goods (determined by ρ_i), and the parameter determining import adjustment costs φ_{mi} . As expected, the presence of adjustment costs means that non-oil consumption and investment goods distributors must solve a dynamic cost minimization problem (with the Lagrangian from the problem facing investment distributors determining the price of new investment goods P_{It} that appears in the household's budget constraint).

2.4 The Oil Market

We assume that the home country is simply endowed with an exogenous flow supply of oil each period of Y_O . In our benchmark model, we assume that the pre-tax price of oil relative to the domestic output good ψ_t is determined by an exogenous stochastic process. With both domestic oil supply and the oil price determined exogenously, foreign oil production $Y_{O_t}^*$ must adjust endogenously to clear the world oil market:

$$Y_O + Y_{O_t}^* = O_{Y_t} + O_{C_t} + O_{Y_t}^* + O_{C_t}^*. \quad (14)$$

Thus, the sum of the home and foreign oil production equals the sum of home and foreign consumption (by firms and households). The relative price of oil is assumed to follow a simple AR(1) of the form:

$$\log(\psi_t) = \rho \log(\psi_{t-1}) + \epsilon_t. \quad (15)$$

The persistence parameter ρ is set arbitrarily close to one, so that the log of the relative price of oil effectively follows a random walk.

In Section 6, we conduct sensitivity analysis in an alternative framework in which the oil price is determined endogenously to equate supply and demand in the world oil market. In this case, the foreign production level $Y_{O_t}^*$ in equation (14) is assumed to follow an exogenous stochastic process (as discussed below, an AR(1) in the growth rate).

2.5 Fiscal Policy

We assume that a fixed share g of the domestic output good Y_{DT} is purchased by the government (n.b., the import content of government purchases is zero). These government purchases G_t have no direct effect on household utility. Given the availability of lump-sum taxes and the Ricardian structure of our model, it is convenient to assume that lump-sum transfers T_t are adjusted each period to balance the government receipts and revenues, so that:

$$P_{Dt}G_t + T_t = \tau_{ov}P_{Ot}(O_{Yt} + O_{Ct}) + \tau_{os}P_{Dt}(O_{Yt} + O_{Ct}) \quad (16)$$

2.6 Resource Constraints for Nonoil Goods, and Net Foreign Assets

The resource constraint for the goods sector of the home economy can be written as:

$$Y_t = C_{Dt} + I_{Dt} + G_t + M_t^* + \phi_{It}, \quad (17)$$

recalling that M_t^* denotes the home country's exports, and ϕ_{It} resources that are lost due to costs of adjusting investment.

The evolution of net foreign assets can be expressed as:

$$\frac{e_t P_{Bt}^* B_{Ft+1}}{\phi_{Bt}} = e_t B_{Ft} + e_t P_{Mt} M_t^* - P_{Mt} M_t + P_{Ot} (O_{Yt} + O_{Ct} - Y_{Ot}). \quad (18)$$

This expression can be derived from the budget constraint of the households after imposing the government budget constraint and the definition of firm profits.

3 Solution Method and Calibration

We solve the model by log-linearizing the equations around the model's steady state. To obtain the reduced-form solution of the model, we use the numerical algorithm of Anderson and Moore (1985), which provides an efficient implementation of the method proposed by Blanchard and Kahn (1980) (see also Anderson (1997)).⁶

⁶The steady state around which we linearize depends on the relative level of technology in each country, which we initialize to unity (so that per capita income in each country is identical in the steady state, though GDP may differ across countries due to population differences). We evaluated the robustness of our solution procedure by using a nonlinear Newton-Raphson algorithm that does not rely on linearization around an initial steady state, and found that the results were nearly identical to those reported.

The model is calibrated at a quarterly frequency. The parameter values for the home economy under our benchmark calibration are listed in Table 1; parameters for the foreign economy are identical except for the trade share parameters (which are determined by balanced trade given relative population sizes).

The discount factor β is 0.99. The utility functional parameter σ is set equal to 1. The parameter determining the degree of habit persistence in consumption \varkappa is set to 0.8. We choose $\chi = 10$, implying a Frisch elasticity of labor supply of $1/5$, which is considerably lower than if preferences were logarithmic in leisure, but well within the range of most empirical estimates. The utility parameter χ_0 is set so that employment comprises one-third of the household's time endowment. The population size parameter ζ is set so that U.S. output comprises one fourth of world output.

The production function parameter ρ_v is set to -2, implying an elasticity of substitution between capital and labor of $1/2$. Thus, capital and labor are less substitutable than the unitary elasticity case implied by the Cobb-Douglas specification. We set the parameter for the cost of adjusting investment, $\phi_i = 4$, close to the value used by Christiano, Eichenbaum, and Evans (2005). The depreciation rate of capital $\delta = 0.025$ is consistent with an annual depreciation rate of 10 percent. Given that we set the government share of output to 0.18, the quasi-capital share parameter ω_k is chosen so that the U.S. investment to output equals an empirically-realistic value of 0.20.

Our calibration of ω_{oy} and ω_{oc} is determined by the overall oil share of output, oil end-use ratios in consumption and production, and the level of oil taxes. The overall oil share of output is set equal to 0.03 for the home country. From end-use data provided by the Energy Information Administration we set the share of oil demand in consumption and production to $1/3$ and $2/3$, respectively. The level of oil imports for the home country is set at 50% of total demand. These estimates are based on data for the United States obtained from British Petroleum, and reflect

expenditures and local production of both oil and natural gas for 2003.⁷ Ad-valorem taxes on oil are set at 10 percent, while specific oil taxes are set at 20 percent. These values were chosen using data on oil taxes in the U.S. from OPEC and from the International Energy Agency.

Our choices for the parameters that determine the short and long-run elasticity of substitution for oil (ρ_{oy} , ρ_{oc} , φ_{oy} and φ_{oc}) are motivated by the regression results summarized in Appendix A. The parameter ρ_{oy} is chosen so that the elasticity of substitution between oil and the other factors of production is 0.5. Similarly, the parameter ρ_{oc} is chosen so that the elasticity of substitution between oil and other inputs into the households' consumption basket is 0.5. The adjustment cost parameters φ_{oy} and φ_{oc} are set so that in reaction to a permanent oil price increase the half life of the response of oil demand is 10 years. Our estimates for the short-run and long-run price elasticities for oil demand are within the wide range of estimates reported in surveys such as Dahl and Sterner (1991). Also, as in most studies we find a markedly lower short-run elasticities than long-run elasticities.⁸

Turning to the parameters determining trade flows, the parameter ω_{mc} is chosen to match the estimated average share of imports in total U.S. consumption of about 9 percent (using NIPA data), while the parameter ω_{mi} is chosen to match the average share of imports in total U.S. investment of about 38 percent. This parameterization implies a nonoil-goods import-to-GDP ratio for the home country of about 12 percent. Given that trade is balanced in steady state, and that the oil import share for the home country is 1.5 percent of GDP, the goods export share is 13.5 percent of GDP.

We assume that $\rho_c = \rho_i = 2$, consistent with a long-run price elasticity of

⁷This calibration choice reflects the close substitutability of oil and natural gas as energy inputs, and the high correlation between movements in oil prices and natural gas prices. Over the period 1970-2005 we computed the correlation between oil crude prices and gas prices to be 0.9 (using data from NYMEX).

⁸n.b., in our framework the elasticity of substitution between oil and the other factors of production coincides with the long-run price elasticity of demand for oil.

demand for imported consumption and investment goods of 1.5.⁹ While this is higher than most empirical estimates using macro data, we emphasize that the presence of adjustment costs translates into a lower relative price sensitivity in the short-run. In particular, we set the adjustment cost parameters $\varphi_{mc} = \varphi_{mi} = 10$, implying a price-elasticity still well below unity after a year.¹⁰ We choose a small value (0.0001) for the financial intermediation cost ϕ_b , which is necessary to ensure the model has a unique steady state.

The calibration of most parameters in the foreign country matches the calibration for the home country, with the exception of the parameters governing relative output sizes and trade flows. The foreign population size parameter is set so that foreign output comprises three fourth of world output. The trade parameters are adjusted to imply balanced trade in steady state with the home country.

4 Model Simulations

4.1 Benchmark Calibration

Figure 1 shows responses of the home country (the United States) to a 50 percent permanent rise in the relative price of oil under our benchmark calibration. The shock induces a progressive fall in U.S. oil demand as both households and firms substitute away from this more costly input (upper left panel). With a pre-tax oil

⁹Hooper and Marquez (1995) surveyed trade elasticity estimates obtained from aggregate data. They reported import price elasticities in a range of 1 to 1.5 for the United States. McDaniel and Balistreri (2003) gave a good summary of the industry-level estimates of trade elasticities. They concluded that the estimated values increase with the level of data disaggregation

¹⁰Hooper, Johnson, and Marquez (2000) found significant differences between short-run and long-run trade elasticities for the United States when using aggregate trade data. The survey of industry-level estimates by McDaniel and Balistreri (2003) argued that available studies provide definitive evidence that long-run estimates are higher than short-run. Gallaway, McDaniel, and Rivera (2003) provided long-run estimates for 309 industries that range from 0.52 to 4.83; their short-run estimates were up to five times smaller.

elasticity of demand of one-half and roughly 80 percent passthrough of the pretax oil price to the aftertax price, oil demand drops roughly 20 percent in the long-run.¹¹ As noted above, our calibration implies that the half-life of the adjustment to this new long-run level is about 10 years.

The progressive decline in oil use has effects on output, the expenditure components, and the real interest rate that are qualitatively similar to those of a persistent decline in productivity growth. Thus, gross output declines gradually in response to lower oil use and a falling capital stock, while consumption contracts due to a reduction in household permanent income (however, in contrast to the case of a productivity shock, the long-run consumption decline exceeds that of gross output, since energy imports rise when expressed in units of the domestic good). Investment spending falls reflecting that a lower path of oil usage pushes down the current and future marginal product of capital. The decline in the marginal productivity of capital induces a modest fall in domestic real interest rates (notwithstanding a transient initial rise due to habit persistence in consumption).

Turning to the implications for the external sector, the rise in oil prices causes a sharp and immediate deterioration of the overall (nominal) trade balance equal to about 1 percentage point of GDP. In the short-run, the nonoil trade balance is unresponsive to the shock, so that the deterioration in the overall nominal trade balance essentially reflects the deterioration in the oil component due to higher prices (with the volume of oil imports nearly fixed). In the longer-term, the overall trade deficit narrows and even shifts into surplus roughly ten years after the occurrence of the shock.

Several forces provide a catalyst for eventual trade balance improvement. First, oil import volumes gradually decline. Importantly, because the United States produces half of its oil demand in the pre-shock steady state and U.S. oil supply is

¹¹While the fall in energy demand due to the direct effect of higher prices is amplified by an endogenous decline in output and consumption, the latter effect is fairly small quantitatively.

assumed to be price inelastic, the decline in oil imports is much larger than the fall in total oil demand in percentage terms: thus, while oil demand falls around 20 percent in the long-run, oil imports fall more than 40 percent. Second, a decline in the domestic terms of trade spurs an improvement in U.S. real nonoil exports, while retarding nonoil imports (as the relative price of imports rises, and exports fall). We provide an extensive discussion of the factors affecting the direction and magnitude of this terms of trade effect below. Finally, the fall in domestic consumption and investment also retards real nonoil imports (although the effect on net exports is somewhat muted due to a decline in foreign investment spending).

