

The Real Interest Rate, the Real Oil Price, and US Unemployment Revisited

Spyros Andreopoulos*

*Department of Economics, University of Bristol
8 Woodland Road, Bristol BS8 1TN
United Kingdom*

24 December 2006

Abstract

The time series evidence on the relationship between unemployment and the real prices of capital and energy is re-examined for US data. In contrast to previous studies, results indicate that the real interest rate matters little, if at all, for equilibrium unemployment. Using a Markov Switching vector autoregressive method proposed by Psaradakis, Ravn, Sola (2005) [*JApplEconometrics* 20(5), pp. 665-683] to investigate time-varying Granger causality, the paper shows that the real rate helps forecast unemployment during NBER expansions only. Granger causality from the oil price to unemployment occurs in recessions. The results support the view that the price of crude induces at least some recessions, while not being a regular feature of the US business cycle.

Keywords: Unemployment; Real Interest Rate; Oil Price; Granger Causality; US Recessions; Markov Chain; Regime Switching; Structural Instability

JEL classification: C32, E24, E32

* This paper is based on Chapter 2 of my PhD dissertation at Birkbeck College, University of London. For encouragement and advice I am indebted to my supervisor, Ron Smith. Further, I wish to thank Martin Sola for useful comments and for making available his GAUSS code. All mistakes are my own.

1. Introduction

Motivated by the analyses of Phelps (1994) and Carruth et al. (1998), this paper re-examines the time series evidence on the relationship between real input prices – the real price of oil and the real interest rate – and unemployment in the United States. Several objectives are pursued. First, I update the empirical results in Carruth et al. (1998) to include data up to 2005:4. Second, I use new econometric techniques which do justice to the structural instability encountered in the relationships of interest. This allows for a more precise characterisation of the role each input price plays in US aggregate dynamics. Finally, the paper seeks to contribute to the literature on the oil price-macroeconomy nexus.

The main results are as follows.

- In contrast to previous studies in the literature it is shown that the real interest rate does granger cause unemployment, albeit almost exclusively during NBER expansions
- The real price of oil granger causes unemployment during recessions only
- The real interest rate matters little, if at all, for equilibrium unemployment
- Equilibrium unemployment increases with the real price of oil, corroborating earlier results in the literature

The next section motivates the paper by discussing the relevant literature. Section 3 presents the data and revisits the long run relationship between unemployment and real input prices. Section 4 then turns to Granger Causality (GC). I present new evidence on GC between unemployment and the real rate as well as the price of oil, and show that there is time variation in both granger causal relationships. Section 5 contains the main results of the paper. It investigates time-varying GC by means of a Markov Switching VAR. Section 6 concludes.

2. Unemployment and Real Input Prices: Theory and Evidence

A dominant, though not incontrovertible, paradigm in macroeconomics is the theory of the natural rate of unemployment. First proposed by Friedman (1968) and Phelps (1968), natural rate theory posits that in the medium run – the horizon at which the effects of nominal rigidities and misperceptions have petered out – unemployment reverts to an equilibrium which is independent of the level of aggregate demand in the economy. In turn, the stylised fact that most of the variance of unemployment is accounted for by shifts in its mean rate between decades rather than fluctuations around the mean over the cycle (e.g. Layard et al. 1991, Lindbeck 1993) suggests that the natural rate may itself be variable.

Theories which endogenise the natural rate, most prominently Phelps' (1994) “structuralist” approach, and the literature termed “equilibrium” unemployment theory (e.g. Pissarides 2000), have highlighted a host of potential candidates which could induce shifts in the natural rate. Structuralist authors in particular have stressed the importance of real input prices, that is, the prices of capital and energy, as important determinants of medium term unemployment, see Phelps (1994) and Carruth et al. (1998). How do movements in real input prices affect the number of jobs available in the economy? Structuralist theories use variants of efficiency wage models of the labour market to show that higher real prices for inputs will depress employment: when an input price increases, a larger share of output must be spent on this factor, which results in a real wage squeeze. But in an efficiency wage economy a reduction in the real wage is incentive compatible only if the unemployment rate increases; otherwise, workers will withhold effort, or will be more prone to quit.

A great deal of empirical research has gone into testing the structuralist view of the aggregate labour market. Bianchi and Zoega (1998) and Papell et al. (2000) provide

evidence from OECD economies which supports the hypothesis that unemployment is subject to infrequent shifts in its mean rate.¹ Their results show that in many cases the timing of the shifts coincides with major changes in oil prices or real interest rates. More direct evidence on the role of real input prices for the dynamics of unemployment is provided by Phelps (1994) and Carruth et al. (1998). Phelps (1994, ch. 17) provides pooled cross-section time-series evidence indicating that increases in the world real interest rate and the world real oil price increase unemployment in OECD economies. On the other hand, Carruth et al. (1998) – henceforth CHO – estimate an equilibrium correction model for United States unemployment. They show that a long run (cointegrating) relationship operates between unemployment, the real rate of interest, and the real price of oil for most of the second half of the previous century.

I begin by revisiting, and updating, the cointegration evidence presented in CHO. My conjecture is that the sample period CHO use to estimate their model may provide a distorted picture of the importance of the two input prices for US unemployment. This is because almost half of their 1954:2 to 1995:2 sample, the years 1954-1973, witnessed very little variation in the real price of oil as the nominal price was virtually fixed and inflation was moderate.² Indeed, when I estimate a cointegrating relationship for the period 1970:1 to 2005:4, I find that the real rate of interest is not statistically significant in the equilibrium equation. On the other hand, the real price of oil is highly significant. These results are robust to the inclusion of the real wage in the empirical model.

Following CHO, I then turn to the granger causal relationship between unemployment and real input prices. While a large literature investigates Granger causality (GC) between the real economy and the price of crude oil – mainly focussing on output

¹ See also Coakley et al. (2003).

² There were three quarter-of-a-dollar increases in the nominal price of oil, along with a few smaller increases and decreases. The US inflation rate over the period 1954 to 1973 did not exceed 6 percent and was below 3 percent most of the time.

growth – little research effort has gone into GC between the real rate of interest and the macroeconomy. CHO examine GC between the real interest rate (as well as the real price of oil) and unemployment and find little evidence that the real rate contributes to forecasting unemployment. Further, their results are particularly sensitive to the sample period considered.

This sensitivity is a further conceptual point of departure for the present paper. Motivated by structural instability in the granger causal relationship between real input prices and unemployment – on which detailed evidence is also provided here – I propose modelling time varying GC between the three magnitudes by means of the Markov Switching Vector Autoregression (MS VAR) of Psaradakis et al. (2005). The parameters of this VAR are designed to be time varying, and change to reflect the presence or absence of GC. I use this technique to investigate whether the real rate helps forecast unemployment, conditional on the real price of oil.

The model is successfully applied to more than half a century of US quarterly data (1953:2-2005:4). The main result is that the real interest rate does forecast unemployment. Interestingly, comparison of the time path for the model's states with the NBER business cycle chronology indicates that it does so almost exclusively during periods of expansion in the US economy. A further result from the MS VAR is that the real rate plays a quantitatively negligible role for unemployment in the long run: the sum of the VAR coefficients on the real rate in the unemployment equation is minute. This result reinforces the cointegration evidence regarding the lack of importance of the real interest rate in long run unemployment dynamics, arrived at for the period 1970-2005.

While focusing on unemployment, this paper also contributes more generally to the literature on the oil-price macroeconomy relationship. Crucially from a methodological

point of view, the Psaradakis et al. (2005) MS VAR model I employ combines two of the most popular econometric approaches in the literature on aggregate dynamics: Granger Causality and Markov Switching. Thus, since the pioneering contribution of Hamilton (1983), GC constitutes the main vehicle of research on oil prices and the macroeconomy. On the other hand, Hamilton's (1989) MS technique has established itself as perhaps the single most important empirical tool for business cycle analysis. To mention the contribution which is most relevant to the present work, Raymond and Rich (1997) use a MS model to study the influence of oil prices on business cycle phases and transitions in the US.

Returning to my results, the MS VAR indicates that the real rate of interest and the real price of oil play complementary roles over the business cycle with respect to forecasting unemployment. While the real rate granger causes unemployment during upswings, the price of crude helps forecast unemployment *when the real rate does not*, that is, in recessions. I also provide evidence that the granger causal relationship between the real price of energy and unemployment may be present for two specific recessionary episodes only: the downturn associated with the first OPEC shock (1973-75), and the one following Iraq's invasion of Kuwait in 1990/1.

Overall, the results are consistent with the view that the price of crude is not a regular influence on the US business cycle, as argued in Raymond and Rich (1997) among others. More generally, however, the GC and cointegration findings I present both indicate that the real price of crude does matter a great deal for US economic performance, something which has been called into question in the recent empirical debate on oil prices and the macroeconomy (e.g. Hooker 1996, Barsky and Kilian 2004). A suggestive interpretation of my results is that the real rate influences unemployment over the business cycle but has little influence on the *equilibrium* unemployment rate;

the real price of energy on the other hand generates at least some recessions, and probably induces the natural rate of unemployment to ratchet up.³ This paper therefore provides some important qualifications to the existing evidence on US unemployment dynamics, as well as to the more general literature on input price-macroeconomy interactions. At the same time, it adds to the growing econometric literature which emphasises nonlinearities in the operation of the aggregate labour market, e.g. Neftci (1984), Acemoglu and Scott (1994) and more recently Coakley et al. (2003) and Hamilton (2005).

Which functional form is appropriate for the oil price? Several recent studies have stressed that the relationship between oil and the macroeconomy has come to be characterised by asymmetries or nonlinearities, e.g. Hamilton (1996).⁴ These authors use complicated filters to extract certain components of the oil price such as the net change over a certain preceding time period. In the present paper, the real price of oil enters in (log) levels throughout, without any further adjustment. While the motivation for this treatment stems primarily from economic theory, my choice is vindicated by the empirical results.⁵ Indeed, the fact that the price of crude is shown to matter for the US economy without the need for complicated and essentially arbitrary transformations is a further insight offered here.

3. On the Long Run Relationship Between Unemployment and Input Prices

3.1 *The Data*

To explore the relationship between the real interest rate, the real price of energy, and the unemployment rate, I use quarterly US data from 1953:2 to 2005:4 (211

³ I use “equilibrium” unemployment and “natural rate” synonymously in this paper.

⁴ Hamilton (2003) provides a very useful survey and interpretation of this literature.

