NBER WORKING PAPER SERIES

OIL PRICES AND THE TERMS OF TRADE

David K. Backus
Mario J. Crucini

Working Paper 6697
http://www.nber.org/papers/w6697

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 1998

Oil Prices and the Terms of Trade
David K. Backus and Mario J. Crucini
NBER Working Paper No. 6697
August 1998
JEL No. F41, F32, E32

ABSTRACT

The combination of substantial terms of trade variability and unstable correlation patterns of trade prices with output and trade volumes has led some to suggest a break in the link between trade volumes and prices. We find that oil accounts for much of the variation in the terms of trade over the last twenty five years and its quantitative role varies significantly over time. And since our dynamic general equilibrium model predicts that the economy responds differently to oil supply shocks than to other shocks, changes in their relative importance help to account for the unstable correlations in the data.

David K. Backus
Stern School of Business
New York University
44 West 4th Street
New York, NY 10012
and NBER
dbackus@stern.nyu.edu

Mario J. Crucini
Department of Economics
Ohio State University
401 Arps Hall
1945 North High Street
Columbus, OH 43210-1172

1. Introduction

Relative price movements remain the central issue in international macroeconomics, with their source a continuing object of speculation and debate. We study one such price, the terms of trade, which we define as the ratio of the import deflator to the export deflator in national income and product accounts. In four countries for which quarterly data are available for four decades, we find that the terms of trade has been highly variable relative both to real output and to the predictions of existing dynamic general equilibrium models of trade (Backus, Kehoe, and Kydland 1994, for example, or Stockman and Tesar 1995). This variability, moreover, has been substantially greater over the last twenty years than the first twenty. We also find that the correlation of terms of trade movements with real output and trade volumes varies widely over time and across countries.

We show that a large part of the variability of the terms of trade is associated with extreme movements in oil prices. Accordingly, we study the terms of trade and its correlation with other variables in a setting in which events affecting the production of oil interact with the production and sale of other goods. The structure blends Backus, Kehoe and Kydland's (1994) model of two large industrial countries with Crucini and Kahn's (1996) treatment of international trade in intermediate inputs into production. In our theoretical economy, two industrialized countries produce imperfectly substitutable goods using capital, labor, and oil. These two goods are then aggregated into final consumption and investment goods. A third country produces oil but consumes final goods.

This structure mimics in a simple way the interaction between large industrialized countries and largely non-industrial oil producers. The result allows us to study the responses of trade flows and real output to changes in oil production and industrial productivity. As one might expect, the volatility and comovements of the terms of trade depend on the nature and source of disturbances. We consider three different disturbances: a domestic productivity shock, a foreign productivity shock, and an oil supply shock. As in earlier work, productivity shocks alone produce little in the way of terms of trade volatility. The reason, ironically, is that the model generates insufficient variability in export and import quantities. Since quantities and prices are linked (the relative price equals the marginal rate of substitution), lack of variability in the quantities is inherited by the prices unless we adopt unrealistically low substitution elasticities.

Two features of the model help us to account for terms of trade variability. First, the supply of oil becomes partly exogenous, leading to dramatic price changes. We think this fits our experience in the 1970s quite well. Second, the technology exhibits little opportunity to substitute oil for capital and/or labor. As a result, relatively small variations in the quantity of oil are associated with substantial variation in its relative price.

The same mechanism helps to account for the unstable correlation between the terms of trade and output. In the model, fluctuations in the terms of trade and output generated by productivity shocks are positively correlated, while those generated by oil shocks are negatively correlated. A domestic productivity disturbance leads to an increase in domestic output and a fall in the relative price of the domestic good — an increase in the terms of trade. Oil supply disruptions, on the other hand, lead to output reductions in countries that use imported oil as an input into production and therefore an inverse relation between the terms of trade and output. Unstable correlations are the result, in this setting, of changes over time in the composition of shocks hitting the economy. We elaborate on these and other features of the model in the remainder of the paper.

2. International Business Cycles

We start by reviewing properties of business cycles in eight developed countries between 1955 and 1990, with an emphasis on prices and quantities of traded goods. We then turn to the price of oil. Quarterly national income and product account data are from the Organization for Economic Cooperation and Development's (OECD's) *Quarterly National Accounts.* Other sources are described in Appendix A. Unless otherwise noted, statistics are computed after application of the Hodrick-Prescott (1997) filter. Three variables are transformed in a way that merit special mention. We measure net exports as the ratio of net exports to output, both measured in nominal terms. We measure the terms of trade as the ratio of implicit deflators for imports and exports. Finally, we compute a trade ratio of export to import quantities.

2.1 Business Cycles Across Countries

We begin with a review of some of the features of international business cycles that have been highlighted in the existing literature. Table 1 reports standard deviations, correlations with output, and international correlations of various macroeconomic aggregates.

We see that trade variables are uniformly more volatile than aggregate output and are often more volatile than investment. The standard deviation of output ranges from 0.89 in France to 2.17 in Germany while the standard deviation of investment ranges from 2.64 to 5.63 in the same two countries. Trade variables are substantially more volatile than output, with the standard deviation of the trade ratio ranging from a low of 3.66 in France to a high of 9.94 in Australia. The terms of trade is also quite volatile, with the standard deviation ranging between 2.44 for Canada to 5.68 for Japan. Note that the terms of trade is uniformly less volatile than the trade ratio, though some of the differences are small.

The next three columns of Table 1 report correlations between key international variables. The correlation of net exports and output is robustly negative, ranging from –0.05 for Germany to –0.68 for Italy. However, the correlations of the terms of trade with output and net exports vary widely across countries. The terms of trade is approximately acyclical, in the sense that its correlation with output is small on average, but it ranges between 0.38 for Italy and –0.30 for Australia. The correlation of the terms of trade and net exports ranges from –0.66 for Italy to 0.28 for the US.

The final three columns report the correlations of each country's output, consumption, and investment with the same US variable. We see that all of these macroeconomic aggregates tend to be positively correlated, with output more strongly correlated across countries than consumption or investment.

2.2 Business Cycles Across Periods

The dramatic oil price increases of the early 1970's and 1980's and subsequent collapse are salient features of international price behavior in the postwar economy. Given the net trade exposure of industrialized countries in energy and oil, it would be surprising if some dimension of international business cycles were not affected by the rapid emergence of these shocks. We search for evidence of this by examining business cycles across different time periods.

In Table 2 we divide the sample into three periods. We refer to the middle period, 1973:1 to 1986:4, as the OPEC regime, when the most dramatic shifts in oil prices occurred. Many of the statistics change significantly as we move from the early postwar period to the OPEC years, but not all the changes are in the same direction. Beginning with output, volatility rose in France, the United Kingdom, and the United States, but fell in Germany and Japan. Consumption volatility rose in Italy, Japan, the United Kingdom, and the United States, but fell in Australia and Germany. Investment volatility rose in Australia, France, Italy, and the United States, but fell in Germany and Japan.

The trade variables display a more consistent pattern across the two samples. The standard deviation of the trade ratio is about the same in both periods. The terms of trade, however, exhibits much more volatility during the OPEC period, increasing substantially in five of the eight countries and rising in absolute terms in every country. The correlation of output and net exports remains robustly negative, though it declines in France while rising in Australia and Canada. The correlation of the terms of trade and net exports is not significantly different in the OPEC period. The correlation of the terms of trade and output falls significantly in four countries.

Cross-country correlations indicate greater conformity of business cycles with the United States during the OPEC period than before. This is particularly evident in output correlations, which rise for all countries but France.

The focus of the remainder of the paper will be to explain three observations about the

OPEC period: (i) greater volatility of the terms of trade without an associated increase in the volatility of the trade ratio; (ii) smaller correlations of the terms of trade with output (particularly among oil importers); and (iii) increased cross-country correlations of output. Increased variability of the terms of trade can be traced directly to the price of oil, which we examine next.

2.3 Oil Prices

The basic facts relating to the oil price shocks of the 1970's are familiar to most of us: the nominal price of a barrel of crude oil jumped by over 300 percent from $2.7 in 1973 to $11.2 in 1974, rose another 200 percent to 34.4 in 1981 before falling to about half this price by the late 1980's. Given the importance of oil as an internationally traded commodity and the volatility of its price, oil shocks could potentially explain virtually all of the terms of trade variation from the early 1970's to mid–1980's. We explore that possibility in this section.

Our monthly nominal oil price index is a simple average of US dollar prices in three major markets: Brent, Dubai, and West Texas. The series is converted to a quarterly frequency by averaging. The resulting price index is then divided by the quarterly U.S. GDP deflator to arrive at a constant dollar measure of oil prices.

Figure 1 plots the oil price index against the terms of trade for Canada and the United States with the oil price series normalized so that its standard deviation matches that of the terms of trade for the country against which it is plotted. Two observations can be made directly from the figure. First, the terms of trade is remarkably volatile, shifting on many occasions by 20% or more in a few years. The most dramatic shifts in the terms of trade are synchronous with changes in the relative price of oil. The second obvious feature of the series is the negative correlation between the two relative prices in Canada and the positive correlation in the United States. These patterns almost certainly reflect the fact that the United States is a net importer of oil while Canada is a net exporter.

Table 3 reports the net trade share for all fuels for each country in the sample at three points in time: 1970, 1975, and 1987. The share is computed as the ratio of the nominal trade balance in fuels divided by the nominal average level of exports plus imports in a given year, expressed as a percentage. We see that countries with little or no domestic oil production have large negative trade balances in fuel. The most obvious example is Japan, which in 1970 had a net fuel share of trade of –20.2%.