From an accounting perspective, adjustments in both the oil and non-oil components of the nominal trade balance play a quantitatively important role in contributing to the eventual improvement in the overall balance. As indicated by Figure 1, the overall trade balance improves about one percentage point of GDP between the immediate aftermath of the shock and the end of the fifteen year horizon shown, and about 2/3 of this improvement is attributable to the oil balance. Nevertheless, given our benchmark calibration of the elasticity of oil demand of 1/2, the oil balance remains in substantial deficit even in the long-run. Thus, intertemporal balance requires that adjustment in the nonoil balance also play a sizeable role, which is achieved largely through the worsening of the home terms of trade.¹²

4.2 The Oil Elasticity of Substitution and the Discount Factor

The magnitude of the effects of oil shocks on the terms of trade, the composition of trade adjustment, and its speed depend critically on the elasticity of substitution between oil and other factor inputs. To illustrate this sensitivity, Figure 2 contrasts

¹²Laffer and Agmon (1978) observed that the trade balance of oil importing countries rebounded relatively quickly following the first oil crisis of the 1970s. Furthermore, the non-oil trade balance played a prominent role in reducing the overall trade gap.

responses under our benchmark calibration to the 50 percent oil price hike with responses derived under two alternative calibrations of this long-run elasticity. One alternative imposes a long-run elasticity of unity, consistent with a Cobb-Douglas production function over the factor inputs (the dashed line, referred to as the “Cobb-Douglas specification” below). A second alternative imposes a long-run elasticity of .05, much lower than our benchmark, and close to the Leontief specification adopted in many simple theoretical models (the dotted line, referred to as the “near Leontief” specification below). In each case, the adjustment cost parameter on oil is changed so that the half life of adjustment to a permanent shock remains 10 years (while other parameters are unchanged).

The upper left panel compares the response of total oil demand under the three specifications. Unsurprisingly, oil demand eventually shows a much larger decline under the Cobb-Douglas specification than under the benchmark. With a pre-tax oil elasticity of demand of unity (and again, roughly 80 percent passthrough of the pre-tax oil price to the after-tax price), oil demand drops roughly 40 percent in the long-run. Under the near-Leontief specification, oil demand shows a fairly small contraction, even in the long-run. Conversely, the contractionary effects on domestic demand are larger when it is difficult to substitute away from oil, reflecting the larger reduction in permanent income, and a greater decline in investment (since the marginal product of capital falls by a more substantial magnitude).

Turning to the implications for the trade balance, the responses under the near-Leontief case are qualitatively the same as under our benchmark. From a quantitative perspective, the much smaller decline in oil use translates into a smaller fall in oil imports, accounting for the much larger deterioration of the nominal oil balance apparent in Figure 2. Given the requirement of intertemporal balance, the permanent deterioration of the oil component of the trade balance in turn requires a larger offsetting improvement of the nonoil component than under our benchmark calibration; thus, the non-oil balance improves over 1 percentage point of GDP after

25 years, roughly two times as much as under the benchmark. This greater improvement of the non-oil balance is brought about by a decline of the domestic terms of trade that is twice as large as under the benchmark. In our general equilibrium framework, the larger terms of trade decline reflects the interplay of the greater shortfall on the domestic oil balance, and the correspondingly greater wealth effect abroad.

The responses under the Cobb-Douglas specification exhibit some striking qualitative differences from specifications with lower long-run oil elasticities. In particular, the domestic terms of trade remain flat, the oil balance moves into eventual surplus, and the non-oil balance essentially shows no response. Mechanically, while the 50 percent rise in the oil price reduces oil demand around 40 percent in the long-run – reflecting that the elasticity of oil demand with respect to the pretax price is somewhat below unity due to specific taxes – this produces a much sharper decline in oil imports. In particular, because the United States meets a substantial fraction of its oil demand with domestic production, and the latter is price-inelastic, the volume of oil imports falls by roughly 70 percent in the longer-run, which generates the eventual improvement in the nominal oil balance seen in Figure 2. The long-run improvement in the oil balance allows the non-oil balance to move into deficit in the longer term, and the domestic terms of trade to even improve slightly. While the terms of trade response in the short-run is fairly sensitive to factors that determine the wealth effect abroad (as we show below), the fact that domestic nominal oil imports eventually fall implies a much smaller wealth effect abroad than in cases with a lower long-run price elasticity. Thus, a permanent oil price hike would not depress the terms of trade in the longer run if the oil price elasticity was as high as unity.

In Figure 3, we examine the sensitivity of our results to alternative parameterizations of the discount factor β . The discount factor affects the responses through two related channels. First, it affects the cost of U.S. external borrowing, which is

relevant given that the shock induces an initial trade deficit that is highly persistent. Thus, a low value of the discount factor implies a high steady state real interest rate, which means that the oil price hike has a more sharply depressing effect on U.S. wealth and consumption due to the higher debt servicing costs; and conversely for a high discount factor. Second, it determines how the foreign economy discounts current and future oil receipts, and thus can markedly affect the wealth effects of the shock abroad. Given that foreign oil revenues decline through time as the U.S. substitutes away from oil, a low value of discount factor magnifies the wealth effects on the foreign country.

These observations help account for sensitivity of the trade balance and terms of trade responses to the discount factor depicted in Figure 3. The discount factor has a negligible impact on the oil balance. Nevertheless, given that U.S. debt-servicing costs are higher for low values of the discount factor, intertemporal balance requires a larger and more rapid improvement in the non-oil balance, and correspondingly greater terms of trade deterioration. Moreover, because a low discount factor implies a larger rise in foreign wealth, these effects on the non-oil balance and terms of trade are amplified. Thus, the U.S. non-oil balance under the low discount factor of $\beta = 0.975$ improves by more than twice as much as under the benchmark calibration, while the deterioration of the U.S. terms of trade is about twice as large. Conversely, the terms of trade and non-oil balance exhibit smaller movements than under the benchmark when the discount factor rises to $\beta = 0.995$.

Our foregoing results indicate that the effects of an oil price hike on terms of trade and the non-oil trade balance can vary substantially – even qualitatively – depending on the elasticity of substitution between oil and other inputs, and on the discount factor. Given parameter uncertainty, it is useful to identify regions of a two dimensional space in these parameters that tend to imply a substantial U.S. terms of trade worsening and associated improvement in the non-oil balance, from regions that generate a terms of trade improvement and deterioration of the U.S.

nonoil balance.

We perform this exercise in the contour plot shown in Figure 4. Each of the isoquants identifies pairs of the parameters that imply the same magnitude of deterioration in the U.S. terms of trade, with the magnitude labeled next to the relevant isoquant (the terms of trade response is an average over the first two years following the 50 percent oil shock). Because the terms of trade response turns out to be closely aligned with a response of the nonoil trade balance of the opposite sign, we focus exclusively on the former. Thus, to the northeast of the “0” isoquant, the U.S. terms of trade actually improve in response to the price hike, and there is a corresponding deterioration of the non-oil trade balance (not shown). This region corresponds to a high long-run degree of substitutability between oil and other factors, and a high value of the discount factor. As suggested above, under these conditions, U.S. oil imports eventually fall enough that the oil component of the U.S. trade balance moves into surplus (relative to its initial steady state), so that foreign oil revenues fall in the long-run. Given that the U.S. faces low borrowing costs, and that foreign wealth effects are small or negative (reflecting a long horizon for discounting the future, and the expectation of an eventual fall in revenue), the U.S. non-oil balance runs small but persistent deficits, and the terms of trade show the slight improvement implied by Figure 4. By contrast, the terms of trade deteriorate progressively as the two key parameters are varied in a southwesterly direction. For example, given that the vertical axis corresponds to a long-run elasticity of .05, the point on the vertical axis corresponding to a discount factor of $\beta = 0.99$ yields the near-Leontief case shown in Figure 2. In this region of the parameter space, the large and persistent increase in the oil deficit requires the non-oil balance to eventually shift into a substantial surplus, which is accomplished through a terms of trade deterioration exceeding 5 percent as shown. With the same low elasticity but a lower discount factor, higher U.S. borrowing costs and a larger foreign wealth effect would amplify the required depreciation in the U.S. terms of trade.

4.3 Domestic Oil Production

In our benchmark calibration, the United States meets nearly one half of its oil demand through domestic production in the pre-shock steady state. Figure 5 considers an alternative in which U.S. production is fixed at zero, while other features of our calibration remain unchanged. In addition to highlighting the implications of domestic production for our results, this alternative is useful for assessing how a higher degree of reliance on imported oil – similar to that of Japan or the euro area – might affect the response of trade flows and the terms of trade.

Two features of Figure 5 are especially noteworthy. First, the deterioration of the terms of trade under this alternative calibration is several times larger than in the benchmark; and second, the U.S. nonoil trade balance moves into a much larger surplus. These differences reflect that the oil price hike induces a much smaller percentage decline in real oil imports in the case in which domestic production is zero (as oil imports simply fall by the same percentage as oil demand), so that there is a much larger deterioration of the oil component of the nominal trade balance. Because intertemporal current account balance in turn requires a bigger improvement in the non-oil balance, the terms of trade must deteriorate by a larger amount.

Thus, our analysis suggests that the fact that the U.S. produces a sizeable fraction of the oil it consumes may help partly insulate the U.S. terms of trade from the effects of oil price shocks. By contrast, a country such as Japan might experience considerably more pressure on its terms of trade.¹³

¹³The empirical analysis of Kilian (2006b) found that an exogenous oil price increase would lead to a depreciation relative to the dollar for Italy, France, Germany, and Japan, all countries whose level of local oil production is lower than for the United States. Conversely, for an oil exporter like Canada, Kilian's results predicted an appreciation of the Canadian dollar. Golub (1983) provided an early example of a model linking the exchange rate response to an oil shock with the level of oil imports.

4.4 Specific Taxes on Energy

Figure 6 investigates the implications of varying the specific tax rate on oil from its benchmark value of $\tau_{os} = 0.2$. In particular, we consider an alternative in which the specific tax rate is set to 0.4, and another case in which it is set equal to zero. Clearly, a higher tax rate induces a larger decline in the terms of trade, and correspondingly, a somewhat larger improvement of the nonoil trade balance.

These effects reflect that the higher specific oil tax damps the long-run elasticity of oil demand with respect to the pre-tax oil price. To see this, note that this elasticity can be expressed as:

$$\rho_{oy} \left[\frac{1 + \tau_{ov}}{1 + \tau_{ov} + \tau_{os}} \right] \quad (19)$$

where ρ_{oy} is the elasticity of oil demand with respect to the after-tax price (for firms and households), and the remaining ratio is the percentage response of the after-tax oil price to a change in the pre-tax price. Thus, as the specific tax on oil rises, the pre-tax elasticity of oil demand is lower for any given value of the after-tax elasticity, so that oil imports fall by less in response to an oil price hike. Accordingly, intertemporal current account balance requires a bigger improvement in the non-oil balance in the case with the higher tax rate, and correspondingly, a larger decline in the terms of trade. The effects in the case of a lower tax rate can be understood by a symmetric argument.