⁵ In related work, Dotsey and Reid (1992) find that the distinction between oil price increases and decreases does not matter for unemployment.

observations), where 1953:2 is the earliest date for which medium term interest rates are available. Time series plots of the data are shown in Figure 1 and Figure 2. With the exception of the price of oil, my basic specification uses the same data as CHO: the real interest rate measure is a five-year treasury bill rate less the contemporaneous annual GDP deflator inflation rate;⁶ unemployment is the civilian unemployment rate in per cent and the real price of oil is the log of the nominal oil price minus the log of the GDP deflator. All data has been retrieved from the FRED database.⁷

There are two main reasons for using the five year real rate. From a theoretical point of view, a medium term interest rate such as the five year rate probably captures most adequately the time horizon relevant for employment decisions as these are modelled in structuralist and equilibrium theories of unemployment. In those theories, employees are assets to the firm which yield a flow of dividends per unit time in the form of employment services, i.e. contributions to the firm's profit. Those dividends are discounted by an appropriate rate, and CHO argue that a five-year real interest rate is an adequate proxy for this discount rate. The second motivation is to maintain comparability with results in CHO. The next subsection revisits some of their evidence.

3.2 Cointegration Evidence

The main contribution of CHO is to show the existence of a cointegrating relationship between unemployment and real input prices. Their Error Correction Model indicates the existence of a long run relationship with signs as predicted by the theory: unemployment increases in the long run with both the real compensation for capital and energy. However, their results may be driven by the sample period they consider. Inspection of Figure 2 reveals that the real price of oil exhibited very little variation in

⁶ The inflation rate is calculated as the annualised quarterly growth rate (log difference) of the GDP deflator.

⁷ The series mnemonics are: TB5Y for the nominal five year treasury bill yield; OILPRICE for the spot price of oil (West Texas Intermediate); GDPDEF for the deflator; and UNRATE for the civilian unemployment rate. I aggregate to quarterly by using the middle-month observation.

the period prior to the first OPEC supply shock in the fourth quarter of 1973. Until that date the price of oil was virtually fixed in nominal terms with the exception of a few minor discretionary changes.⁸ Variation in the real price of crude between 1954 and the end of 1973 is mostly due to the erosion of the nearly constant nominal price by a slowly increasing price level.⁹ In other words, during almost half the 1954:2 to 1995:2 sample period upon which the estimates in CHO are based, one of the two drivers of unemployment was almost constant. It is therefore possible that their estimates overstate the importance of the other driver, the real interest rate, as the model will tend to attribute the dynamics of unemployment to the factor that varies rather than to the factor that does not. Even if this is not necessarily the case, the lack of variation in the real price of oil certainly renders the 1954-1973 period an atypical sample from a more contemporary perspective. This makes exploring the more recent experience worthwhile.¹⁰

To examine this hypothesis, I estimate cointegrating relationships for unemployment and the two real input prices for the period 1970:1 to 2005:4.¹¹ Beyond investigating whether the results in CHO might be driven by the sample period considered, such evidence is useful for the additional reason that it updates their findings. Crucially, it will address the question whether energy prices are still important for the US macroeconomy, a subject of intense debate in the empirical literature (see, e.g.,

⁸ Prior to 1973 there have been several small adjustments of a few cents each time in both directions. The biggest increases were in 1953:2 (from \$2.57 to \$2.82), 1957:1 (\$2.82 to \$3.07) and 1970:4 (\$3.31 to \$3.56). See Hamilton (2003) for a historical discussion of the market for crude oil.

⁹ Between 1954 and 1973 US inflation was relatively moderate. GDP deflator inflation reached 6.2 percent only in the third quarter of 1973, while in general remaining much lower than that for most of the period.

¹⁰ In addition, the low degree of variation in the real price of crude will induce large standard errors on the estimated oil coefficients, and thus understate the significance of oil in a statistical sense.

¹¹ The integration properties of the series are as follows (all evidence is based on Augmented Dickey Fuller tests). The unemployment rate marginally rejects the unit root hypothesis at the ten percent level for one lagged difference in the test equation chosen on the basis of the Akaike Information Criterion, but fails to reject for higher lag orders which feature better residual properties for the equation. On the other hand, the tests are unable to reject the unit root null for the real oil price, the real rate, and the real wage (used below), even when a trend term is included in the test equation. I do not report these results here to economise space.

Hamilton 2003 and Barsky and Kilian 2004 for recent surveys). In their investigation of the long run relationship between unemployment and real input prices, CHO use a single-equation technique (Engle and Granger 1987) to test for and estimate cointegration. I shall not do so here. Rather, I use the cointegrating vector autoregressive (VAR) method pioneered by Johansen (1995). This approach has the advantage of econometric efficiency (Gonzalo 1994) as well as modelling transparency. My benchmark is a trivariate VAR in unemployment, the real interest rate, and the real oil price. To check robustness of the results, I also estimate a quadrivariate VAR which includes the real wage. The motivation to use this variable is primarily theoretical: in structuralist models á la Phelps (1994) and CHO, the real wage is the main channel of transmission for input price shocks.¹² From a statistical point of view, including the real wage also makes the VAR more flexible since it adds another margin of adjustment. My real wage measure is the log real compensation per hour in the nonfarm business sector.¹³ I use this wage measure because it is broader than manufacturing wages frequently employed in empirical work. Below, I report the main results. A more detailed discussion of the estimation procedures is available from the author upon request.

For both the trivariate and the augmented VAR, a trend is included which is restricted to enter the cointegrating equation.¹⁴ A lag length of two is chosen for both first-differenced models on the basis of information criteria, sequential likelihood ratio tests and residual properties, although none of the results depends on this choice. Table 1 shows cointegration test results for the two models for the case of one and two lagged differences along with critical values. I rely on the Trace statistic as it is more robust to

¹² Rotemberg and Woodford (1996) provide evidence that the real wage falls following an increase in the price of oil.

¹³ The FRED mnemonic for this series is COMPRNFB.

¹⁴ An unrestricted trend in the first-differenced VAR would induce a (counterfactual) quadratic trend in the levels model.

departures from normality than the Maximum Eigenvalue test (Cheung and Lai 1993), and the residuals do indicate nonnormality. There is strong evidence in favour of a single cointegrating relationship for both the benchmark and the augmented model.

Table 2 reports cointegration estimates for the two specifications. In the quadrivariate model I straight away restrict the real wage not to enter the long run relationship (referred to as “overidentified (1)”); the restriction yields a p-value of 0.48. In both the just identified trivariate and the quadrivariate VAR, the two input prices have the sign predicted by theory. However, while the real price of oil is highly statistically significant, the real rate is not. The trend in the cointegrating equation is also significant, while its sign is negative. This downward drift in the equilibrium relationship is compatible with a decrease in the US natural rate of unemployment over the past twenty years, as has been suggested by several authors recently (see for example Gordon 1998 or Laubach 2001).

In the next step, I impose a set of zero restrictions suggested by these first estimates. In the benchmark VAR, I exclude the real interest rate from the cointegrating equation and impose nonadjustment of the oil price equation to equilibrium deviations. These restrictions are easily accepted by the data, yielding a p-value of 0.34. I impose the same restrictions in the quadrivariate VAR; in addition, I restrict the adjustment coefficient in the real wage equation to zero. Again, the test statistic is well to the left of the critical values, with the p-value being 0.54.¹⁵ A graph of the equilibrium error from the augmented VAR is shown in Figure 3. In both cases therefore, the real interest rate does not seem to exert a significant influence on long run unemployment. On the other hand, the real price of oil is highly significant while its coefficient is sizeable.¹⁶

¹⁵ For both models, imposing the zero restrictions in a sequential fashion rather than simultaneously failed to produce a rejection for any hypothesis considered.

¹⁶ The fact that the coefficient on the real price of oil increases substantially when the real rate is excluded probably suggests that the latter is not totally unimportant. However, the oil price coefficient jumps even

Thus, cointegration techniques indicate no important role for the real interest rate in long run US unemployment dynamics. This finding is robust to the inclusion of the real wage in the analysis, the main margin of adjustment in the underlying theory. The result concerning the lack of statistical significance for the real interest rate thus stands in contrast to the evidence in CHO, who report a significant real interest rate in their long run relationship estimated over the period 1954:2 to 1995:2. My interpretation is that their result may be driven by sample selection.¹⁷ Roughly half of CHO's sample, from 1954 to 1973, exhibited very little variation in the real price of oil. Once a less atypical sample period is considered where both real input prices vary substantially throughout, as is the case here, the real rate loses its significance.¹⁸

On the other hand, there does seem to be a robust relationship between the real price of oil and unemployment as predicted by endogenous natural rate theories such as Phelps (1994) and CHO.¹⁹ Indeed, there is strong evidence in the data of cointegration in the *bivariate* relationship between unemployment and the real price of oil: when I estimate a cointegrating VAR in unemployment and crude, the coefficient of the latter in the estimated cointegrating equation is 1.95, which is very close to the estimates above.²⁰ It is thus likely that the successful forecasting performance of CHO's model is due to the

when exclusion restrictions on the adjustment parameters in the oil and real wage equations are the only restrictions imposed.

¹⁷ I have attempted to further substantiate this claim by estimating the cointegrating relationship for the sample in CHO with the Johansen procedure. The tests showed any number of cointegrating relationships between 1 and 3; and long run estimates under the assumption of a single cointegrating vector were not encouraging. It is therefore unlikely that the difference in the results is due to the estimation method.

¹⁸ CHO report results from three equations, all of which have been estimated from the beginning of their sample period, 1955:4, to various endpoints; this data therefore always includes the period of low variation in the real oil price (see their Table 3, p. 625). They do not report results from any post-1970 estimation.

¹⁹ Interestingly, in related research for the Euro Area over the same sample period I find that it is the real rate of interest which is significant in the cointegrating relationship, while the real price of energy is not.

²⁰ Both Trace and Maximum Eigenvalue statistics strongly suggest the presence of a single cointegrating relationship. These results are available upon request.

link between the price of crude and unemployment rather than any contribution from the real interest rate (see their p. 626/7).²¹

There is a more general issue here, however. The finding regarding the importance of the real price of oil for unemployment deserves to be emphasised because the continued existence of an oil price-macroeconomy nexus is a matter of considerable controversy in the recent empirical literature. This literature has utilised a variety of complicated nonlinear specifications in order to investigate the impact of oil on the US economy. The present paper shows that, at least in the long run, a simple, linear relationship operates between unemployment and the real price of oil, as first proposed by CHO.²²

4. Granger Causality

The results above call into question the existence of a long run relationship between the real interest rate and unemployment. On the other hand, the evidence on a long run relationship between the price of crude and unemployment seems to be quite solid. In order to further elucidate the relationship between unemployment and input prices this section considers Granger causality evidence. Such evidence may be useful as it provides an atheoretical summary of the relationships between the variables in question. Further, GC imposes much fewer requirements on the data than cointegration. The question is simply: do past values of variable X contain information for current values of Y, given the information in past values of Y? In the present context GC is especially interesting since previous results in the literature have been inconclusive, both for the real interest rate and the real oil price.

