Oil Shocks: Shaking the Fundamentals or Our Confidence?*

by

Luís Aguiar-Conraria

Department of Economics, Cornell University

NIPE, Universidade do Minho

Yi Wen

Department of Economics, Cornell University

1 Introduction

Despite 30 years of research since the first oil crises in 1973, how exactly a rise in oil price causes nationwide economic downturns is still a unsolved mystery. Oil as an input for the entire US economy and the manufacturing sector accounted for not more than 1% of total production costs in the early 1970s. If the cost of oil rises (being elastically supplied at the new price level), then even a 100% increase in the oil cost can translate only to $0.01 \times 100 = 1$ percent decrease in output, let alone the

*This is a very preliminary version. Please do not cite.

likely counter effects from factor substitutions. Yet the actual decline in output after the oil shock in 1973, that caused a 80 percent increase in oil price, was eight percent from its peak, more than eight times larger than predicted. Where is the missing multiplier?

Kim and Loungani (1992) and Finn (1995) use a real business cycle model, extended to include energy prices. They conclude that the volatility of the oil price cannot account for more than 20 percent of the volatility in output.

Rotemberg and Woodford (1996) showed that with a traditional one sector model, the oil prices shocks can explain at most one fifth of the drop in output. After modifying the model to allow imperfect competition they found closer matches to the data (in particular when they allowed for collusion between oligopolists).

Wei (2003) argues that the jump in energy prices caused capital obsolescence, which in turn could explain the drop in output. To formalize this idea Wei used a putty-clay neoclassical investment model to analyze the effect of the oil shock on the stock market. The putty-clay model is an extreme case of rigidity of the adjustment of installed capital. Her model not only fails to explain the evolution of the market value of the firms, but can only partially explain the drop in output (about 40% of the drop in output that followed the oil crisis of 1973).

None of these hypotheses can completely explain the deep recession in 1974-75. First of all, the predicted output contraction under the oil shock is at most four percent (Rotemberg and Woodford 1996, Wei 2003), while the actual drop in GDP was eight percent from its peak in 1973 to a trough in 1975. Secondly, all these models predict an immediate permanent drop in output on impact after a permanent increase in oil price, while the actual GDP dropped by only 2 percent on impact in 1973 and the drop continued for nearly 5 quarters until 1975; also, actual GDP completely bounced back to its pre-shock level in 4 years since 1975 while the actual oil price remained essentially the same after the sharp increase in 1973 until a second big oil price increase took place in 1979.

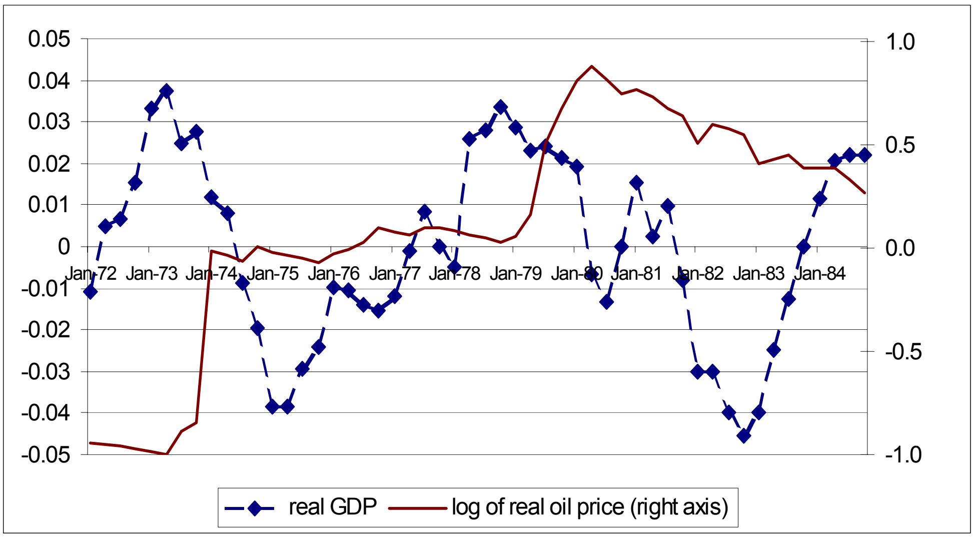

Figure 1: Evolution of GDP and of real oil price

In 1979, a second oil shock hit the US economy. In terms of total increase in oil price, this shock is as big as the first one in 1973. However, a deep recession similar to 1974-75 did not follow the shock immediately. The economy experienced a very mild recession in 1980. If the economy had moved like that during the first shock, it would had continued to slide from 1980 into a much bigger recession in 1981. Instead, a deep recession did not come until 12 quarters later after the shock and there was even a temporary recovery in the middle near the end of 1980 and the beginning of 1981.

Figure 1 summarizes these ideas. In figure 1 we can see the log of real oil price and the movements of GDP[1] in that period of time. Three puzzles are revealed:

- Why are the recessions so deep?
- Why is there a relatively quick rebound of GDP within 2 years of the first oil crises, despite an essentially permanent oil price increase?
- Why is the behavior of the economy after the second oil price shock different from the one that followed the first oil shock?

Although none of these puzzles can be resolved by existing theories, the first puzzle has, nonetheless, draw a substantial amount of attention, while the second puzzle has rarely been addressed by the literature. There are several possibilities to explain the recovery in 1975 to 1979. First, factor substitutions and obsolescence of

[1] We took log of the real GDP series and then we detrended the series using the Hodrick Prescott filter.

energy intensive technologies after oil shocks may explain the full recovery of the US economy after 1975. While plausible, no rigorous models were developed to address this possibility. To our knowledge the only exception is Wei (2003). But Wei's analysis still implies that a permanent oil price increase should cause a permanent decrease in GDP, despite the possibility of new investment in the more energy saving machines in her model. In addition, if the US economy had already switched to energy saving technologies by the end of 1975, then why is there another big recession in 1982 following the second oil price shock? Second, there may have been other positive shocks hit the US economy in 1975 to pull the economy out of recession. No strong empirical evidence exist, however, to show that during the period of 1975 to 1979 there were large enough positive macroeconomic shocks hit the economy. Thirdly, some people have blamed the contractionary monetary policy conducted in 1974 for the deep recession in 1974-75, see for example Bernanke et al. (1997). This theory, however, cannot explain the full rebound of the economy in 1979 (starting in 1975). No monetary policy could have generated that much of output growth to lead the US economy to a full recovery (see Hamilton and Herrera, 2001).

Hence, not only does the magnitude of the 1973-74 recession remain a mystery, but also the full recovery of the economy taking place between 1975 to 1979. None of the theories proposed by Rotemberg and Woodford (1996), Finn (2000)[2] and Wei

[2] Finn (2000) argues that she can explain the dynamics that follows an oil shock. I don't like

(2003), or by others in the literature, can explain these two big puzzles associated with the oil shocks in 1973.