Nevertheless, as suggested by the equation above, the wedge introduced by specific taxes can only change the level of the pre-tax elasticity noticeably if the after-tax elasticity is fairly high. Thus, the effect is quite small quantitatively even for an after-tax oil price elasticity in the range of our benchmark value of 0.5. This analysis suggests that empirically relevant differences in specific tax rates between the United States and the euro area – with the latter close to the 40 percent rate considered in our alternative – probably play little role in contributing to different effects of oil shocks on the terms of trade or external balance across the two regions.

5 Complete vs. Incomplete Markets

As we have seen, oil price shocks have substantial wealth effects that diverge across countries in our benchmark model with incomplete financial markets. It is helpful to contrast these results with an alternative financial structure of complete markets in order to emphasize the pivotal role that such cross-country changes in the distribution of wealth play in driving the dynamics of the terms of trade and non-oil trade balance.

Figure 7 contrasts responses to the permanent oil price shock in our benchmark model with incomplete markets to responses derived from a variant that allows for complete financial markets. The results are strikingly different: in particular, the terms of trade remain unchanged from baseline under complete markets, and the non-oil trade balance also is essentially unaffected.

These disparities reflect that ownership of the profit flow associated with oil production is effectively shared across countries in the complete markets case through insurance transfers. Although oil price hikes associated with an adverse supply shock still have a contractionary effect on consumption – reflecting that the higher factor price lowers productivity – there is no added hit to an oil importer associated with a transfer of wealth across countries. With symmetric wealth effects and labor supply responses across countries, the terms of trade remain nearly unchanged. Even though the oil-importing country runs an overall trade deficit under complete markets, it receives insurance payments that offset the higher nominal oil expenditures: these transfer allow it to avoid having to accrue a surplus on its nonoil balance (as was necessary to achieve intertemporal current account balance in the incomplete markets case).

As a corollary, the structural factors highlighted in Section 4 as influencing the magnitude of the terms of trade and nonoil trade balance responses operate on these variables only insofar as the shock has differential wealth effects across countries. But this hinges on the incomplete market structure. Hence, with these differential

wealth effects (virtually) eliminated under complete markets, variations in these structural factors – including the elasticity of substitution between oil and other inputs, the steady state real interest rate, and the magnitude of domestic oil production of the oil importer – have essentially no effect on the terms of trade or non-oil balance. This is illustrated in Figure 7, which also shows the complete markets response for the case of an extremely low long-run substitution elasticity between oil and other inputs of 0.05, compared with 0.5 in our benchmark. Under complete markets, the response of the terms of trade and the non-oil trade balance are indistinguishable for these very different substitution elasticities; this contrasts sharply with the large divergence in the incomplete markets case that was highlighted in our discussion of Figure 2.

6 An Endogenous Oil Price

We have thus far assumed that the relative price of oil was determined exogenously, and assessed the effects of permanent oil price shocks. This approach has a long precedent in the literature, as it appeared to capture key features of the dramatic OPEC-driven oil price hikes of the 1970s. However, our analysis of the effects of oil shocks on the external balance does not hinge on specifying the relative price of oil as an exogenous process. In this section, we show that similar results obtain in a framework in which the oil price is determined endogenously through the balance of supply and demand in the world oil market.¹⁴

Figure 8 shows responses to a shock that induces a gradual reduction in the foreign flow endowment of oil Y_{Ot}^* . In this framework, the foreign endowment is assumed to follow an AR(1) in the growth rate with an autoregressive parameter of 0.9, and the innovation is scaled so that the relative price of oil rises 50 percent above

¹⁴As emphasized by Barsky and Kilian (2004) and Kilian (2006a) identifying supply and demand disturbances in the oil market is a formidable empirical task. However, Leduc and Sill (2004) suggested that a model-based approach might facilitate this decomposition.

in baseline at its peak under our benchmark calibration. Although the peak effect is identical to that analyzed earlier, it is clear from the figure that the price declines considerably over the long simulation horizon. The gradual decline in the oil price is attributable to falling world oil demand as households and firms substitute away from oil.

Focusing on our benchmark calibration, it is clear from Figure 8 that the qualitative effects of the oil supply shock are identical to those derived under an exogenous shock to oil price (recalling Figure 1). Thus, the higher oil price shifts the trade balance of the oil-importer into persistent deficit. This deficit must be offset by an improvement in the non-oil balance, which is achieved through a deterioration in the terms of trade. The notable difference between the responses in Figure 8 and those in Figure 1 is that the former are quantitatively smaller. This simply reflects that the oil price hike is much less persistent, and the wealth transfer to the oil-importer correspondingly smaller.

Our results showing how key structural parameters affect the response of the terms of trade and the trade balance (Figures 2-5) continue to obtain in the case of an endogenous price response. In fact, the divergence in responses across alternative calibrations of the elasticity of substitution parameter and the domestic production share parameter tends to be much larger than in simulations with an exogenous oil price. To illustrate this, Figure 8 compares the effects of the oil quantity shock under our benchmark calibration in which the oil elasticity of demand is set to 0.5 with alternatives in which the elasticity is set equal to unity and 0.1 (the same alternatives considered in Figure 1). Given that the fall in world oil supply is identical across simulations, the oil price shows a larger and more persistent increase as the long-run oil price elasticity of substitution declines. Accordingly, the terms of trade deteriorate by more (and the non-oil balance improves by more) if the oil price elasticity of demand is low, which mirrors the qualitative pattern in Figure 1. However, because the oil price response is now much larger under the low elasticity

calibration – nearly 100 percent even in the long-run, compared with about a 25 percent rise in our benchmark – the magnitude of the terms of trade deterioration and non-oil trade balance response is greatly amplified. By comparison, under a high elasticity of substitution the oil price remains only slightly above its pre-shock baseline in the long-run, so that the terms of trade and nonoil balance remain nearly unchanged.

7 Allowing for Nominal Rigidities

This section examines the sensitivity of our results to the inclusion of nominal rigidities in both price- and wage-setting. In particular, we consider an alternative version of our model that modifies our flexible price framework along several dimensions that have become standard in the literature on dynamic New-Keynesian models. First, we account for stickiness in the aggregate price of the domestically produced non-oil good (P_{Dt}) by assuming that it is produced by a continuum of monopolistically-competitive firms that set prices for an extended duration in overlapping staggered contracts. In this framework, the domestically produced non-oil good in (7) effectively serves as the factor input to the monopolistic producers (thus, all producers have the same marginal cost). These monopolistically-competitive firms set prices in their domestic currency Calvo-style contracts (so that there is “producer currency pricing” in the export market), with a mean contract duration of four quarters, and full indexation to past prices (as in the model of Christiano, Eichenbaum, and Evans 2005). Second, we allow for nominal wage stickiness by assuming that wages are set in staggered contracts by a continuum of households with differentiated labor inputs. Wages are also set in Calvo-style staggered contracts with a mean duration of four quarters, with full indexation (to past wage inflation in our case). Finally, we assume that monetary policy is determined by an

estimated interest rate reaction function of the form:

$$i_t = (1 - \gamma_i)(\bar{r} + \bar{\pi}) + \gamma_i i_{t-1} + \gamma_\pi(\pi_t^{(4)} - \bar{\pi}) + \gamma_y(y_t - y_{t-1}) + \epsilon_{it}. \quad (20)$$

In the above, i_t is the annualized nominal interest rate, $\pi_t^{(4)}$ is the four-quarter inflation rate of the aggregate output deflator (i.e., $\pi_t^{(4)} = \sum_{j=0}^3 \pi_{t-j}$), \bar{r} and $\bar{\pi}$ are the steady-state real interest rate and the central bank's constant inflation target (both expressed at annual rate). Also, $y_t - y_{t-1}$ is the (annualized) quarterly growth rate of aggregate output. We estimated the parameters of the monetary policy rule using U.S. data from 1983:1-2003:4.¹⁵ Our estimates implied $\gamma_\pi = 0.6$, $\gamma_y = 0.28$, and $\gamma_i = 0.8$.

As seen in Figure 9, the response of the trade balance and its components in the model with nominal rigidities is virtually identical to that in the model with flexible prices and wages, and the response of the terms of trade is also very similar. This similarity reflects that the response of oil demand is mainly driven by the large change in the relative price of energy. Thus, given a similar response of the oil component of the trade balance across the two models, the required adjustment in the nonoil balance and terms of trade turns out to be similar. Importantly, the requisite terms of trade adjustment can occur even in the model with nominal rigidities because the exchange rate is flexible to adjust. As Figure 9 shows an alternative monetary policy that responds to inflation only would yield only minor differences in the responses of the oil and non-oil trade balance relative to our estimated policy rule.¹⁶

¹⁵We estimated the rule using instrumental variables with lags of inflation and output growth as instruments.

¹⁶Leduc and Sill (2004) emphasized that alternative monetary policy rules can lead to markedly different responses of output to oil price increases. In our setup, increasing the long-run elasticity would increase the differences brought about by alternative monetary rules for the response of output and absorption. However, varying the policy rule has a negligible impact on the oil and non-oil trade balance response, because of offsetting movements in the terms of trade.

8 Conclusion

This paper has investigated how a permanent rise in oil prices affects the U.S. trade balance and terms of trade. Our analysis suggests that permanent or highly persistent oil price increases are likely to generate some initial decline the U.S. real exchange rate, allowing an improvement in the nonoil balance to play an important quantitative role in eventually correcting the trade deficit that emerges following an oil shock. As applied to the oil price increases of the past several years, this would suggest that the oil price hikes – if persistent – will serve as a force pushing towards improvement in the nonoil trade balance, as well as in compressing oil imports. Nevertheless, the response of trade flows and the real exchange rate depends on certain structural parameters about which there is considerable uncertainty, especially the long-run price elasticity of oil demand.

Finally, while we have devoted considerable attention to investigating the sensitivity of our model to various changes in structure, a possible limitation of our modeling framework is its representative agent structure. Accordingly, it would be of interest to consider the implications of adopting a more general framework that would allow for departures from Ricardian Equivalence, such as a Blanchard-Yaari formulation with finite horizons. Such an alternative might allow for a more realistic way of capturing the wealth effects of oil shocks, especially to the extent that foreign countries may effectively have a much shorter horizon than implied by our representative agent structure.