Regarding the nexus between oil prices and unemployment, GC has been the main methodological vehicle by which the relationship between energy and the

²¹ This should be particularly the case for the ability of their model to forecast the 1990/91 recession, which was preceded by the sharp increase in oil prices following Iraq's invasion of Kuwait.

²² Most of this literature of course focuses on output growth rather than unemployment. Given Okun's law, the step from unemployment to output growth is a small one, however.

macroeconomy has been investigated empirically, see Hamilton (1983) for a pioneering contribution and Hooker (1996) for more recent work. However, the jury is still out on this issue. Indeed, CHO themselves “... do not view oil price-macroeconomy Granger causality questions as settled ...” (p. 625). The same authors, further, fail to reach a conclusion on the granger causal relationship between unemployment and the real rate as the evidence is very sensitive to both lag length of the VAR employed, and sample period.

This section provides two perspectives on GC. A first perspective is gained by means of tests over various sample periods. Then, I make an attempt to track changes in GC over time. Thus, the following subsection updates and extends some of the evidence presented in CHO on input price-unemployment GC in fixed samples. The second subsection discusses time-varying GC by means of a rolling window procedure.

4.1 Granger Causality in Fixed Samples

I begin this subsection with results from the period 1970-2005, the sample on which the cointegration evidence in the previous section was based. This will enable me to look at evidence from the estimated VECM's as well as unrestricted VARs. I then move on to consider all available observations, i.e. the period 1953:2 to 2005:4, and subperiods thereof.

4.1.1 Sample Period 1970 to 2005

GC test results for the period 1970:1 to 2005:4 are given in Table 3. The left hand side of the table gives results from the unrestricted, levels VARs, while the restricted, cointegrating VARs in first differences are on the right hand side. All these specifications include a linear trend, which is restricted to enter the cointegrating equation for the VECM's. For both cointegrating VARs, all exclusion restrictions discussed above have been imposed. For the trivariate, unrestricted VAR the no GC null

hypothesis is strongly rejected for the price of crude, but not for the real interest rate. For the restricted VAR on the other hand, the data fails to reject the exclusion hypothesis for either variable (in first difference).

Similar results emerge from the quadrivariate VAR which augments the benchmark model with the real wage. The no GC hypothesis for the real wage cannot be rejected for either the unrestricted VAR or the VECM. The inclusion of the wage variable, while not changing the results qualitatively, does however throw into sharp relief the differential importance of the two input prices. When compared to the trivariate model, the p-value for the no GC null increases to almost unity for the real interest rate, while remaining in the vicinity of zero for the price of crude. Once again, neither the real rate nor the real price of oil are significant in first differences, once the long run restrictions are imposed. The results therefore support the interesting conclusion that the forecasting contribution of the oil price stems mainly from the long run relationship: once the dynamics of the system is parameterised as an error correction model, lagged differences of the real price of oil fail to forecast unemployment.

4.1.2 Sample Period 1953 to 2005 and Subsamples

I now use all available data. The full sample period then becomes 1953:2 to 2005:4, and I also consider split sample results. The evidence presented here differs from CHO's in two main aspects. One is that I extend the information set to include the real wage. As elaborated upon previously, the real wage is an important margin of adjustment in the relevant theory. I thus investigate GC in both a trivariate and a quadrivariate context. CHO on the other hand consider the basic trivariate relationship as well as the two bivariate unemployment models it nests. My other difference with CHO is that I employ a different breakpoint around which to split the sample. CHO follow Hooker (1996) and consider 1973:4 as a breakpoint. I use 1979:4 instead, having experimented with the

1973 breakpoint as well. There are several reasons for this choice. First, this breakpoint is more relevant for the MS VAR results considered below. Second, it clarifies whether the evidence is sensitive to the chosen breakpoint. And third, it prevents the two subsamples from becoming too different in length. I experimented extensively with many aspects of model specification such as lag length, the inclusion or otherwise of a deterministic trend, and, as mentioned, the breakpoint. So as not to test the patience of even the most interested reader, I present here a selective summary of this evidence only. It turns out that the evidence is not particularly sensitive to the choice of breakpoint between 1973:4 and 1979:4. The statistical significance of the exclusion restriction for the real price of oil and the real price of capital is not affected substantially by the inclusion of wages, whether there is a trend in the model or not. The presence of the wage variable does matter for the order of the VAR, as the lag length criteria frequently select a more parsimonious model when the price of labour is included. The lag order proved to be the single most important determinant of the GC results. Since, frequently, sequential likelihood ratio tests for the lag length arrive at different conclusions than my preferred information criterion, the Akaike criterion (AIC), I present GC evidence for both 2 and 6 lags.²³ No trend is included in the full sample model. On the other hand I do include a linear trend term in the VAR's of the two subsamples, as this is clearly a more adequate description of the data. In no case did the inclusion of the trend substantially affect the results, however.

GC results are summarised in Table 4. For the full sample, sequential LR tests indicate an optimal lag length of 6, while the AIC selects 2. When the real wage is included, both criteria agree on 2 lags. GC is similar between models, given the lag length. Thus, the exclusion restriction for the oil price is highly significant with 2 lags, with or

²³ The maximum length for the lag selection criteria was set to be 8.

without the real wage. Similarly, the same hypothesis for the real rate is highly nonsignificant in either case. With 6 lags, however, the real rate is substantially closer to statistical significance irrespective of whether the wage has been included. The oil price maintains its significance for the longer lag length.

For the subsample 1953:2 to 1979:4, both AIC and sequential LR tests indicate 6 lags for the benchmark model while they also agree on 7 lags for the quadrivariate specification. This time, neither the inclusion of the wage, nor the lag length seem to matter much: in every case, the no GC null is strongly rejected for the real price of oil, but not for the real rate. This is of course to be expected, as the subsample includes the two OPEC shocks. Things look different for the second subsample (1980:1 to 2005:4). Regarding lag length, the sequential LR tests favour 8 lags, while the AIC 2. The ambiguity is resolved when the real wage is introduced, as both criteria now indicate 2 lags. Here, there is a stark difference between the results depending on the order of the VAR. Thus, the parsimonious models fail to reject no GC for the real rate (and for the oil price when the wage is included).

The models with the longer lag length, by contrast, strongly reject the exclusion null for the real rate, and indeed for the price of crude. The importance of the real rate in this subsample is not surprising as many economists would argue that high real interest rates were the decisive factor behind the deep recession of 1981/2. What is less obvious is why this should show up in the model with the long lags but not in the parsimonious one. The interpretation that the real rate influences unemployment with long lags is not quite satisfactory: the big spike in unemployment (1982:4) follows shortly after the abrupt increase in the real interest rate at the beginning of that year.

I draw two conclusions from this analysis. First, the real price of oil is virtually always highly significant. This finding is in contrast to much of the recent literature on the

energy price-macroeconomy relationship, e.g. Hooker (1996), which has cast doubt on whether the price of crude is still important for the US economy. Second, there is not much evidence that the real rate is granger causal for unemployment, except, possibly, at long lags and for certain periods only.

4.2 Time-Varying Granger Causality

To provide a more detailed account on the presence of structural instability in the granger causal relationships under investigation, and to motivate the use of MS techniques employed later on, this subsection presents results from rolling GC tests. Such evidence seems to be new in the literature. Figure 4 and Figure 5 show the time series of p-values from rolling GC tests for unemployment, obtained from excluding the real rate and the real price of oil, respectively.²⁴

From Figure 4 it is evident that GC between the price of capital and unemployment is far from stable over time, as the real interest rate seems to forecast unemployment over certain periods only. In particular, it does so during 1965-67, 1971-72, 1979, 1981-84 and from 1998 into the new millennium.²⁵ This finding shows why there is so little evidence for the real rate Granger causing unemployment: GC episodes are, actually, rather infrequent. What can we learn from the timing of these episodes?

The shaded areas correspond to recessions in the US economy as referenced by the NBER. Out of the 11 episodes of statistically significant real rate-unemployment GC (listed in Footnote 25), only four coincide somewhat with the recessions of 1969/70,

²⁴ These p-values are obtained from a VAR equation with 2 lags, with the window width set at 28 quarters, estimated by Ordinary Least Squares. The lag length was chosen on the basis of the Akaike criterion for the full-sample VAR. The statistic is asymptotically distributed as F with (2,19) degrees of freedom. Results are only mildly sensitive to the choice of window width. The p-values may only be a poor approximation to the true probabilities due to the small sample size and the asymptotic justification of the tests. Psaradakis et al. (2005) use bootstrapped p-values in a similar exercise.

²⁵ Specifically, the null hypothesis that the real interest rate does not GC the unemployment rate yielded p-values of less than ten percent for the following periods: 65:4-67:1; 69:2; 70:1 (recession: 69:4-70:4); 71:1-72:3; 75:3; 79:2-79:4; 81:4-84:3 (recession: 81:3-82:4); 91:1 (recession: 90:3-91:1); 92:1; 98:4-2003:3 (recession: 2001:1-2001:4) and 2005:2-2005:4.

81/82, 90/91 and 2001, while the remaining 7 are outside recessions. On the other hand, some of the quarters with significant GC are just before or after recessions. The evidence is thus not very clear: to the extent that it does speak, it probably warrants the conclusion that the granger causal relationship between the real rate and unemployment holds mostly outside recessions. At the same time, GC seems to be present in 3 out of the 4 post-1980 recessions.

In Figure 5, the real oil price is excluded from the same VAR equation. Once again, somewhat surprisingly perhaps given the results of the previous section, GC is not stable. The hypothesis that real oil prices do not GC unemployment is rejected for 1965-69, 1974-1978, the first quarter of 1988, and 1991-92 and 1994-97.²⁶ Episodes of GC are thus once more surprisingly infrequent, occurring only six times in the last 45 years. Out of those six episodes, only 2 are associated with recessions (see Footnote 26), the first one after the first OPEC shock in the mid seventies and the other one following the oil price spike prior to Gulf War I. In both cases, GC continues to hold for some time after the recession.