In this paper, we propose a plausible explanation for the two big puzzles associated with the oil price rises in 1973. Our theory does not depend on the quick obsolescence of energy intensive technologies, which we think may take place in a period much longer than 3-5 years due to the high costs involved in developing alternative forms of energy, nor on some unobservable good shocks hitting the US economy in 1975. Our explanation is based on the conventional story of endogenous business cycle due to coordination failure, which is modelled in this paper via externalities and increasing returns to scale. Due to production externalities among firms, a surge in energy price can trigger massive collapse of aggregate activities. The multiplier effect comes from reinforcement of individual firms' actions among each other and it is strong enough to cause a big aggregate recession. Furthermore, interactions among firms' production

---

her specification of energy use (equation (5) of her paper), specially because of her calibration. For example, her calibration implies that if you double your use of your car, the energy you use will be multiplied by 2.7556. It seems too much. Then at some point she stops using oil and starts considering energy in general. So she calibrates the values in such a way that the Steady State energy costs amount for 4.3% of the total output, instead of considering imported oil (about 1%). In her impulse response analysis she considers an increase in the energy cost, and not just oil, and she directly compares her numerical results with Rotemberg and Woodword (1996) who considered the costs of total energy to be 2%.

activities can also give rise to a dynamic accelerator mechanism that not only causes a gradual decline of GDP into a big recession but also a temporary full rebound of aggregate output in a relatively short period of time despite a permanent shock to the energy price.

Thus, the two big puzzles associated with the 1974-75 recession can be fully explained by the oil price increase in 1973 under the multiplier effect and the accelerator effect of increasing returns to scale. The aggregate returns to scale required in our model to match the data is around 1.1, which is consistent with the most recent empirical findings on returns to scale (see Laitner. and Stolyarov (2004)).

To check whether there exist other unobservable shocks hitting the US economy in 1975 that may have pulled the US economy out of the recession, we construct forecasting errors in GDP for the entire period of the 1970s and early 1980s. The forecast errors reflect surprises in the economy that are not explainable by the history of output and other fundamental variables. We feed the forecasting errors into our model to see if they can help explain the recession and the recovery in the 1970s in the absence of oil shocks. By construction, the forecasting errors reflect either the unforecastable fundamental shocks or sunspots. The unobservable fundamental shocks may include technological innovations and monetary policy shocks. The unobservable sunspots is defined as the part of the forecasting errors that is orthorgonal to the fundamentals. We show that neither sunspots nor fundamental shocks except the oil shock in 1973 is responsible for the recession of 1974-75 and the subsequent recovery of 1975-79.

However, we also apply our model to the second oil price shock in 1979. We find that the second oil shock in 1979 alone cannot explain the recovery of 1980-81 and the deep recession in 1982; instead our model shows that some unobservable shocks are largely responsible for the short-lived recovery in 1981 and the deep recession in 1982. Based on our model the oil shock in 1979 would have caused a deep recession in 1981. Our model predict that the dramatic increase in oil price in 1979 could have caused a sharp drop in output and a big recession in 1981 with a full recovery in 1984. This did not happen, however. We find that this is due to some unobservable shocks in 1981 that have neutralized some of the adverse effects of the oil shock in 1979, postponing the big recession in 1981 to 1982; but these shocks may have also caused the big recession in 1982. Therefore, our analysis shows that the 1982 recession may have not been the result of oil price increases in 1979 or in any period between 1979 and 1982.

We estimate the sunspots of the US economy in the 1970s by constructing forecast errors that are unrelated to fundamentals (will be discuss this in more detail). A pessimistic forecast entices agents to cut back consumption spending. In equilibrium, this lower consumption demand leads to lower investment because of expected reduction in future sales.

The rest of the paper is organized as follows. In section 2 we describe the economic model. In section 3 we describe the econometric model that we use to estimate the sunspots. In Sections 4 and 5 we discuss the calibration of the model and present the results. In the last section we conclude.

2 The model

This model is a version of the one-sector indeterminate RBC model studied by Wen (1998) and Benhabib and Wen (2002). This model has also been studied recently by Harrison and Weder (2002) and Xiao (2003). A representative agent in the model chooses sequences of consumption (c), hours (n), capacity utilization (e), and capital accumulation (k) to solve

$$\max \; E_0 \sum_{t=0}^{\infty} \beta^t \left(\log(c_t) - a n_t \right)$$

subject to

$$c_t + [k_{t+1} - (1-\delta_t)k_t] = \Phi_t \left(e_t k_t \right)^{\alpha_k} n_t^{\alpha_n} o_t^{1-\alpha_k-\alpha_n} - p_t o_t, \tag{1}$$

where $\Phi_t \left(e_t k_t \right)^{\alpha_k} n_t^{\alpha_n} o_t^{1-\alpha_k-\alpha_n} = y_t$ is the production function, o_t is oil input, p_t is the real oil price, $e \in [0,1]$ denotes capital utilization rate, and Φ is a measure of production externalities and is defined as a function of average aggregate output which individuals take as parametric:

$$\Phi_t = \left[(e_t k_t)^{\alpha_k} n_t^{\alpha_n} o_t^{1-\alpha_k-\alpha_n} \right]^{\eta}, \quad \eta \geq 0. \tag{2}$$

The rate of capital depreciation is endogenous. In particular, the capital stock depreciates faster if it is used more intensively:

$$\delta_t = \lambda e_t^{\theta}, \quad \theta > 1; \tag{3}$$

which imposes a convex cost structure on capital utilization.[3]

The real oil price, p_t, is exogenously determined by oil supply from outside the economy. Profit maximization implies the demand for oil to be given by,

$$o_t = (1 - \alpha_k - \alpha_n) \frac{y_t}{p_t}.$$

In equilibrium, $y = (ek)^{\alpha_k(1+\eta)} n^{\alpha_n(1+\eta)} o^{(1-\alpha_k-\alpha_n)(1+\eta)}$. Substituting the demand for oil into the production function gives

$$y = (ek)^{\alpha_k(1+\eta)}\, n^{\alpha_n(1+\eta)} \left((1 - \alpha_k - \alpha_n) \frac{y}{p} \right)^{(1-\alpha_k-\alpha_n)(1+\eta)}$$

which implies

$$y = (1 - \alpha_k - \alpha_n)^{\frac{(1-\alpha_k-\alpha_n)(1+\eta)}{[1-(1-\alpha_k-\alpha_n)(1+\eta)]}} \left(\frac{1}{p} \right)^{\frac{(1-\alpha_k-\alpha_n)(1+\eta)}{[1-(1-\alpha_k-\alpha_n)(1+\eta)]}} \left[(ek)^{\alpha_k(1+\eta)}\, n^{\alpha_n(1+\eta)} \right]^{\frac{1}{[1-(1-\alpha_k-\alpha_n)(1+\eta)]}}$$

The budget constraint thus becomes

$$c_t + [k_{t+1} - (1 - \delta_t)k_t] = A_t \left[(e_t k_t)^{\alpha_k(1+\eta)}\, n_t^{\alpha_n(1+\eta)} \right]^{\frac{1}{[1-(1-\alpha_k-\alpha_n)(1+\eta)]}}$$

where $A_t = (\alpha_k + \alpha_n) \left(\frac{1-\alpha_k-\alpha_n}{p_t} \right)^{\frac{(1-\alpha_k-\alpha_n)(1+\eta)}{[1-(1-\alpha_k-\alpha_n)(1+\eta)]}}$ is an index of productivity, which is adversely affected by positive shocks in the oil price.

[3] Also see Greenwood et al. (1988).