References

- Anderson, G. (1997). A Reliable and Computationally Efficient Algorithm for Imposing the Saddle Point Property in Dynamic Models. Federal Reserve Board, Occasional Staff's Studies 4.
- Anderson, G. and G. Moore (1985). A Linear Algebraic Procedure for Solving Linear Perfect Foresight Models. *Economic Letters* 17, 247–52.
- Atkeson, A. and P. J. Kehoe (1999). Models of Energy Use: Putty-Putty versus Putty-Clay. *American Economic Review* 89(4), 1028–1043.
- Backus, D. and M. Crucini (1998). Oil Prices and the Terms of Trade. *Journal of International Economics* 50, 185–213.
- Barsky, R. and L. Kilian (2004). Oil and the Macroeconomy Since the 1970s. *Journal of Economic Perspectives* 18(4), 115–134.
- Baxter, M. and M. Crucini (1995). Business Cycles and the Asset Structure of Foreign Trade. *International Economic Review* 36(4), 821–854.
- Bernanke, B., M. Gertler, and M. Watson (1997). Systematic Monetary Policy and the Effects of Oil Price Shocks. *Brookings Papers on Economics Activity* (1), 91–142.
- Blanchard, O. J. and C. M. Kahn (1980). The Solution of Linear Difference Models under Rational Expectations. *Econometrica* 48(5), 1305–1312.
- 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* 113(1), 1–45.
- Cole, H. L. and M. Obstfeld (1991). Commodity Trade and International Risk Sharing: How Much Do Financial Markets Matter. *Journal of Monetary Economics* 28, 3–24.
- Corsetti, G., L. Dedola, and S. Leduc (2004). International Risk Sharing and the Transmission of Productivity Shocks. CEPR Discussion Papers No. 4746.

- Dahl, C. and T. Sterner (1991). Analysing Gasoline Demand Elasticities: a survey. *Energy Economics* 13(3), 203–210.
- Finn, M. (2000). Perfect Competition and the Effects of Energy Price Increases on Economic Activity. *Journal of Money, Credit, and Banking* 32(3), 400–416.
- Fuhrer, J. C. and G. R. Moore (1995, February). Inflation Persistence. *Quarterly Journal of Economics* (440), 127–159.
- Gallaway, M. P., C. A. McDaniel, and S. A. Rivera (2003). Short-Run and Long-Run Industry-Level Estimates of U.S. Armington Elasticities. *North American Journal of Economics and Finance* 14, 49–68.
- Golub, S. S. (1983). Oil Prices and Exchange Rates. *The Economic Journal* 93(371), 576–593.
- Hamilton, J. (1983). Oil and the Macroeconomy since World War II. *Journal of Political Economy* 91(2), 228–48.
- Hamilton, J. (2003). What Is an Oil Shock. *Journal of Econometrics* 113, 363–98.
- Hamilton, J. and C. Herrera (2004). Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: Comment. *Journal of Money, Credit, Banking* 36(2), 265–86.
- Hooper, P., K. Johnson, and J. Marquez (2000). Trade Elasticities for the G-7 Countries. Princeton Studies in International Economics.
- Hooper, P. and J. Marquez (1995). Exchange Rates, Prices, and External Adjustment in the United States and Japan. pp. 107–68.
- Jones, D. W., P. Leiby, and I. K. Paik (2004). Oil Price Shocks and the Macroeconomy: What Has Been Learned Since 1996. *Energy Journal* 25(2), page 1–32.
- Kilian, L. (2006a). A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries. Manuscript, University of Michigan.

- Kilian, L. (2006b). Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy? Forthcoming, *Review of Economics and Statistics*.
- Kilian, L., A. Rebucci, and N. Spatafora (2007). Oil Shocks and External Balances. Manuscript, University of Michigan.
- Kim, I.-M. and P. Loungani (1992). The Role of Energy in Real Business Cycle Models. *Journal of Monetary Economics* 29(2), 173–189.
- Laffer, A. B. and T. Agmon (1978). Trade, Payments and Adjustment: the Case of the Oil Price Rise. *Kyklos* 31, 68–85.
- 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.
- McDaniel, C. A. and E. J. Balistreri (2003). A Review of Armington Trade Substitution Elasticities. *Integration and Trade Journal* 7(18), 161–173.
- Rotemberg, J. and M. Woodford (1996). Imperfect Competition and the Effects of Energy Price Increases on Economic Activity. *Journal of Money, Credit, and Banking* 28(4), 549–577.
- Schmitt-Grohe, S. and M. Uribe (2003). Closing Small Open Economy Models. *Journal of International Economics* 61(1), 163–185.
- Smets, F. and R. Wouters (2003). An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area. *Journal of the European Economic Association* 1(5), 1124–1175.
- Turnovsky, S. J. (1985). Domestic and Foreign Disturbances in an Optimizing Model of Exchange-Rate Determination. *Journal of International Money and Finance* 4(1), 151–71.
- Wei, C. (2003). Energy, the Stock Market, and the Putty-Clay Investment Models. *American Economic Review* 93(1), 311–23.

Table 1: Calibration of Baseline Model

| Parameter | Used to Determine | Parameter | Used to Determine |
|---|--|-----------------------|--|
| Parameters governing households' behavior | | | |
| $\beta = 0.99$ | discount factor | $\chi = 10$ | labor supply elasticity ^a |
| $\sigma = 1$ | consumption elasticity | $\varkappa = 0.8$ | consumption habits |
| $\phi_i = 4$ | investment adj. cost | $\phi_b = 0.001$ | financial intermediation cost |
| $\rho_{oc} = -2$ | oil elasticity in consumption (0.5) | $\varphi_{oc} = 142$ | oil adj. costs in consumption |
| $\tau_{os} = 0.2$ | specific oil tax | $\tau_{ov} = 0.1$ | ad valorem oil tax |
| $\omega_{oc} = 0.018$ | weight of oil in consumption | | |
| Parameters governing firms' behavior | | | |
| $\delta = 0.025$ | depreciation rate of capital | $\rho_v = -2$ | K-L sub. elasticity (0.5) |
| $\omega_k = 1.61$ | weight of K in value added | $\omega_{oy} = 0.024$ | weight of Oil in production |
| $\rho_{oy} = -2$ | value-added oil sub. elasticity (0.5) | $\varphi_{oy} = 142$ | oil adj. costs in production |
| Parameters governing international trade | | | |
| $\rho_c = -1$ | consumption import sub. elasticity (1.5) | $\omega_{mc} = 0.063$ | weight of imports in non-oil consumption |
| $\rho_i = -1$ | investment import sub. elasticity (1.5) | $\omega_{mi} = 0.43$ | weight of imports in investment |
| $\varphi_{mc} = 10$ | consumption import adj. costs | $\varphi_{mi} = 10$ | investment import adj. costs |

^a The Frisch elasticity is $2/\chi = 0.2$.

Figure 1: A 50% Increase in the Price of Oil

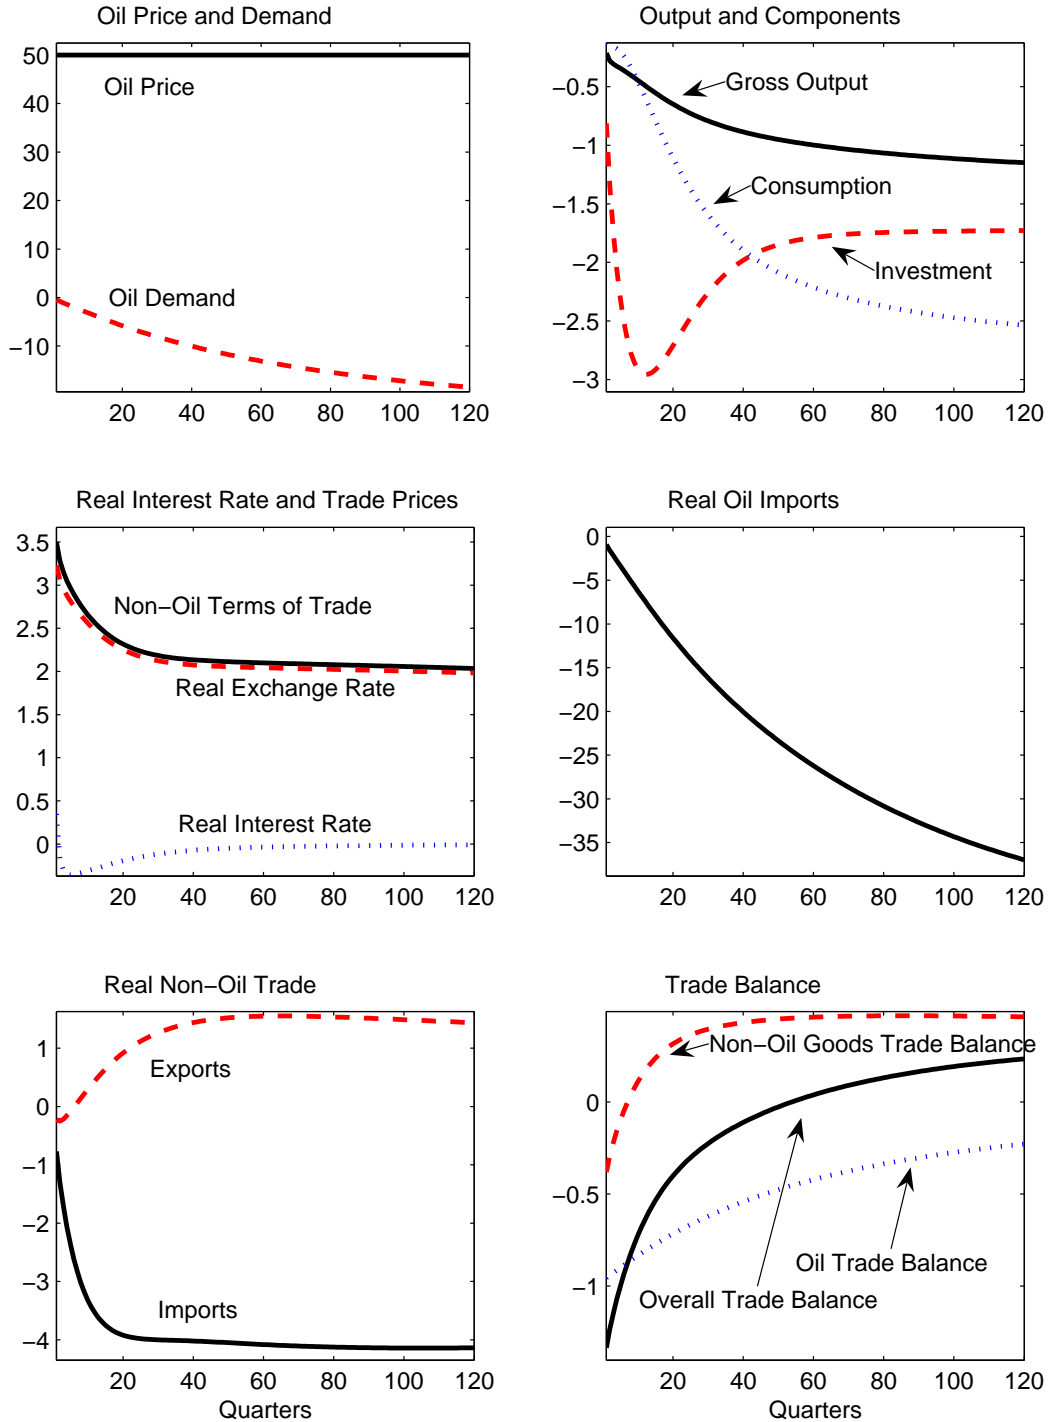


Figure 2: A 50% Rise in the Oil Price: Alternative Price Elasticities of Oil Demand