The first oil price shock altered net fuel trade shares in a predictable way. Between 1970 and 1975, the deficits in fuel trade almost exactly double in France, Germany, Italy, and Japan (all countries that have substantial net import shares of fuel during the entire sample period). The United States' net export share of fuel changes even more, moving from –3.3% to –21.4%. By 1987, the deficits had largely reversed themselves. Some of the reversal was due to the collapse of the relative price of oil in the mid-1980's, but some was also likely due to energy conservation. Other major alterations to fuel trade balances were the increased positive position of fuel in the trade balances of Australia and Canada (moving from approximate balance to 13.3% and from 1.2% to 5.5%, respectively) and the emergence of the United Kingdom as a net exporter of fuel (attributable to the rapid growth of North Sea oil production).

The next two columns of Table 3 report the correlation of each country's terms of trade and the relative price of crude oil in the world market (in terms of US goods) computed using both log-levels of variables and their HP-cyclical components. We find a negative correlation between the terms of trade and oil for the only country that is consistently a net exporter of oil — Canada. The correlation is robustly positive for the countries that were consistently net importers of oil with the exception of France and Germany where the correlation of the cyclical components are close to zero. The correlation for the United Kingdom changes from –0.25 to 0.36 as we move from log-levels to cyclical fluctuations.

The correlations are suggestive of an important role for oil's relative price in the cyclical and secular evolution of the terms of trade. How important? We answer this question by comparing the volatility of the overall terms of trade to an estimate of the non-fuel terms of trade. Our estimate of the non-fuel terms of trade is the overall terms of trade multiplied by the net export share of fuels. The estimate will be crude unless the ratio of the quantity of fuel to non-fuel trade is constant (see the discussion in Appendix B). The standard deviation of the annual terms of trade and our non-fuel terms of trade estimates are shown in the final two columns of Table 3.

As one might anticipate, the smallest adjustment in the volatility of the terms of trade occurs for countries with small or ambiguous trade exposure: rising slightly in Australia from 8.10 to 8.22 percent per year; falling somewhat in Canada from 6.51 to 6.05; and falling in the United Kingdom from 7.63 to 5.23. More dramatic changes are found for the large net importers, where the non-fuel terms of trade is typically one-fourth as volatile as the overall terms of trade. After adjusting for the impact of oil, the volatility of the terms of trade of the smaller countries is comparable or greater than that for the larger countries.

3. The Model

We extend the two-good, two-country, stochastic growth model of Backus, Kehoe, and Kydland (1994) to incorporate a third country that sells oil. Each of the three countries specializes in the production of a single good. The first two goods satisfy the consumption and investment needs of all three countries. We denote them as good a and good b and interpret them as manufactures. The third good is oil (o), which is used to produce manufactures. We use the same indexes to refer to countries as we use for goods.

3.1 The Economic Environment

Consumers in each country maximize the expected value of lifetime utility, given by:

$$E_0 \sum_{t=0}^{\infty} \beta^t U(c_t^j, L_t^j), \beta < 1 \quad j = a, b, o. \tag{1}$$

where consumption, c, is a CES aggregate of the home good, a, and the foreign good, b, and L is leisure.

The production of each manufactured good makes use of labor N, physical capital k, and oil o, as inputs. Output in country j is y^j and is affected by a stochastic productivity variable denoted z^j that is described in detail below.

$$y_t^j = z_t^j F(k_t^j, N_t^j, o_t^j), \quad j = a, b. \tag{2}$$

The key point of our departure from existing models of the macroeconomic impact of oil shocks is that we endogenize the price of oil. We accomplish this by having the world supply of oil, y^o, at any point in time equal the sum of two parts:

$$y_t^o = z_t^o + F(N_t^o) \tag{3}$$

an exogenous component z^o which in practice we treat as OPEC oil supply and an endogenous component, $F(N^o) = N^\alpha$, which we treat as the supply originating from non-OPEC oil producers.

The technology differs in two respects from that of Backus, Kehoe, and Kydland (1994). One is that final consumption and investment goods are separate composites of a and b. The aggregator for investment is

$$i_t^j = i(a_{it}^j, b_{it}^j), \quad j = a, b. \tag{4}$$

where $i(a, b) = [\psi a^{1-\mu} + (1-\psi) b^{1-\mu}]^{1/(1-\mu)}$ and the subscript i identifies the use of a and b in the investment process. The other is that physical capital formation is subject to adjustment costs, as in Baxter and Crucini (1995):

$$k_{t+1}^j = (1-\delta) k_t^j + \phi(i_t^j / k_t^j) k_t^j, \quad j = a, b. \tag{5}$$

with $\phi > 0$, $\phi' > 0$, $\phi'' < 0$.

The resource constraints for the three goods are

$$y_t^a = \sum_{j=a,b,o} a_{ct}^j + \sum_{j=a,b} a_{it}^j \tag{6}$$

$$y_t^b = \sum_{j=a,b,o} b_{ct}^j + \sum_{j=a,b} b_{it}^j \tag{7}$$

$$y_t^o = o_t^a + o_t^b \qquad (8),$$

where a_{ct}^j refers to the amount of good a used in country j for consumption purposes. Note that physical investment does not take place in the oil producing country and we abstract from the use of oil by oil producing countries. Similarly, time allocations satisfy

$$1 - L_t^j - N_t^j = 0, \quad j = a, b, o; \tag{9}$$

the unit endowment is divided between work and leisure.

We solve the model by exploiting the equivalence between competitive equilibria and Pareto optima. An optimum is the solution to the maximization of:

$$\sum_{t=0}^{\infty} \beta^t \left\{ \omega^a U(c_t^a, L_t^a) + \omega^b U(c_t^b, L_t^b) + \omega^o U(c_t^o, L_t^o) \right\} \tag{10}$$

subject to the constraints implied by equations (5)–(9). We approximate the first-order necessary conditions of the planning problem in a neighborhood of the steady state. We specify log-linear stochastic processes for the shocks and calibrate the functional forms of preferences and technology. The resulting linear dynamic rational expectations model is then solved using standard techniques (King, Plosser, and Rebelo (1988), for example).

3.2 Calibration

We explore the impact in the model of changes in the supply of oil. The specifics depend, to some extent, on the choice of parameters. We stay as close as possible to the values used by Backus, Kehoe, and Kydland (1994), but consider alternatives when they seem interesting. The complete set of parameter values is summarized in Table 4. A brief sketch of the logic behind these choices follows.

The functional form of preferences for the industrial countries is exactly as in BKK:

$$U(c, L) = [c^{\theta} L^{1-\theta}]^{1-\gamma} / (1-\gamma) \tag{11}$$

where $0 < \theta < 1$, $\gamma > -1$. Consumption is a CES function of the home and foreign manufactured good: $c = [\psi a^{1-\mu} + (1-\psi) b^{1-\mu}]^{1/(1-\mu)}$. The elasticity of substitution between the two goods is $1/\mu$.

With these preferences, the curvature properties of the utility function are completely determined by three parameters: the intertemporal substitution parameter (γ), the fraction of time spent in the workplace (governed by θ), and the willingness of individuals to substitute between domestic and foreign goods. We set the first two of these parameters, as well as the discount factor, the share of value added going to capital, and the depreciation rate, equal to values suggested by Kydland and Prescott (1982). The elasticity of substitution between domestic and foreign goods (σ) is set at 1.5 for both the consumption and investment aggregator. Costs of adjustment in capital accumulation are essentially set to zero by calibrating $\eta = -(D\phi(i/k)/D^2\phi(i/k))/(i/k)$, the elasticity of the investment-capital ratio with respect to Tobin's "q", to equal to 1000.

The two industrial countries are of equal size, but each consumes a disproportionate amount of its own good. Given that we abstract from non-traded goods, this last feature is needed to ensure that trade shares are not unrealistically high. For simplicity we assume that the oil producer likes a and b equally well so that $a^o = b^o$. The only remaining pieces of information we need to arrive at the steady-state values of quantity ratios are a trade share for the industrialized countries; the ratio of oil use to value added; and the fractions of imports for consumption and investment purposes. We set the first two parameters at 15% and 10% to approximate the share of trade in industrialized countries and the cost share of energy for the United States, respectively. We assume a common home bias for both consumption and investment goods so that the ratio of imports of consumption goods to investment goods equals the ratio of total expenditure on consumption goods to investment goods.

Turning to the oil market we consider features of both the supply and demand. The sensitivity of the demand for oil to changes in industrial country output depend, in part, on how oil enters the production function — one of the most studied and least resolved issues in empirical macroeconomics. Part of the difficulty stems from the different roles played by oil in industrialized countries. Nordhaus (1980) notes that many components of the physical capital stock are engineered in ways that makes substitution possible only when the existing capital is scrapped. A good example is the transportation sector, in which energy use depends largely on the fuel efficiency of the outstanding stock of automobiles. The energy consumption of this sector responded gradually to increases in oil prices during the 1970's, as old vehicles were gradually replaced with more fuel efficient models. Electricity generation is another sector in which infrastructure requirements can make energy substitution costly. Conservation would reduce the end-use demand, but the relationship between physical inputs of oil and other factors of production would be closer to Leontief than Cobb-Douglas. Berndt and Wood (1979) evaluate the conflicting evidence on capital and energy substitutability, in which estimates of complementarity and substitutability have been found. Jorgensen (1986, p 9) concludes that energy and capital are on the borderline between substitution and complementarity.

We follow Kim and Loungani (1992), in nesting capital and oil as a CES function within a Cobb-Douglas production function: $y = zN^{\alpha}[\eta k^{1-\nu} + (1-\eta)o^{1-\nu}]^{(1-\alpha)/(1-\nu)}$ with $0 < \alpha < 1$, $\nu > 0$, $\eta > 0$ and the elasticity of substitution between capital and oil equal to $1/\nu$.

We explore the sensitivity of our results to alternative parameterizations of capital-energy substitutability but our baseline choice of the parameter ν is 11, which translates into an elasticity of substitutability of 0.09, compared to the value 0.7 used by Kim and Loungani (1992). The lower elasticity seems more plausible for an investigation focusing on business cycles, while an elasticity approaching one might be more appropriate for analysis of the secular changes in energy use (Atkeson and Kehoe 1994, for example). We will see, in any case, that a high elasticity of substitution produces strongly counterfactual implications for the time series of prices and quantities in the world oil market. Basically, if other inputs into production are highly substitutable for oil (in a technological sense) it will be next to impossible to match the observed changes in the quantity of oil production that we observe in the data.