The aggregate return to scale in this economy is measured by

$$\frac{\alpha_k(1+\eta)+\alpha_n(1+\eta)}{1-(1-\alpha_k-\alpha_n)(1+\eta)}=\frac{(\alpha_k+\alpha_n)(1+\eta)}{(\alpha_k+\alpha_n)(1+\eta)-\eta},$$

which equals one if $\eta = 0$ and exceeds one if $\eta > 0$. The presence of oil in the production function magnifies the effect of the externality on aggregate returns to scale. To see this, note that in the absence of oil, $(\alpha_k + \alpha_n) = 1$, hence the return to scale is $(1 + \eta)$. In the presence of oil, $(\alpha_k + \alpha_n) < 1$, hence the return to scale is larger than $(1 + \eta)$, assuming $(\alpha_k + \alpha_n) > \frac{\eta}{1+\eta}$.

The model can be solved by log-linearizing the first order conditions around the steady state as in King *et al.* (1988). It is well known that with mild externalities this model possesses multiple dynamic equilibria around a unique steady state. In particular, equilibrium output and capital stock in the model follow the following dynamic process (circumflex variables denote percentage deviations from their steady state values):

$$\begin{pmatrix}\hat{y}_t\\ \hat{k}_t\end{pmatrix}=M\begin{pmatrix}\hat{y}_{t-1}\\ \hat{k}_{t-1}\end{pmatrix}+R_1E_{t-1}\hat{A}_t+R_2\hat{A}_{t-1}+\begin{pmatrix}1\\ 0\end{pmatrix}\Theta_t, \tag{4}$$

where M is a full-rank coefficient matrix with both eigenvalues lying inside the unit circle on the complex plane; and Θ_{t+1} is a one-step ahead forecasting error of output defined as

$$\Theta_t=\hat{y}_t-E_{t-1}\hat{y}_t, \tag{5}$$

which satisfies

$$E_t \Theta_{t+1} = 0 \text{ for all } t.$$

The forecast error is the source of indeterminacy and sunspots in this model. Assuming that the major fundamental shocks to the US economy in the 1970s are oil shocks, the forecast error can then be further decomposed into two orthogonal *i.i.d.* processes:

$$\Theta_t = \zeta \varepsilon_t + v_t,$$

where $\varepsilon_t \equiv \hat{A}_t - E_{t-1}\hat{A}_t$ is an innovation in the fundamentals, i.e. shock to the oil price, and v_t is a measure of confidence (shock to non-fundamentals) and is often named sunspot in the literature (see e.g., Benhabib and Wen, 2001). Notice that both the coefficient, ζ, and the variance of the sunspot, σ_v^2, are free parameters in the model. If the forecast error is assumed to be uncorrelated with the fundamentals, then $\zeta = 0$. Otherwise ζ can be either positive or negative.

In the following section, we propose a method to measure the confidence index series, v_t, in the forecast error of output using data from the US economy. Our empirical measure is consistent with our theoretical model. Namely, we construct a time series of out of sample forecast errors for GDP. And we extract an orthogonal component in this forecast error series by projecting it on contemporaneous fundamental shocks, such as the oil price, the Solow residual, and government shocks. This orthogonal component in the forecast error series is our measure of sunspots.

Based on our simulation of the model, we show that "animal spirits", or sunspots, played a nontrivial role in shaping the recessions that followed the oil price shocks.

3 Empirical Sunspots: An Out-of-Sample Forecasting Error Approach

To simulate the model derived in the previous section we need to estimate the expectational errors of equation 5. Due to the indeterminacy of the model we cannot pin down particular equilibrium values for the variables. We have to rely on some external source to form estimate agents' forecasts. We then feed the model with these estimated expectational shocks.

In the literature two different ways to estimate these expectations have been proposed. Maybe, the closest in spirit to ours is the one proposed by Oh and Waldman (1990). As a measure of false forecasts, Oh and Waldman use revisions of the government's *Index of Leading Indicators*. Then, using simple OLS regressions, they check if these revisions have an effect in future economic performance. They conclude that these effects are significant (although they interpret them in the light of strategic complementarities and not self-fulfilling prophecies as we do here). From the perspective of our model, the biggest merit of this approach is that the forecasts for a given period are done based on the information available only in earlier periods (as we will see this is not as trivial as it seems). Unfortunately two drawbacks of this method are apparent. First, from an empirical perspective, the error of a particular forecast maybe due to sunspots, but may also be due to some fundamental shock that occurred after the forecast was produced. Second, from a theoretical perspective, the government faces a credibility problem. If the announcements have a significant impact on future activity, then government officials have an incentive to lie, and hence the announcements are not credible.

More recently Harrison and Weder (2003) and Xiao (2003) used a different method to estimate sunspots. They run a VAR that includes a measure of consumers' confidence (Xiao considers a Consumer Sentiment Index, while Harrison and Weder use an interest rate spread), real GDP/GNP and several measures of fundamentals. The residuals from the consumers' confidence regression is then extracted and used as a measure of sunspots. This method violates the spirit of the model, as given by equation 5. As we can see in that equation the forecast of a variable value on time $t+1$ should be based only on information available at time t. Since Harrison and Weder (2003) and Xiao (2003) use in-sample residuals, each residual depends on the entire sample. Consider, for example, that we have data from 1950 to 2003, and we apply their method to our model. Then the sunspot of 1973 would depend on the future value of the variables, say the values at 2003. This would mean the events of 9/11/2001 were relevant to explain the evolution of the American economy in 1973 or the Great Depression in the 1930s.

Trying to take the model implications as seriously as possible, we propose a new way to estimate the sunspots. The idea of rational expectations is that the agents use all the data available and, with the help of an economic model, they produce optimal forecasts. As we have already discussed when we have indeterminacy, the model cannot help to pin down the agents' optimal forecasts. In this case a sensible thing for a rational agent to do is to collect all the data available and use a good econometric model to compute out-of-sample forecasts. Since the out of sample forecast errors may be contaminated with shocks to the fundamentals, these need to be further decomposed between sunspots and fundamentals' shocks.

In particular, using a rolling window, we produce out-of-sample forecasts for the GDP between 1973:1 and 1984:4. The econometric model that we use, and the data we have are described in the next subsections. After estimating the forecasts we compute a series of forecast errors for the period we are studying (so we estimate Θ_t of equation 5). Since we want to estimate the sunspots, and the sunspots, by definition, are nonfundamental, we need to purge the fundamental shocks from the forecasting errors. We do this by regressing these on contemporaneous variables, which we believe to represent the fundamental shocks and extracting an orthogonal component. These *purified* forecast errors will be our measure for sunspots. To *clean* the forecasting errors we use variables that can arguably represent fundamental shocks, for example oil prices, labor productivity, money supply, government spending.

3.1 Principal Components Combination

An increasing popular method, which uses many explanatory variables is the Principal Components Regression (PCR), which was applied by Sargent and Sims (1977). More recently, this method has been successfully applied to US Macroeconomic data — Stock and Watson (1998, 1999, and 2002), Bernanke and Boivin (2003). This literature has been growing, and some nice asymptotic results have already been derived — see Stock and Watson (1998 and 2002) and Bai and Ng (2002). This method is becoming popular because Stock and Watson showed, with Monte Carlo simulations and with economic data, that this technique performed significantly better than several other competing techniques. Aguiar-Conraria and Hong (2003), based on the literature of Forecasts Combination, proposed a modification of the PCR method, and showed, in an application to inflation forecasting, that the modified PCR — called Principal Components Combination (PCC) — allowed for some significant accuracy gains.

In this paper, measuring the accuracy by the mean square forecast error, the use of PCC, when compared to the PCR, leads to an increase in accuracy of 20%.