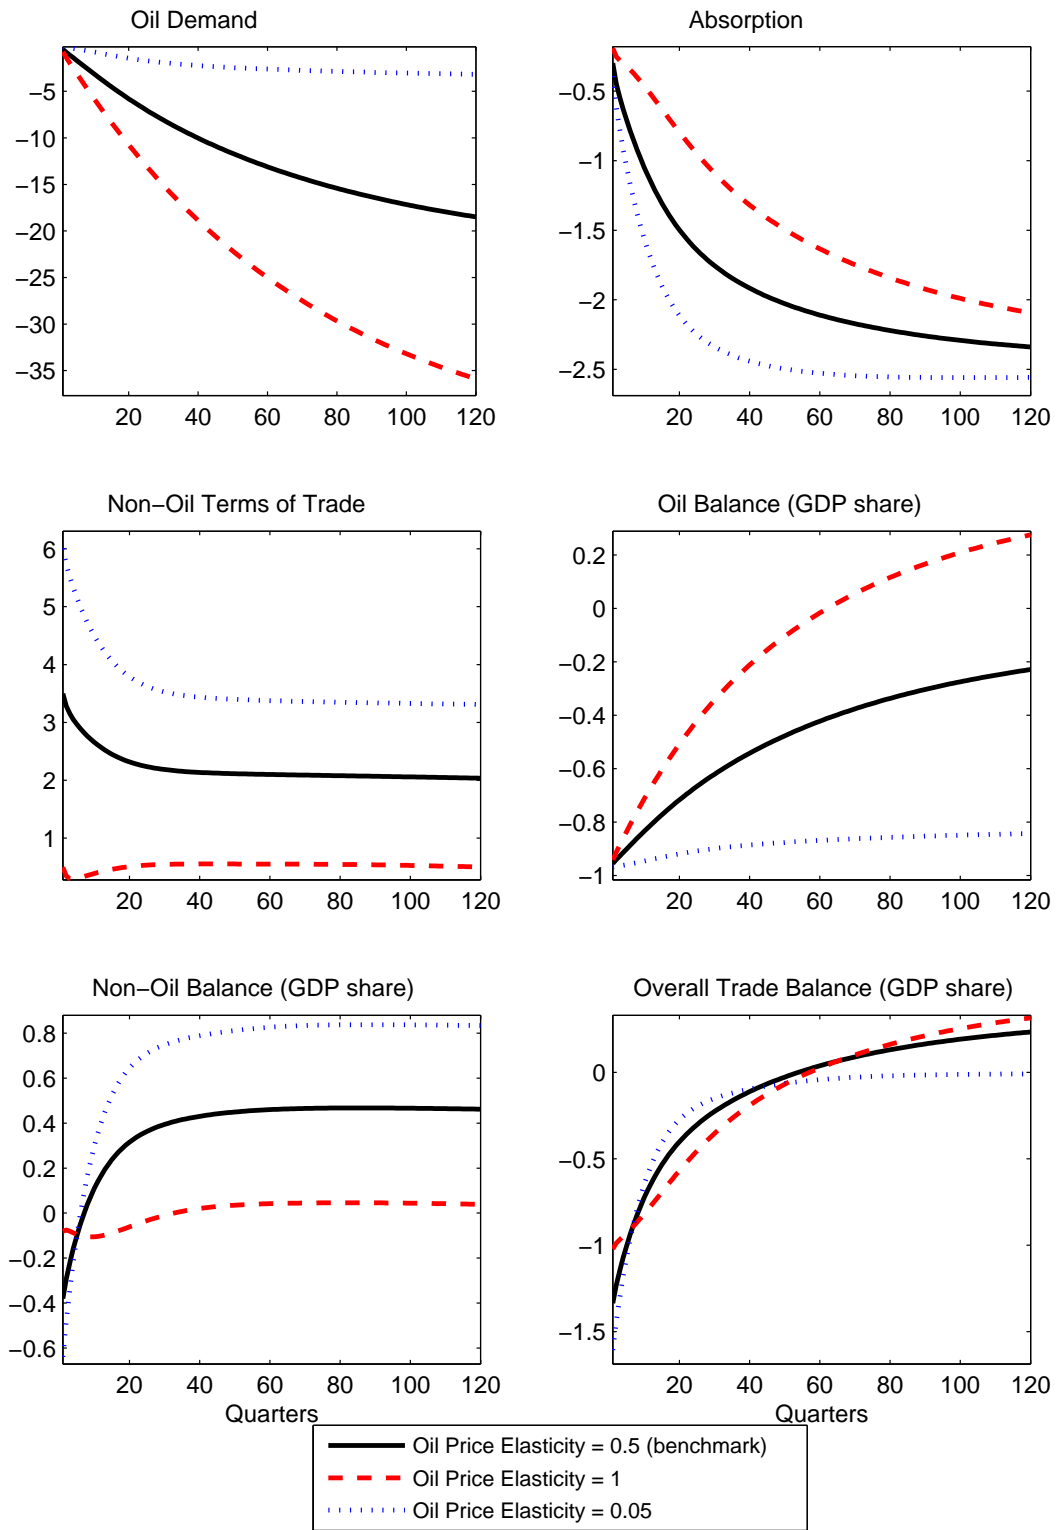


Figure 3: 50% Rise in the Oil Price: Alternative Discount Factors

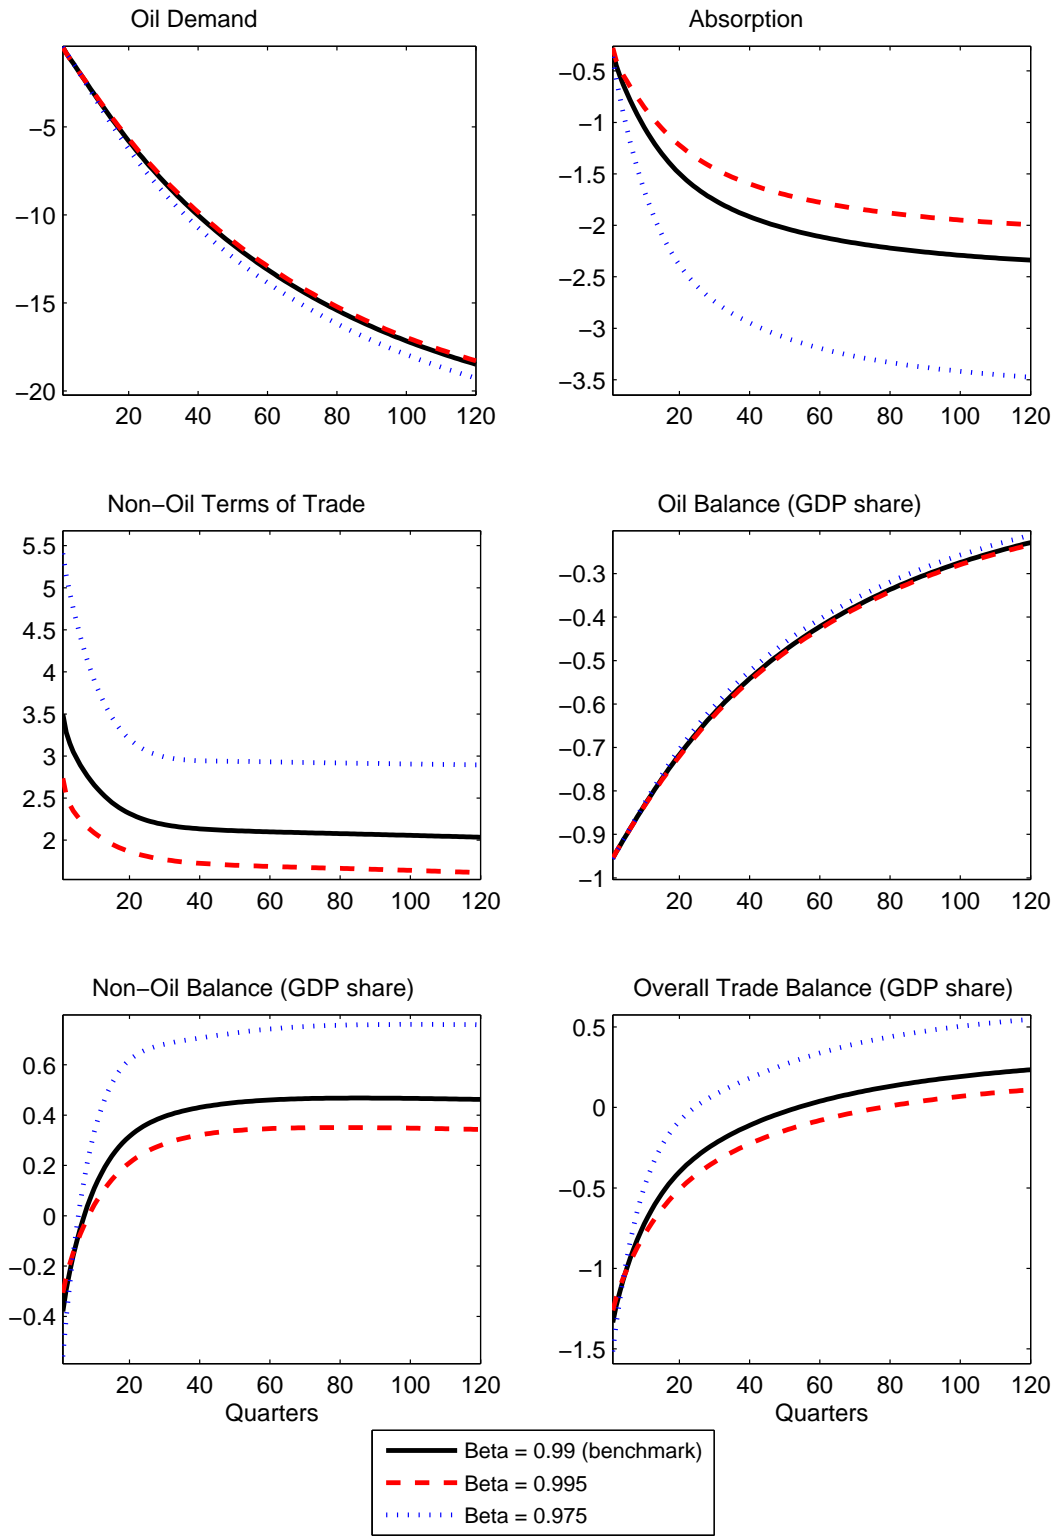


Figure 4. Contours of the 2-Year Average Response of the Terms of Trade to a 50% Permanent Increase in the Price of Oil (baseline calibration)




Figure 5. 50% Rise in the Oil Price: Alternative Levels of Domestic Oil Production

