ME
7/9/07
1:00 PM

The Macroeconomic Effects of Oil Price Shocks: Why are the 2000s so different from the 1970s?

Olivier J. Blanchard **Jordi Galí** *

May 28, 2007. Preliminary.

Abstract

We characterize the macroeconomic performance of a set of industrialized economies in the aftermath of the oil price shocks of the 1970s and the last decade, focusing on the differences across episodes. We examine four different hypotheses for the mild effects on inflation and economic activity of the recent increase in the price of oil: (a) good luck (offsetting shocks) (b) smaller share of oil in production, (c) more flexible labor markets, (d) improvements in monetary policy. We conclude that all four have played an important role.

* MIT and NBER, and CREI, UPF and NBER respectively. We thank Davide Debortoli for excellent research assistance, the NSF and the Banque de France Foundation for financial assistance.

Introduction

Since the 1970s, and at least until recently, macroeconomists have viewed changes in the price of oil as as an important source of economic fluctuations, as well as a paradigm of a global shock, likely to affect many economies simultaneously. Such a perception is largely due to the two episodes of low growth, high unemployment and high inflation that characterized most industrialized economies in the mid and late 1970s. Conventional accounts of those episodes of stagflation blame them on the large increases in the price of oil triggered by the Yom Kippur war in 1973, and the Iranian revolution of 1979, respectively.[1]

The events of the past decade, however, seem to call into question the relevance of oil price changes as a significant source of economic fluctuations. The reason: Since the late 1990s, the global economy has experienced two oil shocks of sign and magnitude comparable to those of the 1970s but, in contrast with the latter episodes, GDP growth and inflation have remained relatively stable in much of the industrialized world.

Our goal in this paper is to shed light on the nature of the apparent changes in the macroeconomic effects of oil shocks, as well as on some of its possible causes. Disentangling the factors behind those changes is obviously key to assessing the extent to which the episodes of stagflation of the 1970s can reoccur in response to future oils shocks and, if so, to understanding the role that monetary policy can play in order to mitigate their adverse effects.

One plausible hypothesis is that the effects of the increase in the price of oil proper have been similar across episodes, but have coincided in time with large shocks of a very different nature (e.g. large rises in other commodity prices in the 1970s, high productivity growth and world demand in the 2000s). That coincidence could significantly distort any assessment of the

1. Most undergraduate textbooks make an unambiguous connection between the two oil price hikes of 1973-1974 and 1979-1980 and the period of stagflation that ensued. See e.g. Mankiw (2007, p. 274).

impact of oil shocks based on a simple observation of the movements in aggregate variables around each episode.

In order to evaluate that hypothesis one must isolate the component of macroeconomic fluctuations associated with exogenous changes in the price of oil. To do so, we identify and estimate the effects of an oil price shock using structural VAR techniques. We report and compare estimates for different sample periods and discuss how they have changed over time. We follow two alternative approaches. The first one is based on a large VAR, and allows for a break in the sample in the mid 1980s. The second approach makes use of a number of rolling bivariate VARs, including the price of oil and one other variable at a time. The latter approach allows for a gradual change in the estimated effects of oil price shocks, without imposing a discrete break in a single period.

Two conclusions emerge clearly from this analysis: First, there were indeed other adverse shocks at work in the 1970s; the price of oil explains only part of the stagflation episodes of the 1970s. Second, and importantly, the effects of a given change in the price of oil have changed substantially over time. Our estimates point to much larger effects of oil price shocks on inflation and activity in the early part of the sample, i.e. the one that includes the two oil shock episodes of the 1970s.

We then focus on the potential explanations for these changes over time. We consider three hypotheses, not mutually exclusive:

First, real wage rigidities may have decreased over time. The presence of real wage rigidities generates a tradeoff between stabilization of inflation and stabilization of the output gap. As a result, and in response to an adverse supply shock and for a given money rule, inflation will generally rise more and output will decline more, the slower real wages adjust. A trend towards more flexible labor markets, including more flexible wages, could thus explain the smaller impact of the recent oil shocks.

Second, changes in the way monetary policy is conducted may be responsible for the differential response of the economy to the oil shocks. In particular, the stronger commitment by central banks to maintaining a low and stable rate of inflation, reflected in the widespread adoption of more or less explicit inflation targeting strategies, may have led to an improvement in the policy tradeoff that make it possible to have a smaller impact of a given oil price increase on both inflation and output simultaneously.

Third, the share of oil in the economy may have declined sufficiently since the 1970s to account for the decrease in the effects of its price changes. Under that hypothesis, changes in the price of oil have increasingly turned into a sideshow, with no significant macroeconomic effects (not unlike fluctuations in the price of caviar).

To assess the merits of the different hypotheses we proceed in two steps. First, we develop a simple version of the new-Keynesian model where (imported) oil is both consumed by households and used as a production input by firms. The model allows us to examine how the economy's response to an exogenous change in the price of oil is affected by the degree of real wage rigidities, the nature and credibility of monetary policy, and the share of oil in production and consumption. We then look for more direct evidence pointing to the relevance and quantitative importance of each of those hypotheses. We conclude that all three are likely to have played an important role in explaining the different effects of oil prices during the 1970s and during the last decade.

The paper is organized as follows. Section 1 gives a short summary of our paper fits in the literature. Section 2 presents basic facts. Section 3 presents the results from multivariate and bivariate VARs. Section 4 presents the model. Section 5 uses the model to analyze the role of real rigidities, credibility in monetary policy, and the oil share. Section 6 concludes.

1 Relation to the Literature

Our paper is related to many strands of research:

Bruno and Sachs (1985) were the first to analyze in depth the effects of oil prices of the 1970s on output and inflation in the major industrialized countries. They explored many of the themes of our paper, the role of other shocks, the role of monetary policy, the role of wage setting.

In a series of contributions, Hamilton (in particular Hamilton (1983, 1996)) showed that most of US recessions were preceded by increases in the price of oil, suggesting an essential role for oil price increases as one of the main cause of recessions. The stability of this relation has been challenged by a number of authors, in particular Hooker (1996). Our findings that the effects of the price of oil have changed over time is consistent with the mixed findings of this line of research.

The hypothesis that the stagflations of the 1970s were largely due to factors other than oil has been explored by Barksy and Kilian (2002) among others, who blame monetary policy. Similarly, Bernanke, Gertler and Watson (1997) argue that much of the decline in output and employment was due to the rise in interest rates, resulting from the Fed's response to the higher inflation induced by the oil shocks. Given the recent experience, that explanation would seem insufficient: given that monetary policy is generally viewed as being more responsive to inflation currently than it was in the 1970s, the Bernanke et al. hypothesis would, on its own, imply a larger decline in output and employment in response to the recent adverse oil shocks, and this at odds with the evidence.

Looking at the effects of the price of oil, Rotemberg and Woodford (1997) argued that it was difficult to explain the sheer size of these effects in the 1970s. They argued that something else was going on, namely an endogenous increase in the markup of firms, leading to a larger decrease in output. Finn (2000) showed that the effects of the relevant size could be generated in a perfectly competitive RBC model with variable capital utilization. To explain the depth of the effects of the 1970s and not in the 2000s, we focus instead on the role of real wage rigidities, and the decline in these rigidities over time, an explanation we find more convincing than changes in either the behavior of markups or capacity utilization over time. In following this line, we build on our earlier work on the implications of real wage rigidities and their interaction with nominal price stickiness (Blanchard and Gali 2007).

Our paper is also closely related to the recent research on the “Great Moderation”, the decrease in output fluctuations over the last 30 years (for example, Blanchard and Simon (2001), Stock and Watson (2003)). Our findings that the dynamic effects of oil shocks have considerably decreased is more consistent with an interpretation based on changes in the propagation mechanism than in decreases in the underlying shocks themselves.

The literature on possible changes over time in the effects of oil shocks, to which the present paper seeks to make a contribution, is, as far as we know, almost non-existent. We are aware of two papers which are closely connected with ours and share some of our findings, though they differ in their empirical approach and focus. Hooker (2002) analyzes empirically the changing weight of oil prices as an explanatory variable in traditional Phillips curve specification for the US economy. He finds that pass-through from oil to prices has become negligible since the early eighties, but cannot find evidence for a significant role of the decline in energy intensity, the deregulation of energy industries, or changes in monetary policy as a factor behind that lower pass-through. De Gregorio, Landerretche, and Neilson (2007) provide a variety of estimates of the degree of pass-through from oil prices to inflation, and its changes over time, for a large set of countries. In addition to estimates of Phillips curves along the lines of Hooker (2002), they also provide evidence based on rolling VARs, as we do in the present paper, though they use a different specification, and focus exclusively on the effects on inflation. They provide some evidence pointing to different explanations, including a change in the response of the exchange rate (in the case of non-US countries), and the virtuous effects of being in a low inflation environment. Many of our findings are consistent with theirs or complement them in different ways. In addition, we provide a theoretical framework that can be used to quantify the role of alternative factors behind the observed improvement inthe policy tradeoff.

2 Basic Facts

Figure 1 displays the evolution of the price of oil since 1960 (all figures are at the end of the paper). More specifically, it shows the quarterly average price of a barrel of West Texas Intermediate, measured in U.S. dollars.[2] The figure shows clearly how a long spell of stability came to an end in 1973, triggering a new era characterized by large and persistent fluctuations in the price of oil, punctuated with occasional sharp run-ups and spikes, and ending with the prolonged rise of the past few years.

Figure 2 displays the same variable, now normalized by the U.S. GDP deflator, and measured in natural logarithms (multiplied by 100, so that its variations can be interpreted as percent changes). That transformation gives us a better sense of the magnitude of the changes in the real price of oil. As the figure makes clear, such changes have often been very large, and concentrated over a relatively short period of time.

We start by presenting some evidence on four large oil shock episodes, listed in Table 1. Each of those episodes has involved a cumulative change in the (log) price of oil above 50 percent, sustained for more than four quarters.

2. The description of the stylized facts discussed below is not altered significantly if one uses alternative oil price measures, such as the PPI index for crude oil (used e.g. by Hamilton (1983) and Rotemberg and Woodford (1996)) or the price of imported crude oil (e.g. Kilian (2006)).

Our four episodes start, respectively, in 1973, 1979, 1999, and 2002, though the largest price changes do not coincide with the starting date. For convenience we refer to those episodes as O1, O2, O3 and O4, respectively. Note that our criterion leaves out the price rise of 1990 (triggered by the Gulf War), due to its quick reversal.

TABLE 1: Postwar Oil Shock Episodes

	run-up period	50% rise date	max log change ($)	max log change (real)
OIL1	1973:3-1974:1	1974:1	104 %	98 %
OIL2	1979:I-1981:1	1979:3	94 %	74 %
OIL3	1999:1-2000:4	1999:3	91 %	87 %
OIL4	2002:1-2005:3	2003:1	112 %	104 %

Table 1 lists, for each episode, (i) the run-up period, (ii) the date at which the cumulative log change attained the 50 percent threshold (which we use as a benchmark date below), and (iii) the percent change from trough to peak (measured by the cumulative log change), both in nominal and real terms. The duration of our selected episodes ranges from 5 quarters (O1) to 15 quarters (O4). Interestingly, the size of the associated nominal price rise is similar across episodes, around 100 percent. A similar characterization emerges when we use the cumulative change in the real price of oil (with the price normalized by the GDP deflator), except for O2 where the rise is smaller because of the high rate of inflation during that episode. In short, the four episodes involve oil shocks of a similar magnitude. In particular, the numbers do not seem to justify a characterization of the two recent shocks as being substantially milder in size than the shocks of the 1970s.

In spite of their relatively similar magnitude, these four oil shock episodes have been associated with very different macroeconomic performances. Figures 3 and 4, which show respectively the evolution of CPI inflation and the unemployment rate in the U.S. over the period 1960-2005, provide a visual illustration. Each figure shows, in addition to the variable displayed, four shaded areas representing our four oil shock episodes, as well as a vertical line marking the date at which the 50 percent price increase is attained. Note that the timing of O1 and O2 coincide with a sharp increase in inflation, and mark the beginning of a large rise in the unemployment rate. In each case, both inflation and unemployment reached a peak a few quarters after the peak in oil prices (up to a level of 12.3% and 15.8%, respectively, in the case of inflation, 8.8% and 10.6% for the unemployment rate). The pattern of both variables during the more recent oil shock episodes is very different. First, while CPI inflation shows a slight upward trend during both O3 and O4, the magnitude of the changes involved is much smaller than that observed for O1 and O2, with the associated rises in inflation hardly standing out relative to the moderate size of fluctuations shown by that variable since the mid-1980s. Second, the variation in the unemployment rate during and after O3 and O4 is much smaller in size than that observed in O1 and O2. The timing is also very different: while O1 and O2 lead a sharp rise in unemployment, the latter variable keeps declining during the length of the O3 episode, with its rebound preceding O4. Furthermore, after a persistent (though relatively small) increase, unemployment starts declining in the midst of O4, i.e. while the price of oil is still on the rise.

Figures 5 and 6 provide related evidence for each of the G7 countries as well as for three aggregates (the G7, the euro-12, and the OECD countries).[3] More specifically, Figure 5a displays, for each country and episode, the average rate of inflation over the 8 quarters following each episode's bench-

3. We use quarterly data from OECD's Economic Outlook Database. Inflation is the annualized quarter-to-quarter rate of change in the CPI.

mark date (at which the 50% threshold oil price rise is reached) *minus* the average rate of inflation over the 8 quarters immediately preceding each run-up. Note that the increase in inflation associated with O1 is typically larger than the one for O2. The most striking evidence, however, pertains to O3 and O4, which are typically associated with a change in inflation in their aftermath of a much smaller size than that following O1 and O2.[4] Figure 5b, which displays the same information but averaging the inflation change for O1-O2 and O3-O4, makes the same point in a more dramatic way.

The evidence on output across episodes resembles that on inflation. This is shown in Figure 6a and 6b, which report for each country and episode (or averages of two pairs of episodes in the case of Figure 6b) the cumulative GDP gain or loss over the 8 quarters following each episode's benchmark date, relative to a trend given by the cumulative GDP growth rate over the 8 quarters preceding each episode. The pattern resembles closely that shown for inflation: O1 and O2 are generally associated with GDP losses that are much larger than those corresponding to O3 and O4 (with the latter involving some small GDP gains in some cases). When averages are taken over pairs of episodes, as in Figure 6b, the pattern becomes uniform, pointing once again to much larger output losses during and after the oil shocks of the 1970s.

Is there a systematic relation, across countries and oil shock episodes, between the size of cumulative GDP losses and the increase in inflation? In the face of shocks of similar magnitude (as is the case here), the presence of a time-invariant tradeoff between output and inflation stabilization would imply that episodes characterized by smaller change in inflation would be also associated with larger output losses. The scatterplot shown in Figure 7 shows that this is not the case. On the contrary, smaller rises in inflation

4. Even for Canada and Germany, the largest change in inflation occurs in either O1 or O2.

tend to accompanied by smaller output losses, with the correlation between the two being highly significant. Figure 7 makes clear that an important component of that correlation results from the difference between the early episodes (marked by grey rhombuses, and involving large inflation rises and output losses) and the late ones (marked by pink squares, and characterized by small inflation rises and small output losses—or even output gains).

The evidence presented above is consistent with the hypothesis that the macroeconomic effects of oil price shocks have become smaller over time, being currently almost negligible (at least in comparison with their effects in the 1970s). But it is also consistent with the hypothesis that other (non-oil) shocks have coincided in time with the major oil shocks, either reinforcing the adverse effects of the latter in the 1970s, or dampening them during the more recent episodes. In order to sort out those possibilities we turn next to a more structured analysis of the co-movements between oil prices and other variables.

3 Estimating the Effects of Oil Price Shocks using Structural VARs

In this section we provide more structural evidence on the macroeconomic effects of oil price shocks, and changes over time in the nature and size of those effects. We first provide evidence for six countries, the U.S., France, Germany, the U.K., Italy, and Japan, using a 6-variable VAR. Then we turn to a more detailed analysis of the U.S. evidence using a longer sample period, and to a battery of bivariate VARs.

In all the estimation results reported below, we identify an oil shock as the reduced form innovation to the (log) nominal oil price, measured in U.S.

dollars. In other words, we interpret unexpected variations in the nominal price of oil as exogenous relative to the contemporaneous values of the remaining macroeconomic variables included in the VAR. One may well question this identification assumption:[5] If the economy under consideration is large (say, the US), or even it is small but its output is highly correlated with world output, the price of oil may respond contemporaneously to developments in the domestic economy. We have therefore explored two alternative approaches. The first was to use the average price of oil in the first month of each quarter, which is arguably less likely to vary in response to output and price developments taking place throughout the quarter. The second was to assume that the price of oil responds to unexpected movements in the quantity variables (output, and employment) within the quarter, and that quantity variables in turn do not respond to the price of oil within the quarter. The basic results emphasized below are invariant to either of these two alternative approaches.[6]

3.1 VAR Evidence for the Six Countries

For each of the six countries, we estimate a VAR containing six variables: the price of oil (expressed in log differences), CPI inflation, GDP deflator inflation, wage inflation, and the log changes in GDP and employment. The data are taken from OECD's Economic Outlook database. The frequency of our data is quarterly and the sample period is 1970:1-2006:4. Our three inflation measures are quarter-to-quarter, expressed in annualized terms. Each equation in our VAR includes four lags of the six variables above, a

---

5. Rotemberg and Woodford (1996), who rely on the same identification assumption when studying the effects of oil shocks on the US economy, restrict their sample period to end in 1980 on the grounds that variations in the price of oil may have a significant endogenous component after that date.
6. In a recent paper, Kilian (2006) decomposes movements in prices into movements due to supply shocks, oil-specific demand shocks, and world demand shocks. He finds that nearly all of the one-period ahead reduced form innovation in the price of oil is due to oil specific demand shocks, nearly none to world demand shocks.

constant term and a quadratic trend fitted measure of productivity growth.

Figures 8a-8f display the estimated impulse response functions (IRFs) for the different variables of interest to an oil price shock where, as discussed above, the latter is identified as the innovation in the oil price equation. Estimates are reported for two different sample periods: 1970:1-1983:4 and 1984:1-2006:4. The break date chosen corresponds roughly to the beginning of the Great Moderation, as identified by several authors (e.g. McConnell and Pérez-Quirós (2000). Note that each subperiod contains two of the four large oil shock episodes identified in the previous section.

Confidence intervals covering a 70% probability, and obtained using a Monte Carlo procedure, are shown on both sides of the point estimates. The estimated responses of GDP and employment are accumulated and shown in levels. The size of the shock is normalized so that it raises the price of oil by 100 percent on impact (the estimated standard deviations of oil price innovations for the two subsamples are nearly identical.) In all cases, the real price of oil shows a near-random walk response (not shown here), i.e. it jumps on impact, hovering around the new plateau subsequently.

The estimates for the U.S., shown in Figure 8a, fit pretty well the conventional wisdom about the effects of a rise in oil prices. For the pre-1984 period, CPI inflation shifts up immediately and so do, GDP inflation and wage inflation, even though with the last two only do so with a one-period lag. Output and employment (the latter with a lag) decline persistently. Most relevant for our purposes, the responses of the same variables in the post-1984 period are considerably more muted, thus suggesting a weaker impact of oil price shocks on the economy.

The estimates for France and the U.K. show a pattern very similar to that of the U.S.. In the case of France, the contrast between the early and the late periods is particularly strong, both in terms of the size and the persistence of the effects. For all three countries, the responses of inflation and GDP to a 100 percent increase in the price of oil are roughly of the same order as the change in those variables following the four large oil shocks analyzed in the previous section.

Some of the estimated responses for Germany and Italy fit conventional wisdom less well. The inflation measures in Germany hardly change in response to the rise in oil prices in either period, though the impact on output and employment is more adverse in the pre-1984 period. In the case of Italy, there is barely any employment response in the pre-1984 period. Still, for both countries the sign of most of the responses accord with conventional wisdom, and the responses are smaller in the post-1984 period.

The story is different for Japan. The sign of many of the responses to the rise in oil prices is the opposite to that found in the previous countries. The effect on inflation is weak and of no clear sign in both periods. There is a (significant) rise in output and employment in the post-1984 period.

In short, except for Japan, most of the responses fit conventional wisdom rather well: An increase in the price of oil leads to more wage and price inflation, and to a decrease in employment and output for some time. In all cases however, the effects on both inflation and activity are considerably weaker in the second subsample than in the first.