An advantage of this approach is that an infinite number of variables can be used to produce the forecasts. This econometric method is consistent with the idea that, in the absence of a good economic model, the agents should gather all the information available at a certain point in time to form their expectations.

After collecting the out of sample forecast errors there is still the need to remove the contemporaneous fundamental shocks from the forecasting errors. Again this is done out of the sample. Call ε_t the forecasting error at time t and x_t a vector with variables representing fundamentals. Suppose we have 100 observations available, then using the first, say, 50 observations we regress ε on X, and so we have $\varepsilon = X\hat{\beta}+e$. Now we clean the 51^{st} forecasting error: $S_{51} = \varepsilon_{51} - X_{51}\hat{\beta}$, where S stands for sunspot. We do the same for the 52^{nd} observation using a rolling window (so the regression sample goes from the 2^{nd} to the 51^{st} observation), and so forth.

We now briefly explain the PCC method that we use to produce forecasts. For a more detailed explanation, and description of the advantageous of this method the reader can consult the paper mentioned above.

Assume that we have N stationary explanatory variables (already demeaned) arranged in a matrix X, then, using univariate regressions it is possible to produce N forecasts of the dependent variable y, which can be combined using the PCR approach:

1. project y onto the space spanned by each of the N explanatory variables: $z_n = x_n \left(x_n' x_n\right)^{-1} x_n' y$, for $n = 1, 2, ..., N$,
2. create a new matrix of explanatory variables: $Z = (z_1, ..., z_N)$,
3. find the eigenvectors u_i of $Z'Z$ associated with positive eigenvalues. Let u_1 be the eigenvector associated with the highest eigenvalue, u_2 with the second highest,

and so on,

4. use as new regressors the variables Zu_A associated with the A highest eigenvalues.

The choice of A (the number of factors to be used) is usually done using a modified version of the Bayes Information Criterion (BIC), proposed by Bai and Ng (2002). We also allow for lagged values of these regressors. The number of lags is chosen using the traditional BIC.

3.2 Our Data

Since we only want to use out-of-sample forecasts, and we are trying to apply the model to the 70s we have to be relatively parsimonious with the choice of variables, because not many observations are available. To produce the real GDP forecasts we use 21 explanatory variables and two lagged values of each. So overall we have 63 explanatory variables, which include gross private investment, real consumption, interest rates, productivity, unemployment, oil prices, etc. All data is taken from the FRED II database and spans from 1950:1 and 1984:4. Monthly variables were transformed in quarterly data by taking averages. All variables that are not in the form of ratios or rates, were logaritmized. A time trend from non-stationary series was removed using the Hodrick Prescott filter.

We use the data from 1950:1 to 1961:2 to produce the real GDP forecast for 1961:3, then 1950:2 to 1961:3 to produce the forecast to 1961:4, and so on until 1984:4. Since the number of observations is not very big, and to avoid over-fitting, we impose a limit of 1 factor, and allow for a maximum of two lags of for the factor, (two) lagged values of real GDP are also used in the regression[4]. So we have forecasts that span 1961:3 and 1984:4[5]. We use ten years of data to "purify" each forecast error (and estimate the sunspot of that period) — so the first sunspot value will be for 1971:4. To do this we regress 10 years of forecasting errors on some contemporaneous variables, namely oil prices, labor productivity, money supply and government spending. Several other variables and combinations were tried. The results were remarkably robust to any changes. Figure 2 displays our estimated sunspots series.

4 Calibration

We calibrate the model using the following parameter values and assumptions. The parameter values are standard for quarterly models: $\alpha_k = 0.36 * .99$ and $\alpha_n = (1 - 0.36) * 0.99$, implying the share of the oil sector to be $1 - \alpha_k - \alpha_n = 0.01$.

[4]We also considered the possibility of more factors and of allowing more laggs. The out of sample mean square errors became larger, suggesting the possibility of over-fitting.

[5]By definition the forecast errors must be serially uncorrelated and mean zero. We used the Breusch-Godfrey Serial Correlation LM Test to test for correlation up to the order four. The p-value found was 0.7, showing no signs of serial correlation. Testing if the mean of the forecast errors is zero produces a p-value of 0.95.

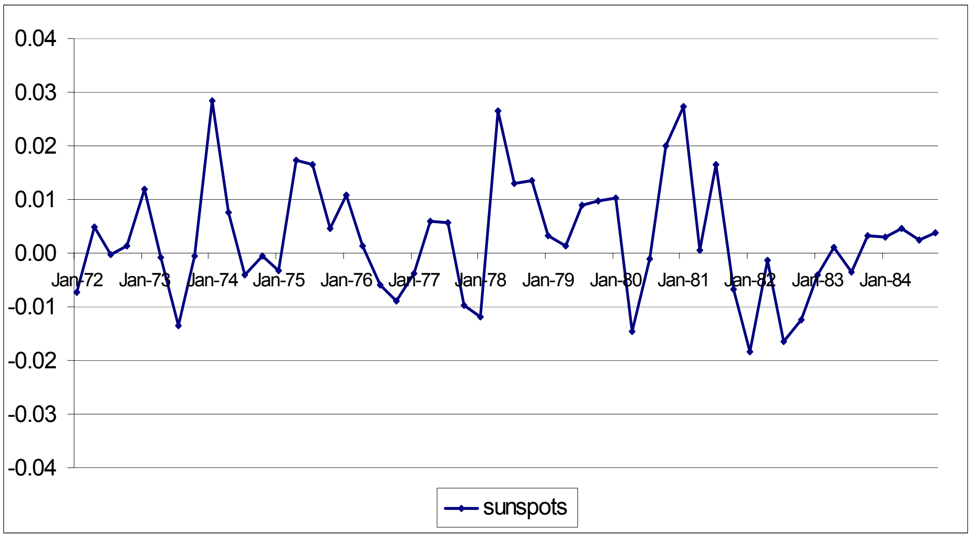

Figure 2: Estimated Sunspots

For the discount factor we assume $\beta = 1.03^{-\frac{1}{4}}$. Given these calibration values, the model exhibits indeterminacy for $\eta > 0.0948$.We will consider several cases for η considering the possibility of Constant Returns to Scale (CRS), Increasing Returns to Scale (IRS) and no indeterminacy, and finally IRS and indeterminacy.

We also have to assume a functional form for $E_t p_{t+1}$. Hamilton (1983, 1996), Burbidge and Harrison (1984), Gisser and Goodwin (1986), and others have argued that oil prices were exogenous to the US economic activity, at least, until mid 1980s. Following this literature we assume that $E_t p_{t+1} = f(p_t, p_{t-1}, ...)$. More precisely,we assume that the logarithm of real oil price follows a random walk: $E_t p_{t+1} = p_t$. Using data, up to the end of 1984, for the real oil price series, based on the Augmented

Dickey Fuller test, we concluded that the hypothesis of a unit root could not be rejected. Running simple autoregressive processes of several orders we concluded that all the coefficients were statistically insignificant, both individually and jointly, except the first lag of the real oil price (even the constant was insignificant). Running a simple AR(1) process, with no constant we found the estimated coefficient to be 0.996. We tested for serial correlation of the residuals using Breusch-Godfrey LM test of order 4. No evidence of serial correlation was found. So our assumption of a pure random walk for the oil prices matches the evidence from the data, implying that past information will not be helpful to predict futures changes in the oil price.