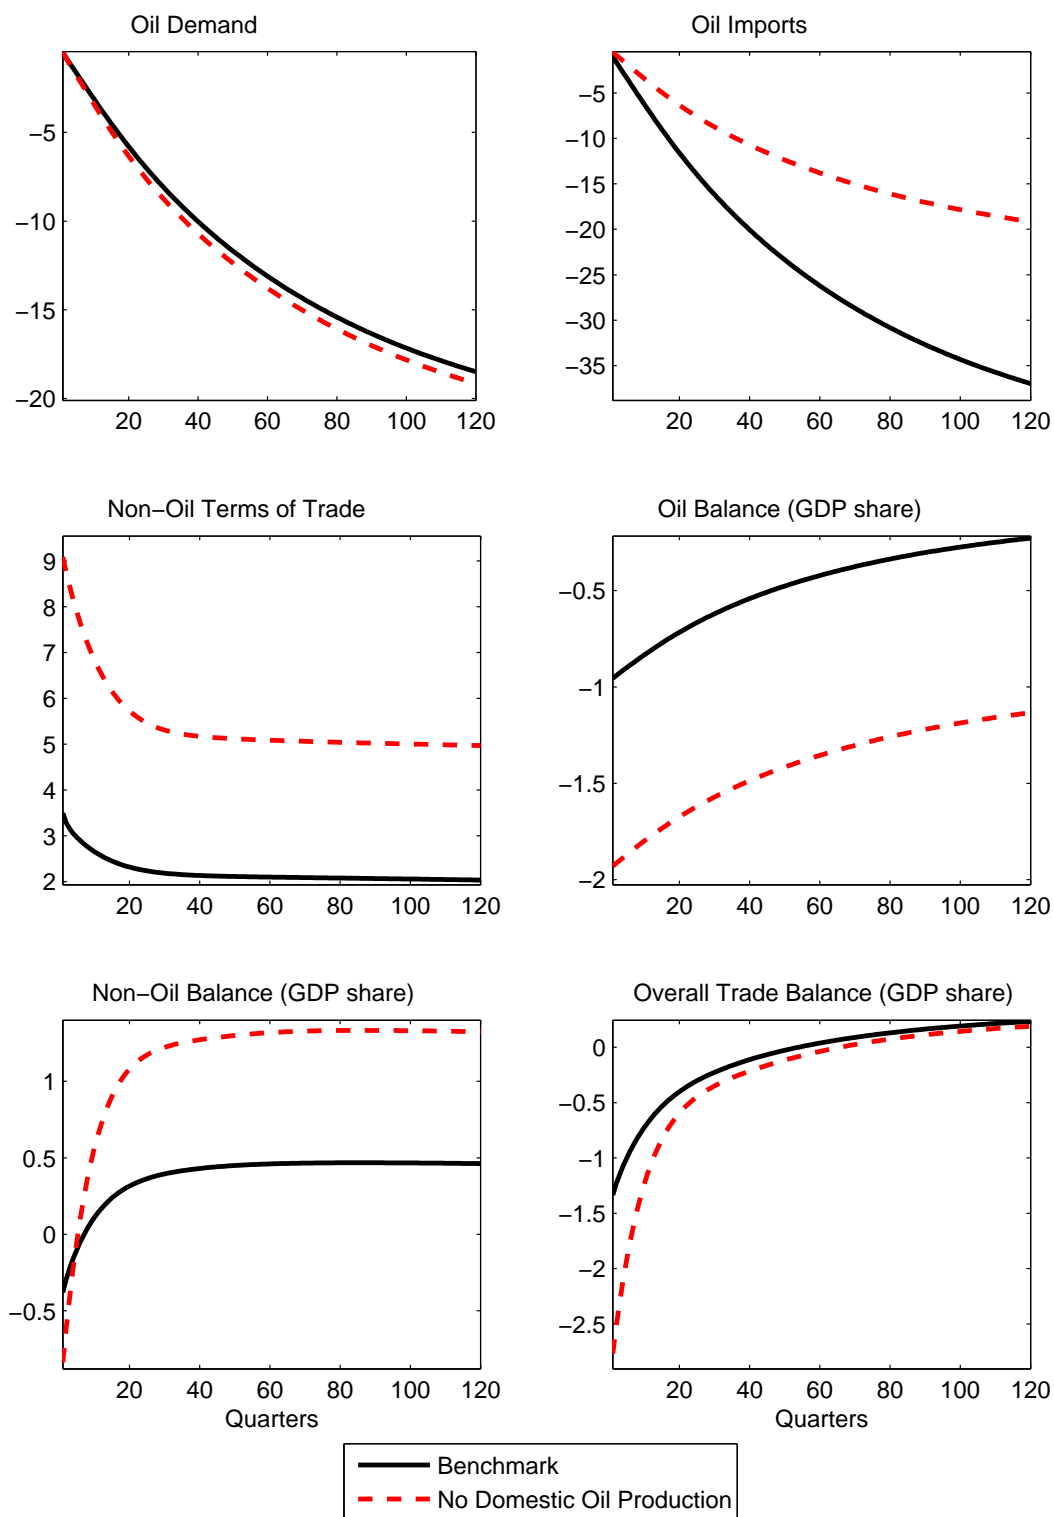


Figure 6. 50% Rise in the Oil Price: Alternative Levels of Specific Oil Taxes

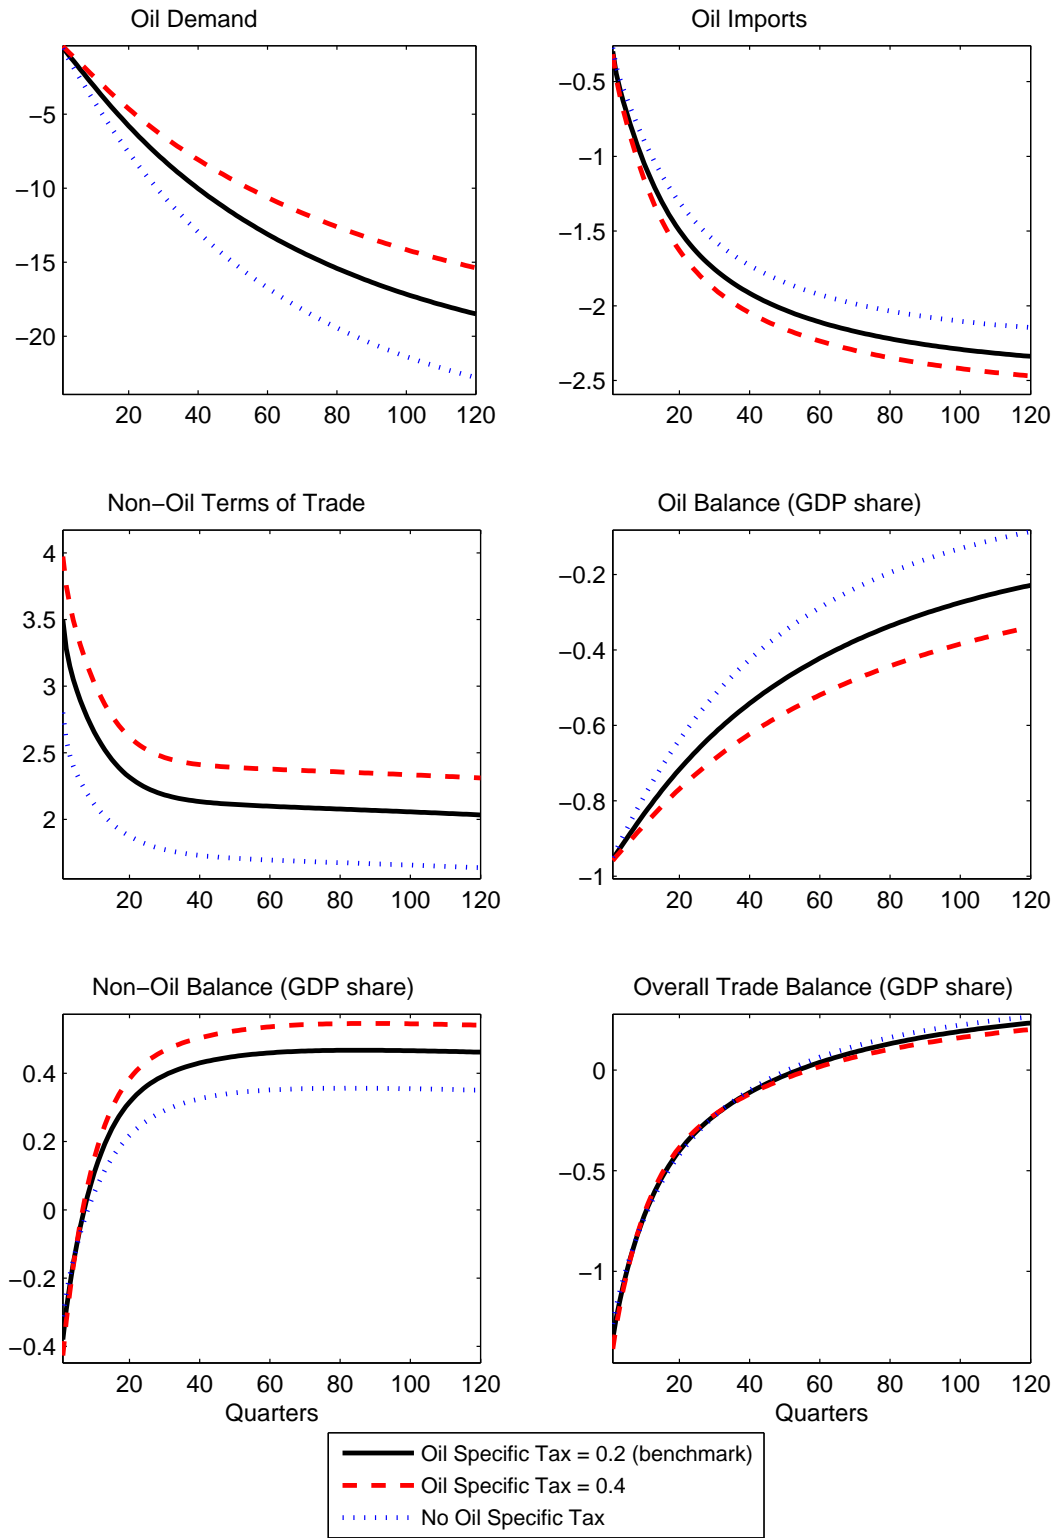


Figure 7. 