Oil supply behavior introduces some new issues. Much of the literature on the supply-side of the oil market bases supply or pricing decisions on optimal behavior given a particular market structure. See, for example, the excellent survey by Cremer and Salehi-Isfahani

(1991). We assume that OPEC supply decisions are taken as exogenous by agents in the two industrialized countries and the non-OPEC oil producer. To see how this assumption allows us to map existing data on OPEC and non-OPEC oil supply into our model consider the linearized version of our 'supply' function:

$$\hat{y}_t^o = \phi^o \hat{z}_t^o + (1 - \phi^o)\alpha \hat{N}_t^o \tag{12}$$

where ϕ^o is OPEC's market share which we treat as constant in the deterministic steady state. We associate z_t^o with OPEC oil production data, and allow $(1 - \phi^o)\alpha \hat{N}_t^o$ to be determined by the endogenous choice of effort by non-OPEC producers. Our formulation makes clear what an oil shocks is: a change in the supply of oil by OPEC, holding fixed the remaining exogenous variables (in our model home and foreign productivity). Note that there is no guarantee that we will match — or even come close to — either the total amount of oil produced in the world economy or its world relative price without putting some restrictions on the endogenous part of oil supply.

The empirical counterparts to the two components of world oil supply are displayed in Figure 2. Non-OPEC supply is measured from the axis to the dashed line while OPEC supply is the distance from the dashed to the solid line which is total world production measured in billions of barrels per year. We see that non–OPEC oil supply is quite smooth suggesting that supply is not highly responsive to either changes in OPEC production or the dramatic changes in the relative price of oil. For this reason, we adopt a preference specification for the non–OPEC oil producers consistent with an inelastic labor supply response to changes in real wages:

$$U(c, L) = c^{1-\gamma}/(1-\gamma) + \theta_L v(L). \tag{13}$$

We use the same intertemporal elasticity of substitution in consumption, $1/\gamma$, as for the industrialized countries, but set the elasticity $\xi_{LL} = LD^2v(L)/Dv(L)$ equal to -10 which when combined with the fraction of time spent devoted to market activity matches the upper bound on the labor supply elasticity of prime age males found in Pencavel (1986).

That leaves us with the shocks. Our world economy is driven by three disturbances: home productivity shocks, foreign productivity shocks, and oil supply shocks. Much of our analysis utilizes the quarterly OECD data, but since our oil production data is annual we use simulations from an annual version of the model to evaluate our characterization of the world oil market. Complete parameterization of the model requires that we also specify forcing processes consistent with the assumed length of a decision period. Thus we present evidence on both the quarterly and annual properties of the shocks.

The processes for technology shocks are based on observed properties of Solow residuals. In the quarterly model, we use the same symmetric bivariate first-order autoregressive process as Backus, Kehoe, and Kydland (1994):

$$\begin{pmatrix} z_t^a \\ z_t^b \end{pmatrix} = \begin{pmatrix} 0.906 & 0.088 \\ 0.088 & 0.906 \end{pmatrix} \begin{pmatrix} z_{t-1}^a \\ z_{t-1}^b \end{pmatrix} + \begin{pmatrix} \epsilon_t^a \\ \epsilon_t^b \end{pmatrix}, \tag{14}$$

where z_t^a and z_t^b are Solow residuals computed for the United States and Europe, respectively. The standard deviations of both innovations are 0.008325 and the correlation between them is 0.258. With annual data (computed by taken averages of quarterly productivity levels), we estimate a similar system for the period 1972 to 1989:

$$\begin{pmatrix} z_t^a \\ z_t^b \end{pmatrix} = \begin{pmatrix} 0.858 & 0.092 \\ 0.281 & 0.804 \end{pmatrix} \begin{pmatrix} z_{t-1}^a \\ z_{t-1}^b \end{pmatrix} + \begin{pmatrix} \epsilon_t^a \\ \epsilon_t^b \end{pmatrix}. \tag{15}$$

The standard deviations of the innovations are 0.0161 and 0.0190, respectively, and the correlation between them is 0.67. Following the approach taken for the quarterly model, we use a symmetric matrix A with the same eigenvalues as the matrix estimated in (15) to get:

$$A = \begin{pmatrix} 0.83 & 0.16 \\ 0.16 & 0.83 \end{pmatrix}. \tag{16}$$

In both cases, productivity shocks are persistent and exhibit two sources of positive correlation across countries: the correlation between the innovations and the positive off-diagonal elements of A.

In the absence of quarterly oil production data, we interpolate the annual OPEC production data (using a quadratic interpolation of the annual data from 1961 to 1991) and estimate a first-order autoregressive model for z_t^o:

$$z_t^o = 0.977 z_{t-1}^o + \epsilon_t^o \tag{17}$$

with $\sigma_{\epsilon^o} = 0.0099$. Estimating the same process over the OPEC period (1973:1 to 1986:4) we get almost the same coefficient, 0.971 but the standard deviation of the innovations rises to 0.0114.

Using the original annual series of OPEC production data from 1961 to 1991 we get:

$$z_t^o = 0.882 z_{t-1}^o + \epsilon_t^o, \tag{18}$$

with an innovation standard deviation of 0.0828. Note that the innovations of OPEC oil supply have a standard deviation 5 times larger than the annual productivity measures. We set OPEC's market share, ϕ^o, at 0.3, which is a lower bound on OPEC's share of world oil production.

3.3 Impulse Response Analysis

We use impulse response functions to illustrate the economic mechanisms that might give rise to the unstable correlations of trade prices and quantities documented in Section 2. To emphasize the roles played by different shocks, we trace out the dynamic responses of output, the terms of trade, and the trade balance to three different shocks: a positive innovation to home productivity; a positive innovation to world productivity; and a negative oil supply shock. The first two disturbances are set such that world output rises by 1 percent on impact while the third is set so that world output falls by 1 percent on impact. The magnitudes of the innovations needed to achieve this are: a 1.85% increase in home productivity, a 0.93% increase in productivity in both industrial countries, and a 22.1% decline in OPEC oil supply. Figure 3 presents the time paths of output, the terms of trade, and net exports in the home country (each in a separate column) in response to the three types of disturbances (each in a separate row).

Beginning with a home productivity shock, we see that output rises in the home country, as does its terms of trade, and the trade balance moves into deficit. The increase in output is due to the productivity led boom in which both effort and physical investment move above their steady state levels for an extended period of time. The terms of trade rises for two reasons. First, the home country is more productive and therefore produces more of its good. Because the country is also large it drives down the world relative price of its export (causing a deterioration of its terms of trade). Reinforcing this effect is the fact that an increase in world output drives up the relative price of oil which is an import of the home country (recall that the terms of trade is the import price index divided by the export price index). The response of the trade balance is familiar from the existing international real business cycle research where a persistent productivity led boom cases consumption and investment to rise by a sufficient amount to result in a trade deficit. Evidence on this channel may be found in empirical work by Glick and Rogoff (1995), and Sachs (1981).

The second panel shows the response of the home country to an equal increase in productivity in both industrialized countries. Again output rises but the terms of trade response is muted and the trade balance moves into deficit. The terms of trade effect is muted because the relative price of the two manufactured goods does not change when productivity increases simultaneously in the two industrialized countries. Without oil as an imported input into production and given the symmetry of the two industrial countries, common disturbances across countries would not alter the terms of trade. However, the relative price of oil rises due to increased oil demand and the terms of trade rises on this account. The trade deficit is a consequence of consumption smoothing by the oil producer and the impact of an investment boom in both industrialized countries.

The last panel shows the response of the home country to a reduction in OPEC oil supply of the magnitude needed to reduce world output by 1%. Note that despite the inelastic labor supply of non-OPEC oil producers, the OPEC supply shock must be very large to induce a worldwide recession; OPEC supply must fall by 22.1% to reduce world output by 1%. The decline needed would be even higher with a production function that allowed more substitutability between capital and oil or a more elastic labor supply on the part of non-OPEC suppliers. We see the terms of trade increase very dramatically, rising by 20% on impact (though it decays rapidly). The trade balance moves into deficit initially before turning into a surplus after about one year.

We can see from these experiments that the volatility of the terms of trade and its correlation with output depend crucially on the source of the disturbance. When productivity shocks are the dominant source of economic fluctuations we expect to see a positive correlation of the terms of trade and output for oil importers while the opposite correlation should prevail when oil shocks are dominant. Accompanying the shift to a greater role for oil shocks should be a dramatic increase in the variance of the terms of trade.

4. Oil and Business Cycles

We assess the role of oil in international business cycles in two ways. First, we consider how a change in the variability of oil shocks relative to productivity shocks alters the character of business cycles. Second, we use time series simulations to compare the actual time paths of economic variables to predictions of the model conditioned on the measured sequence of shocks over the period from 1971 to 1989.

4.1 Business Cycle Regimes

We use the term regime to refer to a shift in the time series properties of oil supply which we characterize as involving a change in OPEC's share of world oil production and a coincident change in the volatility of OPEC's oil supply. We refer to the OPEC regime as a time period during which OPEC's market share is 50% and the standard deviation of the innovations to OPEC oil supply is 0.0114; both having risen from benchmark values of 0.30 and 0.0099, respectively. Recall that these last two statistics are from our analysis of the time series properties of OPEC oil supply. The OPEC market shares are based on the data used in Figure 2, which shows a sharp rise in the world share of oil production accounted for by OPEC between 1960 and 1970. Pre-OPEC and post-OPEC regimes mark a return to "normalcy" in the sense that the market share and volatility of oil supply return to baseline parameter settings.