4.1 The model with no indeterminacy

In this section we show that the model with no indeterminacy fails to explain the shape of the recessions in the US in the time period under consideration. First, we show that the model with constant returns to scale fails to explain deepness and the persistence of recessions, then we show that by allowing for increasing returns the model can explain a big drop in GDP, but it still lacks persistence.

4.1.1 Constant Returns to Scale

No increasing returns imply that $\eta = 0$. In this case it is well-known that the model has a saddle path solution and there is no room for expectation driven cycles. It is also known that the dynamics of the model follow very closely the dynamics of the technological shocks (in our case the oil price shocks). So it is no surprise that the evolution of the generated output, that we can observe in figure 3, is almost a mirror of the evolution of the oil prices (see figure 1). So following a permanent increase in the real oil price (first oil price shock), the economy moves to a new steady state at a lower level. Again, after the second oil price the output dynamics follows very closely the dynamics of the (inverse of) real oil price.

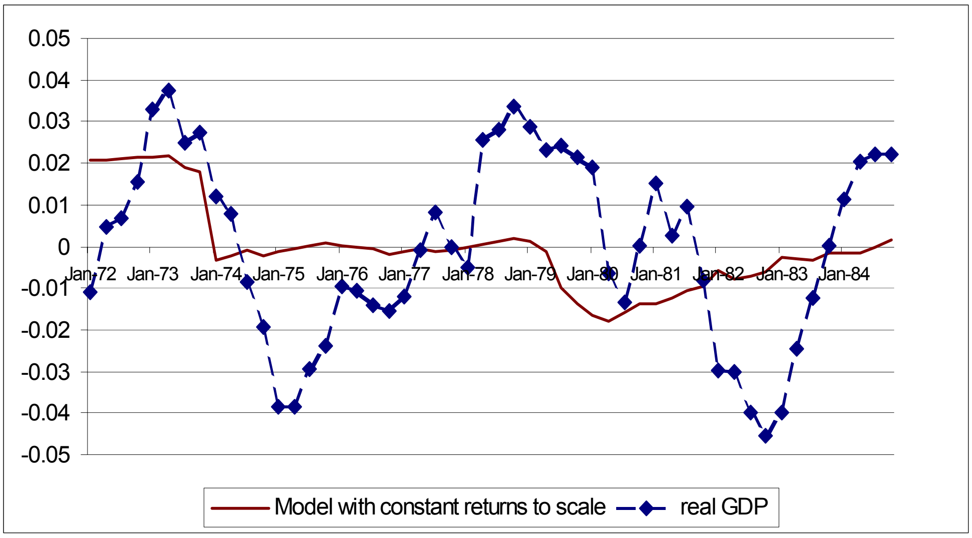

Figure 3: Model with CRS

From the picture, it is quite clear that the RBC model fails to explain the main features of the data, namely it strongly under-predicts the recessions and fails to explain the full rebounce of the economy observed in the end of 1978 and the timing of the recession in the early 1980s.

4.1.2 Increasing Returns to Scale

If we consider increasing returns to scale, but no indeterminacy, the model dynamics does not change qualitatively, but the consequences of the shocks are magnified. For example, with increasing returns of 1.05 ($\eta = 0.05$), the drop of output after the first oil shock is about 3 per cent, and for $\eta = 0.09$ of 6 per cent, which accounts for a reasonable fraction of the first recession. In both situations the model fails to explain the full rebound of the economy that followed the first oil price shock (with the output approching a new steady state at a lower level), and also fails the deepness and the timing of the recession in the 1980s.

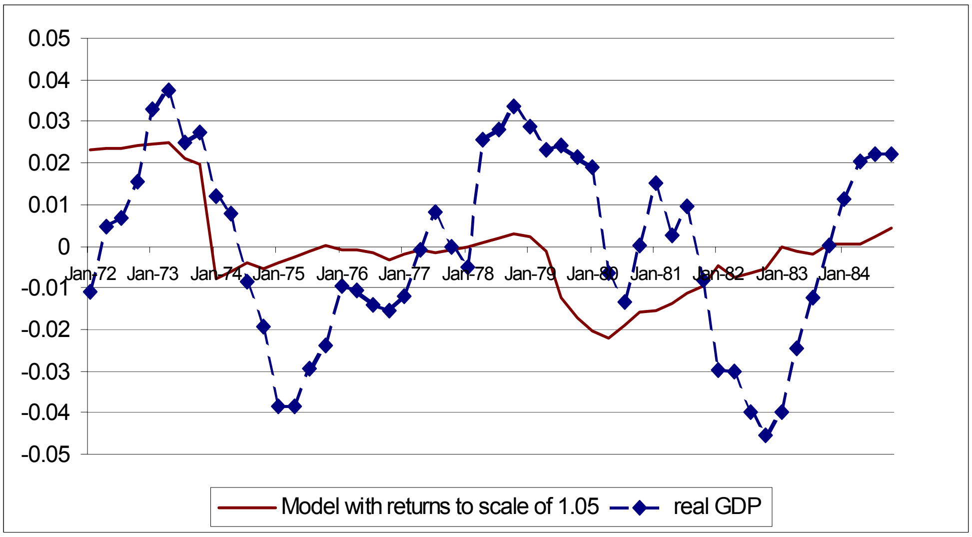

Figure 4: Model with IRS ($\eta = 0.05$)

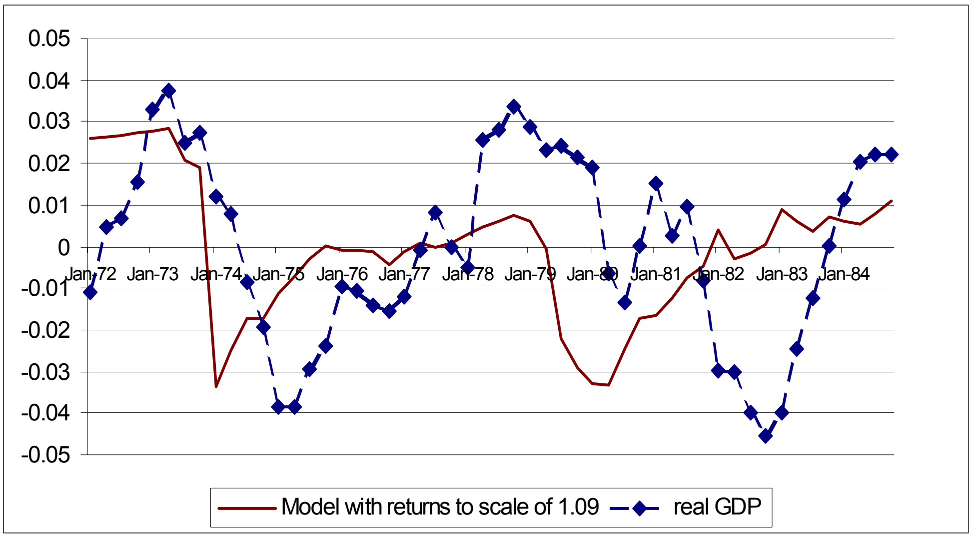

Figure 5: Model with IRS ($\eta = 0.09$)

4.2 The Model with Increasing Returns to Scale and Indeterminacy

The goal of this section is to test if IRS, indeterminacy and sunspots can help explaining the fluctuations in the 70s and early 80s. For that purpose we need first to estimate the value of η. After that we will show indeterminacy is the key to explain the rebound of the economy observed. I simple impulse response analysis will be sufficient to make this point. Then we will run a set of counter-factual experiments to pin down the exact causes of each crises. First we will shut down the forecasting errors ($\Theta_t = 0$) and check how much of the actual output dynamics can be explained by the oil price movements. We will also shut down the oil shocks and check the behavior of the economy when subjecvt only to sunspot shocks. We will argue that while the first oil price shock is sufficient to explain the behavior of the economy in the 1970s, expectational shocks will be determinant in explaining the behavior of the economy in the early 1980s.