3.1.1 A Caveat

As discussed above, and at least for the U.S., France and the U.K, the estimated impulse responses suggest that oil price shocks may have indeed been responsible for much of the large and persistent increases in inflation and the declines in output and employment following the two large oil shocks of the 1970s. Yet, a look at the component of the different variables driven by the identified oil price shocks (not shown here) reveals this is not the case, but also points to a potential reason: the estimated oil price shocks account for only about half of the large, discrete changes in the oil price observed within this period. For example, the residual from the oil price equation in our U.S. VAR equals 49% in 1974:1 against a rise of 85% in the price of oil in the same quarter. We suspect that this is largely due to the use of a small sample including two very large changes in the price of oil (85% is roughly 8 times the standard deviation of the reduced form innovation in the price of oil in the sample), which OLS attempts to explain with the lagged variables included in the regression, creating spurious relations and reducing the size of the fitted residual. When we re-estimate the oil price equation using U.S. data for two alternative longer sample periods 1960:1-1983:4 and 1948:1-1983:1, the fitted residual now accounts, respectively, for a 58% and 68% price increase in 1974:1.

We address this issue in two ways below. First, we focus on the U.S. for which longer quarterly time series are available, and re-estimate the VAR using a longer sample period. Second, and in the context of the rolling VARs introduced below (and for which we are forced to use short samples again), we compute impulse response under the assumption that the oil price follows an exogenous random walk, thus forcing the identified oil price shock to match the (demeaned) change in the oil price period by period by construction.

3.2 Oil Prices and Economic Fluctuations in the U.S.: Evidence based on Longer Samples

Figure 9 shows the responses to an oil price shock based on a six-variable VAR similar to the one described above, estimated using quarterly U.S. data starting in 1960:1 (The estimated impulse responses when we let our sample starts in 1948:1 are almost identical to those in Figure 9). Once again, the shock is normalized to have a 100 percent effect on the oil price on impact. In both cases, the estimated responses for the pre-1984 period show the same qualitative features that characterized the estimates based on the shorter sample.[7] The persistence of the response of the three inflation measures is somewhat more pronounced now and, as a result, so is the contrast with the muted responses that characterize the post-1984 period. It is worth noticing, in particular, the delayed responses of GDP deflator and wage inflation, which contrast with the faster shift upward in CPI inflation. Output and employment also show a gradual downward adjustment in response to the increase in oil prices. In all cases the magnitude of the effect is considerable, and comparable to the changes observed in the wake of the two large oil episodes.

How important are oil shocks in accounting for the observed fluctuations in CPI inflation, output and employment? Figure 10 answers that question using the decomposition associated with the estimated six-variable VAR with data starting in 1960.[8] For each variable and sample period, the figure plots the actual time series, together with the component of the latter that results from putting all shocks except the identified oil price shocks equal to zero. CPI inflation is shown after having its means removed, while in the case of the oil price, GDP and employment, their growth rates have been accumulated, with the resulting series then detrended using an HP-filter (so that the series shown can be interpreted as percent deviations from a slowly moving trend). Table 2 below provides some statistics for the role of oil shocks as a source of fluctuations, including its percent contribution to the volatility of each variable, both in absolute and relative terms.

7. Note that the estimates for the post-1984 period are the same as those showed earlier. We include them for the sake of comparability.
8. Again the results are very similar when we start the sample at 1948.

TABLE 2: The Contribution of Oil Shocks to Economic Fluctuations (1960-2005)

	Conditional Sd (%)			Conditional Sd/Unconditional Sd	
	60:1-83:4	84:1-05:4	Ratio	60:1-83:4	84:1-05:4
Oil	11.0	14.4	1.3	0.67	0.80
CPI Inflation	0.92	0.83	0.90	0.25	0.57
GDP Inflation	0.75	0.21	0.28	0.27	0.31
Wage Inflation	0.71	0.53	0.74	0.27	0.32
GDP	0.50	0.29	0.58	0.29	0.22
Employment	0.64	0.46	0.71	0.35	0.22

The estimated standard deviations of the oil-driven component of the different variables ("conditional standard deviations") show that the volatility of fluctuations caused by oil shocks has diminished considerably for all variables, except for oil itself. That evidence reinforces our earlier IRFs-based findings of a more muted response of all variables to an oil shock of a given size. Thus, the change in the way the economy has responded to oil shocks has contributed to the dampening of economic fluctuations since the mid-1980s, a phenomenon known as the Great Moderation. Interestingly, our estimates suggest that this has been possible in spite of the slightly larger volatility of oil prices themselves.

Our estimates suggest that the relative contribution of oil shocks to fluctuations in all variables has remained roughly unchanged over time, at a level between 1/4 and 1/3. The only exception is CPI inflation for which exogenous oil price shocks have played a larger role in the recent period.

Note that this is consistent with a relatively stable core CPI, with oil price changes being passed through to the energy component of the CPI, and accounting for as much as 1/2 of the fluctuations in overall CPI inflation.

Perhaps most interestingly, however, is the finding, apparent in Figure 10, of the limited role of exogenous oil shocks during the 1973-1974 and 1979-1981 episodes. In particular, only about one third of the observed decline in employment and output during those episodes can be attributed to the oil shocks themselves, according to our estimates. The same is true for CPI inflation: oil shocks generate a spike in that variable, but one that falls short of the observed increase in inflation. Thus, our findings suggest that other shocks may have also played an important role in triggering those episodes. Within our 6-variable VAR, our partial identification approach does not allow us to determine what those additional underlying shocks may have been. Yet, when we replace the price of oil by the PPI index for crude materials in our six-variable VAR, the estimates of the component of CPI inflation, GDP, and employment, driven by exogenous shocks to that broader price index track closely the movements of the actual time series themselves in the pre-1984 period, including the two large oil shock episodes contained in that period, as shown in Figure 11. On the other hand, such broader supply shocks seem to have a very limited role in accounting for the fluctuations in output and employment in the post-1984 period (though a more important one in accounting for variations in CPI inflation, in a way consistent with earlier evidence).

3.3 Evidence Based on Rolling Bivariate Regressions

We have so far analyzed the macroeconomic effects of oil price shocks and their change over time under the maintained assumption of a discrete break sometime around the mid-1980s. While the findings reported above are largely robust to changes in the specific date of the break, some of the potential explanations (discussed below) for the change in the effects of oil price shocks are more likely to have been associated with a more gradual variation over time. This leads us to adopt a more flexible approach, and estimate rolling IRFs to oil price shocks, based on a simple dynamic equation linking a variable of interest to its own lags and the current and lagged values of the change in the (log) oil price. We do this using a moving window of 40 quarters, with the first moving window centered in 1970.

More specifically, letting y_t and p_t^o denote the variable of interest and the price of oil, respectively, we use OLS to estimate the regression:

$$y_t = \alpha + \sum_{j=1}^{4} \beta_j \ y_{t-j} + \sum_{j=0}^{4} \gamma_j \ \Delta p_{t-j}^o + u_t$$

and use the resulting estimates to obtain the implied dynamic response of y_t (or a transformation thereof) to a permanent 100 percent (log) change in the price of oil, thus implicitly assuming in the simulation that Δp_t^o is an i.i.d. process.

In order to check the consistency with our earlier results, we first computed the average IRFs across moving windows within each of the subperiods considered earlier (pre-1984 and post-1984), and found the estimated IRFs (not shown) to be very similar to the ones obtained earlier. In particular, both the inflation variables, as well as output and employment, show a more muted response in the more recent period.

Figures 12a-12e display the rolling IRFs. Several features stand out:

CPI inflation appears to be quite sensitive to the oil shock over the entire sample period, but particularly so in the late 1970s, when inflation is estimated to rise more than 10 percentage points two/three quarters after the 100 percent rise in the oil price. In the more recent period, however, the response is more muted and, perhaps as important, less persistent (in a way consistent with our earlier evidence based on the 6-variable VAR). The evolution over time in the response of GDP deflator inflation to an oil price shock is similar to that that of CPI inflation, but shows a more dramatic contrast, with the response at the end of our sample being almost negligible. The response of wage inflation is rather muted all along, except for its large persistent increases in the late 1970s and early 80s, and a similar spike in the 1990s.

The most dramatic changes are in the responses of output and employment (Figure 12d-e). In the early part of the sample output is estimated to declines as much as 10 percent two years after the 100 log change in the price of oil. The estimated response, however, becomes weaker over time, with the point estimates of that response becoming slightly positive for the most recent period.

The previous evidence thus reinforces the picture that emerged from the earlier evidence, one which strongly suggests a vanishing of oil shocks on macroeconomic variables, and especially on GDP inflation, output and employment. In the remainder we try to uncover some of the reasons why.

4 Modeling the Macroeconomic Effects of Oil Price Shocks: A Simple Framework

We now develop a simple model of the macroeconomic effects of oil price shocks. The purpose is to use the model to help us address two related questions motivated by our earlier findings. First, we want to understand what may account for the more muted response of the economy to oil price shocks in the more recent period. Second, we want to see if we can explain, at least roughly, the magnitude of the responses both in the 1970s and the 2000s.

Our framework is a version of the basic new-Keynesian model with two modifications. First, we introduce oil, and assume it has two uses: It is both consumed by households and used by firms as an input in production. We assume the country is an oil importer, and that the real price of oil (in terms of domestic goods) follows an exogenous process. Second, we allow for real wage rigidities, which, as argued in our earlier work (Blanchard and Galí (2007)), we think are essential to understand the effects of supply shocks, and the policy tradeoff they generate.

We present only the log-linearized equilibrium conditions of the model, relegating their detailed derivation to Appendix 1. Lower case letters denote logs of the original variables. For simplicity we ignore all constants.

4.1 Households

The intertemporal optimality condition of the representative household is given by:

$$c_t = E_t\{c_{t+1}\} - (i_t - E_t\{\pi_{c,t+1}\}) \tag{1}$$

where $c_t \equiv (1-\chi)c_{q,t} + \chi c_{m,t}$ is a consumption index, including both domestically produced goods ($c_{q,t}$) and imported oil ($c_{m,t}$), i_t denotes the nominal interest rate, and $\pi_{c,t} \equiv p_{c,t} - p_{c,t-1}$ is CPI inflation.

The CPI is given in turn by $p_{c,t} = (1-\chi)\, p_{q,t} + \chi\, p_{m,t}$, where $p_{q,t}$ is the price index for domestically produced goods and $p_{m,t}$ is the price of imported oil (expressed in domestic currency). Letting $s_t \equiv p_{m,t} - p_{q,t}$ denote the real price of oil, we can write

$$p_{c,t} = p_{q,t} + \chi\, s_t \tag{2}$$

which implies a direct positive effect of the real price of oil on the CPI, given domestic output prices.

Under the assumption of perfect competition in labor markets (to be relaxed below), the intratemporal optimality condition for the households' problem is given by

$$w_t - p_{c,t} = c_t + \phi\, n_t \tag{3}$$

where w_t is the nominal wage, $p_{c,t}$ is the consumption price index (CPI), n_t denotes employment or hours of work, where the right hand side represents the marginal rate of substitution between consumption and labor, and ϕ is the inverse of the Frisch elasticity of labor supply.

Since we are interested in exploring the consequences of of real wage rigidities, the analysis below is based on a generalization of (3) of the form

$$w_t - p_{c,t} = (1-\gamma)\ (c_t + \phi\ n_t) \tag{4}$$

where the parameter $\gamma \in [0,1]$ can be interpreted as an index of the degree of real wage rigidities. While clearly ad-hoc, (4) is meant to capture in a parsimonious way the notion that real wages may not respond to labor market conditions as fast and fully as implied by the model with perfectly competitive markets. (We have explored the implications of a dynamic version of (4), in which the wage adjusts over time to the marginal rate of substitution. That alternative is more attractive conceptually, and gives richer dynamics. However, it is also analytically more complex, and we have decided to present results using the simpler version above.)

4.2 Firms

Firms are monopolistically competitive, with a production function given by

$$q_t = a_t + \alpha_n n_t + \alpha_m m_t$$

where q_t is (gross) output, a_t is an exogenous technology parameter, m_t is the quantity of oil used in production, and $\alpha_n + \alpha_m \leq 1$.

Cost minimization implies that the firms' demand for oil is given by $m_t = -\mu_t^p - s_t + q_t$, where μ_t^p is the price markup. Using this expression to eliminate m_t in the production function gives a reduced form aggregate production relation

$$q_t = \frac{1}{1-\alpha_m} \ (a_t + \alpha_n n_t - \alpha_m s_t - \alpha_m \mu_t^p) \tag{5}$$

Note that (5) implies that increases in the real price of oil or the markup reduce gross output, given technology and employment.

Prices are staggered à la Calvo (1983), an assumption which yields the following log-linearized equation for domestic inflation $\pi_{q,t} \equiv p_{q,t} - p_{q,t-1}$:

$$\pi_{q,t} = \beta \ E_t\{\pi_{q,t+1}\} - \lambda_p \ \mu_t^p \tag{6}$$

where $\lambda_p \equiv [(1-\theta)(1-\beta\theta)/\theta][(1-\alpha_k)/(1-\alpha_k+\alpha_k\epsilon)]$, where θ denotes the fraction of firms that leave prices unchanged, β is the discount factor of households, ϵ is the elasticity of substitution between domestic goods in consumption, and $\alpha_k \equiv (1-\alpha_m-\alpha_n) \geq 0$.

Combining the cost minimization conditions with the aggregate production function yields the following factor price frontier:

$$\mu_t^p + (1-\alpha_m) \ (w_t - p_{c,t}) + \alpha_k \ n_t = a_t - (\alpha_m + (1-\alpha_m)\chi) \ s_t \tag{7}$$

An increase in the real price of oil which is not offset by higher productivity must be "absorbed" through one or more of the following adjustments: (i) a lower level of employment, (ii) a lower consumption wage, or (iii) a lower markup (and, hence, higher inflation, given (6)).

4.3 Equilibrium

To replace the consumption wage in the previous equation using equation (4), we first need an expression for consumption. Under the assumption of balanced trade, consumption is given by

$$c_t = q_t - \chi \ s_t + \eta \ \mu_t^p \tag{8}$$

where $\eta \equiv \alpha_m/(\mathcal{M}^p - \alpha_m)$, with $\mathcal{M}^p$ denoting the steady state markup. Given aggregate output, an increase in the real price of oil decreases consumption. An increase in the markup lowers the fraction of output that is transferred to foreign oil producers, thus raising income and consumption.

Combining (7) with (4), together with (5) and (8), yields an expression for the markup as a function of employment, productivity, and the real price of oil:

$$\mu_t^p + \Gamma_n \ n_t = - \ \Gamma_s \ s_t + \Gamma_a \ a_t \tag{9}$$

where

$$\Gamma_n \equiv \frac{\alpha_k \gamma + (1-\alpha_m)(1-\gamma)(1+\phi)}{1-(1-\gamma)(\alpha_m - (1-\alpha_m)\eta)} \geq 0$$

$$\Gamma_a \equiv \frac{\gamma}{1-(1-\gamma)(\alpha_m - (1-\alpha_m)\eta)} \geq 0$$

$$\Gamma_s \equiv \frac{\gamma \ (\alpha_m + (1-\alpha_m)\chi)}{1-(1-\gamma)(\alpha_m - (1-\alpha_m)\eta)} \geq 0$$

Using this expression for the markup in (6) gives a version of the new Keynesian Phillips curve for domestic inflation

$$\pi_{q,t} = \beta \ E_t\{\pi_{q,t+1}\} + \lambda_p \Gamma_n \ n_t + \lambda_p \Gamma_s \ s_t - \lambda_p \Gamma_a \ a_t \tag{10}$$

A change in the real price of oil (or a change in technology) generates a tradeoff between stabilization of domestic inflation and stabilization of employment only if real wage rigidities are present, i.e. $\gamma > 0$. In the latter case, an increase in the real price of oil triggers a rise in domestic inflation, unless employment declines.

Note that the characterization of the equilibrium did not require introducing either value added or the value-added deflator. But these are needed to compare the implications of the model to the data.

Letting the value added deflator $p_{y,t}$ be implicitly defined by $p_{q,t} = (1-\alpha_m)p_{y,t} + \alpha_m p_{m,t}$, and rearranging terms gives:

$$p_{y,t} = p_{q,t} - \frac{\alpha_m}{1-\alpha_m} \ s_t \tag{11}$$

thus implying a negative effect of the real price of oil on the value added deflator, given domestic output prices.

The definition of value added, combined with the demand for oil, yields the following relation between value added and output:

$$y_t = q_t + \frac{\alpha_m}{1-\alpha_m} \ s_t + \eta \ \mu_t^p \tag{12}$$

This in turn implies the following relation between value added and consumption:

$$y_t = c_t + \left(\frac{\alpha_m}{1-\alpha_m} + \chi\right) \ s_t \tag{13}$$

An increase in the price of oil decreases consumption given value added both because (imported) oil is used as an input in production, and used as an input in consumption.

Under the approximation $\left(\eta - \frac{\alpha_m}{1-\alpha_m}\right) \mu_t^p \simeq 0$ (which will be good if we have a low net steady state markup $\mathcal{M}^p$ and a small α_m), (5) and (12) imply the following relation between value added and employment:

$$y_t = \frac{1}{1-\alpha_m} \ (a_t + \alpha_n \ n_t) \tag{14}$$

Note that, under this approximation, the relation between value added and employment does not depend on the real price of oil.

Equations (1), (2), (13), (14), and (10) describe the equilibrium dynamics of consumption, the CPI, value added, employment, and domestic inflation, given exogenous processes for technology and the real price of oil, and a description of how the interest rate is determined (i.e. an interest rate rule). In the following subsection we use those conditions to characterize the economy's response to an oil price shock.

4.4 Quantifying the Effects of Oil Price Shocks

Assume that $a_t = 0$ for all t (i.e abstract from technology shocks), and assume that the real price of oil follows an AR(1) process

$$s_t = \rho_s \ s_{t-1} + \varepsilon_t \tag{15}$$

Under this assumption we can summarize the equilibrium dynamics of value added and domestic inflation through the system:

$$\pi_{q,t} = \beta \ E_t\{\pi_{q,t+1}\} + \kappa \ y_t + \lambda_p \Gamma_s \ s_t \tag{16}$$

$$y_t = E_t\{y_{t+1}\} - (i_t - E_t\{\pi_{q,t+1}\}) + \frac{\alpha_m(1-\rho_s)}{1-\alpha_m} \ s_t \tag{17}$$

where $\kappa \equiv (\lambda_p \Gamma_n (1-\alpha_m))/\alpha_n)$.

These two equations must be complemented with a description of monetary policy. Assume an interest rate rule of the form

$$i_t = \phi_\pi \ \pi_{q,t} \tag{18}$$

where $\phi_\pi > 1$. Note that in our model $\pi_{q,t}$ corresponds to core CPI inflation, a variable that many central banks appear to focus on as the basis for their interest rate decisions.

We can then solve for the equilibrium analytically, using the method of undetermined coefficients. This yields the following expressions for domestic inflation and output:

$$\pi_{q,t} = \Psi_\pi \ s_t$$

$$y_t = \Psi_y \ s_t$$

where

$$\Psi_s = \frac{(1-\rho_s)\left(\frac{\kappa \ \alpha_m}{1-\alpha_m} + \lambda_p \Gamma_s\right)}{(1-\rho_s)(1-\beta\rho_s) + (\phi_\pi - \rho_s)\kappa} \ s_t$$

and

$$\Psi_y = \frac{\frac{\alpha_m}{1-\alpha_m}(1-\rho_s)(1-\beta\rho_s) - (\phi_\pi - \rho_s)\varphi_s}{(1-\rho_s)(1-\beta\rho_s) + (\phi_\pi - \rho_s)\kappa} \ s_t$$

Expressions for CPI inflation and employment can be obtained using (2) and (14), respectively:

$$\pi_{c,t} = \Psi_\pi \ s_t + \chi \ \Delta s_t$$

$$n_t = \Psi_y \frac{1-\alpha_m}{\alpha_n} \ s_t$$

With these equations, we can turn to the discussion of three hypotheses on the nature of the changing effects of oil price shocks. In all cases we use the evidence we presented earlier for the U.S. as a benchmark.