The first two columns of Table 5 report the cross-country averages of the individual country statistics that were reported in Table 2. We summarize the salient features of pre-OPEC and OPEC regimes as follows. The volatility of consumption and income are virtually unchanged across the two periods while investment volatility rises modestly. The volatility of the terms of trade rises sharply from 2.75 to 4.27, accompanied by almost no change in the volatility of the trade ratio. The correlation of the terms of trade and output falls from 0.08 to -0.12. The international comovement of output rises from -0.02 to 0.61 while investment comovement increases from 0.07 to 0.32.

Beginning with the pre-OPEC period, we see that the model does a reasonably good job in capturing volatility with the notable exception of the trade variables. Output is more volatile than consumption but less so than investment, both of which we see in international data. The presence of oil combined with exogenous fluctuations in its supply increases the standard deviation of the terms of trade from 0.48 in BKK to 1.71. This prediction as well as the volatility of the trade ratio (at 1.05) remain an order of magnitude below the standard deviations of their empirical counterparts in the pre–OPEC period of 5.64 and 2.75, respectively.

The model correctly predicts the sign of the correlation between output and net exports, output and the terms of trade, and the terms of trade and net exports. Again, the magnitudes are quite far off from their empirical counterparts: net exports are too sharply counter-cyclical, –0.55 in the model compared to –0.25 in the data; the terms of trade are too sharply pro-cyclical, 0.49 in the model compared to 0.08 in the data; and the correlation of the terms of trade and net exports is too low, –0.60 in the model compared to –0.18 in the data. The comovement of international output is about right, –0.02 in the data and 0.09 in the model. As with virtually all dynamic equilibrium models that feature complete markets and physical capital as the mobile factor, the international correlation of consumption is too high and the international correlation of investment is too low (negative in fact).

More interesting to us is that the model predicts a number of significant changes in the character of the international business cycles across the two regimes. The presence of large oil shocks increases the volatility of the terms of trade, investment, and the trade ratio while reducing the correlation of output and the terms of trade. Each of these could be anticipated from the impulse response analysis. They also mimic some of what we see in the data, although the magnitudes differ. The predicted increase in investment volatility is too small, rising from 4.61 to 4.75, while rising from 3.88 to 4.37 in the data. The predicted increase in the volatility of the trade ratio is a bit too large, rising from 1.05 to 1.14, while rising very modestly from 5.64 to 5.72 in the data. The predicted decline in the correlation of the terms of trade and output is dramatic in both the model, falling from 0.49 to 0.22, and in the data where it falls from 0.08 to –0.12. The international comovement of output, consumption and investment is virtually unchanged in the model (except for a modest increase in the case of investment) compared to quite dramatic increases in the data. However, the comovement of international business cycles continues into the post–OPEC period so it is unclear what to make of this shift in terms of oil's role.

4.2 Sensitivity analysis

We have used our model to suggest how output, the terms of trade, and the trade balance would evolve in response to different shocks over the post World War II period. While the model predicted an increase in the volatility of the terms of trade and a reduction in the comovement of output and the terms of trade during the OPEC period, it was less successful in matching the actual magnitudes of business cycle statistics. The question is whether the discrepancies that remain are sensitive to changes in the model's parameter values or theoretical structure.

Table 6 presents a detailed sensitivity analysis. The first two rows of the table present the cross–country averages of business cycle moments for all countries and countries that are consistently net importers (France, Italy, Germany, Japan, and the United States).

We begin our sensitivity analysis by stepping back and considering the moment predictions of the model in the presence and absence of each type of disturbance. Consider first a case in which variation in productivity is the only source of business cycle fluctuations. We see in Table 6 that this parameterization of the model is basically identical to the baseline version we used to describe the pre-OPEC period except that the correlation of output and the terms of trade is higher. The standard deviation of output is 1.59, only slightly lower than in the presence of oil supply shocks. While oil shocks alone can produce substantial variation in output (0.38), they are incapable of replicating the volatility of output that we see in the data, which averages about 1.58 across industrialized countries.

Examination of the other moments makes clear how dramatically the international business cycle is altered as we move from a productivity driven business cycle to a cycle driven only by oil shocks: the volatility of the terms of trade rises from 1.20 to 6.90; the correlation of the terms of trade and output falls from 0.76 to -0.94; the international correlation of output goes from 0.08 to almost unity; and the international correlation of investment goes from -0.73 to 1.00.

The next two experiments consider a larger trade share and a lower substitutability of goods a and b. Increasing the trade ratio makes the aggregator functions for consumption and investment more symmetric across countries. Our intuition is that this will make the model look more like the one sector model in terms of international comovement of consumption and investment thereby increasing consumption correlations while reducing investment correlations across countries. In fact the international correlation of consumption rises from 0.79 in the baseline to 0.97 while the international correlation of investment falls from –0.68 to –0.87.

Altering the substitutability of the two goods would be expected to increase price variability relative to quantity variability based on the relationship between relative prices and quantities in general equilibrium. However, the presence of oil complicates the relationship, since both the terms of trade and quantity ratio are now functions of all three goods. Making goods a and b less substitutable increases the variability of oil demand since output of the two industrialized countries will be more highly correlated in this case (output correlations increase from 0.09 to 0.18). As a result, the ratio of price variability to quantity variability actually falls modestly as we lower the substitution elasticity.

The next two experiments illustrate the effect on the economy of alterations in parameters that govern key features of the oil market, a less elastic demand for oil (approximately a Leontief relationship between oil and capital) and a more elastic oil supply, which we accomplish by using the same preferences for the oil producer that we use for the other two countries which gives rise to a more elastic labor supply.

Our intuition was that the first of these would give rise to changes in business cycles similar to an increase in oil supply shocks. We find that this is the case. The volatility of the terms of trade increases from 1.71 in the pre-OPEC benchmark to 2.68 while the correlation of output and the terms of trade falls from 0.49 to 0.33, and the correlation of the terms of trade and the trade balance falls from -0.60 to -0.91.

In contrast to this case, the more elastic oil supply case mimics quite closely the case without oil supply shocks. The reason for this should be obvious. If oil supply responds endogenously to offset OPEC quantity changes the total world oil supply will reflect mainly fluctuations in demand for oil driven by productivity changes. As such we would expect the relative price of oil to be less volatile and the terms of trade and output to be positively correlated as in the productivity model. In fact the terms of trade volatility drops from the pre-OPEC benchmark of 1.71 to 0.99 while the correlation of the terms of trade and output increases from 0.49 to 0.73. Lowering the energy cost share has qualitatively similar effects.

4.3 Simulations

We conclude our empirical analysis of the model by considering it's ability to match the actual time path of economic variables over the period in which oil supply disturbances were most dramatic. We focus our attention on four aggregates: U.S. GDP, European GDP, the relative price of crude oil, and the world production of oil.

The procedure we use to generate the simulations is as follows. We calibrate an annual version of the baseline model including the forcing processes of U.S. and European productivity and OPEC oil production data as described in the calibration section. We transform the linearized model into first–difference form and compute changes in the endogenous variables as functions of the actual sequences of changes in the three forcing processes.[1] The sequences of growth rates that result are accumulated and then filtered using the Hodrick–Prescott filter with a smoothing parameter of 100.

Figure 4 presents simulation results along with their empirical counterparts. Each column displays the time path of a different economic variable: oil prices, world oil production, U.S. output, and European output. The first row of the figure presents simulations using all three shocks while the second row presents simulations with only productivity shocks.

Beginning with the productivity driven model, in the lower panel of the figure, we see very modest oil price fluctuations and almost no variation in oil production. This is largely independent of the elasticity of oil supply — we get some additional oil price volatility when the preferences of the oil producer are such that labor supply is inelastic, but not much.

Turning to the upper panel which shows the simulations using all three shocks we see that adding oil supply disturbances increases both the oil price and oil production variation dramatically. We also see that the inelastic supply case (our benchmark parameterization) does a better job of matching the magnitude of fluctuations in these markets than does a more elastic supply specification.

More difficult to see in the Figure is much of a quantitative role for oil shocks in explaining the actual path of U.S. and European output beyond what is accounted for by the productivity shocks. Table 7 makes the differences across the simulations transparent by comparing the standard deviations of the sample paths and their RMSE relative to the actual data.

We see that the baseline parameterization that incorporates both shocks does are very good job of matching the standard deviations of the four series over the period 1970 to 1989. The standard deviation of world oil output is predicted to be 4.12 compared to 4.43

in the data while the comparable statistics for the relative price of oil are 27.49 and 28.99. The standard deviations of the simulated output cycles of the U.S. and Europe are also very close to their empirical counterparts. Turning to the RMSE statistics which also take into account the ability of the simulated paths to actually track the data we see similar degrees of success with one notable exception. As is apparent in Figure 4, the model misses the timing of the oil price changes and underpredicts the magnitude of the first oil price shock in 1973.

The model is much less successful when only one type of disturbance is used to generate the simulations. Continuing with the baseline simulation, we see that the productivity model does a reasonable job of matching output variability but a very poor job of matching the volatility of either the relative price or quantity of oil that we observe over the sample period of study. In particular, the productivity driven model can only account for about 1% of the observed relative price variation in oil. The case with only oil shocks has the opposite problem with only modest output variation and almost exactly the right amount of oil price and quantity variation. We see also from Table 7 that moving to a more elastic supply of oil reduces the explanatory power of the model in all dimensions, but particularly in terms of matching the behavior of the oil market.

We conclude this section with a brief discussion of the contribution of oil shocks and productivity shocks to the overall U.S. business cycle. Finn (1995) estimates energy price shocks contribute between 7.47 and 18.75 percent to U.S. output volatility compared estimates of between 2.6% and 12% produced by Kim and Loungani (1992).[2]

Recall that theirs are partial equilibrium models so that the persistence and volatility of energy prices are calibrated to match the actual data. In our model the price of oil is endogenously determined, given the path of productivity and oil supply (both exogenous and endogenous parts). The parameterization that uses an inelastic labor supply for the endogenous part of oil supply matches the volatility of the relative price of oil almost exactly and is therefore the natural one to compare to the variance decompositions performed by Finn and Kim and Loungani. Using this parameterization we estimate that OPEC oil supply changes account for about 17.4% of U.S. output variation, a figure at the high end of existing estimates.