4.2.1 Estimation of the Degree of Returns to Scale

Given our calibration, the model exhibits indeterminacy for $\eta > 0.0948$. The parameter governing the increasing returns, η, is chosen to maximize the correlation between the true GDP and the GDP generated by the model, when we consider the data generator process to be given by equation 4. So we estimate η by maximizing:

$$\max_{\eta > 0.0948} \frac{\sum_{t \in [1971.5, 1984]} (GDP_t \hat{y}_t)}{\sqrt{\sum_{t \in [1971.5, 1984]} (GDP_t)^2 \sum_{t \in [1971.5, 1984]} (\hat{y}_t)^2}} \tag{6}$$

where $\hat{y}_t$ is given by equation 4, and Θ_t is the estimated forecast error at time t (not the sunspot). The estimated η is 0.1001939[6], with a correlation of 0.86. It is interesting to note that this value is im the middle of the range found by Laitner and Stolyarov (2004) in an independent study. These authors estimate the value of increasing returns to scale in the US economy to be in the range $0.09 - 0.11$.

In figure 6 we can see the behavior of the model economy when equation 4 is fed with the forecasting errors.

[6] To estimate the value of η, we did a grid search with $\eta \in [0.0948, 1]$, considering increments of 0.0000001.

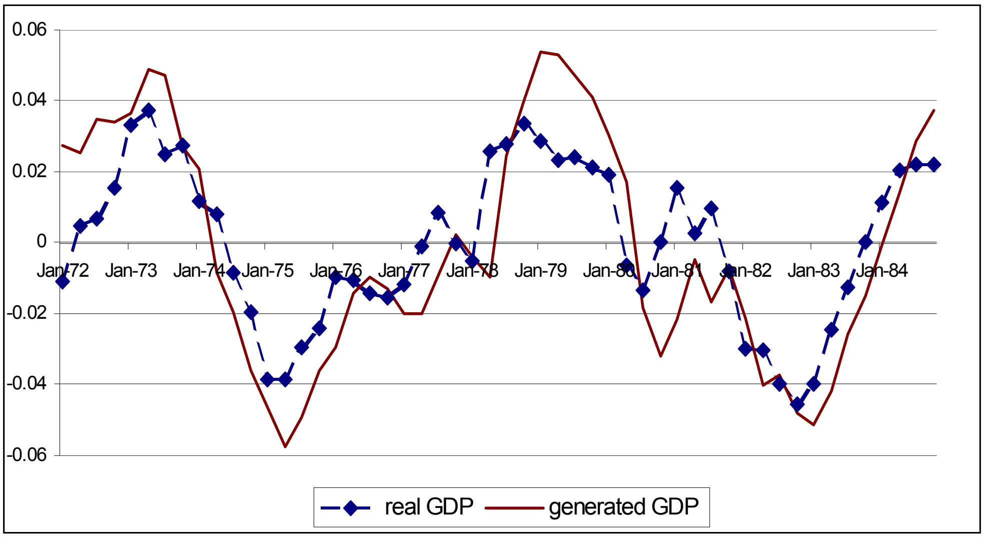

Figure 6: The model with oil shocks and forecasting errors

The model over-predicts the size of the recession that followed the first oil shock, predicting a drop in the output of 10.5 per cent, but is able to predict the expansion of the economy in the late 1970s and the small rebound in 1981.

4.2.2 Impulse Response Analysis

In figure 7 we can see the predicted response to a 100 per cent permanent increase in the oil price. After the shock (in period 2) the economy enters in a recession, with the output continuosly dropping for 6 quarters. In the long run the economy oscillates towards a lower steady state level but before that one can observe a full rebound after 5 years of the initial shock.From this analysis we conclude that increasing returns and indeterminacy is the key to understand the reaction of the economy to the first oil

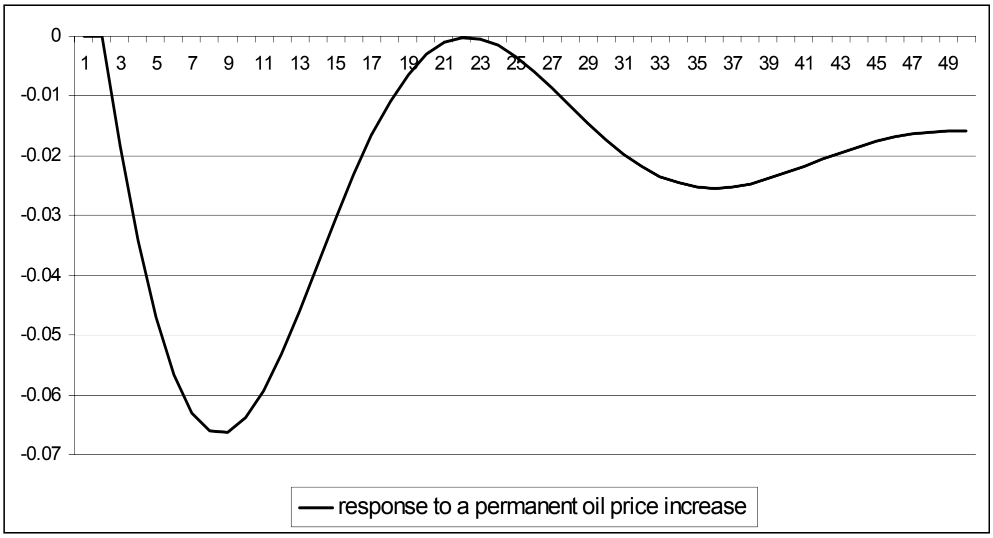

Figure 7: Response to a permanent real shock

shock.

4.2.3 The Model with Oil Price shocks

Now we feed the model with the observed oil prices and, following the impulse response analysis, we keep the forecasting errors off.

The model does a good job in matching the until 1979, being able, as predicted, to explain the deepness of the recession and the full recovery of the economy. Observed in the end of 1978. This model fails to explain the characetristic features of the recession that followed the second oil shock, namely it predicts a big recession in the end of 1980 and not of 1982. It also fails to explain the recovery in the end of 1980.