Overall, what appears to be conspicuous about these results is the absence of the second OPEC shock in the set of periods where real oil prices forecast unemployment. That is, the price of crude does not granger cause unemployment immediately before, or during, the 1980 recession. Further, the second recession of the early eighties, which saw US unemployment reach a postwar high at 10.8 percent in the fourth quarter of 1982, does not belong to that set of periods either. Rather, Figure 4 indicates that it is the real rate which granger causes unemployment in that episode. This suggests that the two most severe recent US recessions may have had different root causes: a steep increase in the

²⁶ The dates for which the null hypothesis that the real price of oil does not GC the rate of unemployment was rejected at levels of less than ten percent are as follows: 65:4-69:2; 74:4-78:2 (recession: 73:4-75:1); 88:1; 90:4-92:2 (recession: 90:3-91:1); 94:3-95:1; 95:3-97:1.

price of oil for the recession of 1973/5,²⁷ and high real interest rates for the 1981/2 recession. I leave more careful investigation of this issue for future research.

To summarise, fixed sample tests, discussed in the previous subsection, consistently, and strongly, reject the no causality null hypothesis for the price of crude. Rolling results on the other hand indicate that there are actually very few periods for which the real oil price granger causes unemployment. Two of those periods, importantly, coincide with the 1973/5 and 1990/1 recessions, suggesting a causal economic interpretation. On the whole, therefore, these results seem to lend further support to the view that the behaviour of the real price of energy, while important for the US macroeconomy, does not constitute a *regular* influence on the business cycle (see, e.g., Raymond and Rich 1997).

Looking at all the evidence adduced so far allows some speculation about the dynamics of unemployment. The real price of capital fails, as cointegration results have established, to enter the long run input price relationship significantly. Further, GC test results indicate, overall, a weak influence of the real rate on unemployment. On the other hand, while the real price of oil granger causes unemployment for a few recessionary episodes only, fixed sample results indicate a strong granger causal relationship irrespective of sample size and period; at the same time, there is a cointegrating relationship between unemployment and the real price of oil. This evidence allows the inference that the real interest rate may have a short run influence only on unemployment, while the real price of oil has long run effects, possibly by inducing recessions which increase the natural rate. I investigate this further in the next section.

²⁷ Barsky and Kilian (2004) are critical towards the view that oil is the cause of the 73/5 US recession. One of their arguments is that the US economy was already in a recession in the fourth quarter of 1973, when the price of crude shot up.

5. Granger Causality in a Markov Switching Framework

The rolling GC procedure discussed in the previous section generated a number of interesting results and allowed some speculation as to the relative importance of real interest rates and real oil prices for the US business cycle. However, the analysis also raised some questions, the most important one perhaps concerning the timing of GC. This is the case especially for the relationship between the real interest rate and unemployment. In particular, it is occasionally not clear whether GC between the real rate and unemployment occurs during, or outside recessionary episodes. This issue is at the heart of the major methodological shortcoming of the rolling window approach (see Psaradakis et al. 2005 – PRS for a discussion). Effectively, the procedure assumes continuous change in the parameters or relationships under investigation. Frequently, however, this assumption may be inappropriate, for example when there are relationships which are conditional on the phase of the business cycle. In such cases, a framework is desirable which can accommodate discrete shifts, and the Markov Switching methodology of Hamilton (1989) presents such a framework. More specifically, I shall employ the Markov Switching VAR (MS VAR) technique pioneered by Psaradakis et al. (2005 – PRS). This modelling procedure accommodates time-varying GC in Hamilton's Markov Switching framework. The next subsection briefly explains the model setup, and the second subsection discusses the results.

5.1 The *Psaradakis/Ravn/Sola Markov Switching VAR Model*²⁸

PRS have proposed modelling time-varying GC by means of a MS VAR, where the unobserved state reflects the existence or otherwise of a GC relationship between two variables of interest.²⁹ For present purposes it is particularly important that time-varying

²⁸ See in particular PRS (2002), pp. 6-8 (the manuscript version).

²⁹ Thus, I postulate a different type of nonlinearity than the recent literature on the oil price – macroeconomy relationship does, see Hooker (1996) and Hamilton (1996, 2003). These papers emphasise

GC is modelled in a transparent and parsimonious fashion. More specifically, in a bivariate VAR there are 2 state variables capturing GC, one for each possible granger causal relationship. Each of these variables can take on the values zero (absence of GC) or unity (existence of GC). The VAR as a whole is then driven by a single state variable with four distinct states reflecting the possible combinations of the two underlying state variables.

When investigating GC of real interest rates and real oil prices for unemployment, I include the real rate as an endogenous variable in the VAR along with unemployment, while the real price of oil is included only on the right hand side as a conditioning variable. This treatment will allow me to economise on parameters compared to a fully-fledged trivariate VAR, while conditioning enables the model to control for the effects of the oil price.

The Markov switching VAR model is formalised as follows:

$$(1) \quad \begin{bmatrix} y_{1,t} \\ y_{2,t} \end{bmatrix} = \begin{bmatrix} \mu_{10} + \mu_{11}(1 - s_{1,t}) \\ \mu_{20} + \mu_{21}(1 - s_{2,t}) \end{bmatrix} + \sum_{\tau=1}^k \begin{bmatrix} \varphi_{10}^{(\tau)} + \varphi_{11}^{(\tau)}(1 - s_{1,t}) & \psi_1^{(\tau)} s_{1,t} \\ \psi_2^{(\tau)} s_{2,t} & \varphi_{20}^{(\tau)} + \varphi_{21}^{(\tau)}(1 - s_{2,t}) \end{bmatrix} \begin{bmatrix} y_{1,t-\tau} \\ y_{2,t-\tau} \end{bmatrix} \\ + \sum_{\tau=1}^h \begin{bmatrix} \zeta_{10}^{(\tau)} + \zeta_{11}^{(\tau)}(1 - s_{1,t}) \\ \zeta_{20}^{(\tau)} + \zeta_{21}^{(\tau)}(1 - s_{2,t}) \end{bmatrix} z_{t-\tau} + \begin{bmatrix} u_{1,t} \\ u_{2,t} \end{bmatrix}, \quad t = 1, \dots, T.$$

y_1 will be the US unemployment rate. y_2 will be a measure of the real interest rate, while the conditioning variable z is the real oil price. $s_{1,t}, s_{2,t} \in \{0,1\}$ are unobservable state variables which indicate whether there is GC between the two endogenous variables or not.³⁰ If, for example, $s_{1,t}$ is “switched on”, that is, equal to unity, then the

asymmetries in the impact of positive and negative oil price changes or use explicitly nonlinear functional forms.

³⁰ The existence of the two states is an assumption that PRS do not test in their application, and neither will this paper. It is worth pointing out that this type of test is nonstandard, as it falls under the class of tests in which an unidentified nuisance parameter is present only under the alternative hypothesis. By now, several tests for the null hypothesis of “no switching” have been proposed in the econometric literature, see in particular Carrasco, Hu, and Ploberger (2004) for a recent example, who also discuss earlier tests.

real rate, y_2 , is granger causal with respect to unemployment (given significance of at least one of the $\psi_1^{(\tau)}$ coefficients). The variable $s_{2,t}$, in turn, reflects “reverse causality,” such that when it equals unity unemployment GC the real interest rate, though I will not focus on this.

Notice that the existence or otherwise of GC must also influence the respective autoregressive coefficient in each equation, as well as the intercept. This is achieved by specifying those coefficients to be a function of the state variable which determines GC as well. The model is thus quite flexible in that the remaining parameters in each equation will change whenever the zero restriction associated with the absence of GC becomes binding. Finally, h and k are positive integers, both equal to 2 in the subsequent application as two lags are used for both the endogenous and the conditioning variables.

To complete the model setup, $\mathbf{u}_t' = [u_{1,t} : u_{2,t}]$ is assumed to be vector white noise with mean zero and independent of the two state sequences $\{s_{1,t}\}$ and $\{s_{2,t}\}$. The alternative states of nature represented by $s_{1,t}$ and $s_{2,t}$ can be summarised by means of a single state variable s_t which takes on four values according to the confluence of states for the underlying state variables:

$$s_t = \begin{cases} 1, & \text{if } s_{1,t} = 1, s_{2,t} = 1, \\ 2, & \text{if } s_{1,t} = 1, s_{2,t} = 0, \\ 3, & \text{if } s_{1,t} = 0, s_{2,t} = 1, \\ 4, & \text{if } s_{1,t} = 0, s_{2,t} = 0. \end{cases}$$

The shock process then has the following state-dependent covariance matrices:

$$\Sigma_\ell = E(\mathbf{u}_t \mathbf{u}_t' | s_t = \ell) = [\sigma_{ij,\ell}], \quad i, j = 0, 1, \quad \ell = 1, \dots, 4.$$

The two state variables are assumed to be first order Markov chains with time invariant transition probabilities defined as

$$p_{ij}^{(\ell)} = \Pr(s_{\ell,t+1} = j | s_{\ell,t} = i), \quad i, j = 0, 1; \quad \ell = 1, 2.$$

PRS also assume that the two sequences are independent. While not unreasonable from an economic point of view – it is not obvious why the timing of GC in the two “directions” should be related – this assumption also conserves 8 parameters as the unrestricted transition matrix would involve 12 parameters (probabilities) whereas under stochastic independence only 4 parameters need to be estimated.

The independence assumption thus implies the following transition matrix, denoted by **P**:

$$\mathbf{P} = \begin{bmatrix} p_{11}^{(1)} p_{11}^{(2)} & p_{11}^{(2)} (1 - p_{00}^{(1)}) & p_{11}^{(1)} (1 - p_{00}^{(2)}) & (1 - p_{00}^{(1)}) (1 - p_{00}^{(2)}) \\ p_{11}^{(2)} (1 - p_{11}^{(1)}) & p_{00}^{(1)} p_{11}^{(2)} & (1 - p_{11}^{(1)}) (1 - p_{00}^{(2)}) & p_{00}^{(1)} (1 - p_{00}^{(2)}) \\ p_{11}^{(1)} (1 - p_{11}^{(2)}) & (1 - p_{00}^{(1)}) (1 - p_{11}^{(2)}) & p_{11}^{(1)} p_{00}^{(2)} & p_{00}^{(2)} (1 - p_{00}^{(1)}) \\ (1 - p_{11}^{(1)}) (1 - p_{11}^{(2)}) & p_{00}^{(1)} (1 - p_{11}^{(2)}) & p_{00}^{(2)} (1 - p_{11}^{(1)}) & p_{00}^{(1)} p_{00}^{(2)} \end{bmatrix}$$

To summarise, GC is modelled here by means of unobservable Markov state variables which move stochastically between “GC” and “no GC” states. The Markov switching setup allows examination of both the probability of being in any of the two states at any time over the sample period, obtained from Hamilton’s (1994) recursive filtering algorithm, and conventional statistical testing of the presence of GC: it is present (for real input prices to unemployment), if at least one of the parameters $\psi_1^{(\tau)}$ ($\tau = 1, \dots, k$) is significantly different from zero, and analogously for the parameters governing the “reverse causality” state variable.