Our novel contribution to the literature that aims to evaluate the relative important of oil price movements for U.S. business fluctuations is in providing evidence that their contribution is more likely to be at the high end of estimates coming out of existing calibration exercises. We arrive at this conclusion based on the observation that such estimates are more plausible because they obtain in parameterizations that produce both the right amount of price and quantity volatility in the world market for crude petroleum.

5. Conclusions

Recently a number of authors have examined the role of trade in intermediate inputs for the character of international business cycles. Most of these studies have focused on the role of intermediate inputs in producing more comovement in international output than found in one sector models (see for example, Ambler, Cardia and Zimmerman (1995) or Costello and Praschnik (1993)). Exceptions are Kouparitsas (1995), who demonstrates that the pattern of North-South specialization gives rise to volatile terms of trade across developing and developed economies, and Boileau (1997), who shows that trade in equipment can add volatility to the terms of trade of industrialized countries.[3]

We found that Hamilton's (1983) view of the U.S. business cycle has important implications for international business cycles too. We documented how changes in the terms of trade in major industrialized countries during the period 1972 to 1987 were driven primarily by the dramatic changes in the relative price of a single commodity — oil. Using a suitably adapted dynamic equilibrium model of international business cycles, we found that unstable relationships between relative prices and quantities are to be expected when oil price shocks play vastly different roles across time periods. Once we control for the sources of shocks driving the terms of trade the comovement and volatility of the terms of trade, output, and the trade balance become less puzzling.

Given the timing of the oil price shocks, the increased volatility in the terms of trade since Bretton Woods seems largely due to the increased volatility in the relative price of oil rather than the increased volatility of nominal or real exchange rates. A direction that we plan to pursue in future research is to relate the properties of the terms of trade to similar properties of real exchange rates documented by Mussa (1986) and others. Although the two relative prices are conceptually and operationally different, preliminary analysis shows that they are positively correlated in the data. The question then is whether the striking change in the variability of real exchange rates between the pre-1972 Bretton Woods fixed exchange rate regime and the more recent floating exchange rate period is related to the similar change in the behavior of oil prices. Perhaps further work will tell us more.

Acknowledgements

We are indebted to two anonymous referees for helpful comments on an earlier version of the paper. Betty Dow at the World Bank provided generous help with oil data.

Appendix

Data Sources

Quarterly macroeconomic aggregates. The quarterly macroeconomic aggregates: gross domestic product, consumption, investment, exports, imports, civilian employment are the same series used in Backus, Kehoe, and Kydland (1995). The countries are Australia, Canada, France, Germany, Italy, Japan, the UK, the US, and Europe (an aggregate constructed by the OECD). The sample periods for the quarterly data are given in note to Table 1.

Annual macroeconomic aggregates and fuel shares of trade. The annual macroeconomic data (US GDP and Economic Community GDP) are taken from the World Bank's *World Tables: Fall 1990 Update*, which is available in machine-readable form as the *STARS database.* Annual fuel imports and exports are taken from the same source. Imports are measured on a customs basis and include freight, while exports are free on board.

Oil production data. From 1970 to 1989, the annual world production data are from Knight-Rider Commodity Research Bureau's *Commodity YearBook*, various issues. The data are found in the table entitled: "World Production of Crude Petroleum, by Specified Countries." From 1973 to 1989, the OPEC crude oil production data are from and the Energy Information Adminstration's *Monthly Energy Review, December 1995.* The data are found in Table 10.1b, "World Crude Oil Production." Figures for 1970-1972 are from the same CRB Table as the world totals, but among OPEC members include only: Libya, Iran, Iraq, Kuwait, Indonesia, Saudi Arabia, and Venezuela (the Monthly Energy Review also included Gabon, Qatar, United Arab Emirate which together accounted for about 10% of OPEC production in 1973). The CRB totals are scaled up by factor 1.1 to make the series comparable across periods.

Oil price data. The annual data are nominal price data from the World Bank's *Commodity Trade and Price Trends.* The index is constructed using quantity weights for a larger number of markets than is the monthly index described below. The sample runs from

1955 to 1990. The monthly nominal oil price data were provided by Betty Dow of the World Bank. The index is a simple average of the U.S. dollar prices for crude in three major international markets: Brent, Dubai, and West Texas. The sample is montly from 1950 to 1995. The series is converted to quarterly taking the average of the monthly index within the quarter. Both price series are converted to real terms using the quarterly U.S. CPI-U inflation index.

Estimates of the Non-Fuel Terms of Trade

The estimates of the non-fuel terms of trade are constructed from the overall terms of trade and the shares of non-fuel exports and imports as follows. Let P_t^m be the price index for imports in the current period, consisting of both fuel and non-fuel commodities with prices and quantities: P_t^{mf}, P_t^{mnf}, Q_t^{mf}, and Q_t^{mnf}, respectively. We can write the current price of imports relative to a base period as:

$$\begin{aligned}\frac{P_t^m}{P_0^m} &= \frac{P_t^{mf}Q_t^{mf} + P_t^{mnf}Q_t^{mnf}}{P_o^{mf}Q_t^{mf} + P_o^{mnf}Q_t^{mnf}} \\ &= \frac{P_t^{mnf}}{P_0^{mnf}}\frac{1}{S_t^{mnf}}\frac{P_0^{mnf}Q_t^{mnf}}{P_0^{mf}Q_t^{mf} + P_0^{mnf}Q_t^{mnf}} \\ &= \frac{P_t^{mnf}}{P_0^{mnf}}\frac{1}{S_t^{mnf}} \cdot Q_t^m\end{aligned}$$

where S_t^{mnf} is the share of non-fuel imports relative to total imports (in current value terms) and Q_t^m is the ratio of the quantity of non-fuel imports to the quantity of total trade (both evaluated at base period prices). Repeating this process for the export price index and taking the ratio of the import price index to the export price index to the right-hand-side we have:

$$\frac{P_t^{mnf}}{P_t^{xnf}} = \frac{P_t^m}{P_t^x}\frac{S_t^{mnf}}{S_t^{xnf}}\frac{Q^m}{Q^x}$$

Given the data at hand, we must ignore the time variation in the physical quantity of fuel to non-fuel trade in our estimate setting $Q_t^i = Q^i, i = m, x$.

References

S. Ambler, E. Cardia, and C. Zimmermann, 1995, Intermediate goods and the international business cycle, mimeo, Centre de Recherche sur l'Emploi et les Fluctuations Economiques, Universite du Quebec.

Atkeson, A. and P. J. Kehoe, 1994, Putty-clay capital and energy, Federal Reserve Bank of Minneapolis Staff Report.

Backus, D. K., P. J. Kehoe, and F. E. Kydland, 1994, Dynamics of the trade balance and the terms of trade: the J-curve?, American Economic Review, 84, 84-103.

Backus, D., P. Kehoe, and F. Kydland, 1995, International business cycles: theory and evidence, in Thomas F. Cooley, ed., Frontiers of business cycle research, (Princeton University Press, Princeton.)

Bidarkota, P. and M. J. Crucini, 1997, Commodity prices and the terms of trade, Ohio State University Working Paper 98-01, December.

Baxter, M. and M. J. Crucini, 1995, Business cycles and the asset market structure of foreign trade, International Economic Review 36, 821-854.

Berndt, E. R. and D. O. Wood, 1979, Engineering and econometric interpretations of energy-capital complementarity, American Economic Review 69, 3, 342-354.

Boileau, M., 1997, Trade in capital goods and the volatility of net exports and the terms of trade, mimeo, University of Toronto.

Costello, D. M. and J. Praschnik, 1993, Intermediate goods and the transmission of international business cycles, mimeo, University of Western Ontario.

Cremer, J. and D. Salehi-Isfahani, 1991, Models of the oil market, (Harwood Academic Publishers, New York.)

Crucini, M. J., and J. Kahn, 1996, Tariffs and aggregate economic activity: lessons from the Great Depression, Journal of Monetary Economics 38, 427-468.

Finn, M., 1995, Variance properties of Solow's productivity residuals and their cyclical implications, Journal of Economic Dynamics and Control 19, 1249-1281.

Glick, R. and K. Rogoff, 1995, Global versus country-specific productivity shocks and the current account, Journal of Monetary Economics 35, 159-192.

Hamilton, J., 1983, Oil and the macroeconomy since World War II, Journal of Political Economy 91, 228-248.

Hodrick, R., and E. Prescott, 1997, Post-war U.S. business cycles: an emprical investigation, Journal of Money, Credit, and Banking 29, 1-16.

Hogg, R. V. and E. A. Tanis, 1983, Probability and statistical inference, (Macmillan Publishing, New York.)

Kim, I. and P. Loungani, 1992, The role of energy in real business cycle models, Journal of Monetary Economics 29, 173-189.

King, R. G., C. I. Plosser and S. T. Rebelo, 1988, Production, growth, and business cycles: I. The basic neoclassical model, Journal of Monetary Economics 21, 195-232.

Kouparitsas, M., 1997, North–south terms of trade: an empirical investigation, Federal Reserve Bank of Chicago, Working Paper No. 97-5, September.

Kouparitsas, M., 1995, North-south asset markets and business cycle transmission, mimeo.

Kydland, F. and E. Prescott, 1982, Time to build and aggregate fluctuations, Econometrica 50, 1345-1370.

Jorgensen, D. W., 1986, The great transition: energy and economic ehange, The Energy Journal 7, 3, 1-13.

Mussa, M., 1986, Nominal exchange rate regimes and the behavior of real exchange rates: evidence and implications, Carnegie-Rochester Conference Series on Public Policy 25, 117-214.