5 Three Hypothesis on the Changing Effects of Oil Price Shocks

In order to assess quantitatively the potential for oil price shocks to generate significant macroeconomic fluctuations, we first need to calibrate our model. We assume the following parameter values:

The time unit is a quarter. We set the discount factor β equal to 0.99. The Calvo parameter θ is assumed to take a value of 0.75. We choose the elasticity of output with respect to labor, α_n , equal to 0.7. We assume $\phi = 1$, thus implying a unitary Frisch labor supply elasticity.

As discussed in previous sections, changes in the volatility of the real price of oil are unlikely to lie behind the changes in the size of the effects of oil shocks. Thus, for simplicity, we assume an unchanged process for the real price of oil. Based on an estimated AR(1) process over the period 1960:1-2006:4 we set $\rho_s = 0.97$ and $var(s_t) = (0.11)^2$. Also, and unless otherwise noted, we set the shares of oil in production and consumption (α_m and χ) to equal 0.015 and 0.023, respectively, which correspond to their values in 1973.

Most of the parameters above are kept constant across all the simulations presented below. The exceptions, as well as our treatment of the remaining parameters, varies depending on the hypothesis being considered in each case.

5.1 On the Role of Real Wage Rigidities and its Changes

In the framework above, the presence of some rigidity in the adjustment of real wages to economic conditions is a necessary ingredient in order to generate significant fluctuations in measures of inflation and economic activity. Figure 13 illustrates this point by showing the range of volatilities of domestic and CPI inflation (annualized, and expressed in percent), employment, and value added, implied by our calibrated model under the assumption of perfectly competitive labor markets ($\gamma = 0$). The left-hand graph represents the policy frontier for the standard deviations of domestic inflation and value added, as ϕ_π varies from 1 to 5, a range of values that cover the empirically plausible set (conditional on having a unique equilibrium). The right-hand graph does the same for CPI inflation and employment. Two features stand out:

First, note that the slope of the frontier is positive. This should not be surprising: In the absence of real wage rigidities, there is no tradeoff between inflation and value added stabilization. Hence, a policy that seeks to stabilize domestic inflation more aggressively, also stabilizes value added. In fact, one can reduce the volatility of domestic inflation, value added and employment as much as desired by choosing ϕ_π to be arbitrarily large, as suggested by the figure (this is what we called the "divine coincidence in an earlier paper). Under the assumed rule, on the other hand, CPI inflation faces a lower bound to its volatility, since it is affected directly by any change in the price of oil, in proportion to the share of oil in the consumption basket.

Second, the standard deviations for the four macro variables implied by our simulation are well below their corresponding estimated conditional standard deviations, reported in Table 2. That volatility gap is particularly dramatic for GDP and employment.

The introduction of real wage rigidities alters that picture substantially. Figure 14 makes this clear, by plotting the policy frontier under flexible wages shown in the previous figure (though with a different scale), together with the corresponding policy frontiers generated under the assumptions $\gamma = 0.6$ and $\gamma = 0.9$. Several results are worth pointing out:

First, the tradeoff generated by the presence of real wage rigidities is apparent in the negative relationship between inflation volatility on the one hand and GDP (or employment) volatility on the other.

Second, while the introduction of real wage rigidities raises the volatility of all variables (for any given ϕ_π), the model's predictions still fall short of matching the (conditional) standard deviations of CPI inflation and employment in our two samples, represented by the two crosses.

Yet, that shortcoming notwithstanding, the figure also makes clear that a moderate reduction in the degree of real wage rigidities can account for a substantial improvement in the policy tradeoff and hence on a simultaneous reduction in the volatility of inflation, employment and GDP resulting from oil price shocks.

To what extent that reduction in the degree of real wage rigidities may have been a factor behind the more muted effects of oil shocks in recent years? We rely again on the bivariate rolling VAR approach used earlier to try to answer that question, by seeking evidence of faster wage adjustment in recent years. In particular, we use that approach to estimate the responses of wage inflation, the real consumption wage, the wage markup, and the unemployment rate. We define the wage markup as the gap between the (log) consumption wage $w_t - p_{c,t}$ and the (log) marginal rate of substitution $c_t + \phi n_t$, with $\phi = 1$, as in our baseline calibration. In response to a rise in the real price of oil, we would expect that gap to become negative in the presence of real wage rigidities, which in turn should be associated with a rise in unemployment.

Figures 15a-c display the relevant IRFs, representing the estimated response of each variable to a permanent, 100 percent increase in the dollar price of oil. The response of the consumption wage and its evolution is shown Figure 15a. Note that the consumption wage tends to decline in response to the adverse oil price shock. Furthermore, while the size of the initial decline and subsequent pattern shows some variability, no tendency towards a larger downward adjustment is apparent. This is confirmed in the graphs corresponding to the real wage in Figure 16, which displays the average IRFs for the period 1960:1-1983:4 and 1984:1-2005:4, respectively, for a number of labor market variables. But it is clear that such evidence cannot be interpreted as suggesting that the degree of real wage rigidities has remained unchanged, since the response of the marginal rate of substitution may differ over time, as our earlier evidence on the changing response of employment might suggest. Figure 15b, which displays the response of the wage markup (shown from two different perspectives), suggests that the size of that response has declined substantially over time, from an initial impact near 15 percent to a slightly negative short term response in recent years. Such a pattern is qualitatively analogous to that observed for unemployment, as displayed in Figure 15c. Both observations are corroborated by the corresponding average IRFs shown in Figure 16.

5.2 Changes in Monetary Policy

A number of studies (e.g. Clarida, Galí, and Gertler (2000)) have provided evidence of a stronger interest rate response to variations in inflation over the past two decades, relative to the 1960s and 1970s. It should be clear, however, from the simulations of our model presented above that, other things equal, the stronger anti-inflationary stance should have reduced the volatility of inflation, while increasing that of GDP and employment. It cannot explain—at least by itself—the evidence of lower volatility of both inflation and economic activity in response to oil price shocks.

In addition to that change in behavior, captured by the literature on empirical interest rate rules however, there is widespread agreement that central banks' commitment to keep inflation low and stable has also become more credible over the past two decades, thanks to improved communications, greater transparency, the adoption of more or less explicit quantitative inflation targets and, ultimately, by the force of deeds. In this section we use the framework developed above to study the role that such an improvement in credibility may have had in accounting for the reduced impact of oil shocks.

We model credibility as follows. As in our baseline model we assume that the central bank follows an interest rate rule

$$i_t = \phi_\pi \pi_{q,t}$$

The public, however, perceives interest rate decision to follow

$$i_t = \phi_\pi (1-\delta)\pi_{q,t} + v_t$$

where $\{v_t\}$ is taken by the public to be an exogenous i.i.d monetary policy shock, and $\delta \in [0,1]$ can be interpreted as a measure of the credibility gap.[9]

In addition to the above actual and perceived policy rules, the model is exactly as the one developed above, with the dynamics of value added, domestic inflation, and the real price of oil summarized by (15)-(17). Solving the model for domestic inflation and value added gives:

$$\begin{aligned} \pi_{q,t} &= a s_t + b v_t \\ y_t &= c s_t + d v_t \end{aligned}$$

where a, b, c, and d are given by:

$$\begin{aligned} a &= \frac{(1-\rho_s)\ \left(\kappa\alpha_m(1-\alpha_m)^{-1} + \varphi_s\right)}{(1-\rho_s)(1-\beta\rho_s) + (\phi_\pi(1-\delta) - \rho_s)\kappa} > 0 \\ b &= -\ \frac{\kappa}{1+\phi_\pi(1-\delta)\kappa} < 0 \\ c &= \frac{\alpha_m(1-\alpha_m)^{-1}(1-\rho_s)(1-\beta\rho_s) - (\phi_\pi(1-\delta)-\rho_s)\varphi_s}{(1-\rho_s)(1-\beta\rho_s) + (\phi_\pi(1-\delta) - \rho_s)\kappa} \\ d &= -\ \frac{1}{1+\phi_\pi(1-\delta)\kappa} \end{aligned}$$

Imposing $v_t = \delta\phi_\pi\pi_{H,t}$ into the solution (so that the central bank actually adheres to its chosen rule) we get

$$\pi_{q,t} = \frac{a}{1 - b\delta\phi_\pi}\ s_t$$

9. Below we restrict ourselves to values of δ greater than $1 - \phi_\pi^{-1}$, in order to guarantee a unique equilibrium.

Employment is given by:

$$\begin{aligned} y_t &= c\ s_t + d\phi_\pi \delta\ \pi_{q,t} \\ &= \left(c + \frac{da\phi_\pi \delta}{1 - b\delta\phi_\pi}\right)\ s_t \end{aligned}$$

Figure 17 displays the volatility frontiers associated with $\delta = 0$ and $\delta = 0.5$. In both cases we restrict ϕ_π to values above 2, in order to guarantee a unique equilibrium. We set γ equal to 0.9. There are two points worth noting:

First, allowing for both real wage rigidities and poor credibility, the model's predictions come closer but still fall somewhat short of matching the (conditional) standard deviations of CPI inflation and employment in our two samples, represented by the two crosses. Given the primitive nature of the model, this may not be overly worrisome.

Second, credibility gains can improve the tradeoff facing policymakers significantly. The quantitative gains, however, do not seem sufficient to account, by themselves, for the observed decline in macro volatility in the face oil shocks, documented earlier in the paper. But they show that improved credibility may certainly have contributed to that decline.

Figures 18a-c provides some evidence of the changes in the Fed's response to oil price shocks, as well as an indicator of potential changes in its credibility. The rolling IRFs displayed are based on estimated bivariate VARs with the price of oil and, one at a time, a measure of inflation expectations over the next 12 months from the Michigan Survey, the 3-month Treasury Bill rate, and the real interest rate (measured as the difference between the previous two variables).

First, and most noticeable, the response of expected inflation to an oil price shock of the same size has shrunk dramatically over time, from a rise of about 6 percentage points in the 1970s, to about 2 percentage points since the mid-1980s, and has remained surprisingly stable after that.

Second, while the strength of the response of the nominal interest rate does not seem to have changed much across sample periods, the shrinking response of expected inflation implies that the response of the real rate to an oil price shock has become stronger over time. In fact, the real rate appears to decline significantly in response to an oil price shock in the 1970s, an observation consistent with the (unconditional) evidence in Clarida et al. (2000), and which in the context of the model above could be a source of multiplicity of equilibria. Perhaps surprisingly, over the most recent period, the real interest rate shows very little change in response to an oil price shock, which is somewhat at odds with estimated interest rate rules for the Volcker-Greenspan period.

Once again, Figure 19 summarize the differences in the average patterns across the two sample periods used throughout the paper.

Overall, we see the analysis and evidence presented above as suggesting that, while reduced sensitivity of inflation expectations to oil price shocks during the recent episodes may be partly the result of a less inflation-accommodating Fed policy, the simultaneous observation of a much smaller decline in employment and GDP documented earlier in this paper suggests that an enhanced anti-inflation credibility may also have played a role.

5.3 Declining Oil Shares

A third hypothesis is that the share of oil in consumption and in production is much smaller today than it was in the 1970s. To examine the possible impact of those changes we simulate two alternative versions of our model, with α_m and χ calibrated using 1973 and 1997 data on the share of oil in production costs and consumption expenditures (see Appendix 2 for details of construction). In light of that evidence we choose $\alpha_m = 1.5\%$

and $\chi = 2.3\%$ (1973 data) for the 1970s and $\alpha_m = 1.2\%$ and $\chi = 1.7\%$ (based on data for 1997) for our two calibrations.

Figure 20 displays the policy frontiers for the two calibrations, keeping the index or real wage rigidities unchanged at $\gamma = 0.9$ (and $\delta = 0$). The conclusion is similar to those reached for the other two candidate explanations. The reduction in the oil shares in consumption and production cannot account for the full decline in volatility, but it clearly accounts for part of it. (The values of α_m and χ in 1977, thus after the first but before the second oil shock were 1.8% and 3.6% respectively. This suggests that, other things equal, the second oil shock should have had larger effects than the first. As we saw earlier, the opposite appears to be however true.)

6 Concluding Comments

We have reached five main conclusions:

First, that the effects of oil price shocks coincided in time with large shocks of a very different nature. Given our partial identification strategy, we have not identified these other shocks. We have given some evidence that increases in other commodity prices were important in the 1970s. We have not identified the other shocks for the 2000s.

Second, that the effects of oil price shocks have changed over time, with a steadily smaller effects on prices and wages, as well as on output and employment.

Third, that a first plausible cause for these changes is a decrease in real wage rigidities. Such rigidities are needed to generate the type of large stagflation in response to adverse supply shocks such as took place in the 1970s. We have shown that the response of the consumption wage to the marginal rate of substitution, and thus to employment, appears to have increased over time.

Fourth, that a second plausible cause for these changes is increased credibility of monetary policy. We have offered a simple formalization of lack of credibility and its effect on the volatility frontier. We have shown that the response of expected inflation to oil shocks has substantially decreased over time.

Fifth, that a third plausible cause for these changes is simply the decrease in the share of oil in consumption and in production. The decline is large enough to have quantitatively significant implications.

Despite the length of the paper, we are conscious however of the limitations of our arguments. Some of the evidence, for example, the IRF evidence for Japan, does not fit our story. The model we have developed is too primitive in many dimensions, and its quantitative implications must be taken with caution. The development of a richer model, at least with respect to the specification of production, and of real wage rigidities, and its estimation, seem natural next steps to check the conclusions reached above.

Appendix 1: A New-Keynesian Model for an Oil-Importing Economy

The present appendix describes the model used in Section 4 and derives the equilibrium conditions underlying the simulations in the main text.

Households

We assume a continuum of identical infinitely-lived households. Each household seeks to maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \; U(C_t, N_t)$$

where period utility is given by

$$U(C_t, N_t) \equiv \log C_t - \frac{N_t^{1+\phi}}{1+\phi}$$

and

$$C_t \equiv \Theta_\chi \; C_{m,t}^{\chi} \; C_{q,t}^{1-\chi}$$

where $C_{m,t}$ denotes consumption of (imported) oil, $C_{q,t} \equiv \left(\int_0^1 C_{q,t}(i)^{1-\frac{1}{\epsilon}} \; di\right)^{\frac{\epsilon}{\epsilon-1}}$ is a CES index of domestic goods, N_t denotes employment or hours worked,. and $\Theta_\chi \equiv \chi^{-\chi}(1-\chi)^{-(1-\chi)}$.

The period budget constraint, conditional on optimal allocation of expenditures among different domestic goods (not derived here) is given by:

$$P_{q,t}C_{q,t} + P_{m,t}C_{m,t} + Q_t^B \; B_t = W_t N_t + B_{t-1} + \Pi_t$$

where $P_{q,t} \equiv \left(\int_0^1 P_{q,t}(i)^{1-\epsilon} \; di\right)^{\frac{1}{1-\epsilon}}$ is a price index for domestic goods, $P_{m,t}$ is the price of oil (in domestic currency), and W_t is the nominal wage. Q_t^B is the price of a one-period nominally riskless domestic bond, paying one unit of domestic currency. B_t denotes the quantity of that bond purchased in period t. For simplicity we assume no access to international financial markets.

The optimal allocation of expenditures between imported and domestically produced good implies

$$P_{q,t}C_{q,t} = (1-\chi)\ P_{c,t}C_t$$

$$P_{m,t}C_{m,t} = \chi\ P_{c,t}C_t$$

where $P_{c,t} \equiv P_{m,t}^{\chi}\ P_{q,t}^{1-\chi}$ is the CPI index. Note that χ corresponds, in equilibrium, to the share of oil in consumption. Note that $P_{c,t} \equiv P_{q,t}\ S_t^{\chi}$, where $S_t \equiv \frac{P_{m,t}}{P_{q,t}}$ denotes the real price of oil (or the terms of trade),expressed in terms of domestically produced goods Taking logs,

$$p_{c,t} = p_{q,t} + \chi\ s_t$$

where $s_t \equiv p_{m,t} - p_{q,t}$ is the log of the real price of oil (measured in terms of domestic goods).

Furthermore, and conditional on an optimal allocation between the two types of goods, we have $P_{q,t}C_{q,t} + P_{m,t}C_{m,t} = P_{c,t}C_t$, which can be plugged into the budget constraint. The resulting constraint can then be used to derive the household's remaining optimality conditions. The intertemporal optimality condition is given by:

$$Q_t^B = \beta\ E_t \left\{ \frac{C_t}{C_{t+1}} \frac{P_{c,t}}{P_{c,t+1}} \right\}$$

Under the assumption of perfect competition in labor markets (to be relaxed below), the household's intratemporal optimality condition is given by

$$\frac{W_t}{P_{c,t}} = C_t\ N_t^{\phi} \equiv MRS_t$$

which is the perfectly competitive labor supply schedule. The log-linearized version of the previous two equations, found in the text, are given by:

$$c_t = E_t\{c_{t+1}\} - (i_t - E_t\{\pi_{c,t+1}\} - \rho) \tag{19}$$

$$w_t - p_{c,t} = c_t + \phi \, n_t \tag{20}$$

where we use lower-case letters to denote the logarithms of the original variables, and where $\pi_{c,t} \equiv p_{c,t} - p_{c,t-1}$ represents CPI inflation.

Firms

Each firm produces a differentiated good indexed by $i \in [0,1]$ with a production function

$$Q_t(i) = A_t \; M_t(i)^{\alpha_m} N_t(i)^{\alpha_n}$$

where $\alpha_m + \alpha_n \leq 1$.

Independently of how prices are set, and assuming that firms take the price of both inputs as given, cost minimization implies that firm i's nominal marginal cost $\Psi_t(i)$ is given by:

$$\Psi_t(i) = \frac{W_t}{\alpha_n(Q_t(i)/N_t(i))} = \frac{P_{m,t}}{\alpha_m(Q_t(i)/M_t(i))} \tag{21}$$

Letting $\mathcal{M}_t^p(i) \equiv P_{q,t}(i)/\Psi_t(i)$ denote firm i's gross markup, we have

$$\mathcal{M}_t^p(i) \; S_t M_t(i) = \alpha_m \; Q_t(i) \; \frac{P_{q,t}(i)}{P_{q,t}}$$

.

Let $Q_t \equiv \left(\int_0^1 Q_t(i)^{1-\frac{1}{\epsilon}} \, di\right)^{\frac{\epsilon}{\epsilon-1}}$ denote aggregate gross output. It follows that

$$M_t = \frac{\alpha_m \; Q_t}{\mathcal{M}_t^p \; S_t} \tag{22}$$

where we have used the fact that $Q_t(i) = (P_{q,t}(i)/P_{q,t})^{-\epsilon} \; Q_t$ (the demad schedule facing firm i), and defined $\mathcal{M}_t^p$ as the average gross markup, weighted by firms' input shares.

Taking logs and ignoring constants

$$m_t = -\mu_t^p - s_t + q_t$$

which can be plugged back in to the (log linearized) aggregate production function to yield

$$q_t = \frac{1}{1-\alpha_m} \; (a_t + \alpha_n n_t - \alpha_m s_t - \alpha_m \mu_t^p) \tag{23}$$

Consumption and Gross Output

Note that in an equilibrium with balanced trade (and hence $B_t = 0$) the following relation holds:

$$P_{c,t} C_t = P_{q,t} Q_t - P_{m,t} M_t$$

Taking logs and using the relations between the different price indexes, we can derive

$$c_t = q_t - \chi \; s_t + \eta \; \mu_t^p \tag{24}$$

where $\eta \equiv \frac{\alpha_m}{\mathcal{M}^p - \alpha_m}$, with $\mathcal{M}^p$ denotes the steady state markup.

Combining (23) and (24), and invoking the fact that $\left(\frac{\alpha_m}{\mathcal{M}^p - \alpha_m} - \frac{\alpha_m}{1-\alpha_m}\right) \mu_t^p \simeq 0$ for plausibly low values of α_m and the net markup measures $\mathcal{M}^p - 1$ and μ_t^p, we can write

$$c_t = \frac{1}{1-\alpha_m} \; a_t + \frac{\alpha_n}{1-\alpha_m} n_t - \left(\frac{\alpha_m}{1-\alpha_m} + \chi\right) \; s_t \tag{25}$$

Gross Output, Value Added, and the GDP Deflator

The GDP deflator $P_{y,t}$ is implicitly defined by

$$P_{q,t} \equiv (P_{y,t})^{1-\alpha_m} \ (P_{m,t})^{\alpha_m}$$

Taking logs and using the definition of the terms of trade s_t

$$p_{y,t} = p_{q,t} - \frac{\alpha_m}{1-\alpha_m} \ s_t$$

Value added (or GDP), Y_t, is then defined by

$$\begin{aligned} P_{y,t} Y_t &\equiv P_{q,t} Q_t - P_{m,t} M_t \\ &= \left(1 - \frac{\alpha_m}{\mathcal{M}_t^p}\right) P_{q,t} Q_t \end{aligned}$$

which can be log linearized to yield

$$\begin{aligned} y_t &= q_t + \frac{\alpha_m}{1-\alpha_m} \ s_t + \eta \ \mu_t^p \\ &= \frac{1}{1-\alpha_m} \ (a_t + \alpha_n \ n_t) \end{aligned}$$

where the last equality uses the approximation invoked above.