50% Rise in the Oil Price: Complete and Incomplete Markets

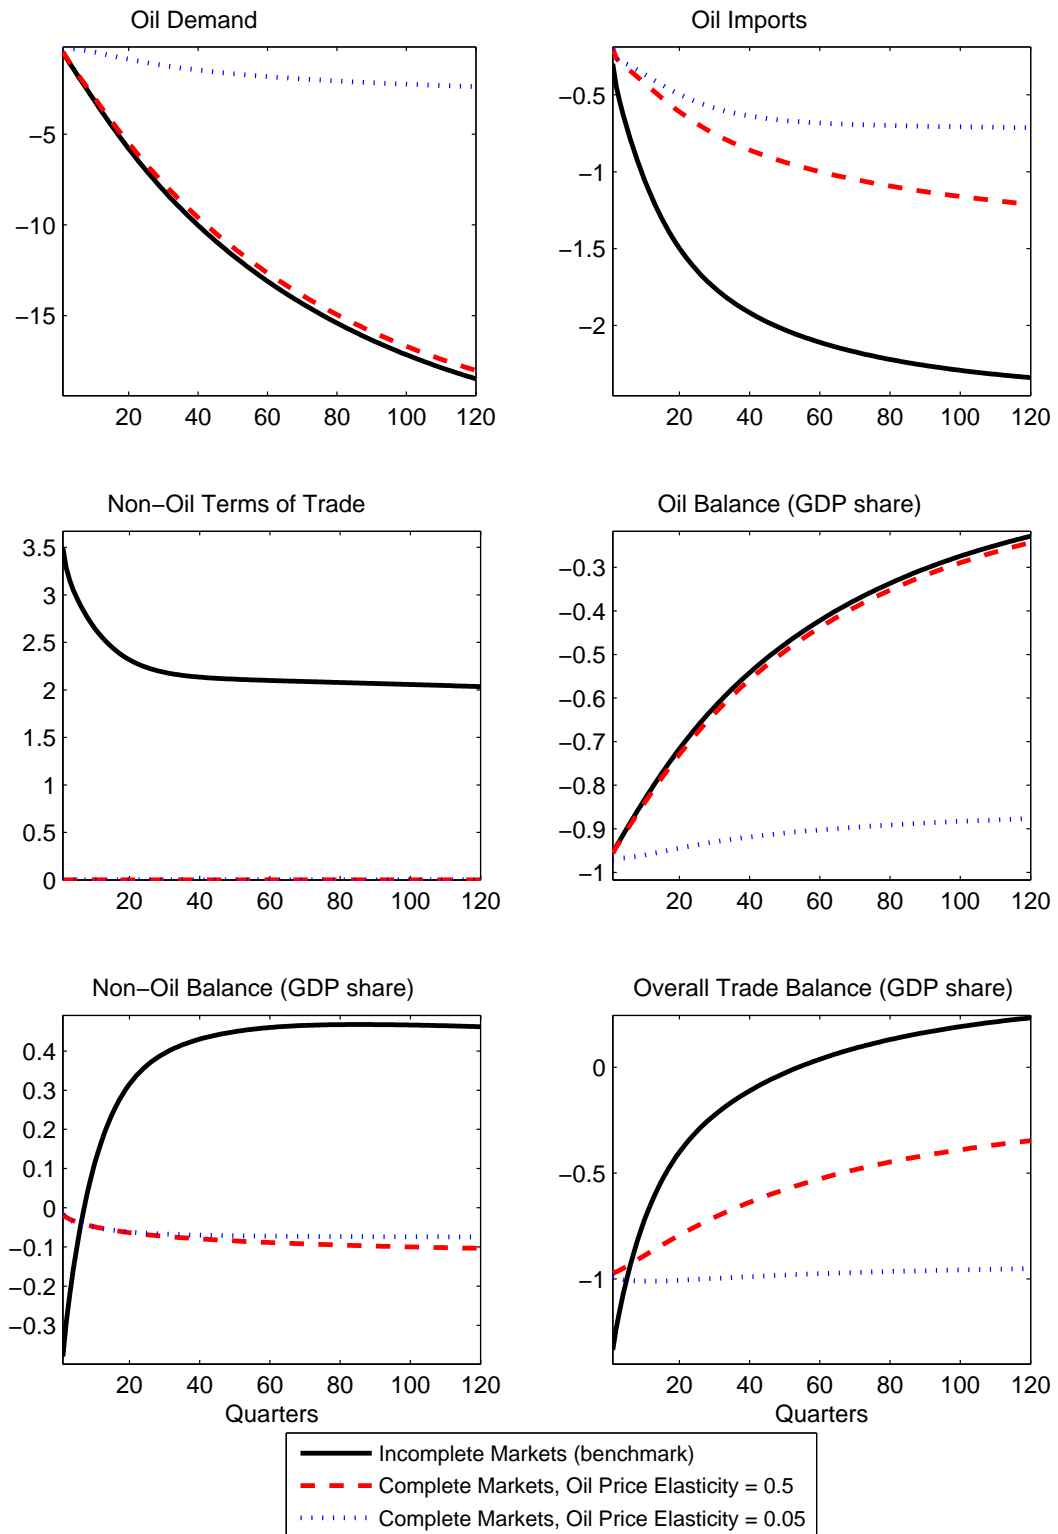


Figure 8. Contraction in Foreign Oil Supply

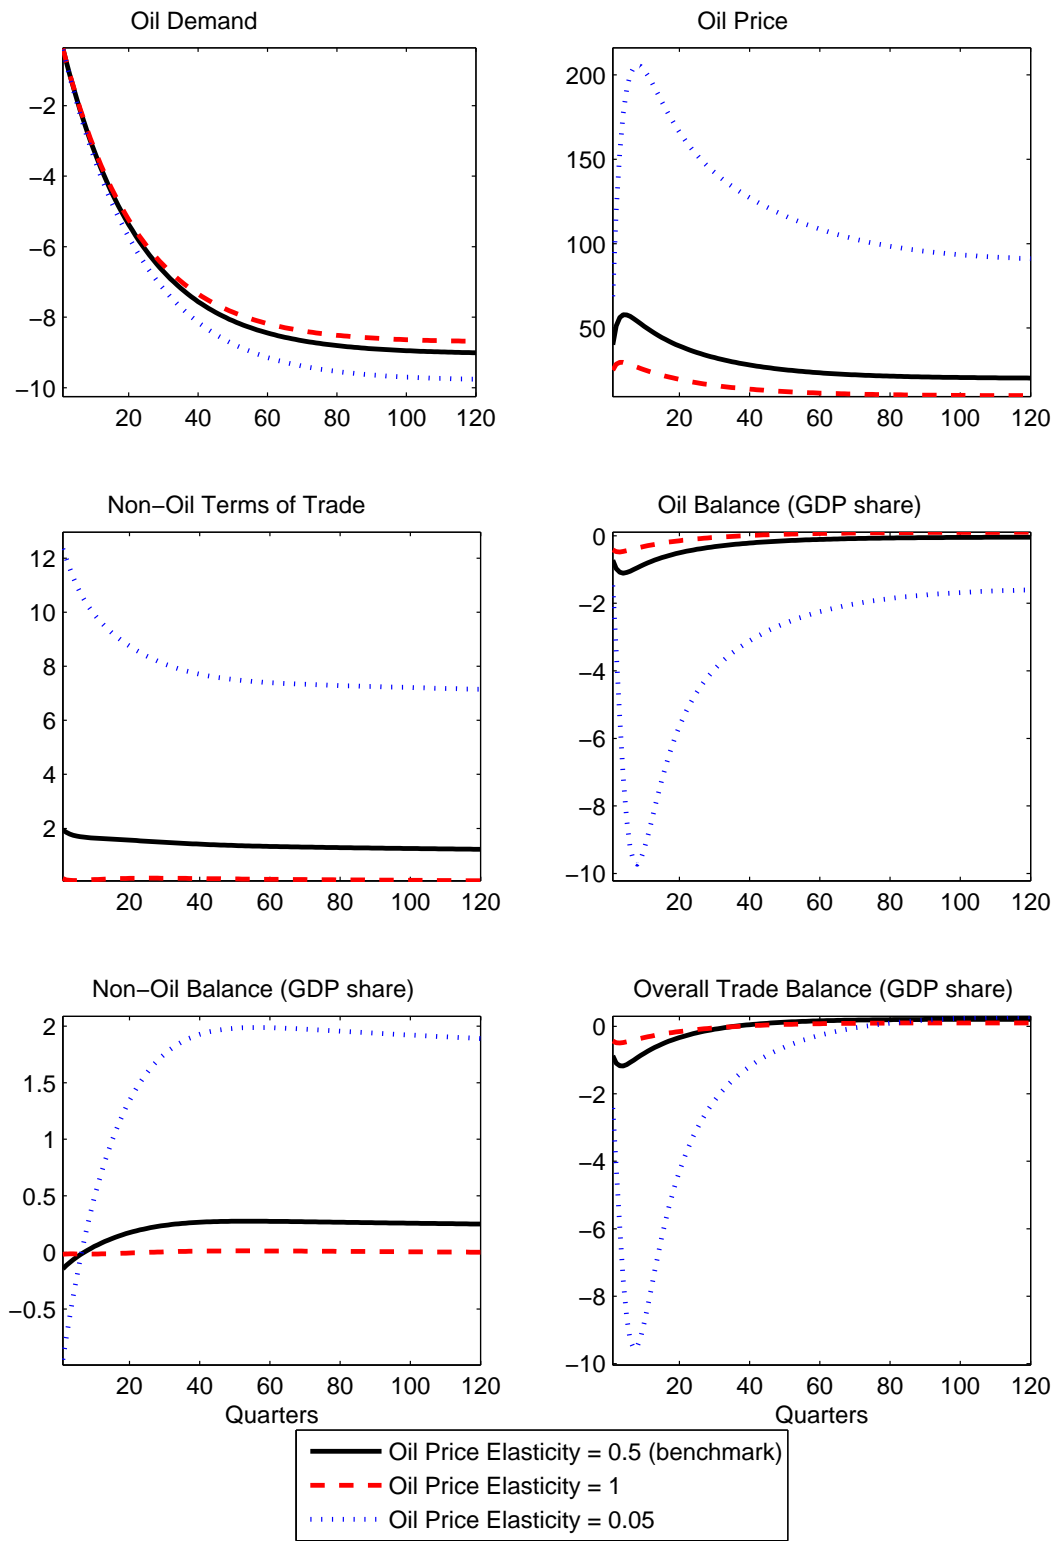
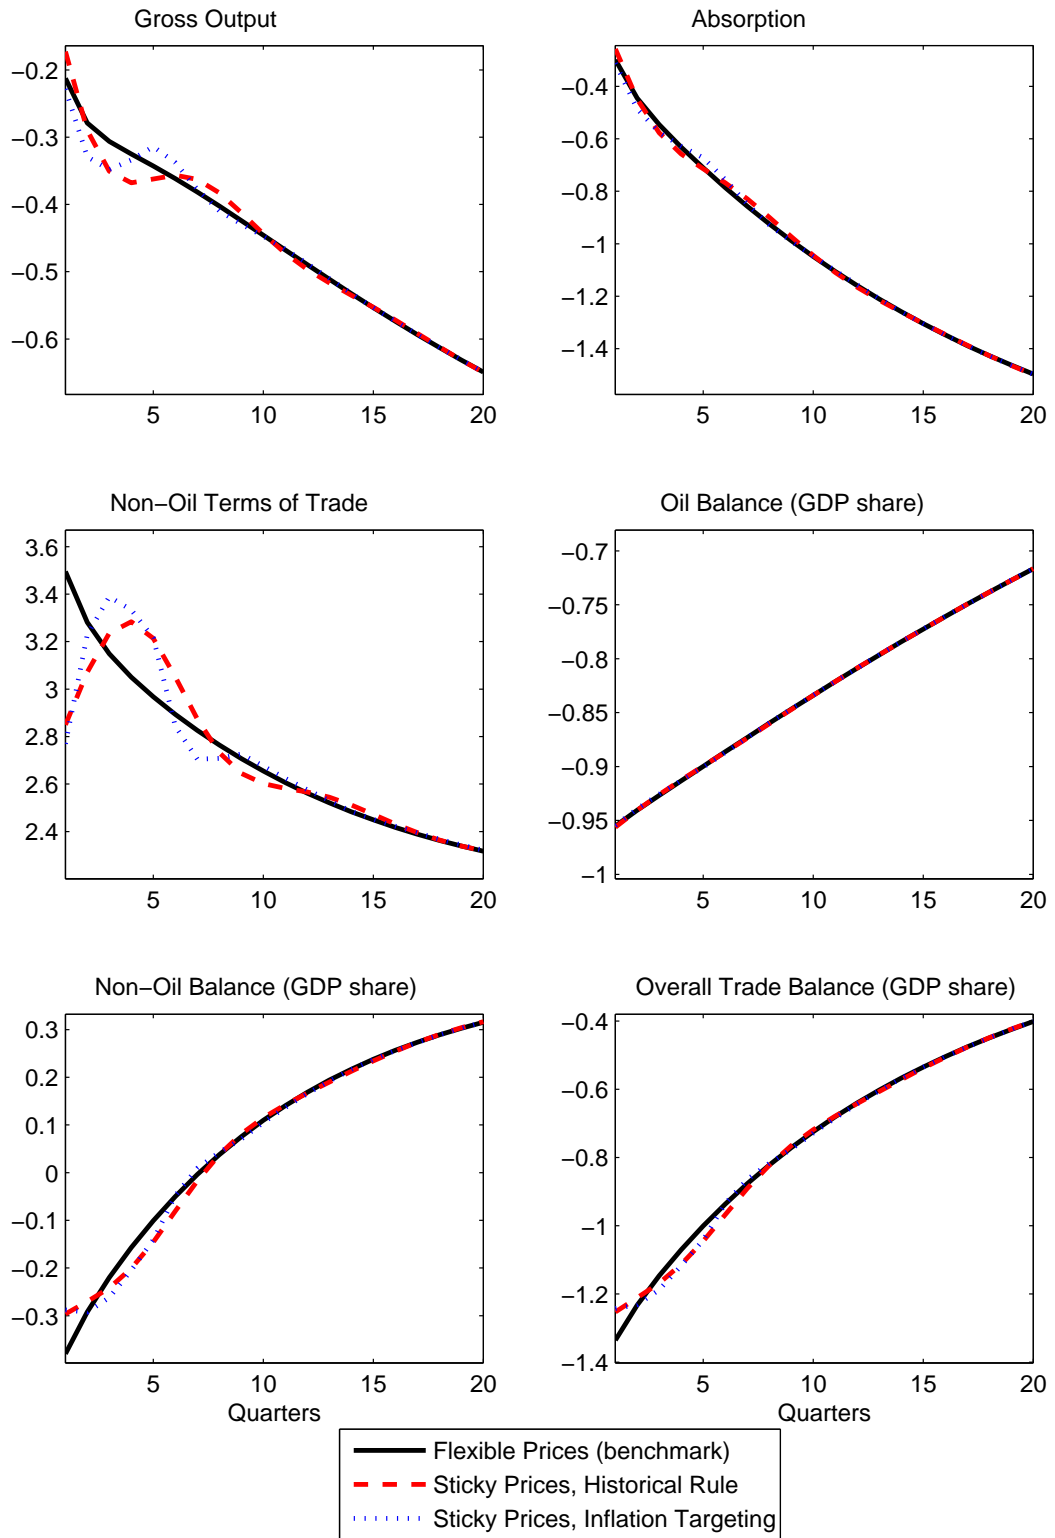