Nordhaus, W. D., 1980, Oil and economic performance in industrialized countries, Brookings Papers on Economic Activity 2, 341-399.

Pencavel, J., 1986, Labor supply of men: a survey, in: Orley Ashenfelter and Richard Layard, eds., Handbook of labor economics (North-Holland, Amsterdam).

Rotemberg, J. J., and M. Woodford, 1996, Imperfect competition and the effects of energy price increases on economic activity, Journal of Money, Credit, and Banking 28, 549-577.

Sachs, J., 1981, The current account and macroeconomic adjustment in the 1970's, Brookings Papers on Economic Activity 12, 201-268.

Stockman, A. C. and L. L. Tesar, 1995, Tastes and technology in a two-country model of the business cycle: explaining international comovements, American Economic Review, 85, pp. 168-185.

Endnotes

1. One limitation of our analysis is that we do not adjust the computation of the Solow residual to account for the mismeasurement induced by variation in oil use. The presumption that oil and capital are not highly substitutable in the short–run would tend to limit the quantitative impact of such adjustments. We currently do not have energy use available on an international basis so we leave this as a topic for future research.

2. Both of these studies assume perfect competition, for a quantitative analysis of oil price shocks on the U.S. economy under imperfect competition see Rotemberg and Woodford (1996).

3. In a separate paper Kouparitsas (1997) demonstrates how variation in the price of primary commodities (fuel and non–fuel primary commodity price indices) relative to manufactures account for most of terms of trade variation in developing countries and much of it in industrialized countries. Bidarkota and Crucini (1997) show that changes in developing country's terms of trade are typically dominated by fluctuations in the world price of a single primary commodity.

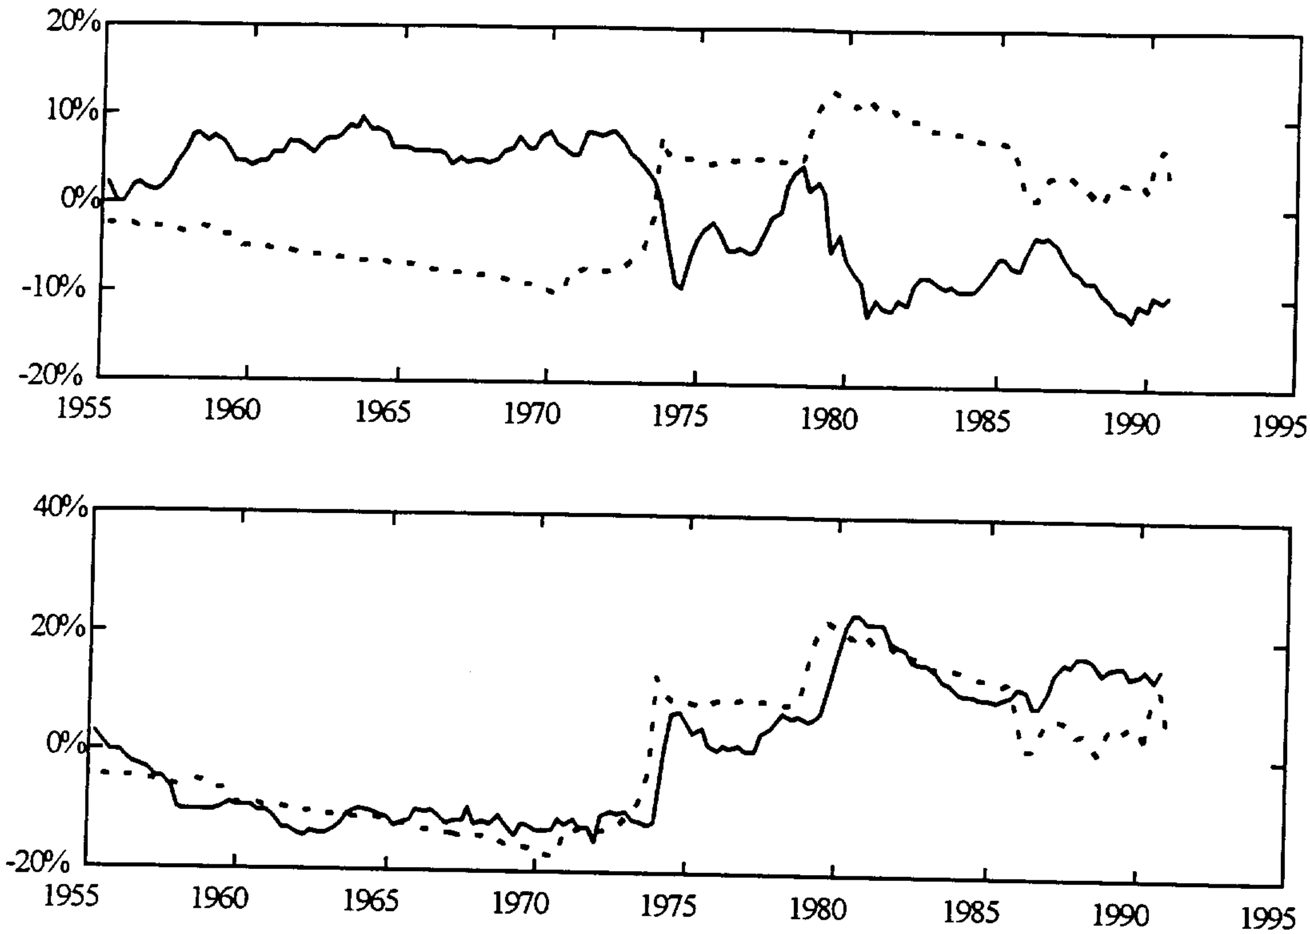

Fig. 1. The figure presents the terms of trade expressed as the ratio of the implicit deflator for imports to exports (solid lines) for Canada (upper panel) and the United States (lower panel). The dashed line is the international price of crude petroleum expressed in constant U.S. dollars (see the data appendix for the details of its construction) and normalized such that the sample standard deviation matches the standard deviation of the terms of trade with which it is plotted.

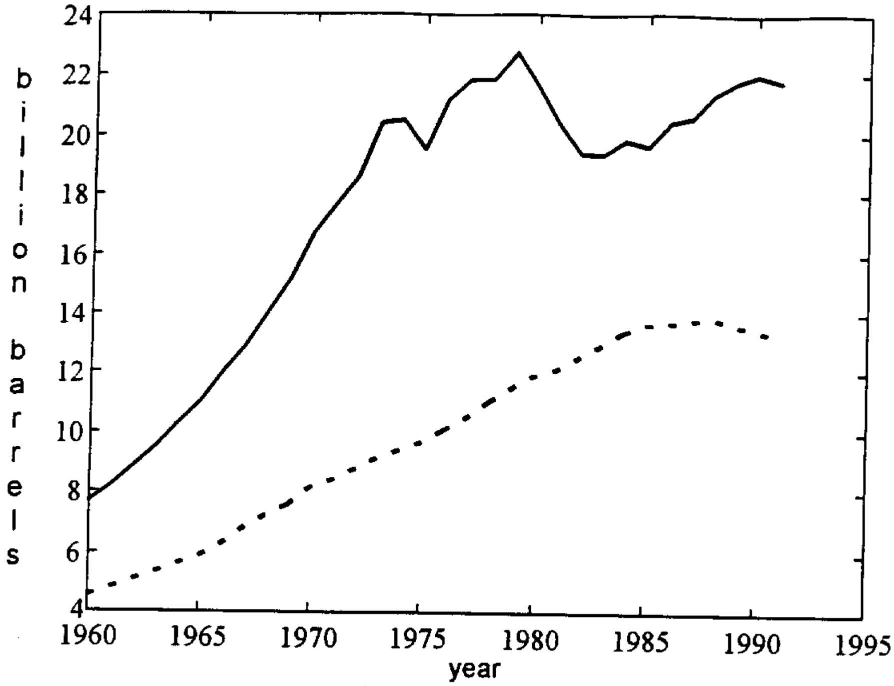

Fig. 2. The figure presents the annual production volume of crude petroleum for all major producers (solid line) and for non-OPEC countries (dashed line).

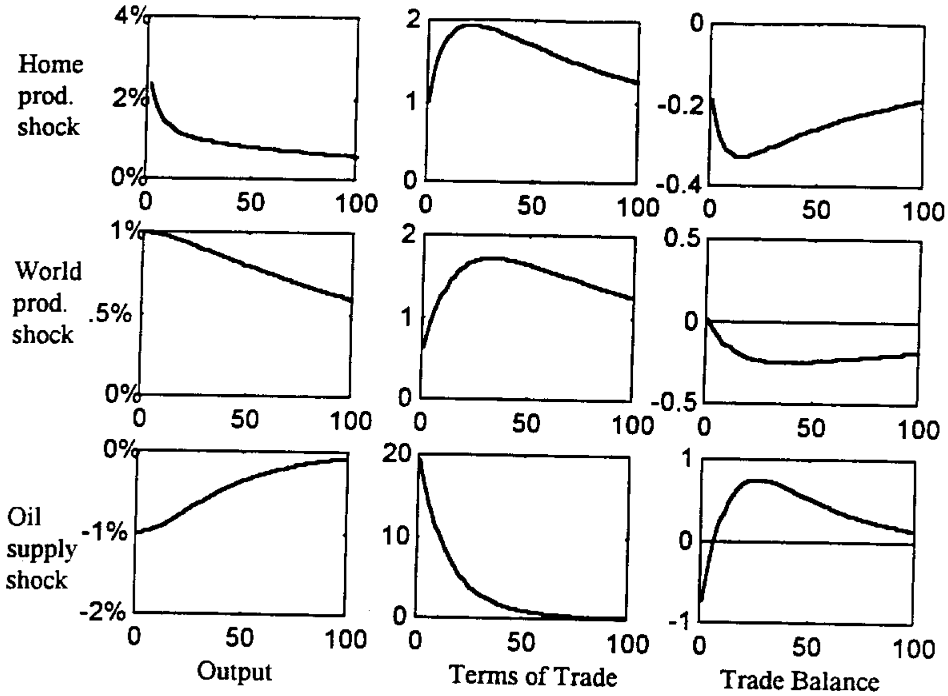

Fig. 3. The figure presents the impulse responses of the home country to a 1% increase in home productivity (upper panel), a 1% increase in both home and foreign productivity (middle panel), and an oil supply reduction of sufficient size to lower world output by 1% (lower panel).