Estimation is performed by Maximum Likelihood under the assumption that the conditional probability density function of the vector process $\mathbf{y}_t' = [y_{1,t} : y_{2,t}]$ given past realisations $\{\mathbf{y}_{t-1}, \dots, \mathbf{y}_1\}$ of itself and of the scalar sequences $\{z_{t-1}, \dots, z_1, s_t, s_{t-1}, \dots, s_0\}$ is Gaussian. That is, write the VAR conditional on s_t being ℓ , as

$$(2) \quad \mathbf{y}_t = \boldsymbol{\mu}_\ell + \mathbf{A}_{1,\ell} \mathbf{y}_{t-1} + \mathbf{A}_{2,\ell} \mathbf{y}_{t-2} + \mathbf{B}_{1,\ell} z_{t-1} + \mathbf{B}_{2,\ell} z_{t-2} + \mathbf{u}_{t,\ell}.$$

Collect the parameters in a the matrix

$$\mathbf{\Pi}_\ell' \equiv [\boldsymbol{\mu}_\ell \quad \mathbf{A}_{1,\ell} \quad \mathbf{A}_{2,\ell} \quad \mathbf{B}_{1,\ell} \quad \mathbf{B}_{2,\ell}]$$

and variables in the vector

$$\mathbf{x}_t' \equiv [1 \quad \mathbf{y}_{t-1}' \quad \mathbf{y}_{t-2}' \quad \mathbf{z}_{t-1}' \quad \mathbf{z}_{t-2}'].$$

Then the conditional density of the t -th observation is

$$(3) \quad f(\mathbf{y}_t \mid s_t = \ell, s_{t-1}, \dots, s_0, \mathbf{y}_{t-1}, \dots, \mathbf{y}_1, \mathbf{z}_{t-1}, \dots, \mathbf{z}_1; \boldsymbol{\theta}) \\ = (2\pi)^{-n/2} |\boldsymbol{\Sigma}_\ell^{-1}|^{1/2} \exp\{-(1/2)(\mathbf{y}_t - \mathbf{\Pi}_\ell' \mathbf{x}_t)' \boldsymbol{\Sigma}_\ell^{-1} (\mathbf{y}_t - \mathbf{\Pi}_\ell' \mathbf{x}_t)\}$$

where $\boldsymbol{\theta}$ is a vector of parameters.

The likelihood function emerges as a byproduct of Hamilton's filter algorithm and is maximised numerically by means of a quasi-Newton algorithm.³¹ In addition to the "filtered" probabilities that result from the Hamilton procedure, so-called "smoothed" probabilities (which use information from the entire sample) of the state sequences have been computed using Kim's (1994) smoothing algorithm.

5.2 Results

5.2.1 Timing and Pattern of Granger Causality

Maximum Likelihood estimates for the bivariate VAR with two lags, conditional on two lags of the real price of oil, are given in Table 5.³² The coefficient estimates demonstrate that the model generally fits the data quite well. It is particularly successful in explaining unemployment, as 10 out of 12 parameters in the unemployment equation are statistically significant, but less so for the real rate, where only 3 out of 12 estimated

³¹ Details of the filter are well-known. See, e.g., Hamilton (1994).

³² The starting values have been obtained from the OLS estimates of the corresponding linear VAR, estimated over the whole sample. The starting values for the ten "hidden" parameters were set equal to zero ("agnostic priors"). To check robustness, I also used the obtained maximum likelihood estimates as starting values for another round of iterations. These converged instantaneously to the same results. In general, results were, as is commonplace for this class of models, relatively sensitive to the choice of starting values. I thus experimented with many different sets of starting values to ensure that a global maximum has been found.

parameters are significant. The real rate is mainly driven by its own autoregressive dynamics, as well as by the price of oil.

Crucially, $\psi_1^{(1)}$ and $\psi_1^{(2)}$, the coefficients that govern GC for the real interest rate in the unemployment equation of the system, are significantly different from zero. Indeed, in both cases the coefficients are precisely estimated and rejections are strong, demonstrating that the most important links postulated by the model are in place: there is, periodically, a granger causal relationship between the real interest rate and unemployment. On the other hand, $\psi_2^{(1)}$ and $\psi_2^{(2)}$ – the coefficients that indicate how the two lags of the unemployment rate enter the interest rate equation – are not significantly different from zero, suggesting the absence of “reverse causality,” that is, of feedback from unemployment to the real interest rate.

The time path of GC is shown in Figure 6 and Figure 7 which depict, respectively, filtered and smoothed probabilities for the state where the real rate granger causes unemployment ($\Pr(s_t = 1) + \Pr(s_t = 2)$); the shaded areas are again NBER-dated recessions. One observes that, contrary to the results from the constant parameter, fixed sample VARs discussed in the previous section, there is GC from the real interest rate to unemployment. The MS VAR also improves upon the rolling GC results in that the periods where GC does occur are, for the most, very clearly delimited. Thus, the time path of the probabilities indicates that a granger causal relationship between the real rate and unemployment, while unstable, exists over most of the sample period.

Two issues are worth commenting on here. In the first place, it appears that GC was substantially more volatile before 1980. After 1980, and with the exception of the two brief spikes around the early eighties recessions, the path of probabilities becomes more

settled.³³ This is certainly a different picture from the one obtained with the rolling GC procedure, where the episodes of causality are usually very brief, and often involve a mere downward spike in the p-value. The stabilisation in the granger causal relationship after 1980 may be due to a smaller variance of shocks or to a change in monetary policy behaviour or both, though I am reluctant to speculate about this question.

A perhaps more important issue concerns the timing of GC with respect to the US business cycle. Thus, the MS VAR results show that there exists GC from the real rate to unemployment for some time prior to a recession, which then ceases abruptly with the advent of the downturn. This pattern can be observed for the recessions of 1957/58, 1969/70, 1973-75, 1980, 1981/82, 1991, and probably for 2001 as well. The MS VAR methodology thus clarifies the ambiguity present in the rolling window results of the previous section. Recall from the rolling GC procedure that the real rate helps forecast unemployment in eleven distinct episodes, seven of which were outside recessions while the other four coincided, to some extent, with recessions. The results here, however, indicate that GC from the real rate to unemployment occurs almost exclusively during expansions. The MS approach therefore does succeed in providing a more accurate delimitation of the regimes and hence enables better identification of the granger causal episodes.

Of course, this result calls for an explanation. Why would the real rate forecast unemployment during expansions only? A tentative interpretation could be that the real rate induces the downturn (or at least correlates with factors that induce the downturn) –

³³ It is therefore possible that the transition probabilities have changed pre- and post 1980, i.e. there may have been a structural break in the transition matrix. In principle this may be testable by estimating the model for each subsample. Implementation, however, is unlikely to be feasible due to lack of degrees of freedom. Another possibility is that there were larger and/or more frequent shocks hitting the system in the pre – 1980 period. In that case, there may have been a structural break in the covariance matrices. This would be compatible with evidence in Sims and Zha (2006). These authors argue that a model with switching variances fits US data best in a system with a commodity price, the federal funds rate, output, unemployment, consumer prices and a money stock, estimated on monthly data.

this would explain why it forecasts unemployment just up to the recession; yet why the real rate should lose the forecasting ability during the recession itself is not obvious. Indeed, it is not clear whether the observed pattern is a monetary or a real phenomenon. Is the forecasting ability of the real rate over the business cycle due to monetary policy and its influence on aggregate demand? Or are these results just another manifestation of business cycle asymmetries in the real economy, the likes of which have been detected by many authors? To an important extent the answer to these questions depends on what drives the (medium term) real rate. If it is central bank interest rate policy and expected inflation, then the roots of the nonlinearity may be monetary in nature. If, on the other hand, the behaviour of medium term real rates is determined primarily by real factors such as, for instance, the expected profitability of capital, then the answer may be found in the real sector of the economy. Answering these questions is, however, beyond the scope of this paper.

5.2.2 Parameter Estimates

It is also well worth exploring some of the parameter estimates of the model. I start with the coefficients of the real rate in the unemployment equation. The coefficient for the first lag is negative, -0.094, while the second lag is almost equal in magnitude, but has the opposite sign (0.092). The *sum* of the two coefficients – which quantifies the long run impact of the real rate on unemployment – is, therefore, minute: -0.002. Although the restriction that the coefficients are of opposite signs but equal in magnitude is rejected by a Likelihood Ratio test,³⁴ the implication is that the real interest rate has a negligible influence on unemployment in the long run.³⁵

³⁴ The maximised log likelihood from the restricted model is 182.586 compared to 188.471 of the unrestricted model. The corresponding LR statistic is 11.77 which is asymptotically chi squared with one degree of freedom; the hypothesis is rejected, as the one percent percentile is 6.635.

³⁵ The fact that the sum of the real rate coefficients is negative is not very worrying given its absolute size. In any case, the long run effect is probably more accurately identified in the cointegrating VAR where the real rate is positive, if insignificant in the equilibrium relationship.

To see this, assume that GC is present ($s_{1,t} = 1$) and impose the restriction of equal but opposite signs for the two parameters of the real interest rate in the unemployment equation of the VAR (1) above (that is, set $-\psi_1^{(1)} = \psi_1^{(2)}$).³⁶ Then, and with obvious notation, from $U_t = \mu_{10} + \phi_{10}^{(1)}U_{t-1} + \phi_{10}^{(2)}U_{t-2} - \psi_1^{(1)}R_{t-1} + \psi_1^{(1)}R_{t-2} + \dots$, one obtains the long run or steady state solution for the unemployment rate (denoted \bar{U}) as $\bar{U} = \mu_{10}/(1 - \phi_{10}^{(1)} - \phi_{10}^{(2)}) + [(\zeta_{10}^{(1)} + \zeta_{10}^{(2)})/(1 - \phi_{10}^{(1)} - \phi_{10}^{(2)})]\overline{OILP}$,³⁷ where \overline{OILP} is the real price of oil in steady state. In other words, the equilibrium rate of unemployment does not depend on the real rate at all.

The overall evidence presented in this paper thus indicates that the real rate matters little, if at all, for the long run dynamics of unemployment. Recall from the cointegration results that the real rate, while correctly signed, is not significant in any of the long run relationships. While the latter result was arrived at for the sample period 1970-2005, the findings from the MS VAR presented here reinforce this conclusion for the full sample for which data is available.