Note that combining the above expressions for consumption and value added we can obtain the following relation between the two

$$c_t = y_t - \left(\frac{\alpha_m}{1-\alpha_m} + \chi\right) \ s_t$$

Price setting

Here we assume that firms set prices in a staggered fashion, as in Calvo (1983). Each period only a fraction $1-\theta$ of firms, selected randomly, reset prices. The remaining firms, with measure θ, keep their prices unchanged. The optimal price setting rule for a firm resetting prices in period t is given by

$$E_t \left\{ \sum_{k=0}^{\infty} \theta^k \ \Lambda_{t,t+k} \ Q_{t+k|t} \ \left(P_t^* - \mathcal{M}^p \ \Psi_{t+k|t} \right) \right\} = 0 \tag{26}$$

where P_t^* denotes the price newly set at time t, $Q_{t+k|t}$ and $\Psi_{t+k|t}$ are respectively the level of output and marginal cost in period $t+k$ for a firm that last set its price in period t, and $\mathcal{M}^p \equiv (\epsilon/(\epsilon-1)$ is the desired gross markup, which also corresponds to the gross markup in the zero inflation perfect foresight steady state. Log-linearization of (26) around a zero inflation steady state, combined with the (log-linearized) law of motion for the domestic price level yields the following equation for domestic inflation, $\pi_{q,t} \equiv p_{q,t} - p_{q,t-1}$:

$$\pi_{q,t} = \beta \ E_t\{\pi_{q,t+1}\} - \lambda_p \ \widehat{\mu}_t^p \tag{27}$$

where $\widehat{\mu}_t^p \equiv \mu_t^p - \mu^p$ denotes the (log) deviation of the average markup from its desired level, and $\lambda_p \equiv \frac{(1-\theta)(1-\beta\theta)}{\theta} \frac{1-\alpha_k}{1-\alpha_k+\alpha_k\epsilon}$.

Note that the (log) average price markup is given by

$$\begin{aligned} \mu_t^p &= p_{q,t} - (w_t - mpn_t) \\ &= q_t - n_t - (w_t - p_{c,t}) - \chi s_t \end{aligned}$$

Appendix 2. Computation of the Oil Share

We think of the U.S. economy as having two sectors, an oil-producing sector and a non-oil producing sector. We define the oil producing sector as the sum of the "oil and gas extraction" sector (NAIC code 211) and the "petroleum and coal" sector (NAIC code 324). ("petroleum refineries", a subsector of "petroleum and coal" is available only for benchmark years, the last available one being 1997. It represents 85% of the gross output of the "petroleum and coal" sector.) We define the non-oil producing sector as the rest of the economy.

To compute relevant numbers for 2005, we use data from the IO tables from the BEA site.

In 2005, "oil and gas extraction" output was \$227b, imports were \$223b, for a total of \$450b. Of this total, \$5b was for domestic final uses, \$440b was for intermediates, of which \$259 went to "Petroleum and coal", and \$181b went to the non-oil sector. Petroleum and coal output was \$402b, imports were \$65b, for a total of \$467b. Of this total, \$167 was for domestic final uses, \$279b for intermediates to the non-oil producing sector.

In 2005, total US value added was \$12,455b. Value added by "Oil and gas" was \$12b, value added by "Petroleum and coal" was \$12b, so value added in the non oil-producing sector was \$12,431b.

These numbers imply a value for χ of (181+279)/(12,431+181+279)= 3.5%, and an estimate of α is (5+167)/(12,431+181+279)= 1.3%.

The shares obviously depend very much on the price of oil. The same computation for the benchmark year of 1997 (which allows us to use "petroleum refining" rather than "Petroleum and coal" together) gives 1.7% and 1.2% respectively.

For the years 1973 and 1977, sectors are classified according to industry number codes. We construct the oil-producing sector as the sum of of "crude petroleum and natural gas" (1977 industry number 8) and "petroleum refining" (1977 industry number 31). The same steps as above yield $\chi = 2.3\%$ and $\alpha = 1.5\%$ in 1973, and $\chi = 3.6\%$ and $\alpha = 1.8\%$ in 1977.

REFERENCES

Barsky, Robert, and Lutz Kilian (2002): "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," *NBER Macroeconomics Annual 2001*, May, 137-183.

Bernanke, Ben, Mark Gertler, and Mark Watson, (1997) "Systematic Monetary Policy and the Effects of Oil Shocks," *Brookings Papers on Economic Activity*, 1997-1, 91-157.

Blanchard, Olivier, and John Simon (2001) "The Long and Large Decline in US Output Volatility," *Brookings Papers on Economic Activity*, 2001-1, 135-174.

Blanchard, Olivier, and Jordi Gali (2007) "Real wage rigidities and the New-Keynesian model", *Journal of Money, Credit, and Banking*, forthcoming

Bruno, Michael, and Jeffrey Sachs, *Economics of Worldwide Stagflation*, Harvard University Press: Cambridge Massachusetts, 1985.

Calvo, Guillermo (1983) "Staggered Contracts in a Utility-Maximizing Framework", *Journal of Monetary Economics*, September, 12, 383-398

Castelnuovo, Efrem, Sergio Nicoletti-Altimari, and Diego Rodriguez-Palenzuela (2003): "Definition of Price Stability, Range and Point Inflation Targets: The Anchoring of Long Term Inflation Expectations," in *Background Studies for the ECB's Evaluation of its Monetary Policy Strategy*,

Clarida, Richard, Jordi Gali, and Mark Gertler (2000) "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory", *Quarterly Journal of Economics*, 115-1, February, 147-180

De Gregorio, José, Oscar Landerretche, and Christopher Neilson (2007): "Another Passthrough Bites the Dust? Oil Prices and Inflation," mimeo.

Erceg, Christopher, Dale Henderson and Andrew Levin (2000), "Optimal monetary policy with staggered wage and price contracts", *Journal of Monetary Economics*, 46, 281-313

Finn, Mary G. (2000): "Perfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money Credit and Banking*, vol. 32, no. 3., 400-416.

Hamilton, James (1983); "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, April 1983, 228-248

Hamilton, James (1996), "This is What Happened to the Oil Price Macroeconomy Relationship," *Journal of Monetary Economics*, vol 38, no2, October, 215-220.

Hooker, Mark A. (1996) "What Happened to the Oil Price Macroeconomy Relationship?" *Journal of Monetary Economics*, 38, 195-213.

Hooker, Mark A. (2002): "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime," *Journal of Money, Credit and Banking*, vol. 34, no. 2, 540-561.

Kilian, Lutz (2006), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," CEPR Discussion Paper 5994.

Levin, Andrew T., Fabio M. Natalucci, and Jeremy M. Piger (2004): "The Macroeconomic Effects of Inflation Targeting," Federal Reserve Bank of St. Louis Review, vol. 86, no. 4, 51-80.

Mankiw, N. Gregory, *Macroeconomics*, 6th edition, Worth Publishers, 2006.

McConnell, Margaret, and Gabriel Perez-Quiros (2000) "Output Fluctuations in the United States; What Has Changed Since the Early 1980s?", *American Economic Review*, 90-5, December, 1464-1476.

Rotemberg, Julio and Michael Woodford (1996): "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal* *of Money, Credit and Banking*, vol. 28, no. 4, 549-577.

Stock, James, and Mark Watson (2003) "Has the Business Cycle Changed and Why?," *Federal Reserve Bank of Kansas City*, 9-56.

Figure 1. Oil Price (West Texas Int., $ per barrel)

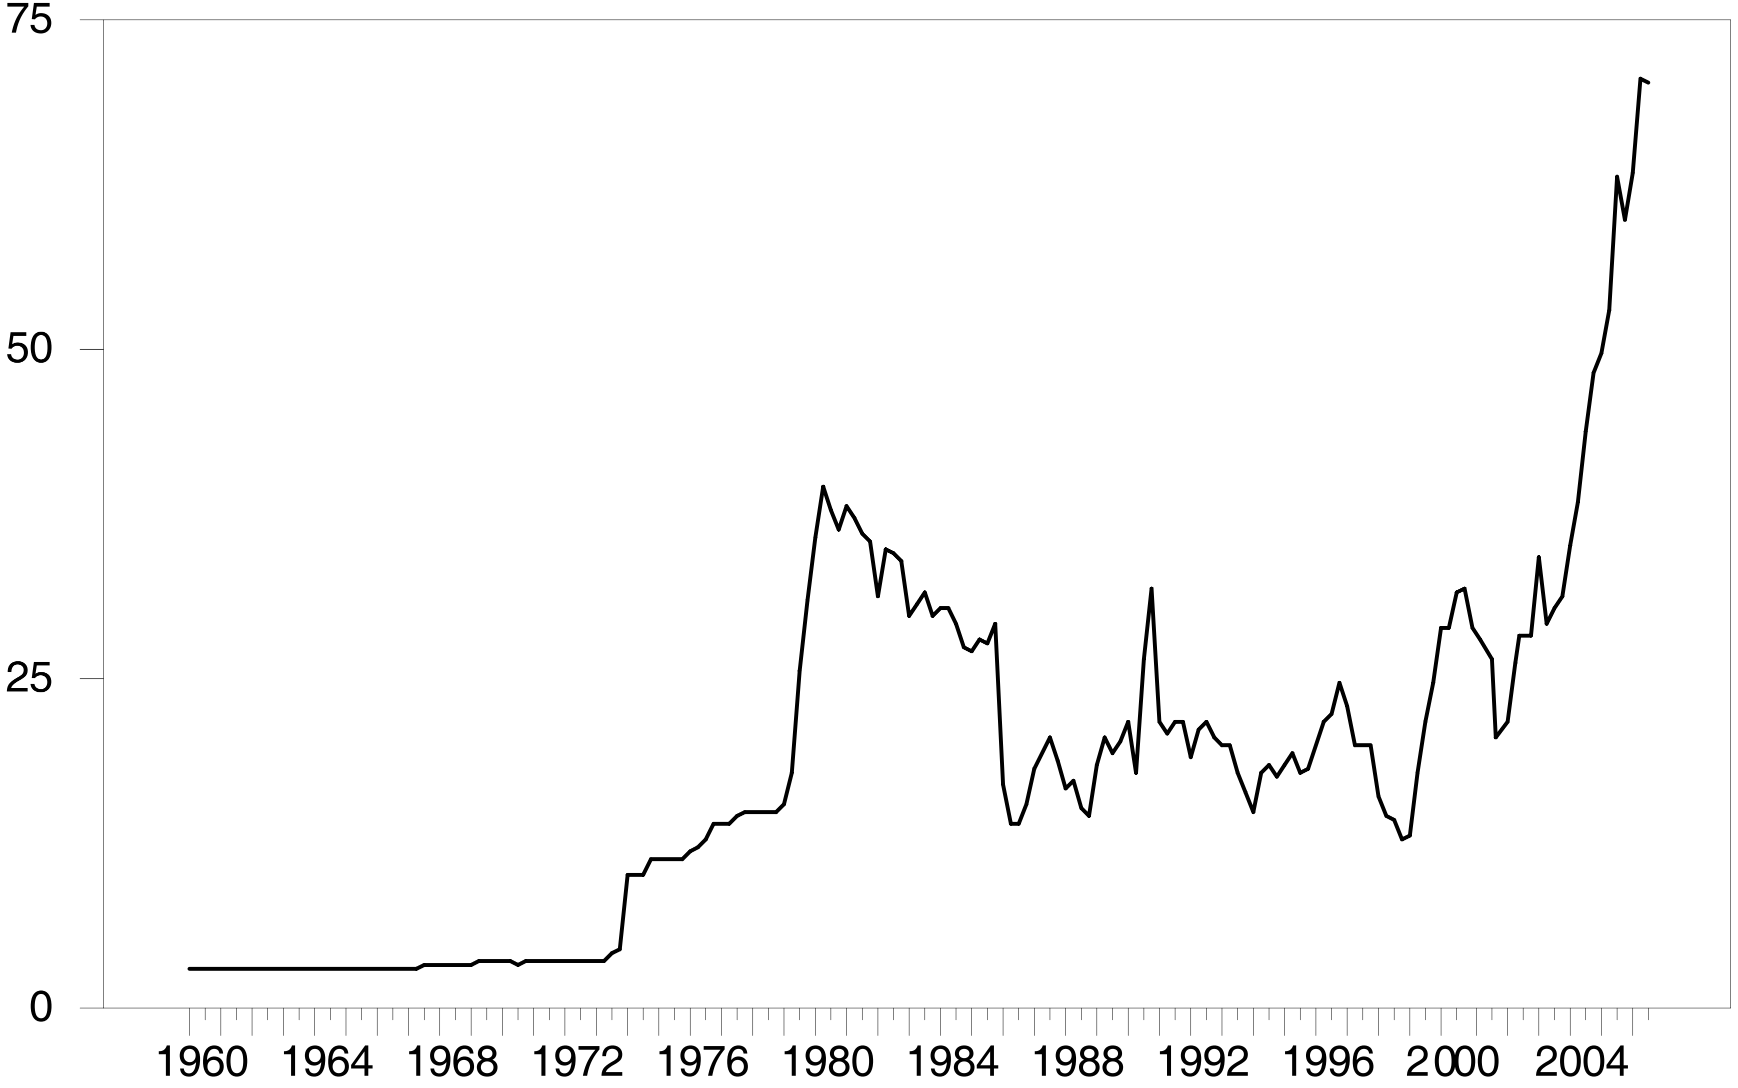

Figure 2. Log Real Oil Price (1960=100)