Figure 9. 50% Rise in the Oil Price: Flexible vs. Sticky Prices



A Appendix: Estimation of the Oil Demand Equation

Abstracting from taxes, the log-linearized behavioral equation determining aggregate oil demand can be expressed as:

$$\begin{aligned} \hat{O}_t = & \varepsilon_a \left(\frac{O_Y}{O} \hat{Y}_t + \frac{O_C}{O} \hat{C}_t \right) - \frac{\varepsilon_o}{1 + \varepsilon_o \varphi_o} \left(\hat{P}_{O,t} - \frac{O_Y}{O} \hat{P}_{Dt} - \frac{O_C}{O} \hat{P}_{Ct} \right) \\ & - \frac{\varepsilon_o \varphi_o}{1 + \varepsilon_o \varphi_o} \left(\frac{O_Y}{O} \hat{Y}_{t-1} + \frac{O_C}{O} \hat{C}_{t-1} - \hat{O}_{t-1} \right) + u_t. \end{aligned} \quad (21)$$

In equation (21), a “hat” denotes a variable’s percentage deviation from its steady state value. Steady state values are distinguished by the omission of a time subscript. To derive the equation above, we imposed the restrictions $\varepsilon_o = \frac{1+\rho_{oy}}{\rho_{oy}} = \frac{1+\rho_{oc}}{\rho_{oc}}$ and $\varphi_o = \varphi_{oy} = \varphi_{oc}$, by which the price elasticity of oil demand is equalized across the two end uses in our model at each point in time. The parameter ε_o is the absolute value of the long-run price elasticity of oil demand. The variable u_t is a stochastic shock to oil demand; although omitted from our model description, this shock may be derived from a taste shock entering into equations (7) and (10), respectively. Our model imposes that the coefficient on contemporaneous activity ε_a equals one, but we estimate it as a free parameter as a test of our specification.

Equation (21) can be manipulated to yield:

$$\hat{O}_t = \varepsilon_a \left(\frac{O_Y}{O} \hat{Y}_t + \frac{O_C}{O} \hat{C}_t \right) - \frac{\frac{\varepsilon_o}{1+\varepsilon_o \varphi_o}}{1 - \frac{\varepsilon_o \varphi_o}{1+\varepsilon_o \varphi_o} L} \left(\hat{P}_{O,t} - \frac{O_Y}{O} \hat{P}_{Dt} - \frac{O_C}{O} \hat{P}_{Ct} \right), \quad (22)$$

where L represents the lag operator (we have suppressed the shock u_t for simplicity). Thus, oil demand depends on a measure of activity $\left(\frac{O_Y}{O} \hat{Y}_t + \frac{O_C}{O} \hat{C}_t \right)$ and on a distributed lag of a relative price measure for oil (a ratio of the oil price to a weighted average of the deflators for domestic production and consumption). The sum of the lag coefficients equals the long-run elasticity ε_o . The parameter φ_o determines the half life of the response of oil demand to a permanent increase in the relative price of oil (keeping the activity measure constant).

To control for endogeneity, our econometric approach follows Fuhrer and Moore (1995). As a statistical model, we use a three equation VAR(4) that includes oil demand, the activity measure, and the relative price measure.¹⁷ The demand equation (21) replaces the oil equation in that VAR system. We estimate this system using maximum likelihood.

As a measure of oil, demand we take total petroleum consumption (product supplied) from the Energy Information Administration. We construct the activity measure as a weighted average of consumption and GDP from the National Income and Product Accounts (NIPA).¹⁸ For the relative price measure, we use the crude oil producer price index series from the Bureau of Labor Statistics, and the consumption and GDP deflators from NIPA. We run all regression equations using log linearly detrended quarterly data. The estimation sample covers the period from the second quarter of 1948 to the fourth quarter of 2005.

The regression results for the oil demand equation are summarized in Table 2. The estimated elasticity for the activity measure is not significantly different from 1, which is the value imposed by our theoretical model. The oil price elasticity is estimated at 0.28. Our baseline calibration value of 0.5 is within one standard deviation. The estimated adjustment cost parameter φ_o is 139.2, which implies a half life of 7 years for the response of oil demand to a permanent price increase. Our baseline half life of 10 years lies within one standard deviation of the estimate of φ_o .

The Durbin Watson statistics suggests serial correlation in the regression residuals of the oil demand equation. To address the potential misspecification, we modify

¹⁷We chose a lag length of 4 to avoid serial correlation in the residual.

¹⁸The weights reflect the share of oil demand in consumption and production for our baseline calibration, respectively 1/3 and 2/3. Gross output is the theoretically preferred production measure implied by our model. However, because deriving a gross output measure requires estimating the share of oil imports directed towards production, we used GDP in our benchmark specification. As a robustness check, we constructed a proxy for gross output, and found little difference from the results reported below.

the adjustment cost function for oil demand in production (equations 8) as follows:

$$\varphi_{OYt} = \left[1 - \frac{\varphi_{oy1}}{2} \left(\frac{\frac{O_{Yt}}{V_t}}{\frac{O_{Yt-1}^A}{V_{t-1}^A}} - 1 \right)^2 - \frac{\varphi_{oy2}}{2} \left(\frac{\frac{O_{Yt-1}}{V_{t-1}}}{\frac{O_{Yt-2}^A}{V_{t-2}^A}} - 1 \right)^2 \right]. \quad (23)$$

We change the adjustment cost function for oil demand in consumption (equation 11) analogously.

Upon log-linearizing, the aggregate oil demand equation now takes the form:

$$\hat{O}_t = \varepsilon_a \left(\frac{O_Y}{O} \hat{Y}_t + \frac{O_C}{O} \hat{C}_t \right) - \frac{\varepsilon_o}{1 + \varepsilon_o (\varphi_{o1} + \varphi_{o2})} \left(\hat{P}_{O,t} - \hat{P}_{D,t} \right) + \quad (24)$$

$$\begin{aligned} & - \frac{\varepsilon_o \varphi_{o1}}{1 + \varepsilon_o (\varphi_{o1} + \varphi_{o2})} \left(\frac{O_Y}{O} \hat{Y}_{t-1} + \frac{O_C}{O} \hat{C}_{t-1} - \hat{O}_{t-1} \right) \\ & - \frac{\varepsilon_o \varphi_{o2}}{1 + \varepsilon_o (\varphi_{o1} + \varphi_{o2})} \left(\frac{O_Y}{O} \hat{Y}_{t-2} + \frac{O_C}{O} \hat{C}_{t-2} - \hat{O}_{t-2} \right). \end{aligned} \quad (25)$$

Following the same estimation procedure as described above we obtain estimates summarized in Table 3.

This alternative regression specification reduces the residual autocorrelation, as indicated by a lower Durbin-Watson statistic. The coefficient estimates ε_a and ε_o are not statistically significantly different from those in the simpler specification. Again, we fail to reject the restriction imposed by our theoretical model that ε_a is 1. The estimates for the adjustment cost parameters φ_{o1} and φ_{o2} imply a half life of 10 years for the response of oil demand to a permanent price increase, as in our baseline calibration. Given the short length of the time series and the slow adjustment of oil demand, it is difficult to estimate the long-run elasticity precisely. Nonetheless, the regression does very well at capturing the variation in oil demand, as suggested by the high R^2 statistic.

Table 2: Regression results for baseline oil demand equation*

| | Coefficient | Std. Error | Prob. |
|-----------------|----------------------------|------------|-------|
| ε_a | 1.05 | 0.05 | 0.00 |
| ε_o | 0.28 | 0.27 | 0.30 |
| φ_o | 139.2 | 95.0 | 0.14 |
| $R^2 = 0.98$ | Durbin-Watson stat. = 2.40 | | |

* See equation (21).

Table 3: Regression results for augmented oil demand equation*

| | Coefficient | Std. Error | Prob. |
|-----------------|----------------------------|------------|-------|
| ε_a | 1.04 | 0.05 | 0.00 |
| ε_o | 0.45 | 0.42 | 0.29 |
| φ_{o1} | 81.9 | 40.3 | 0.04 |
| φ_{o2} | 21.3 | 11.6 | 0.07 |
| $R^2 = 0.98$ | Durbin-Watson stat. = 2.05 | | |

* See equation (24).