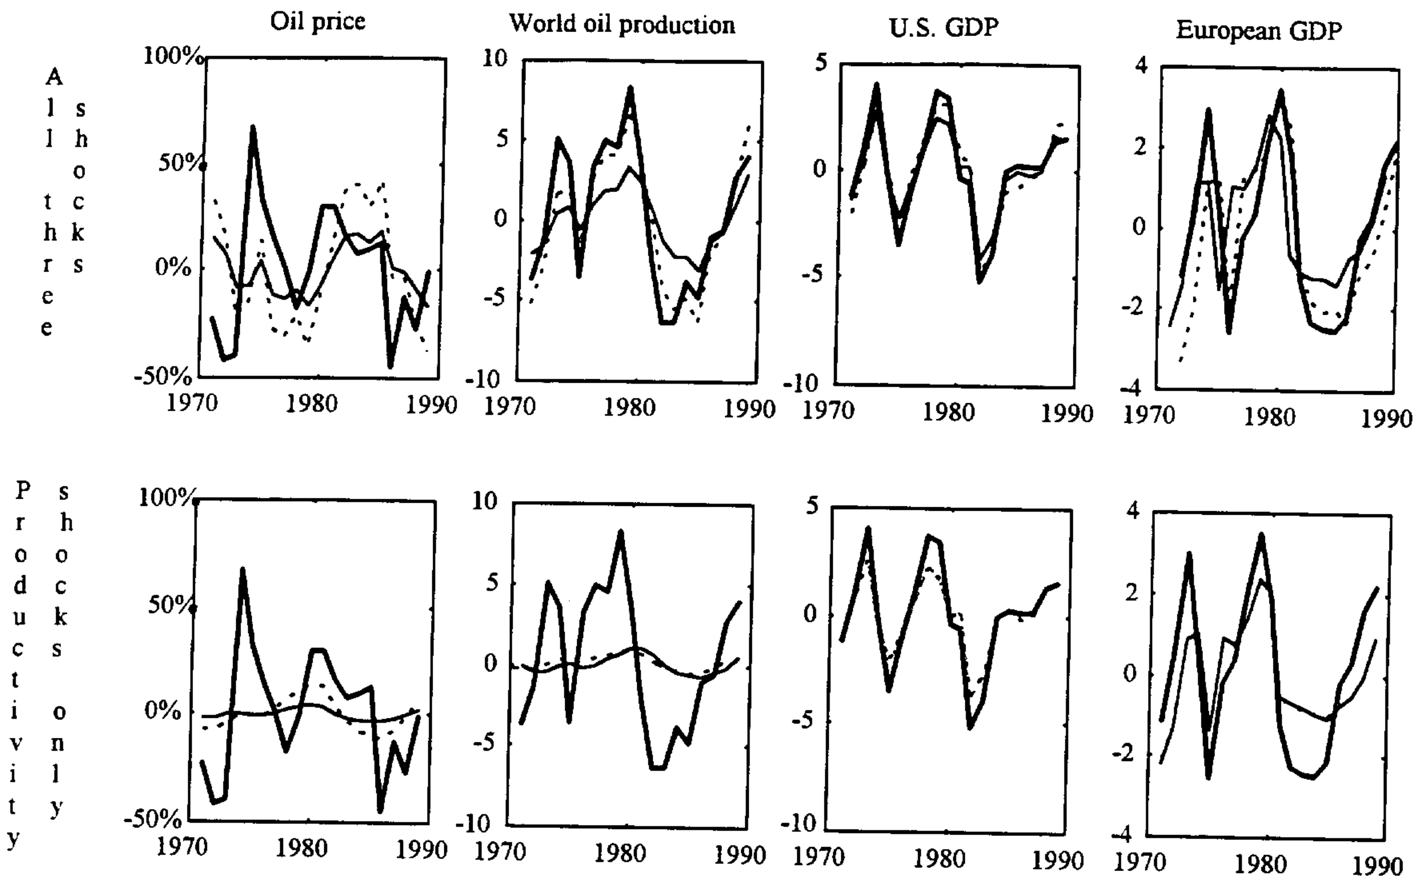

Fig. 4. Simulation results from the model using actual productivity changes in the U.S. and Europe and oil supply shocks (OPEC production changes) as inputs. The upper panel presents the results with all three shocks simultaneously while the lower panel shows the model predictions when only productivity disturbances (for the U.S. and Europe) are used in the simulation. All variables are percentage deviations from HP trends. The thick solid lines are actual sample paths. Simulations are performed over the period 1971-1989 using the baseline calibration with inelastic labor supply (dashed line) and elastic labor supply (thin solid line).

Table 1
Business Cycle Statistics for Eight OECD Countries

	Standard deviation					Cross-correlation of:					
Country	y	c	i	x/m	p	(y,nx)	(y,p)	(p,nx)	(y,y^{us})	(c,c^{us})	(i,i^{us})
Australia	1.54	1.08	3.79	9.94	5.21	-0.19	-0.30	-0.07	0.32	-0.10	0.19
Canada	1.53	1.36	4.54	4.60	2.44	-0.43	-0.11	0.06	0.70	0.53	0.11
France	0.89	0.88	2.64	3.66	3.50	-0.30	-0.14	-0.51	0.43	0.40	0.25
Germany	2.17	2.30	5.63	3.90	2.61	-0.05	-0.09	-0.00	0.37	0.23	0.34
Italy	1.69	1.32	3.28	4.89	3.50	-0.68	0.38	-0.66	0.40	0.03	0.31
Japan	1.60	1.35	3.94	7.29	5.68	-0.23	-0.09	-0.50	0.21	0.28	0.21
United Kingdom	1.49	1.61	3.62	3.94	2.64	-0.25	0.22	-0.54	0.46	0.35	0.31
United States	1.70	1.22	5.29	6.73	2.90	-0.30	-0.08	0.28	1.00	1.00	1.00

Notes: Statistics are based on Hodrick–Prescott–filtered data. Variables are y, real gross national production; c, real consumption; i, real investment; x, real exports; m, real imports; p, ratio of import price index to export price index; nx nominal exports less nominal imports divided by nominal gross national product. Except for the net export ratio all variables are transformed to logarithms before filtering. The sample periods are: Australia, 1960:1 to 1990:3; France, 1970:1 to 1990:3; Germany, 1968:1 to 1990:3; Italy, 1970:1 to 1990:2; Japan, 1955:2 to 1990:3; United Kingdom, the United States and Canada, 1955:1 to 1990:3. Source: OECD, *Quarterly National Accounts*.

Table 2
Business Cycle Statistics by Subperiod

Country	Standard deviation: y	c	i	x/m	p	Cross-correlation of: (y,nx)	(y,p)	(p,nx)	(y,y^{us})	(c,c^{us})	(i,i^{us})
Panel A: Pre-OPEC Period											
Australia	1.60	1.25	2.94	11.33	4.99	-0.38	-0.30	-0.05	-0.13	0.05	0.25
Canada	1.49	1.37	4.51	4.60	1.56	-0.60	-0.24	0.16	0.57	0.64	0.28
France	0.59	0.72	1.60	2.03	2.57	0.30	-0.47	-0.79	0.20	0.37	0.85
Germany	3.72	4.29	9.21	4.81	1.97	0.11	-0.10	0.52	-0.15	-0.13	-0.08
Italy	1.31	0.61	1.44	4.17	3.91	-0.72	0.77	-0.57	-0.64	-0.16	-0.63
Japan	1.85	1.26	4.49	7.97	3.29	-0.24	0.06	-0.63	-0.17	0.06	-0.21
United Kingdom	1.27	1.23	3.18	4.06	2.32	-0.31	0.47	-0.39	0.18	0.29	0.08
United States	1.34	0.90	3.67	6.13	1.39	-0.12	0.44	0.27	1.00	1.00	1.00
Panel B: OPEC Period											
Australia	1.61	0.93 –	4.19 +	8.55 –	5.19	0.04 +	-0.30	-0.25	0.56 +	-0.18	0.18
Canada	1.69	1.44	4.72	5.02	3.39 +	-0.25 +	-0.02	-0.01	0.80 +	0.47	-0.02 –
France	1.00 +	1.01	2.99 +	4.09 +	3.97	-0.37 –	-0.20	-0.53	0.44	0.42	0.19 –
Germany	1.64 –	1.41 –	4.39 –	3.73	2.80 +	-0.19	-0.23	-0.14 –	0.80 +	0.61	0.63 +
Italy	1.92	1.55 +	3.83 +	5.44	3.54	-0.67	0.31 –	-0.69	0.46 +	0.01 +	0.35 +
Japan	1.42 –	1.59 +	3.63 –	7.26	7.88 +	-0.20	-0.28 –	-0.50	0.62 +	0.43	0.58 +
United Kingdom	1.80 +	2.05 +	3.89	3.66	3.12 +	-0.15	0.00 –	-0.69	0.59 +	0.35 +	0.45 +
United States	2.21 +	1.63 +	7.36 +	8.03 +	4.28 +	-0.42	-0.24 –	0.37	1.00	1.00	1.00

Table 2
Business Cycle Statistics by Subperiod
(continued)