A further interesting aspect of the parameter estimates is that the real price of oil is statistically significant only during the periods where the real interest rate does not enter the unemployment equation, that is, when the exclusion restriction of noncausality applies: $\zeta_{10}^{(1)}$ and $\zeta_{10}^{(2)}$ are not significant, while $\zeta_{11}^{(1)}$ and $\zeta_{11}^{(2)}$ are.³⁸ Recall from the VAR (1) that whenever there is GC from the real rate to unemployment (that is, when $s_{1,t} = 1$), lag τ of the real oil price enters the unemployment equation with a parameter

³⁶ Of course, this is strictly speaking not true since the hypothesis of equal but opposite signs was rejected statistically. Given that size of the coefficient sum is very small, however, it should be a good first approximation.

³⁷ For simplicity, and because there does not seem to be any feedback from the unemployment rate on the real rate (recall that the coefficients of unemployment in the interest rate equation are insignificant) I ignore the second equation here.

³⁸ In general, the estimated coefficients of the oil price variables exhibit large standard errors in both equations. I attribute this to the structural breaks in the oil price during the 1970's, as well as its high volatility in the eighties and nineties.

$\zeta_{10}^{(\tau)}$. These parameters are insignificant for both lags. Hence, when the real rate is granger causal for unemployment – during expansions – the real oil price is not. On the other hand, when the real rate is *not* granger causal for unemployment, $s_{1,t} = 0$ and the real price of oil enters with parameter $\zeta_{10}^{(\tau)} + \zeta_{11}^{(\tau)}$, of which only the latter part is statistically significant. Thus, during recessions, when the real rate does not granger cause unemployment, the real oil price does. Notice also that the sum of the (statistically significant) coefficients on the real price of oil, $\zeta_{11}^{(1)} + \zeta_{11}^{(2)}$, is large and positive (0.323), which indicates a substantial quantitative effect of the price of crude on long run unemployment.

These results mesh well with both the cointegration and particularly the rolling GC findings discussed in the previous sections. Rolling GC tests showed that the price of oil forecast unemployment, among other periods, in two recessionary episodes following major increases in the real price of oil: the OPEC I price hike and the Gulf War I spike.³⁹ The MS VAR above thus provides another piece of evidence to suggest that the oil price is not a regular feature of the US business cycle. In addition, variables which are more cyclical in nature such as the real rate may be more important for fluctuations of unemployment at business cycle frequencies. In the language of natural rate theory, real rates may be more helpful in explaining movements of unemployment around the natural rate, with the price of oil possibly being more useful in explaining the – less frequent – shifts of the natural rate itself. This may explain the large impact of the price of crude on unemployment indicated by the point estimates.

³⁹ The finding that oil granger causes unemployment when the real rate fails to do so, i.e. during recessions, need not, of course, imply that oil explains *every* US downturn. The present estimates are perfectly compatible with the view, say, that energy price hikes (granger) caused only the 1973-75 and 1990/91 recessions, as the rolling GC evidence suggests.

Finally, a brief look at equation diagnostics also indicates that the MS VAR fits the data well.⁴⁰ Thus, Portmanteau tests for autocorrelation in the residuals fail to reject the no autocorrelation null for all lag orders the statistic was computed for.⁴¹ This is in stark contrast with the same statistics from the benchmark linear VAR with constant parameters. Here, the no autocorrelation null is strongly rejected for each and every lag order it was computed for. I interpret these results as further evidence in favour of the MS VAR model.

To summarise, three main results have emerged from the above analysis. First, the evidence suggests that the real interest rate plays a quantitatively negligible role in long run US unemployment dynamics. Second, GC from the real rate to unemployment occurs mainly during upswings in economic activity. Third, the real price of oil is not a regular influence on the US business cycle, as the price of energy is granger causal for unemployment in recessions only.

6. Conclusion

Theories seeking to explain the dynamics of unemployment and its natural rate have placed a great deal of emphasis on real input prices (Phelps 1994, CHO). This paper contributes to this literature both by updating existing time series evidence as well as offering new empirical results on the unemployment-real input relationship. The findings presented here are of more general interest also because of their relevance for the debate on the oil price-macroeconomy nexus.

I first revisit some of the evidence in CHO regarding the long run relationship between unemployment and the real prices of energy and capital. These authors estimate a cointegrating relationship between the three magnitudes for the period 1954 to 1995 and

⁴⁰ These results are available upon request.

⁴¹ The statistic was computed for lag orders 1-10, 20, and 30. The null hypothesis is always “no autocorrelation” up to lag h .

find that both input prices are statistically significant as well as correctly signed. I show that this result may be the product of sample selection, as the period 1954 to 1973 saw very little variation in the real price of oil. When estimated over the period 1970 to 2005, a period in which both input prices exhibit substantial variation, the real interest rate turns out to be insignificant in the cointegrating relationship.

The paper then turns to Granger causality. GC is shown to exhibit structural instability, with the ability of both input prices to forecast unemployment being confined to certain periods in the sample. Motivated by the temporal structural instability encountered, and in order to further elucidate GC between unemployment and the real rate, a Markov Switching VAR proposed by Psaradakis et al. (2005) is estimated. In agreement with the aforementioned cointegration evidence, the MS VAR indicates that the real rate has little quantitative influence on unemployment in the long run. The results also show that the forecasting ability of the real rate for unemployment is limited almost exclusively to periods of expansion in the US economy. Further, oil prices and real rates seem to play complementary roles as the former helps forecast unemployment only when the latter does not. Given the timing of GC for the real rate-unemployment relationship, the economically significant implication is that the real price of oil granger causes unemployment in recessions.

Several important conclusions emerge. First, the evidence suggests that the real interest rate plays no important role for long run US unemployment dynamics. On the other hand the real price of oil does matter for US unemployment, both in the long run and in the short run, and is likely to be an important factor behind at least some recessions. This means that the price of crude is still important for the US macroeconomy. Indeed, much of the evidence on the oil price-unemployment nexus presented here is based on simple, linear models, in contrast to the recent literature in this area which has

introduced a raft of complicated nonlinear specifications. Finally, the evidence is consistent with the view, advocated in Raymond and Rich (1997) and elsewhere, that the real price of oil does not constitute a regular feature of the US business cycle.

The broad picture emerging from this paper is consistent with the structuralist view of unemployment dynamics and direct (Phelps 1994, CHO) as well as indirect empirical evidence (e.g. Bianchi and Zoega 1998, Papell et al. 2000), whereby infrequent shocks in input prices shift the equilibrium unemployment rate. This paper qualifies this view by suggesting that the factors behind at least some of those shifts are probably energy prices, with little role for real interest rates. A stylised, if somewhat speculative, summary of the relationship between unemployment and input prices suggested by this paper is that the price of oil is responsible for some of the big ratchets in the natural rate at sub-business cycle frequencies, while the real interest rate contributes to the business cycle frequency movements around the natural rate.

Several open questions remain. In the first place, the result that the real rate forecasts unemployment during expansions in the US economy, but ceases to do so once recessions arrive, demands some economic intuition. The influence of the price of oil on US recessions also warrants further investigation, and the MS VAR framework utilised in this paper is a promising vehicle for such an analysis. In future research I intend to provide more evidence on GC with alternative functional forms for the oil price by using some of the specifications that have been proposed recently in the empirical literature. More generally, the MS VAR framework can be used to investigate the sources of recessions by placing likely culprits in competition with one another. Such an exercise would essentially constitute an attempt at “recession accounting.”

References

- Acemoglu, D. and A. Scott (1994): "Asymmetries in the Cyclical Behaviour of UK Labour Markets." *Economic Journal* 104, November, pp. 1303-1323.
- Barsky, R. and L. Kilian (2004): "Oil and the Macroeconomy Since the 1970s." *NBER Working Paper* 10855, Cambridge, Mass.
- Bianchi, M. and G. Zoega (1998): "Unemployment Persistence: Does the Size of the Shock Matter?" *Journal of Applied Econometrics* 13(3), pp. 283-304.
- Carrasco, M., L. Hu and W. Ploberger (2004): "Optimal Test for Markov Switching." Working Paper, University of Rochester.
- Carruth, A., M. Hooker, and A. Oswald (1998): "Unemployment Equilibria and Input Prices: Theory and Evidence from the United States." *Review of Economics and Statistics* 80, pp. 621-628.
- Cheung, Y.W. and K.S. Lai (1993): "Finite Sample Sizes of Johansen's Likelihood Ratio Tests for Cointegration." *Oxford Bulletin of Economics and Statistics* 55, pp. 313-328.
- Coakley, J., Fuertes, A.-M., and G. Zoega (2003): "Evaluating the Persistence and Structuralist Theories of Unemployment from a Nonlinear Perspective." *Studies in Nonlinear Dynamics and Econometrics* 5(3), Article 2.
- Dotsey, M. and M. Reid (1992): "Oil Shocks, Monetary Policy, and Economic Activity." *Federal Reserve Bank of Richmond Economic Review* 78, pp. 14-27.
- Engle, R. and C. Granger (1987): "Cointegration and Error Correction: Representation, Estimation, and Testing." *Econometrica* 55, pp. 251-276.
- Friedman, M. (1968): "The Role of Monetary Policy." *American Economic Review* 58, pp. 1-17.

- Gonzalo, J. (1994): "Five Alternative Methods of Estimating Long-Run Equilibrium Relationships." *Journal of Econometrics* 60, pp. 203-233.
- Gordon, R. (1998): "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU." *Brookings Papers on Economic Activity* 2, pp. 297-346.
- Hamilton, J. (1983): "Oil and the Macroeconomy Since World War II." *Journal of Political Economy* 91, pp. 228-248.
- Hamilton, J. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." *Econometrica* 57, pp. 357-384.
- Hamilton, J. (1994): *Time Series Analysis*. Princeton, New Jersey. Princeton University Press.
- Hamilton, J. (1996): "This is What Happened to the Oil Price – Macroeconomy Relationship." *Journal of Monetary Economics* 38, pp. 225 – 230.
- Hamilton, J. (2003): "What is an Oil Shock?" *Journal of Econometrics* 113, pp. 363-398.
- Hamilton, J. (2005): "What's Real About the Business Cycle?" *NBER Working Paper* 11161. Cambridge, Massachusetts.
- Hooker, M. (1996): "What Happened to the Oil Price – Macroeconomy Relationship?" *Journal of Monetary Economics* 38, pp. 195 – 213.
- Johansen, S. (1995): *Likelihood-Based Inference in Cointegrated Vector Autoregressive Models*. Oxford, Oxford University Press.
- Kim, C.-J. (1994): "Dynamic Linear Models with Markov-Switching." *Journal of Econometrics* 60, pp. 1-22.
- Laubach, T. (2001): "Measuring the NAIRU: Evidence from Seven Economies." *Review of Economics and Statistics* 83(2), pp. 218-231.