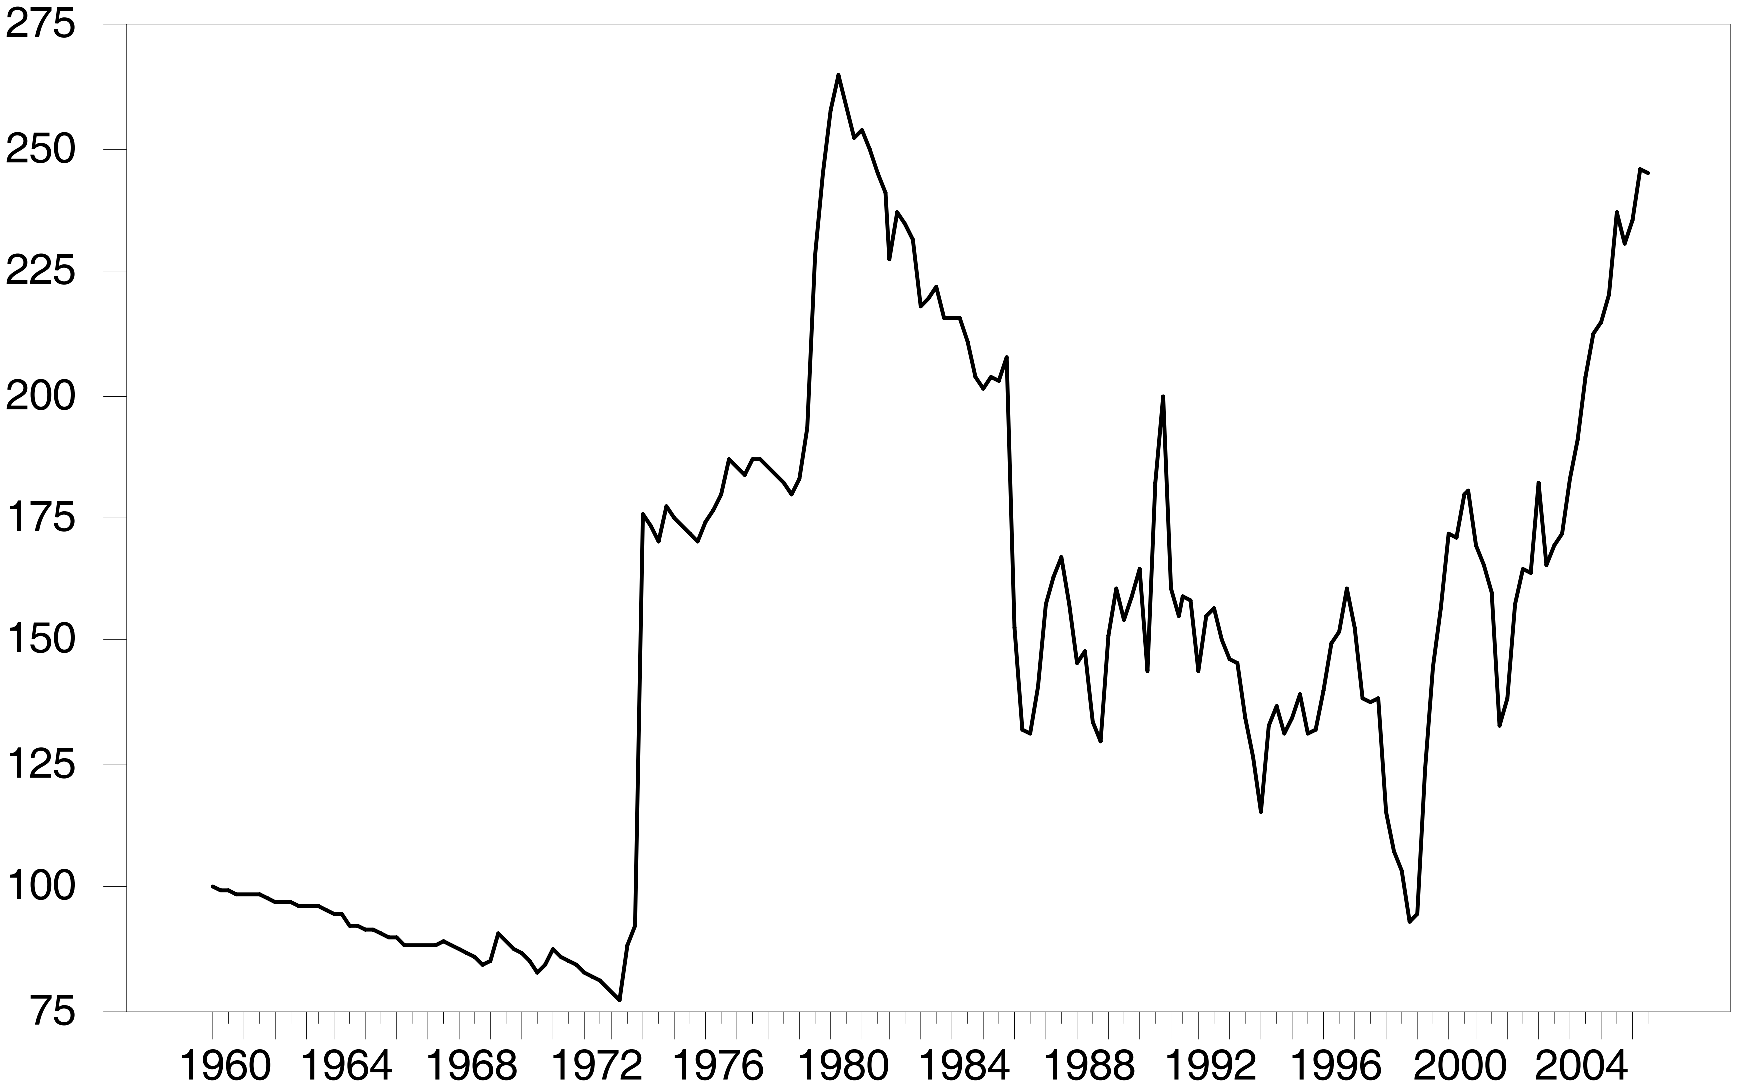

Figure 3: Oil Shocks and CPI Inflation

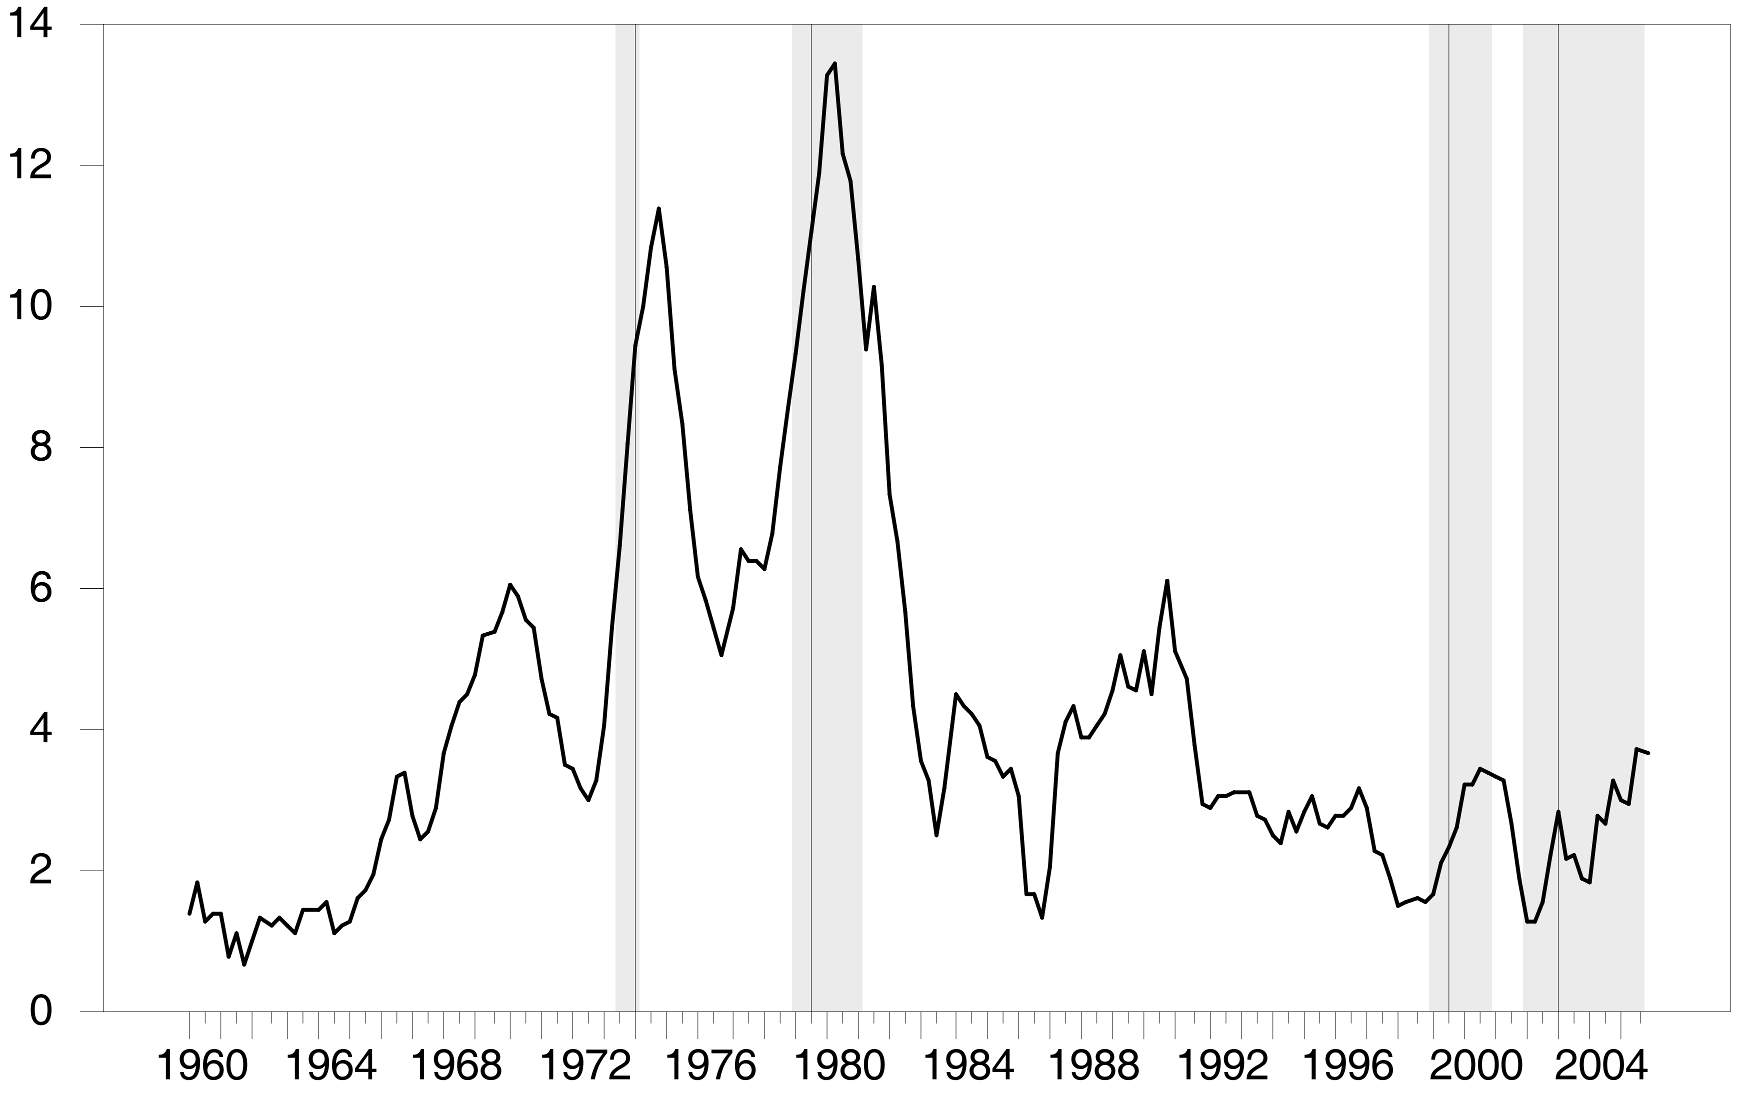

Figure 4: Oil Shocks and the Unemployment Rate

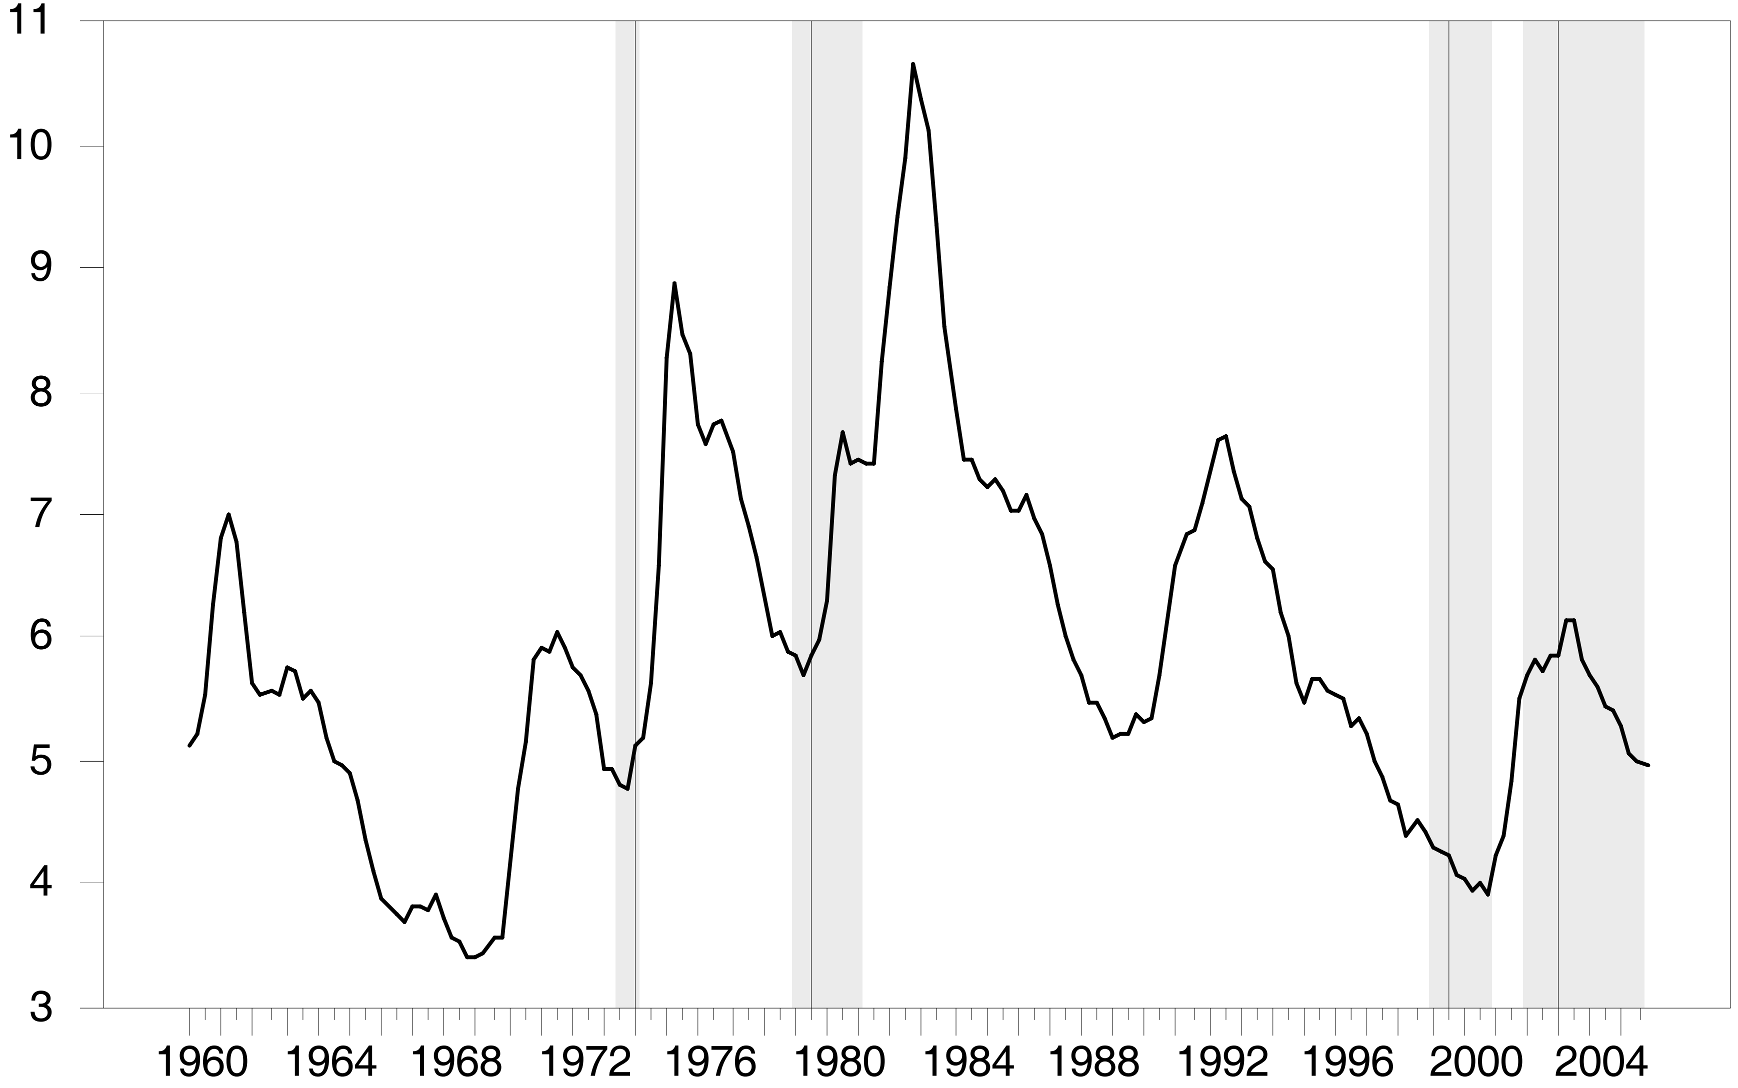

Figure 5a

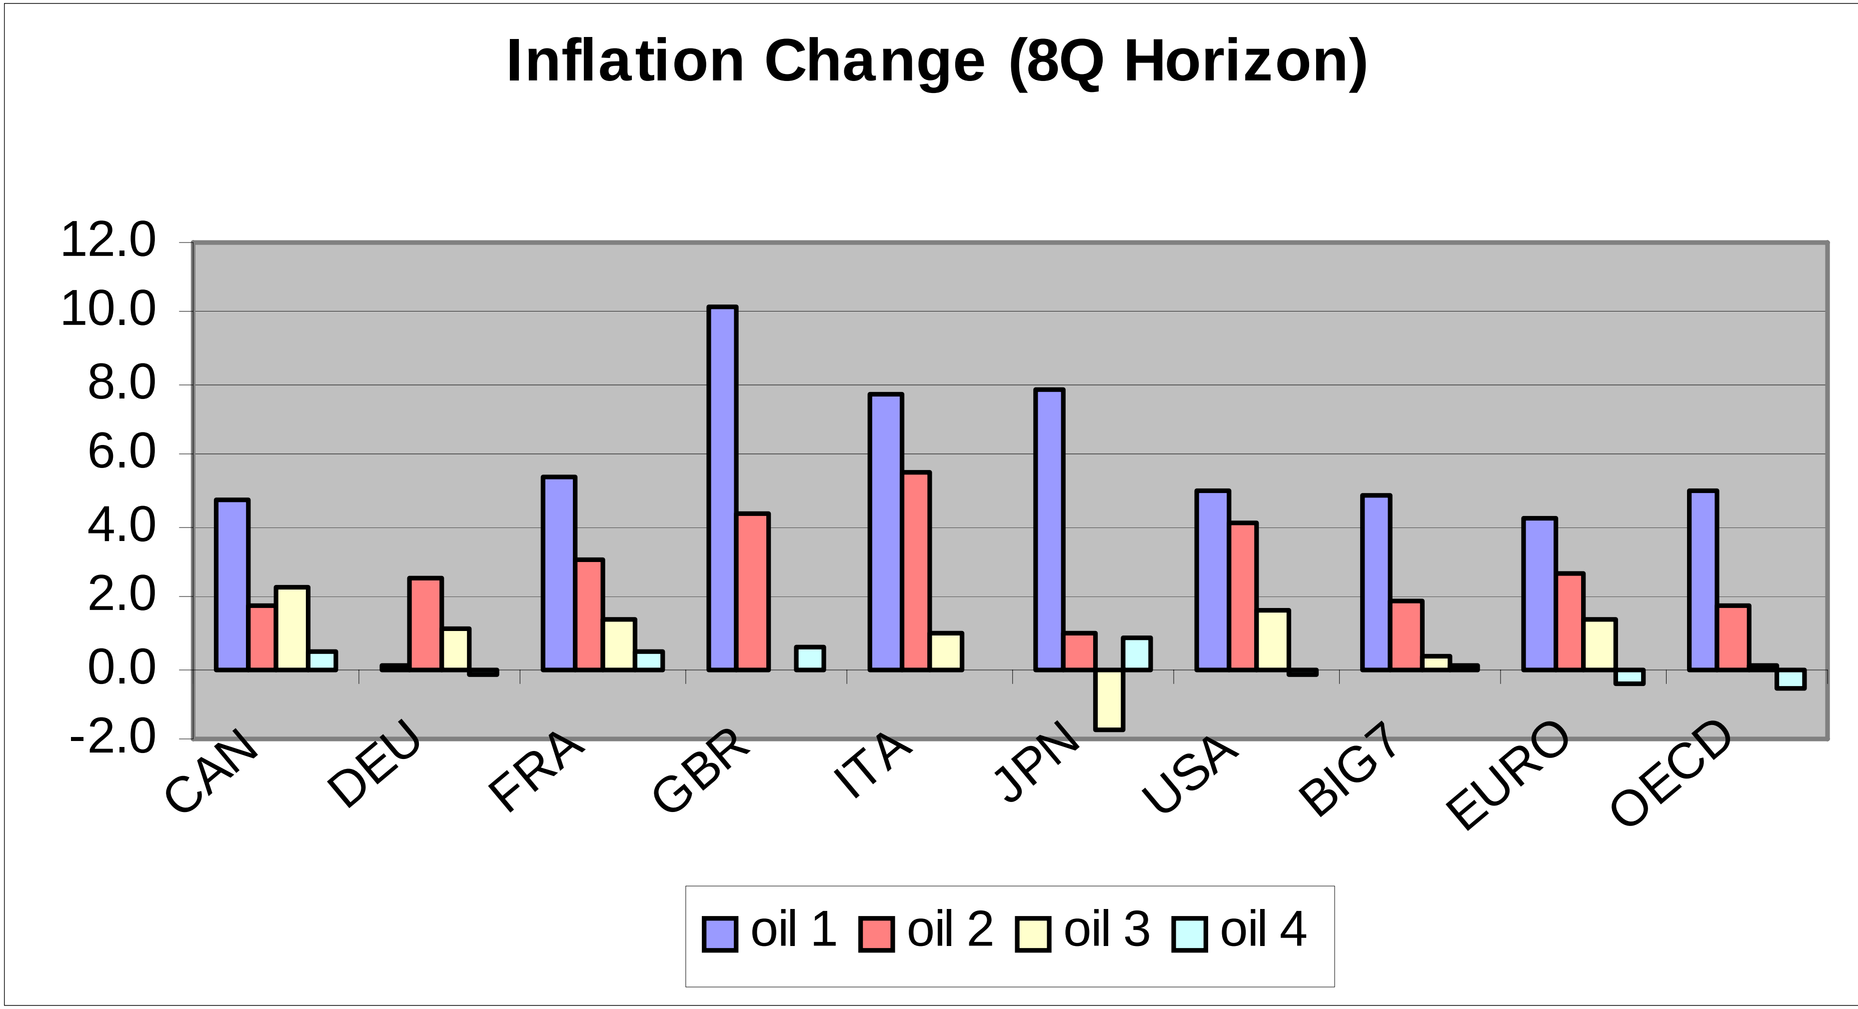

Inflation Change (8Q Horizon)
12.0
10.0
8.0
6.0
4.0
2.0
0.0
-2.0
CAN
DEU
FRA
GBR
ITA
JPN
USA
BIG7
EURO
OECD
oil 1
oil 2
oil 3
oil 4