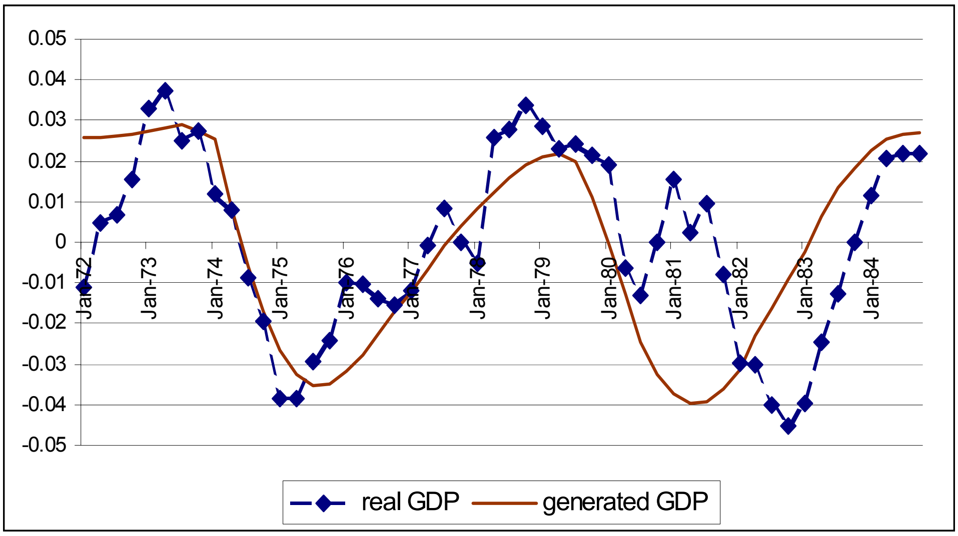

Figure 8: The model fed with oil prices

4.2.4 The model with forecasting errors and sunspots

In figure 9 we shut down the oil shocks and feed the model with forecasting errors only. As it can be seen the model is unable to match the data until 1979 but, after tthat it can predict the recovery in early 1980s and the slump that followed.

This model is unable to fully predict the recession in the 1970s. Even the predicted small slump is not robust to the cleaning of the forecast errors. If we feed the model with the estimated sunspots, instead of the forecasting errors we get figure 10. figure 10 suggests that even the small slump in 1974-75 predicted in figure 9 is due to fundamental shocks and not pure sunspots[7]. On the other hand the predicted behavior

[7] Actually, if we clean the forecasting errors using nothing but the oil prices the results would be quite similar. This suggests that the relevant fundamental shocks are indeed the oil price shocks.

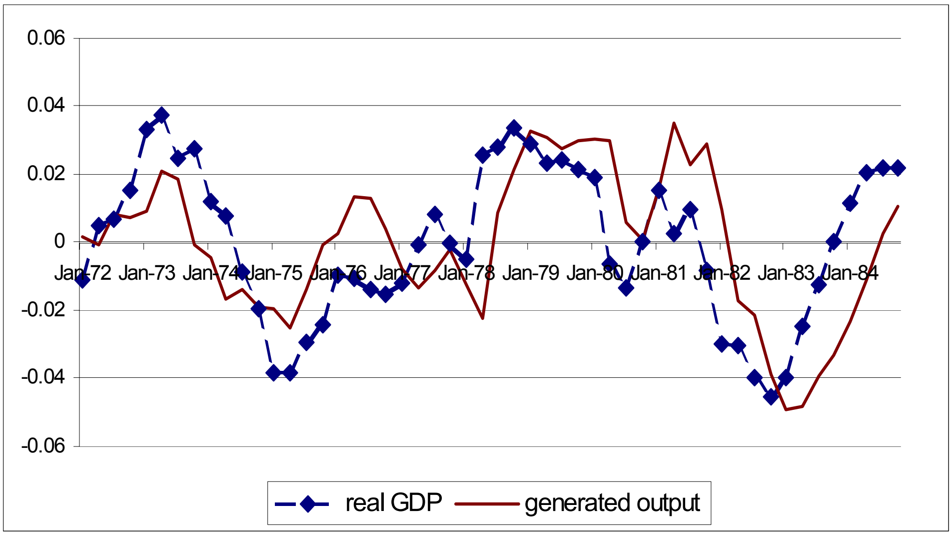

Figure 9: The model with forecasting errors

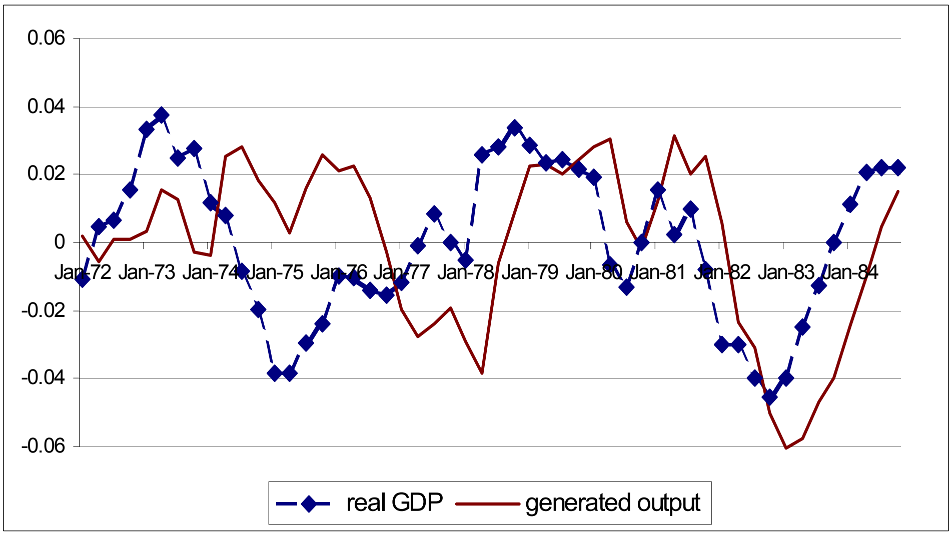

Figure 10: The model with sunspots

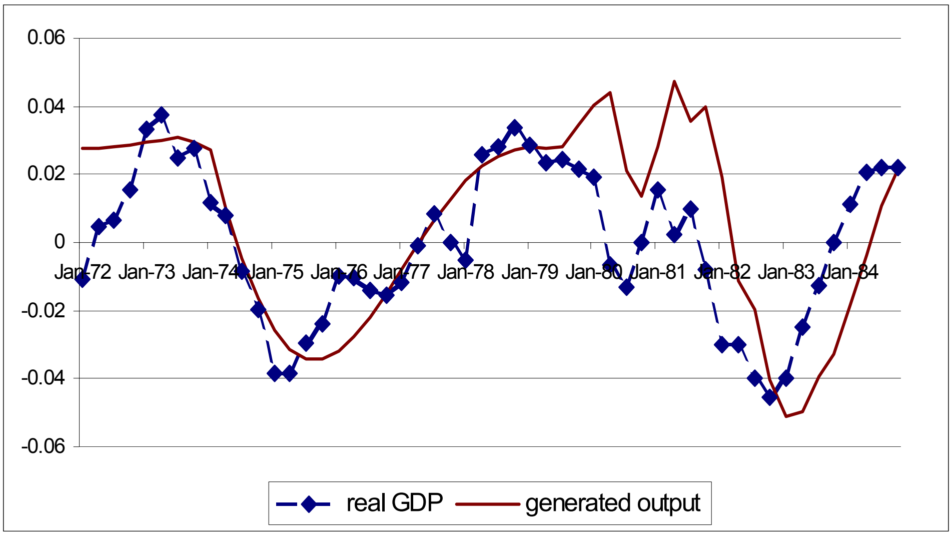

Figure 11:

for the economy after the second oil shock is qualitatively correct. This suggests that sunspots may have played a non trivial role in shaping the economy after the second oil shock.