- Layard, R., Nickell, S., and R. Jackman (1991): *Unemployment. Macroeconomic Performance and the Labour Market*. Oxford, Oxford University Press.
- Lindbeck, A. (1993): *Unemployment and Macroeconomics*. Cambridge, Massachusetts, MIT Press.
- Neftci, S. (1984): "Are Economic Time Series Asymmetric over the Business Cycle?" *Journal of Political Economy* 92, 307-328.
- Papell, D., Murray, C. and H. Ghiblawi (2000): "The Structure of Unemployment." *Review of Economics and Statistics* 82(2), pp. 309-315.
- Phelps, E. (1968): "Money Wage Dynamics and Labor Market Equilibrium." *Journal of Political Economy* 76(4), pp. 678-711.
- Phelps, E. (1994): *Structural Slumps – The Modern Equilibrium Theory of Unemployment, Interest, and Assets*. Harvard University Press, Cambridge, Massachusetts.
- Pissarides, C. (2000): *Equilibrium Unemployment Theory*. Cambridge, Massachusetts. MIT Press.
- Psaradakis, Z., M. Ravn, and M. Sola (2002): "Markov-Switching Causality and the Money-Output Relationship." Manuscript, Universidad Torcuato di Tella.
- Psaradakis, Z., M. Ravn, and M. Sola (2005): "Markov-Switching Causality and the Money-Output Relationship." *Journal of Applied Econometrics* 20(5), pp. 665-683.
- Raymond, J. and R. Rich (1997): "Oil and the Macroeconomy: A Markov State-Switching Approach." *Journal of Money, Credit, and Banking* 29(2), pp. 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.

Sims, C. and T. Zha (2006): “Where there Regime Switches in US Monetary Policy?”

American Economic Review 96 (1), March, pp. 54-81.

Figures

Figure 1: United States Unemployment and Real Interest Rate, 1953:2-2005:4

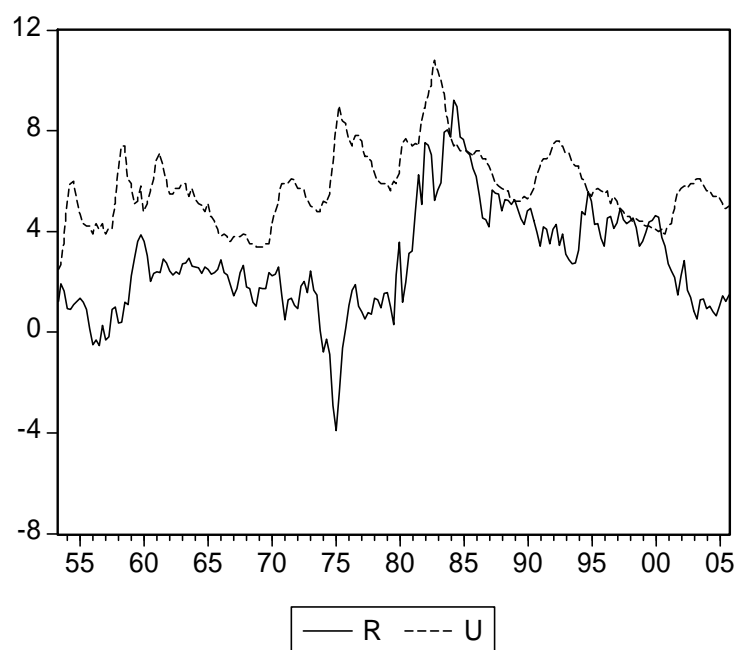
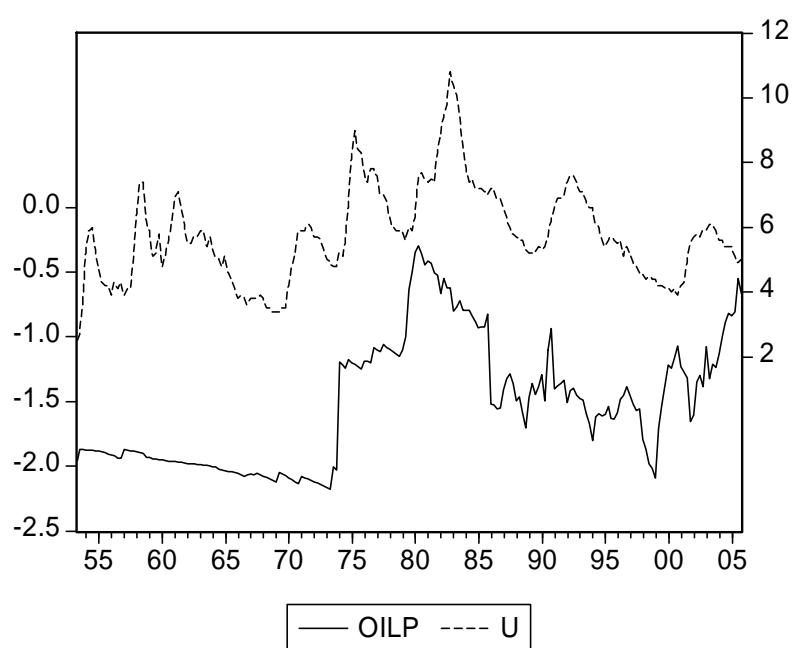
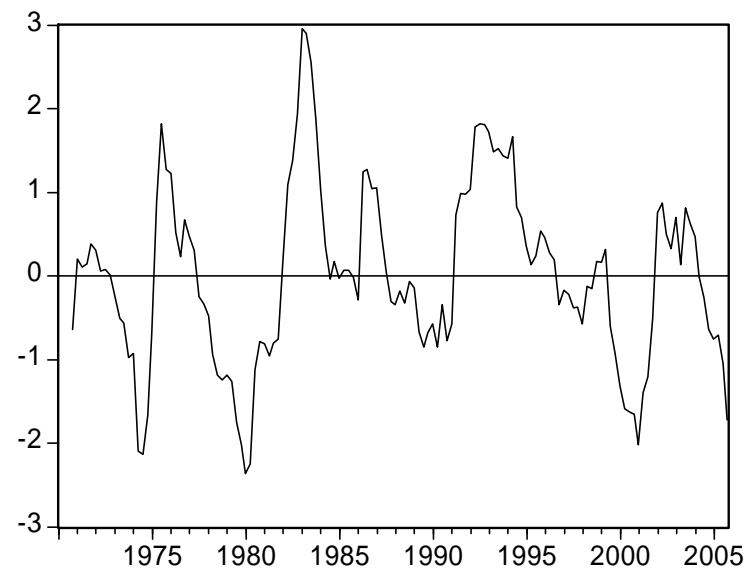


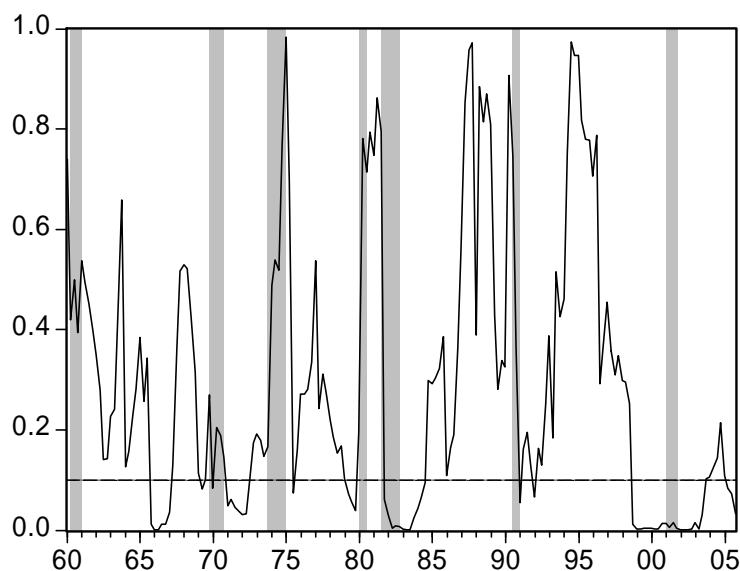
Figure 2: United States Unemployment and Real Price of Oil, 1953:2-2005:4



**Figure 3: Equilibrium Error of Cointegrating Relationship; Augmented VAR,
1970:1-2005:4**



**Figure 4: Rolling Granger Causality Test Results; P-Values for Null Hypothesis
“Real Interest Rate not Granger Causal for Unemployment” (Shaded Areas:
NBER Recessions)**



**Figure 5: Rolling Granger Causality Test Results; P-Values for Null Hypothesis
“Real Oil Price not Granger Causal for Unemployment” (Shaded Areas: NBER
Recessions)**

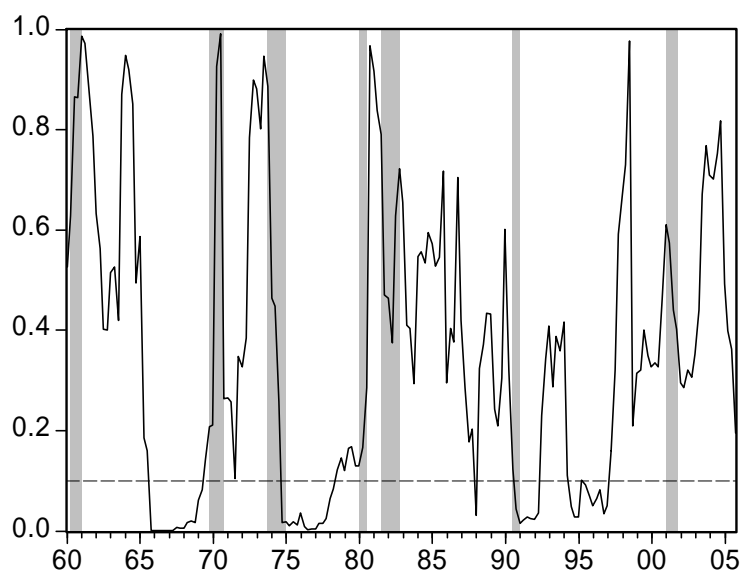


Figure 6: Filtered Probabilities of Real Interest Rate Being Granger Causal for Unemployment (Shaded Areas: NBER Recessions)