Figure 5b

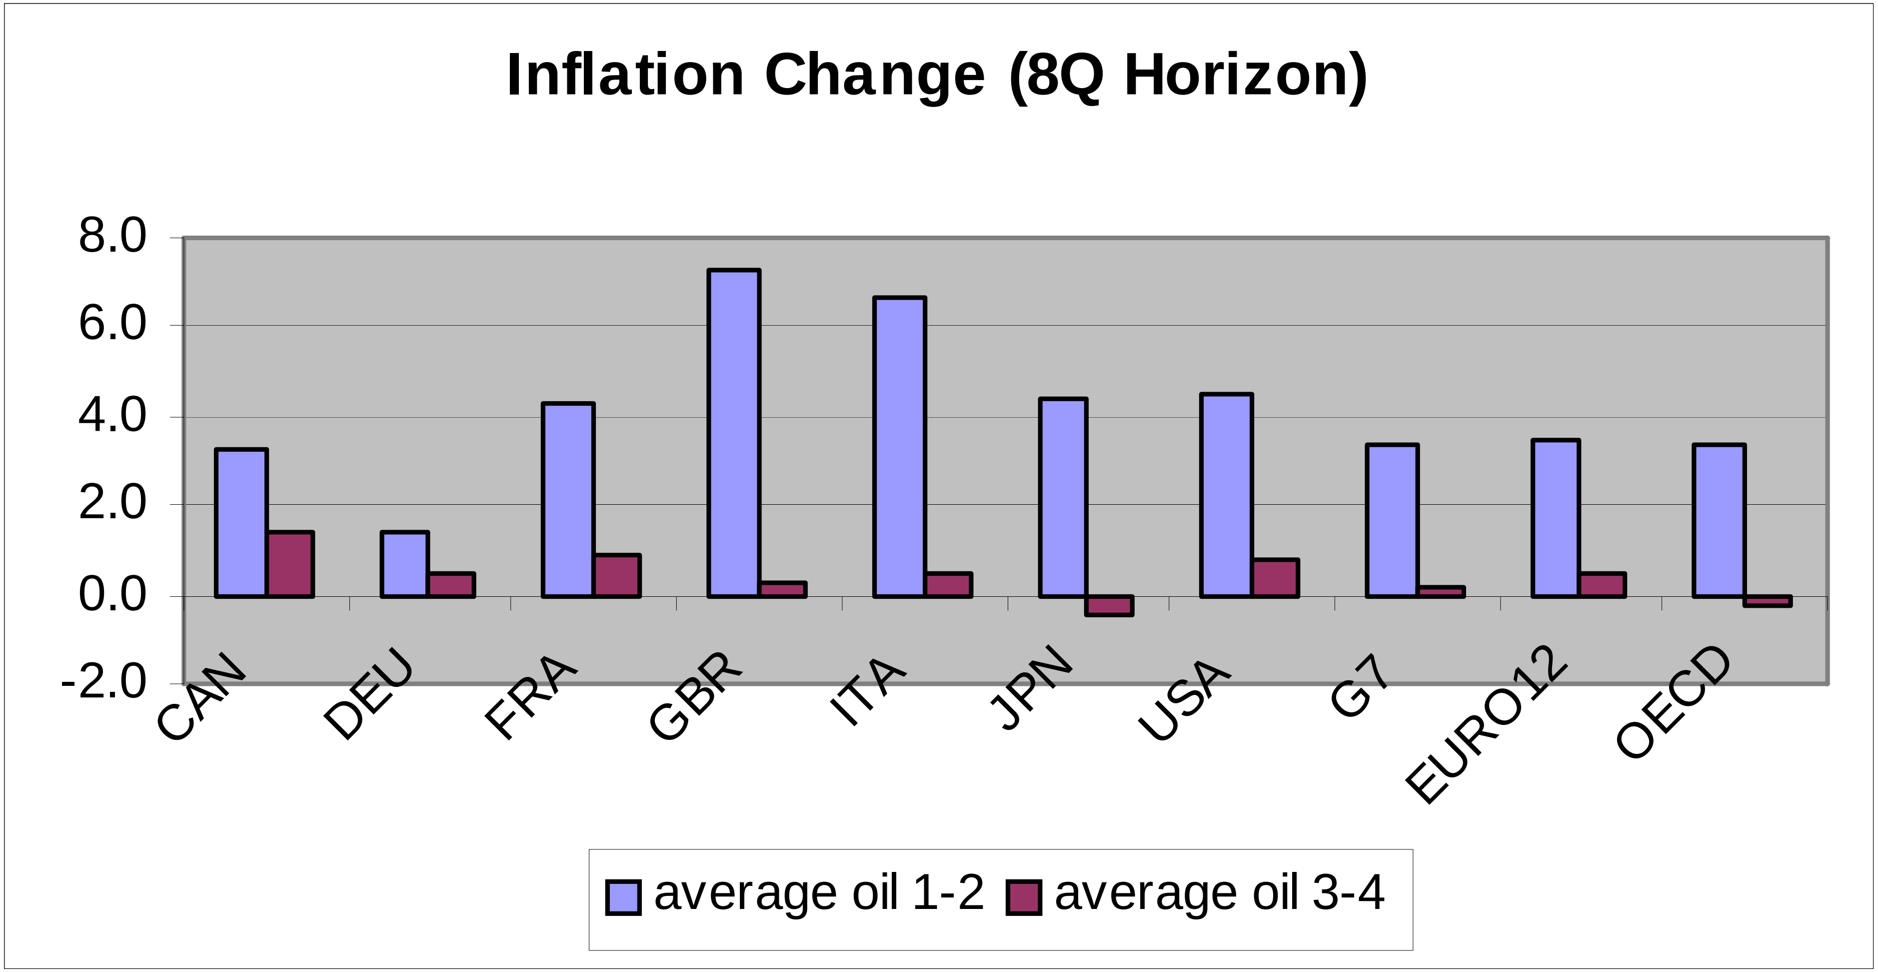

Inflation Change (8Q Horizon)
8.0
6.0
4.0
2.0
0.0
-2.0
CAN
DEU
FRA
GBR
ITA
JPN
USA
G7
EURO12
OECD
average oil 1-2
average oil 3-4

Figure 6a

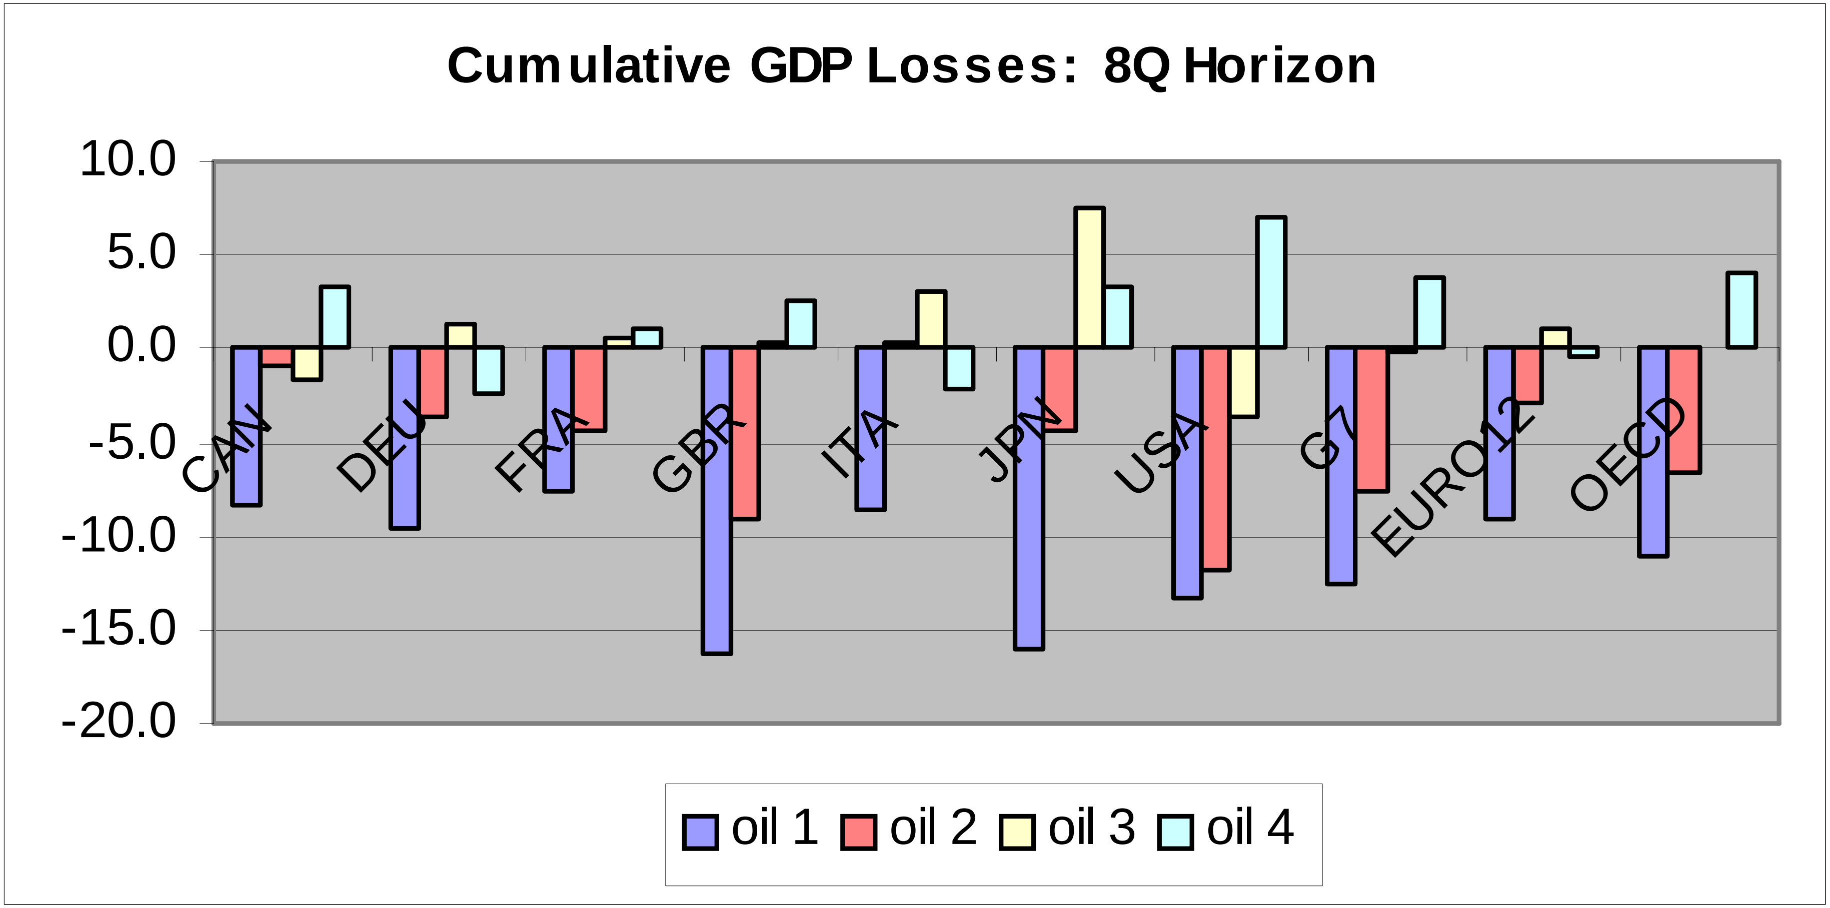

Cumulative GDP Losses: 8Q Horizon
10.0
5.0
0.0
-5.0
-10.0
-15.0
-20.0
CAN
DEU
FRA
GBR
ITA
JPN
USA
G7
EURO12
OECD
oil 1
oil 2
oil 3
oil 4

Figure 6b

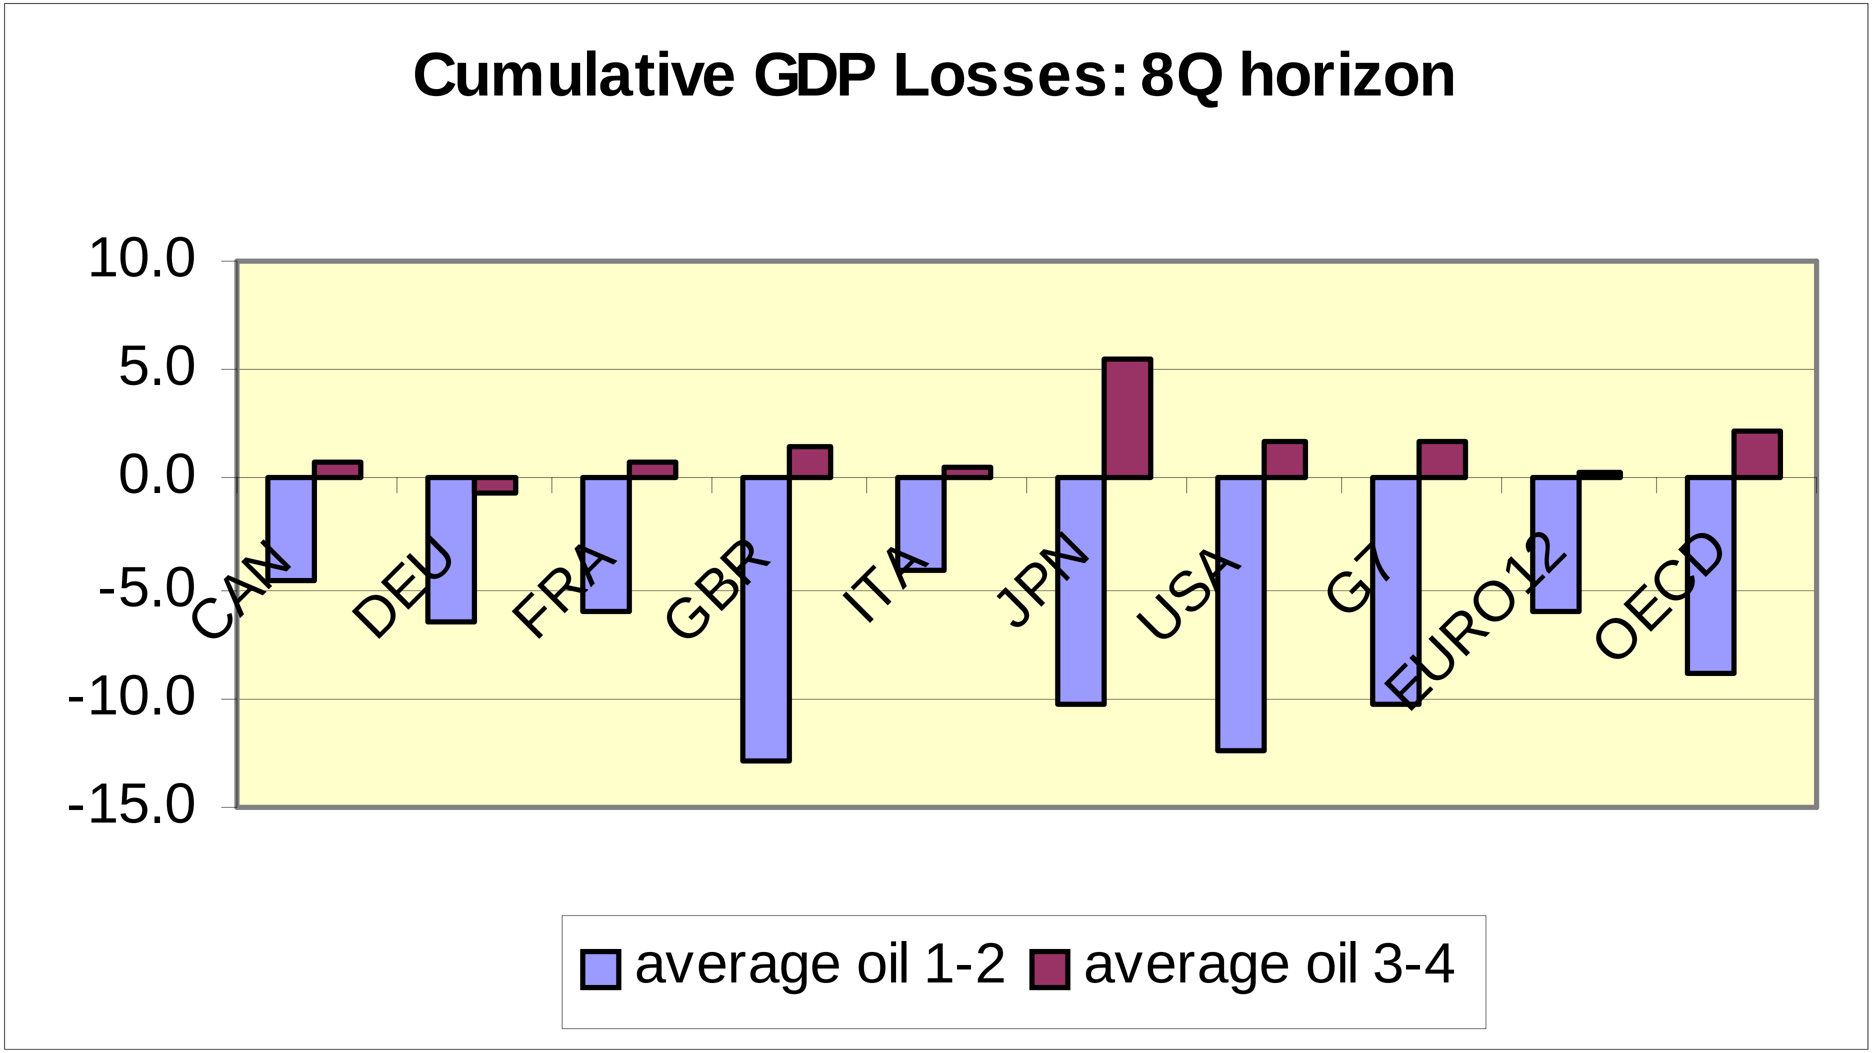
Cumulative GDP Losses: 8Q horizon
10.0
5.0
0.0
-5.0
-10.0
-15.0
CAN
DEU
FRA
GBR
ITA
JPN
USA
G7
EURO12
OECD
average oil 1-2
average oil 3-4