4.2.5 The Model with Oil Shocks and Sunspots

The previous results suggest that while the oil price evolution seems sufficient to explain the chracteristics of the first recession, sunspots are necessary to understand the second recession. We now do two more experiments. In figure 11 we shut down the oil price movements after the first jump in 1973 and, simultaneously we shut down the sunspots until 1979, but switch them on after that. In figure 12 we do the same experiment without ever switching the oil price movements off.Both figures reveal

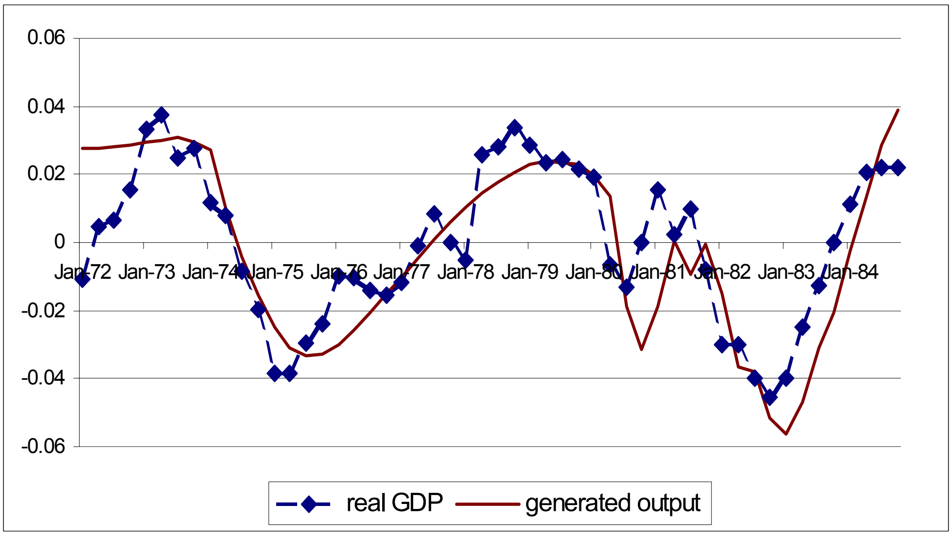

Figure 12:

what we suspected, although they second one seems to fit the data better, sunspots played an important role in explaining the small expansion in the end of 1980 and in explaning the delayed big recession.

5 Tentative conclusions and directions for future research

In this paper we have shown that a standard neoclassical growth model with mild increasing returns to scale, and with indeterminacy can explain, using nothing but the evolution of oil prices can explain a big chunk of the depression that followed the oil price increase in 1973. It can explain not only the big depression that followed the oil price increase, but also the its dynamics. Namely the continuous drop of output for more than one year and the recovery that followed.

We have also shown that animal spirits may have played an important role explaining the recession that followed the oil price increase in the end 1970s. Sunspots were important to explain the delayed recession and its deepness.

So far we have not included policy variables in the model, like fiscal and monetary policies. Including variables representing these policies will allow us to be check if our claim that these policies are not important to explain the recessions is correct, and we will also be able to study the best policy responses to a real shock, like an oil price shock.

References

Aguiar-Conraria, L., and Y. Hong, 2003, Forecasting in data-rich environments, working paper.

Azariadis, C. 1981, Self-fulfilling prophecies, Journal of Economic Theory 25, 380-396.

Azariadis, C. and R. Guesnerie, 1986, Sunspots and cycles, Review of Economic Studies 53(5), 725-738.

Bai, J. and S. Ng, 2002, Determining the number of factors in approximate factor models, Econometrica 70, 191–221.

Benhabib, J. and Farmer, R., 1994, Indeterminacy and increasing returns, Journal of Economic Theory 63, 19–41.

Bernanke, B and Boivin, J., 2003, Monetary policy in a data-rich environment, Journal of Monetary Economics 50, 525–546.

Bernanke, B.S., M. Gertler, M. Watson, 1997, Systematic monetary policy and the effects of oil price shocks, Brookings Papers on Economic Activity, 91–142.

Cass, D. and K. Shell, 1983, Do sunspots matter?, Journal of Political Economy 91, 193-227.

Chauvet, M. and J.-T. Guo, 2003, Sunspots, animal spirits, and economic fluctuations, Macroeconomic Dynamics 7, 140–169.

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

Finn, M., 2000, Perfect competition and the effects of energy price increases on economivc activity, Journal of Money Credit, and Banking, 400–417.

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

Hamilton, J., 2003, What is an oil shock?, Journal of Econometrics 113, 363–398.

Hamilton, J. and A. Herrera, 2001, "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy," Journal of Money, Credit, and Banking, forthcoming

Harrison, S. and M. Weder, 2003, Did sunspot forces cause the Great Depression? Working Paper, Columbia University.

Hooker, M., 1996, What happened to the oil price–macroeconomy relationship?, Journal of Monetary Economics, 38, 195–213.

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

Laitner, J. and D. Stolyarov, 2004, Aggregate returns to scale and embodied technical change: Theory and measurement using stock market data, Journal of Monetary Economics, 51, 191–234.

Leduc, S. and S. Keith, 2004, A quantitative analysis of oil-price shocks, systematic monetary policy, and economic downturns, Journal of Monetary Economics, 51, 781–808.

Oh, S. and Waldman, M., 1990, The macroeconomic effects of false announcements, The Quarterly Journal of Economics 105, 117–34.

Raymond, J. and R. Rich, 1997, Oil and the macroeconomy: a Markov state-swititching approach, Journal of Money, Credit, and Banking 29, 193–213.

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, 549–577.

Sargent, T. and Sims, C., 1977, Business cycle modeling without pretending to have too much a priori economic theory, in C. Sims, New Methods in Business Cycle Research, Minneapolis: Federal Reserve Bank of Minneapolis, 45–109.

Shell, K. 1977, Monnaie et allocation intertemporelle, Mimeo, Seminaire d'Econometrie Roy-Malinvaud, Centre National de la Recherche Scientifique, Paris.

Shell,K. 1987, Sunspot equilibrium, in The New Palgrave: A Dictionary of Economics (J. Eatwell, M. Milgate, and P. Newman, eds.), Vol. 4, New York: Macmillan, 549-551.

Shell, K. and B. Smith, 1992, Sunspot Equilibrium, in the New Palgrave Dictionary of Money and Finance (J. Eatwell, M. Milgate, and P. Newman, eds.), Vol. 3, London: Macmillan, 1992, 601-605.

Stock, J., and Watson, M. 1998, Diffusion indexes, NBER working paper.

Stock, J., and Watson, M. 1999, Forecasting inflation, Journal of Monetary Economics 44, 293–335.

Stock, J., and Watson, M. 2002, Macroeconomic forecasting using diffusion indexes, Journal of Business & Economic Statistics 20, 147–162.

Tobin, J., 1980, Stabilization policy ten years after, Brookings Papers on Economic Activity, 19–71.

Wei, C., 2003, Energy, the stock market, and the putty-clay investment model, American Economic Review 93, 311–323.

Wen, Y., 1998, Capacity utilization under increasing returns to scale, Journal of Economic Theory 81, 7–36.

Woodford, M. 1986, Stationary sunspot equilibria in a finance constrained economy. Journal of Economic Theory, 40, 128-137.

Woodford, M. 1991, Self-fulfilling expectations and fluctuations in aggregate demand. In: N. G. Mankiw and D. Romer (eds.), New Keynesian Economics: Coordination Failures and Real Rigidities, Vol. 2. MIT Press, Massachusetts, 77-110.

Xiao W., 2003, Explaining speculative expansions, Working Paper, Department of Economics and Finance, University of New Orleans (forthcoming in B.E. Journal of Macroeconomics).