Country	Standard deviation					Cross-correlation of:					
	y	c	i	x/m	p	(y,nx)	(y,p)	(p,nx)	(y,y^{us})	(c,c^{us})	(i,i^{us})
	Panel C: Post-OPEC Period										
Australia	1.09	0.97	4.92	9.52	5.71	-0.44	-0.36	0.53	0.84	0.74	0.86
Canada	1.01	0.96	4.00	2.66	1.41	-0.51	-0.59	0.18	0.93	0.79	0.84
France	0.60	0.33	1.49	2.79	2.05	0.50	0.89	0.50	0.22	-0.04	0.60
Germany	0.93	0.86	3.20	2.34	2.50	0.08	0.78	-0.16	-0.32	0.02	0.06
Italy	0.55	0.37	1.14	2.59	2.43	-0.70	0.47	-0.46	0.75	0.57	0.53
Japan	0.82	0.63	1.72	2.64	4.86	-0.49	0.58	-0.81	0.02	-0.20	0.23
United Kingdom	1.15	1.36	4.59	4.65	1.79	-0.63	0.51	0.09	0.57	-0.41	0.52
United States	0.91	0.83	1.60	3.36	1.63	0.02	0.50	-0.67	1.00	1.00	1.00

Notes: Each of the three panels are sample statistics computed over different periods. The first sample begins with the first available obersvation and ends at 1972:4, the second sample is from 1973:1 to 1986:4 and the last sample is from 1987:1 to the last available observation. In the second panel a "+" (-) indicates a statistically signficant increase (decrease) in the moment. The variances are tested for shifts at the 5% level using critical values F(n-1,m-1) where n is the number of observations in the OPEC regime (56) and m is the number of observations available before 1973:1 which varies by country (n and m are reverse to test for statistically significant declines in the variances). The correlations are indicated as significantly changed across periods if the 80% confidence intervals for the two periods do not overlap. The intervals are constructed as $[1+\rho-(1-\rho)e^{x}]/[1+\rho+(1-\rho)e^{x}] \geq \rho \leq [1+\rho-(1-\rho)e^{-x}]/[1+\rho+(1-\rho)e^{-x}]$ where $x = 2c/\sqrt{n-3}$, c is the critical value (here 1.282), and n is the number of observations (see Hogg and Tanis (1983), pg. 451). See the notes to Table 1 for variable definitions and exact sample periods.

Table 3
Oil Prices and the Terms of Trade

				Terms of Trade			
	Net fuel trade share (annual data)			Correlation with oil price (quarterly data)		Standard deviation (annual data)	
Country	1970	1975	1987	Level	HP–cycle	Overall	Ex. Energy
Australia	-0.5	3.0	13.3	0.17	-0.60	8.10	8.22
Canada	1.2	2.8	5.5	-0.47	-0.47	6.51	6.05
France	-10.4	-20.6	-9.2	0.49	-0.05	6.08	3.60
Germany	-5.0	-12.5	-6.9	0.45	0.09	7.36	4.95
Italy	-10.1	-22.6	-11.6	0.60	0.25	8.02	5.03
Japan	-20.2	-44.8	-20.4	0.50	0.14	12.77	6.28
United Kingdom	-8.6	-15.9	3.1	-0.25	0.36	7.63	5.23
United States	-3.3	-21.4	-11.5	0.63	0.16	5.99	3.25

Notes: The correlation of the terms of trade and relative price of oil is computed using quarterly data over the OPEC period, 1973:1 to 1986:4. The net fuel trade shares are annual, computed as the ratio of the nominal trade balance in fuel divided by the average level of nominal imports plus exports for the year indicated above each column and expressed as a percentage. The last two columns report the annual standard deviation of the terms of trade in raw terms (ratio of import to export price index) and our estimate adjusted to exclude the impact of relative fuel prices as described in the technical appendix. Both series are filtered using the Hodrick–Prescott (1997) filter with smoothing parameter equal to 100.

Table 4
Benchmark Parameter Values

	Description	Symbol	Parameter Value	
			Quarterly model	Annual model
Preferences	Discount factor	β	0.99	0.96
	Intertemporal substitution	$1/\gamma$	0.5	0.5
	Atemporal substitution	$1/\mu$	1.5	1.5
	Fraction of time spent working	$\bar{N}$	0.3	0.3
Technology	Labor's share	α	0.64	0.64
	Depreciation rate of capital	δ	0.025	0.10
	Cost of adjustment parameter	η	1000	1000
Trade	Industrial country shares	s_x	0.15	0.15
Oil sector	Energy cost–to–value added	ξ	0.10	0.10
	Elasticity of substitution k,o	$1/\nu$	0.09	0.09
	OPEC's share of oil output	ϕ^o	(0.3, 0.5)	0.5
	Labor parameter	α	0.64	0.64
Productivity shocks	Persistence		0.906	0.83
	Spillover		0.088	0.16
	Innovation variance		0.00825	NA
OPEC oil supply	Persistence		0.97	0.88
	Innovation variance		(0.0099, 0.0114)	NA

Note: The standard deviations of innovations to OPEC oil supply differ across regimes as shown within the parentheses with the first entry referring to the non–OPEC regimes and the second entry referring to the OPEC regime.

Table 5
OPEC and the International Business Cycle

	Country Averages		Model Predictions	
Country	1955:3 1972:4	1973:1 1987:1	Pre–OPEC	OPEC
Standard deviation				
Output	1.65	1.66	1.60	1.60
Consumption	1.45	1.45	1.07	1.09
Investment	3.88	4.37	4.61	4.75
Trade ratio	5.64	5.72	1.05	1.14
Terms of trade	2.75	4.27	1.71	3.00
Cross–correlations				
Output and net exports	-0.25	-0.28	-0.55	-0.48
Output and terms of trade	0.08	-0.12	0.49	0.22
Terms of trade and net exports	-0.18	-0.31	-0.60	-0.59
Correlations with U.S. counterpart				
Output	-0.02	0.61	0.09	0.09
Consumption	0.15	0.30	0.79	0.80
Investment	0.07	0.34	-0.68	-0.59

Notes: The country averages are simple averages of the statistics in the first and second panels of Table 2. The model predictions are population moments for HP filtered results.

Table 6
Sensitivity Analysis

	Standard deviation					Domestic comovement			International comovement		
	y	c	i	x/m	p	(y,nx)	(y,p)	(p,nx)	(y,y^{us})	(c,c^{us})	(i,i^{us})
International data											
All countries	1.58	1.39	4.09	5.62	3.56	-0.30	-0.03	-0.24	0.42	0.25	0.25
Oil importers	1.61	1.42	4.16	5.29	3.64	-0.31	0.00	-0.28	0.35	0.23	0.28
Model predictions											
pre–OPEC	1.60	1.07	4.61	1.05	1.71	-0.55	0.49	-0.60	0.09	0.79	-0.68
No oil shocks	1.59	1.08	4.55	1.04	1.20	-0.58	0.76	-0.69	0.08	0.79	-0.73
No productivity shocks	0.38	0.39	3.45	1.19	6.90	0.56	-0.94	-0.79	0.98	1.00	1.00
High trade share	1.52	1.02	7.25	2.31	1.17	-0.33	0.74	-0.39	0.20	0.97	-0.87
Low substitution across goods	1.53	1.10	4.37	1.22	1.74	-0.52	0.58	-0.62	0.18	0.70	-0.65
Inelastic oil demand	1.55	1.14	4.43	1.05	2.68	-0.25	0.33	-0.91	0.05	0.81	-0.77
Elastic oil supply	1.62	1.03	4.76	0.89	0.99	-0.62	0.73	-0.72	0.12	0.77	-0.58
Low energy cost share	1.58	1.10	4.39	1.16	1.47	-0.39	0.77	-0.55	0.09	0.73	-0.73

Note: The parameters that change across cases are: Pre-OPEC benchmark (see Table 4); No oil shocks, $\sigma^2_{\epsilon^o} = 0.8325 \times 10^{-6}$; No productivity shocks, $\sigma^2_{\epsilon^j} = 0.8325 \times 10^{-6}$, $j = a, b$; High trade share, $s_x = 0.5$; Low substitution across goods, $1/\gamma = 0.1$; Inelastic oil demand, $\nu = 100$. Elastic oil supply, uses the same preference specification for the oil producer as for the industrial countries which gives rise to a more elastic labor supply curve. Low energy cost share uses an energy cost share of 0.05.

Table 7
Simulation Results

Variable	Data	Baseline Parameterization: Both shocks	Baseline Parameterization: Productivity shocks	Baseline Parameterization: Oil shocks	Elastic Oil Supply: Both shocks	Elastic Oil Supply: Productivity shocks	Elastic Oil Supply: Oil shocks
Panel A: Standard deviations							
World oil output	4.43	4.12	0.45	4.12	1.92	0.55	1.91
World oil price	28.99	27.49	7.76	28.98	11.95	2.39	12.93
U.S. GDP	2.42	2.21	1.61	1.01	1.81	1.62	0.47
European GDP	1.95	1.98	1.26	0.86	1.49	1.24	0.37
Panel B: Root Mean Squared Errors							
World oil output		1.62	4.35	1.62	2.91	4.37	2.93
World oil price		37.53	26.63	38.25	30.92	28.52	30.88
U.S. GDP		0.77	0.90	2.07	0.74	0.89	2.26
European GDP		1.13	1.25	1.41	1.12	1.23	1.78

Notes: The data and simulations are annual from 1970 to 1989. The GDP data are from the World Bank, *STARS database* with European GDP being an aggregate of each EC member's GDP. Trends are removed from the data and simulations using the Hodrick–Prescott (1997) filter with smoothing parameter set equal to 100. Simulations are produced for a case in which U.S. and European productivity and OPEC oil supply shocks are included, labelled *Both shocks*; a case in which only U.S. and European productivity shocks are included, labelled *Productivity shocks*; and a case in which only OPEC oil supply shocks are included, labelled *Oil shocks*. The parameterization *Elastic Oil Supply* uses the same preference specification for the oil producer as for the industrial countries which gives rise to a more elastic labor supply curve.