Figure 7. Oil Price Shocks and Inflation–GDP Tradeoffs

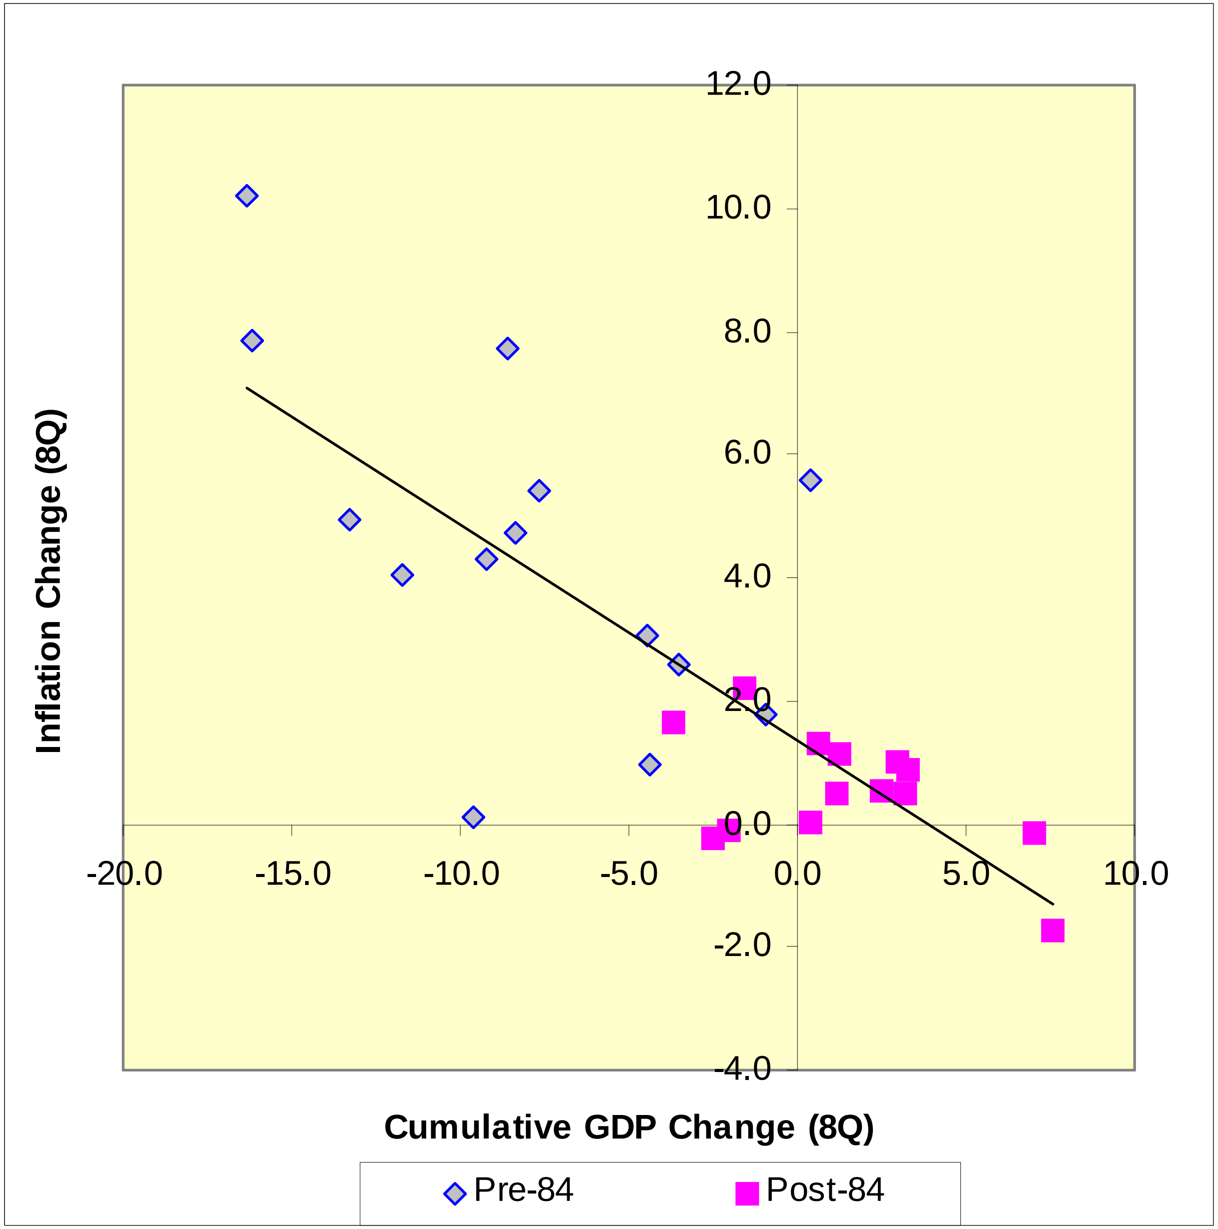

Figure 8a

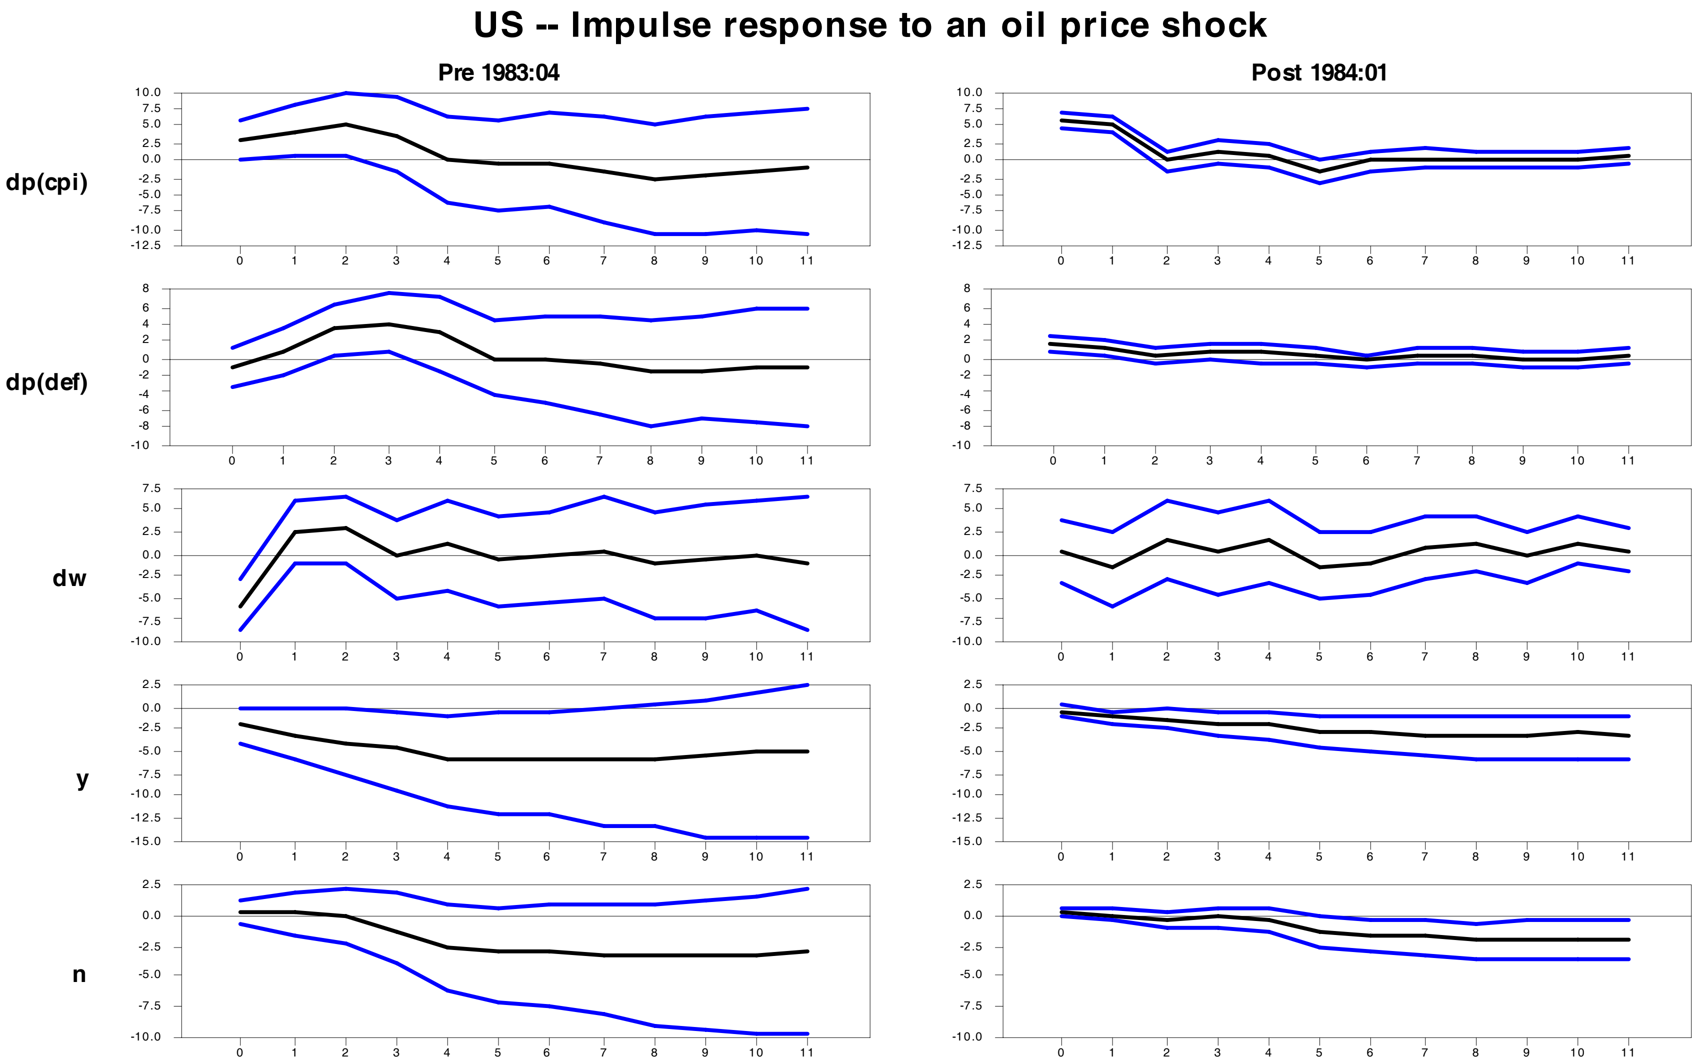
US -- Impulse response to an oil price shock
Pre 1983:04
Post 1984:01
dp(cpi)
dp(def)
dw
y
n

Figure 8b

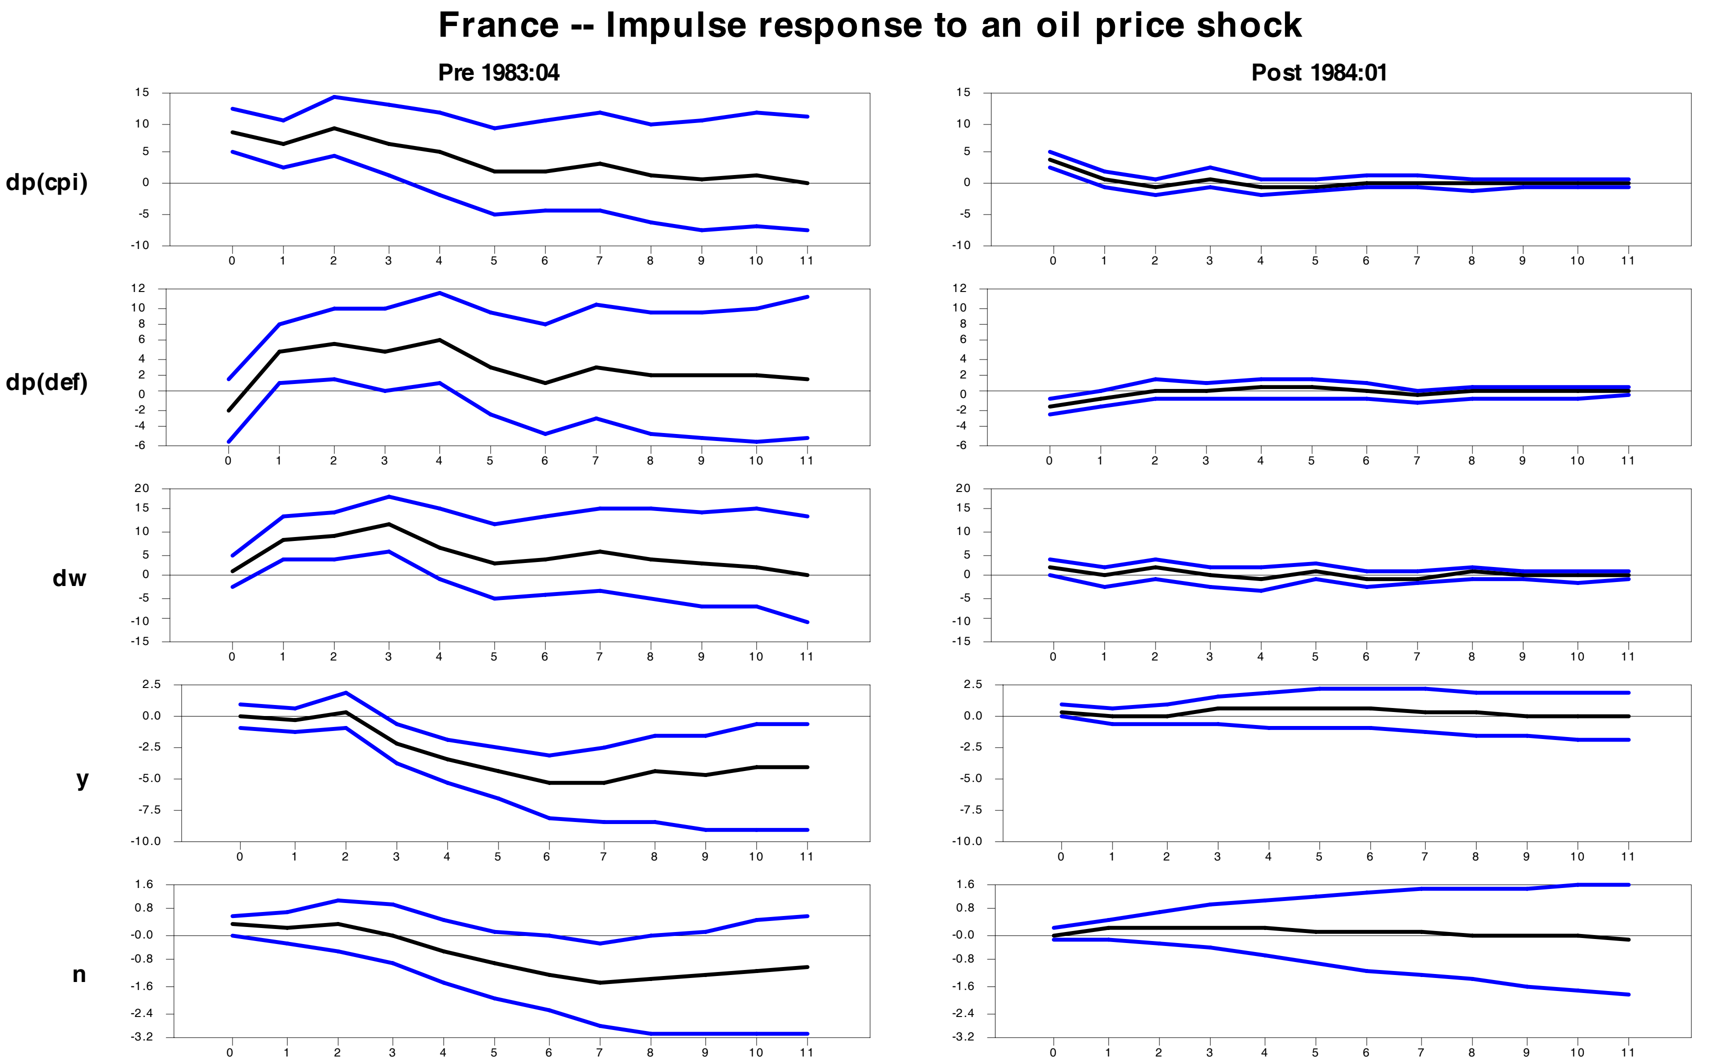
France -- Impulse response to an oil price shock
Pre 1983:04
Post 1984:01
dp(cpi)
dp(def)
dw
y
n

Figure 8c

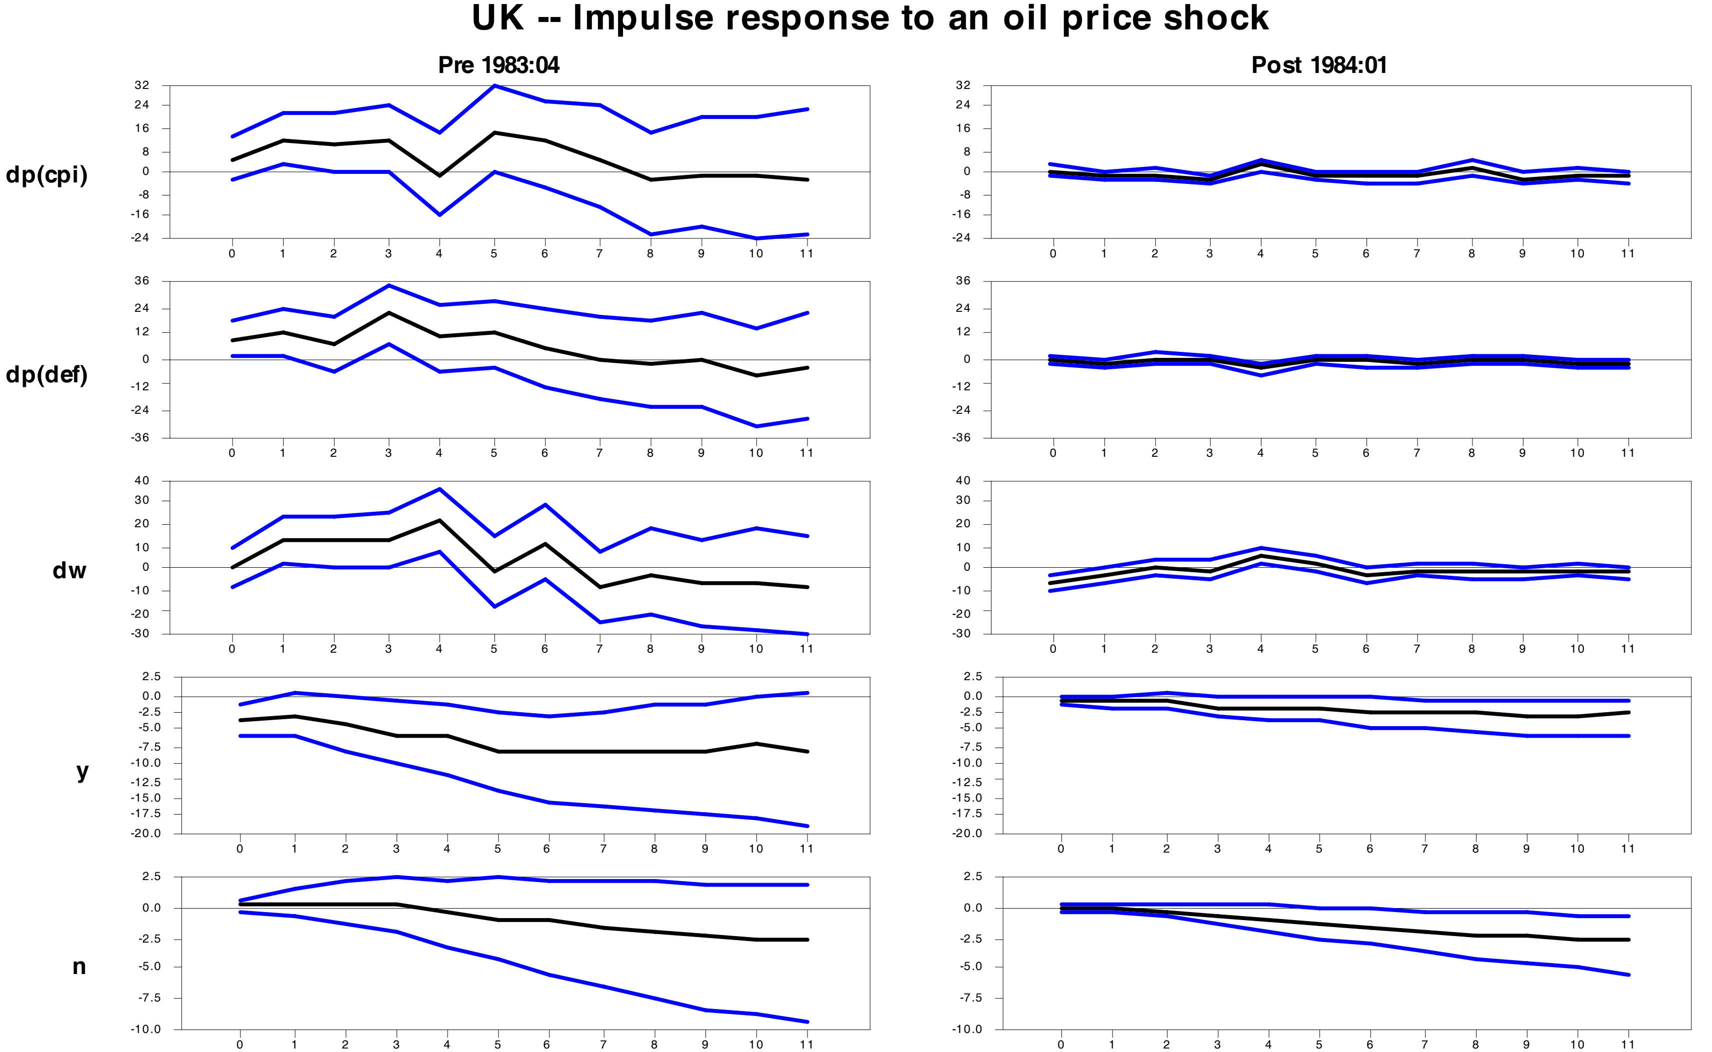
UK -- Impulse response to an oil price shock
Pre 1983:04
Post 1984:01
dp(cpi)
dp(def)
dw
y
n