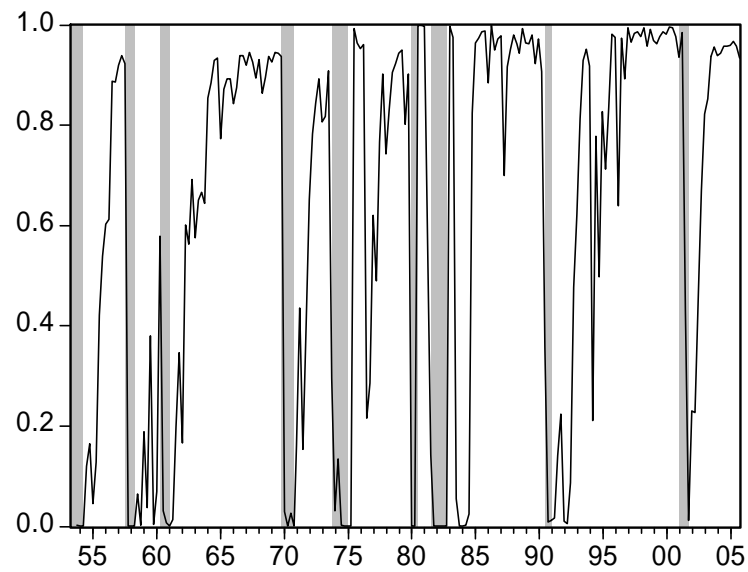
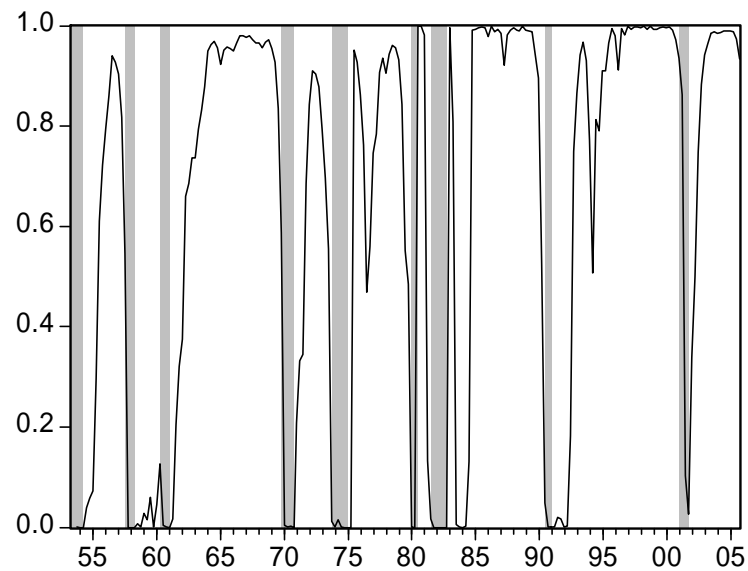


Figure 7: Smoothed Probabilities of Real Interest Rate Being Granger Causal for Unemployment (Shaded Areas: NBER Recessions)



Tables

Table 1: Johansen Trace Tests for Cointegration (1970:1-2005:4)

Benchmark (trivariate) VAR				
<i>Null Hypothesis: cointegrating rank</i>	<i>Trace Statistic</i>		<i>Critical Values</i>	
	1 lagged difference	2 lagged differences	5 percent	1 percent
0	47.21*	49.45**	42.44	48.45
At most 1	22.37	22.51	25.32	30.45
At most 2	8.80	9.48	12.25	16.26
Augmented (quadrivariate) VAR				
<i>Null Hypothesis: cointegrating rank</i>	<i>Trace Statistic</i>		<i>Critical Values</i>	
	1 lagged difference	2 lagged differences	5 percent	1 percent
0	67.76*	70.17**	62.99	70.05
At most 1	41.20	40.99	42.44	48.45
At most 2	15.03	16.03	25.32	30.45
At most 3	3.34	4.11	12.25	16.26

NOTES: Benchmark trivariate VAR: unemployment, real interest rate, real oil price; Augmented quadrivariate VAR: unemployment, real interest rate, real oil price, real wage (see Section 2 for definitions of the variables); *(**) denotes rejection at the five (one) percent level

Table 2: Cointegrating VAR Estimates (1970:1-2005:4)

Benchmark (trivariate) VAR				
	<i>Just identified</i>		<i>Overidentified</i>	
	Coint. Equ.	Adj. Coeff.	Coint. Equ.	Adj. Coeff.
U	1	-0.143 (5.082)	1	-0.138 (5.069)
R	0.111 (1.356)	0.211 (2.913)	0	0.135 (1.897)
OILP	1.552 (3.807)	-0.011 (0.762)	2.050 (4.810)	0
t	-0.016 (3.770)	—	-0.016 (3.527)	—
Augmented (quadrivariate) VAR				
	<i>Overidentified (1)</i>		<i>Overidentified (2)</i>	
	Coint. Equ.	Adj. Coeff.	Coint. Equ.	Adj. Coeff.
U	1	-0.144 (5.070)	1	-0.144 (5.176)
R	0.082 (1.023)	0.213 (2.921)	0	0.152 (2.114)
OILP	1.365 (3.389)	-0.017 (1.136)	1.900 (4.519)	0
W	0	-0.000 (0.172)	0	0
t	-0.015 (3.557)	—	-0.015 (3.440)	—

NOTES: Estimation Method: Maximum Likelihood; Absolute t-statistics in parentheses; U: unemployment rate; R: real interest rate; OILP: real oil price; W: real wage; t: linear trend (see Section 2 for details)

Table 3: Fixed Sample Granger Causality Tests, 1970:1-2005:4

	Statistic	p-value	Statistic	p-value
Trivariate VAR				
	<i>Unrestricted</i>		<i>Restricted (VECM)</i>	
R	0.462	0.794	3.521	0.172
OILP	14.639	0.001	2.265	0.322
Quadrivariate VAR				
	<i>Unrestricted</i>		<i>Restricted (VECM)</i>	
R	0.055	0.973	2.794	0.247
OILP	13.867	0.001	3.236	0.198
W	1.866	0.393	1.774	0.412

NOTES: “Statistic” is the Chi-squared statistic (degrees of freedom equal to the number of lags) resulting from imposing the exclusion restriction on the variable in question; “p-value” is the probability the hypothesis is true; *R* is the real interest rate; *OILP* is the real price of oil; *W* is the real wage (see Section 2 for details); estimation is by Ordinary Least Squares

Table 4: Fixed Sample Granger Causality Tests, Various Samples

Sample Period		2 lags		6 lags	
		Statistic	p-value	Statistic	p-value
1953:2 to 2005:4	<i>Trivariate VAR</i>				
	R	0.111	0.946	9.348	0.155
	OILP	13.412	0.001	13.274	0.039
	<i>Quadrivariate VAR</i>				
	R	0.263	0.876	7.913	0.245
	OILP	18.498	0.000	12.472	0.052
1953:2 to 1979:4	W	9.071	0.011	12.781	0.097
	<i>Trivariate VAR</i>				
	R	1.882	0.390	5.098	0.531
	OILP	6.317	0.042	19.130	0.004
	<i>Quadrivariate VAR</i>				
	R	1.821	0.402	5.604	0.469
1980:1 to 2005:4	OILP	5.728	0.057	14.689	0.023
	W	2.170	0.338	10.204	0.116
	<i>Trivariate VAR</i>				
	R	0.586	0.746	19.883	0.003
	OILP	5.313	0.070	12.718	0.048
	<i>Quadrivariate VAR</i>				
	R	0.697	0.706	17.205	0.009
	OILP	3.616	0.164	12.836	0.046
	W	0.870	0.647	3.256	0.776

NOTES: “Statistic” is the Chi-squared statistic (degrees of freedom equal to the number of lags) resulting from imposing the exclusion restriction on the variable in question; “p-value” is the probability the hypothesis is true; *R* is the real interest rate; *OILP* is the real price of oil; *W* is the real wage (see Section 2 for details); estimation is by Ordinary Least Squares

**Table 5: Markov Switching Vector Autoregression Estimates, United States
1953:2-2005:4**

<i>Unemployment Rate Equation</i>			<i>Real Interest Rate Equation</i>		
	Estimate	Std. Error		Estimate	Std. Error
μ_{10}^*	0.543	0.163	μ_{20}	0.958	0.812
μ_{11}^*	1.165	0.364	μ_{21}	-0.802	0.841
$\varphi_{10}^{(1)*}$	0.799	0.072	$\varphi_{20}^{(1)*}$	0.778	0.109
$\varphi_{11}^{(1)*}$	0.676	0.109	$\varphi_{21}^{(1)}$	0.223	0.151
$\varphi_{10}^{(2)}$	0.126	0.072	$\varphi_{20}^{(2)}$	0.047	0.091
$\varphi_{11}^{(2)*}$	-0.751	0.103	$\varphi_{21}^{(2)}$	-0.223	0.138
$\psi_1^{(1)*}$	-0.094	0.032	$\psi_2^{(1)}$	-0.081	0.266
$\psi_1^{(2)*}$	0.092	0.028	$\psi_2^{(2)}$	0.133	0.251
$\zeta_{10}^{(1)}$	0.095	0.129	$\zeta_{20}^{(1)}$	0.861	0.474
$\zeta_{11}^{(1)*}$	-0.511	0.243	$\zeta_{21}^{(1)*}$	-2.274	0.753
$\zeta_{10}^{(2)}$	0.054	0.140	$\zeta_{20}^{(2)}$	-0.491	0.468
$\zeta_{11}^{(2)*}$	0.834	0.256	$\zeta_{21}^{(2)*}$	1.832	0.766
Markov Probabilities					
$p_{00}^{(1)*}$	0.912	0.031	$p_{00}^{(2)*}$	0.974	0.017
$p_{11}^{(1)*}$	0.851	0.050	$p_{11}^{(2)*}$	0.981	0.013
Covariance Matrices					
$\sigma_{11,1}^*$	0.020	0.004	$\sigma_{11,2}^*$	0.047	0.011
$\sigma_{12,1}^*$	0.220	0.011	$\sigma_{12,2}^*$	0.160	0.016
$\sigma_{22,1}$	0.063	0.071	$\sigma_{22,2}$	0.086	0.045
$\sigma_{11,3}^*$	0.109	0.029	$\sigma_{11,4}^*$	0.238	0.064
$\sigma_{12,3}^*$	1.557	0.059	$\sigma_{12,4}^*$	0.236	0.049
$\sigma_{22,3}$	0.271	0.409	$\sigma_{22,4}^*$	0.236	0.082

Log Lik.: 188.471

NOTES: Estimation method: Maximum Likelihood; * denotes statistical significance at the five percent level or lower