Figure 8d

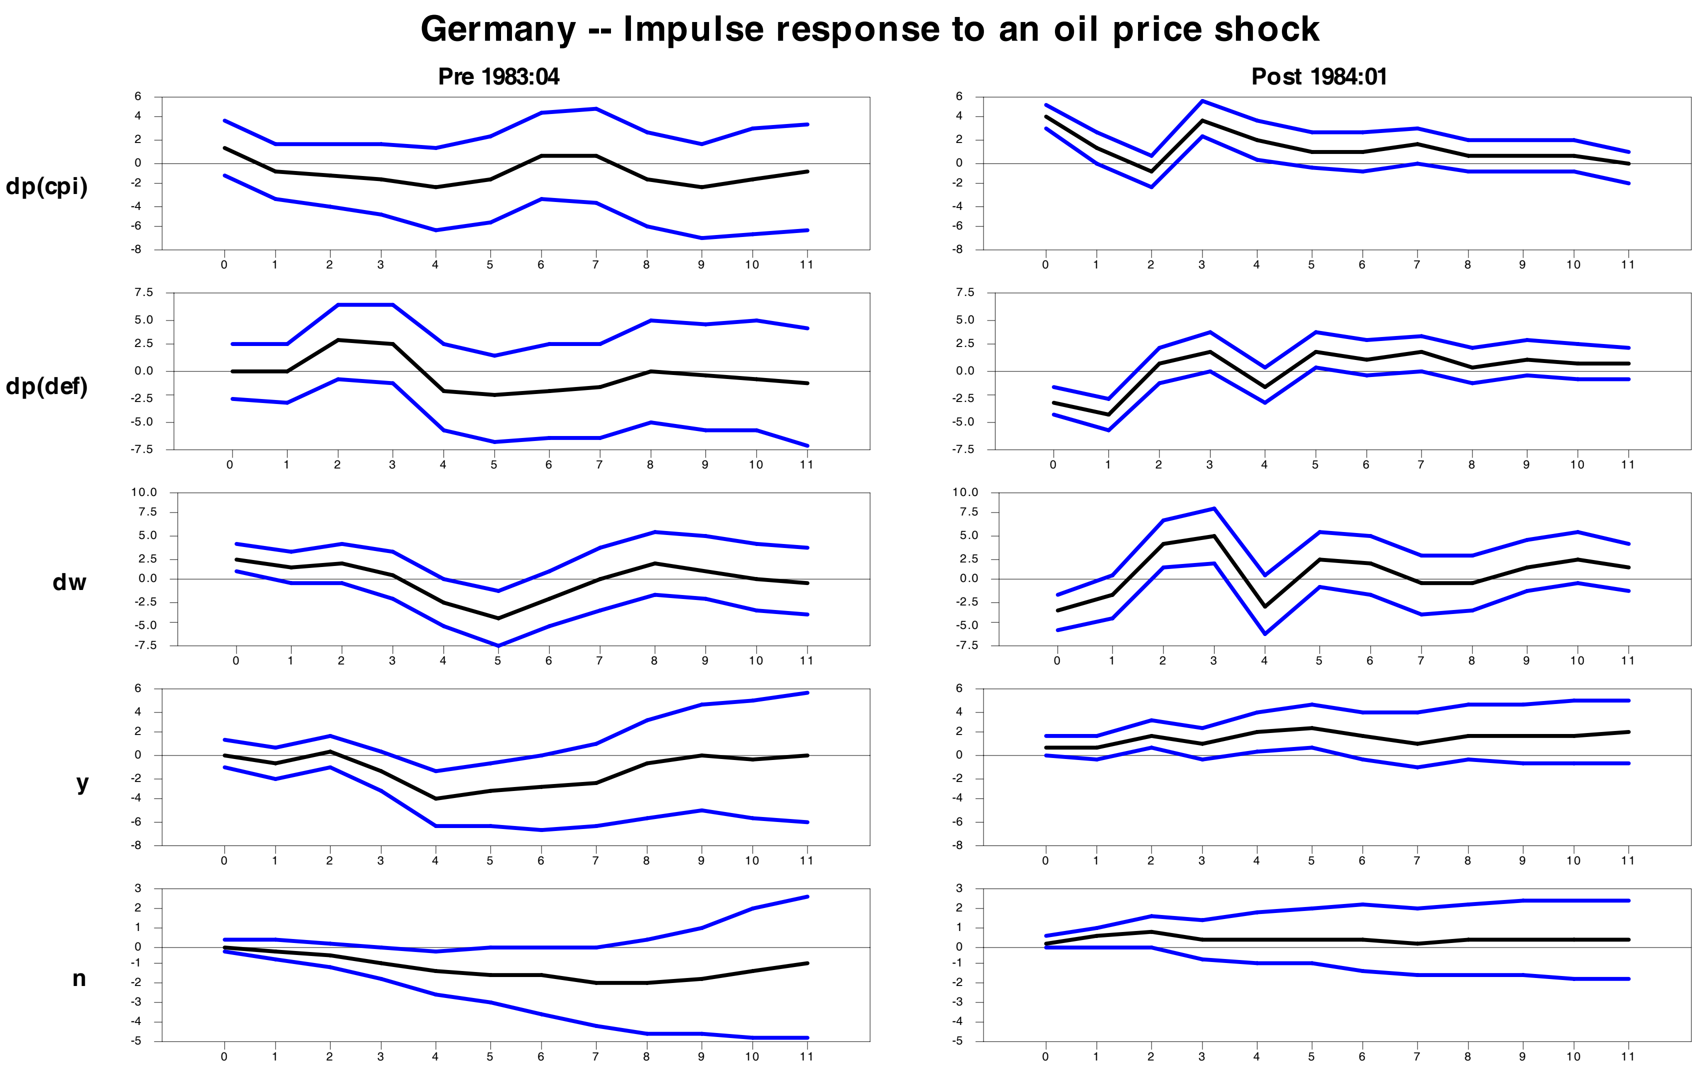
Germany -- Impulse response to an oil price shock
Pre 1983:04
Post 1984:01
dp(cpi)
dp(def)
dw
y
n

Figure 8e

Italy -- Impulse response to an oil price shock

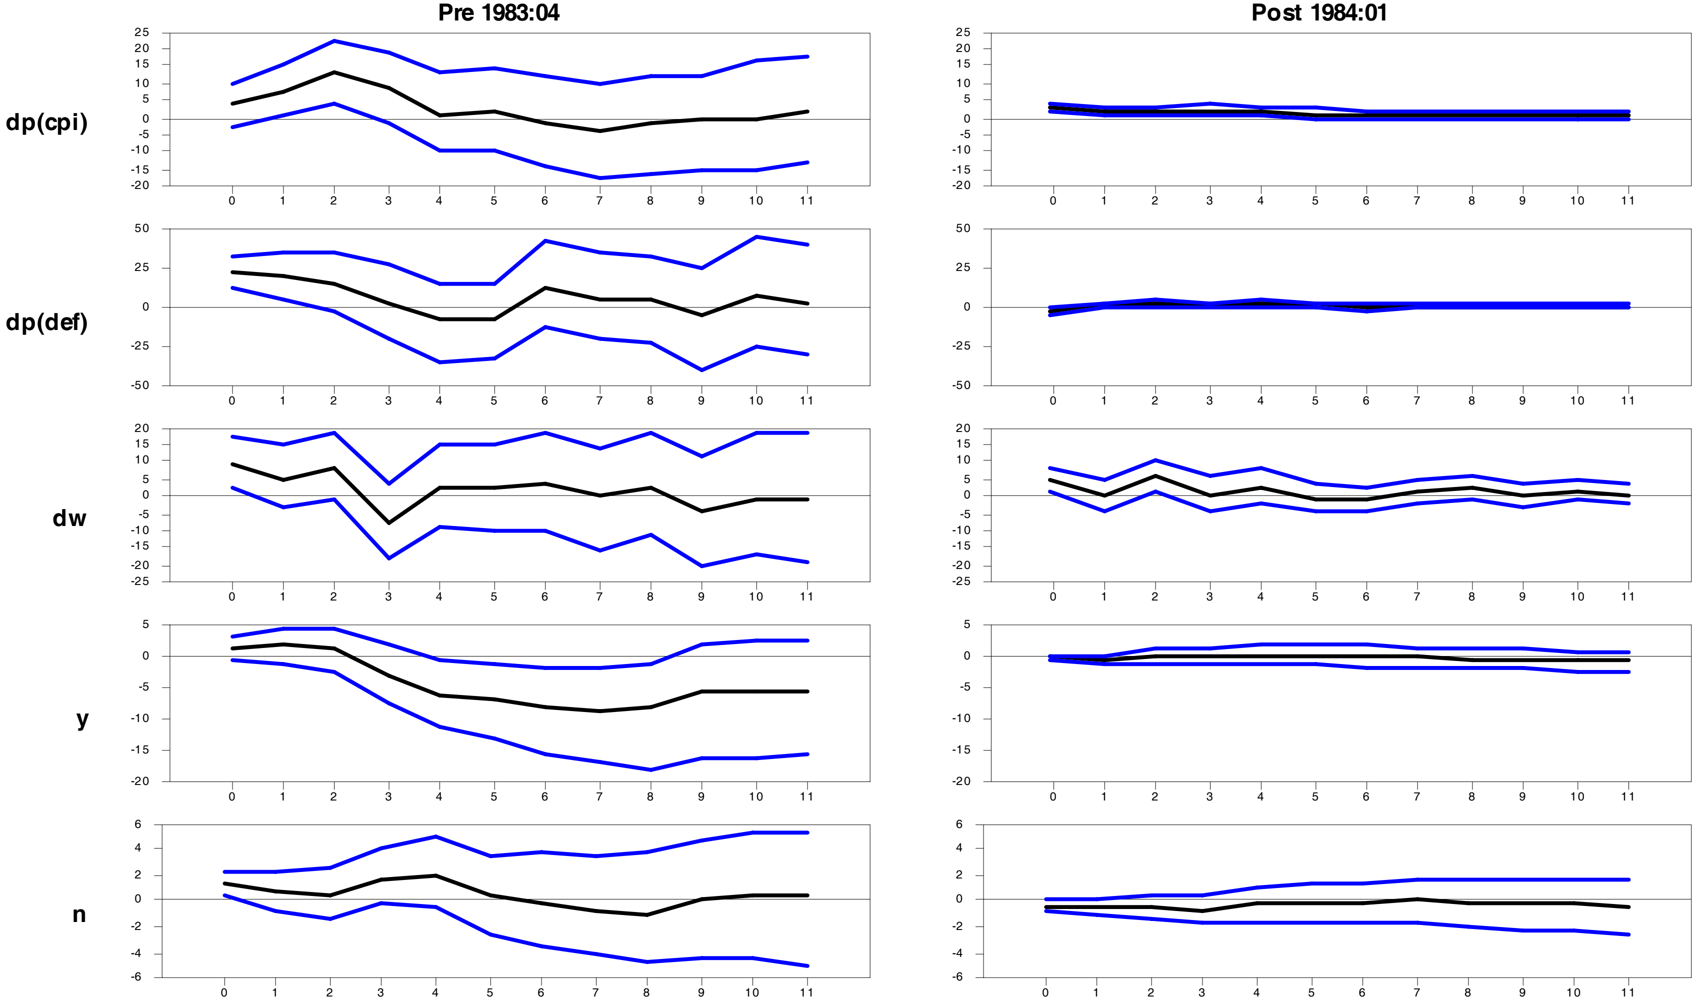

Figure 8g

Japan -- Impulse response to an oil price shock

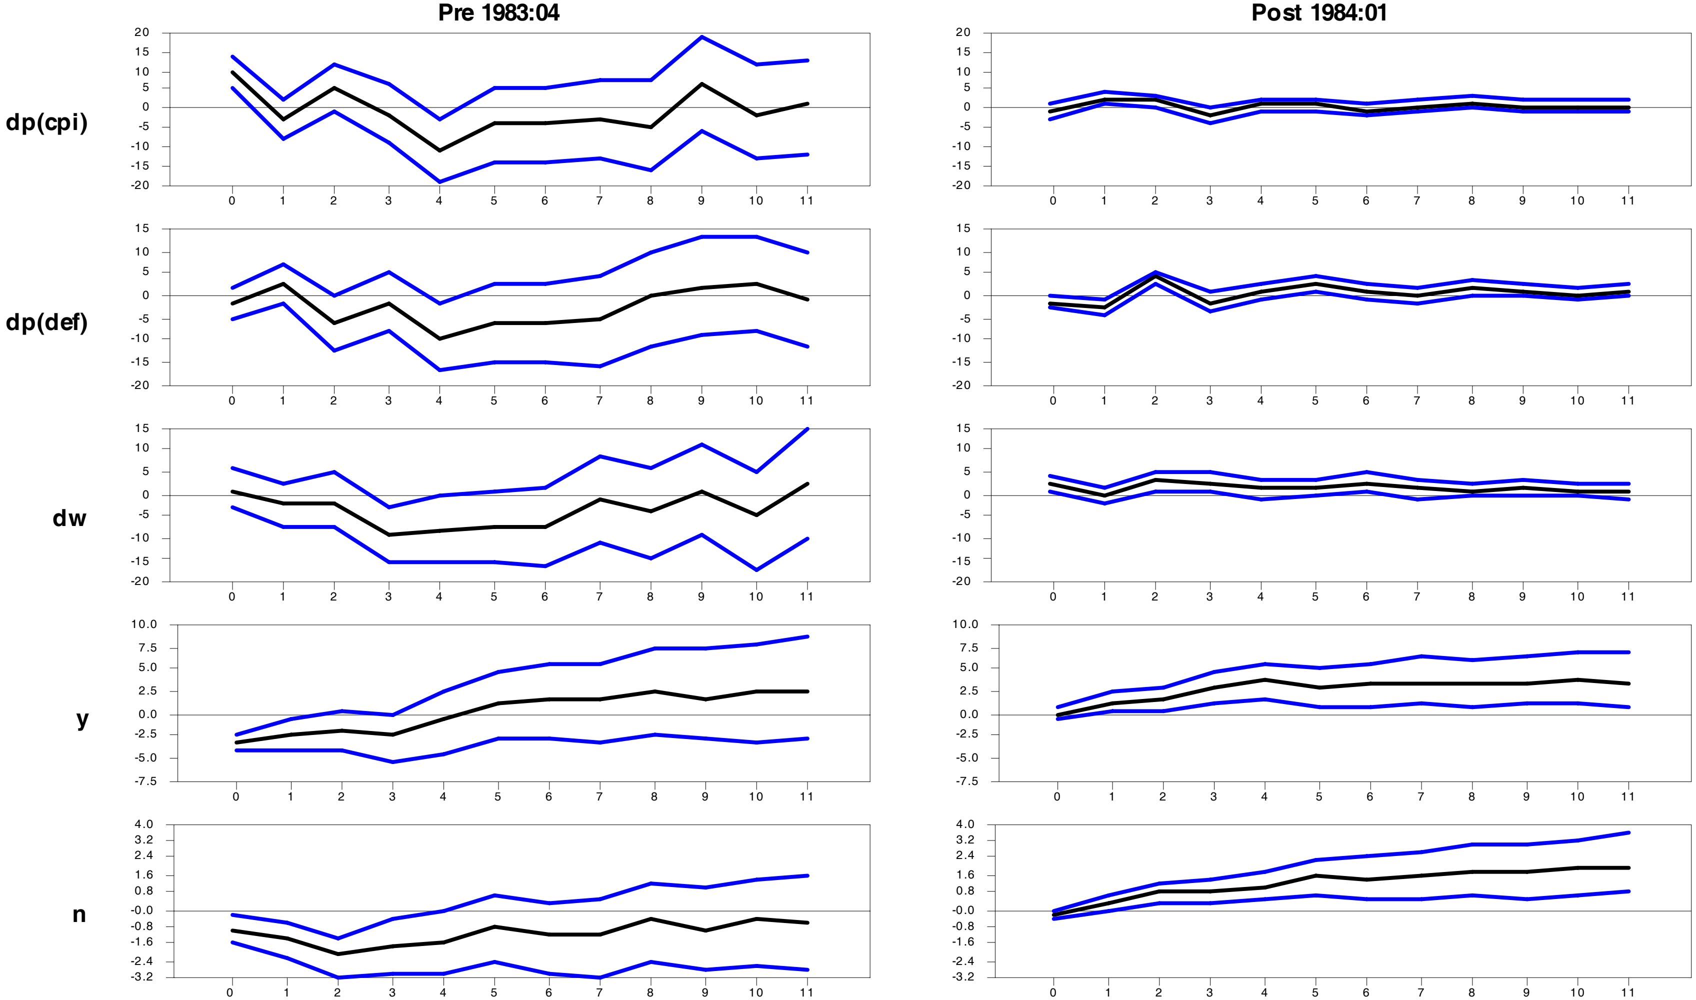

Figure 9

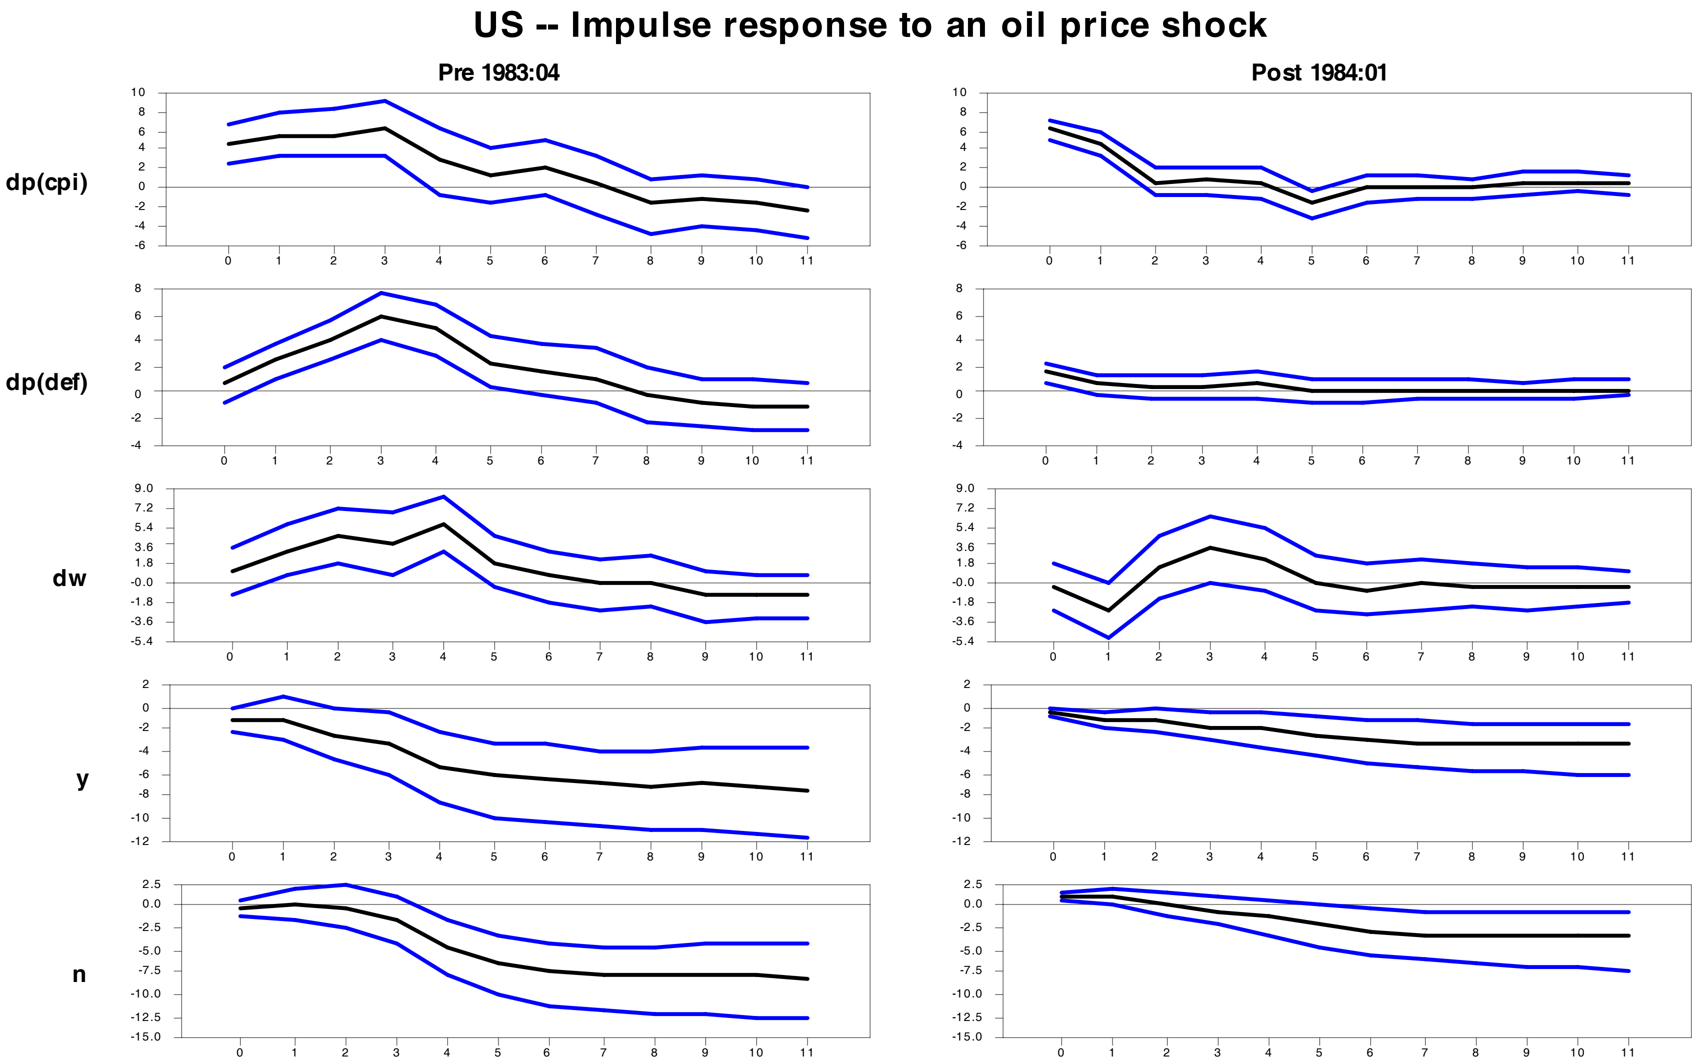

US -- Impulse response to an oil price shock
Pre 1983:04
Post 1984:01
dp(cpi)
dp(def)
dw
y
n

Figure 10

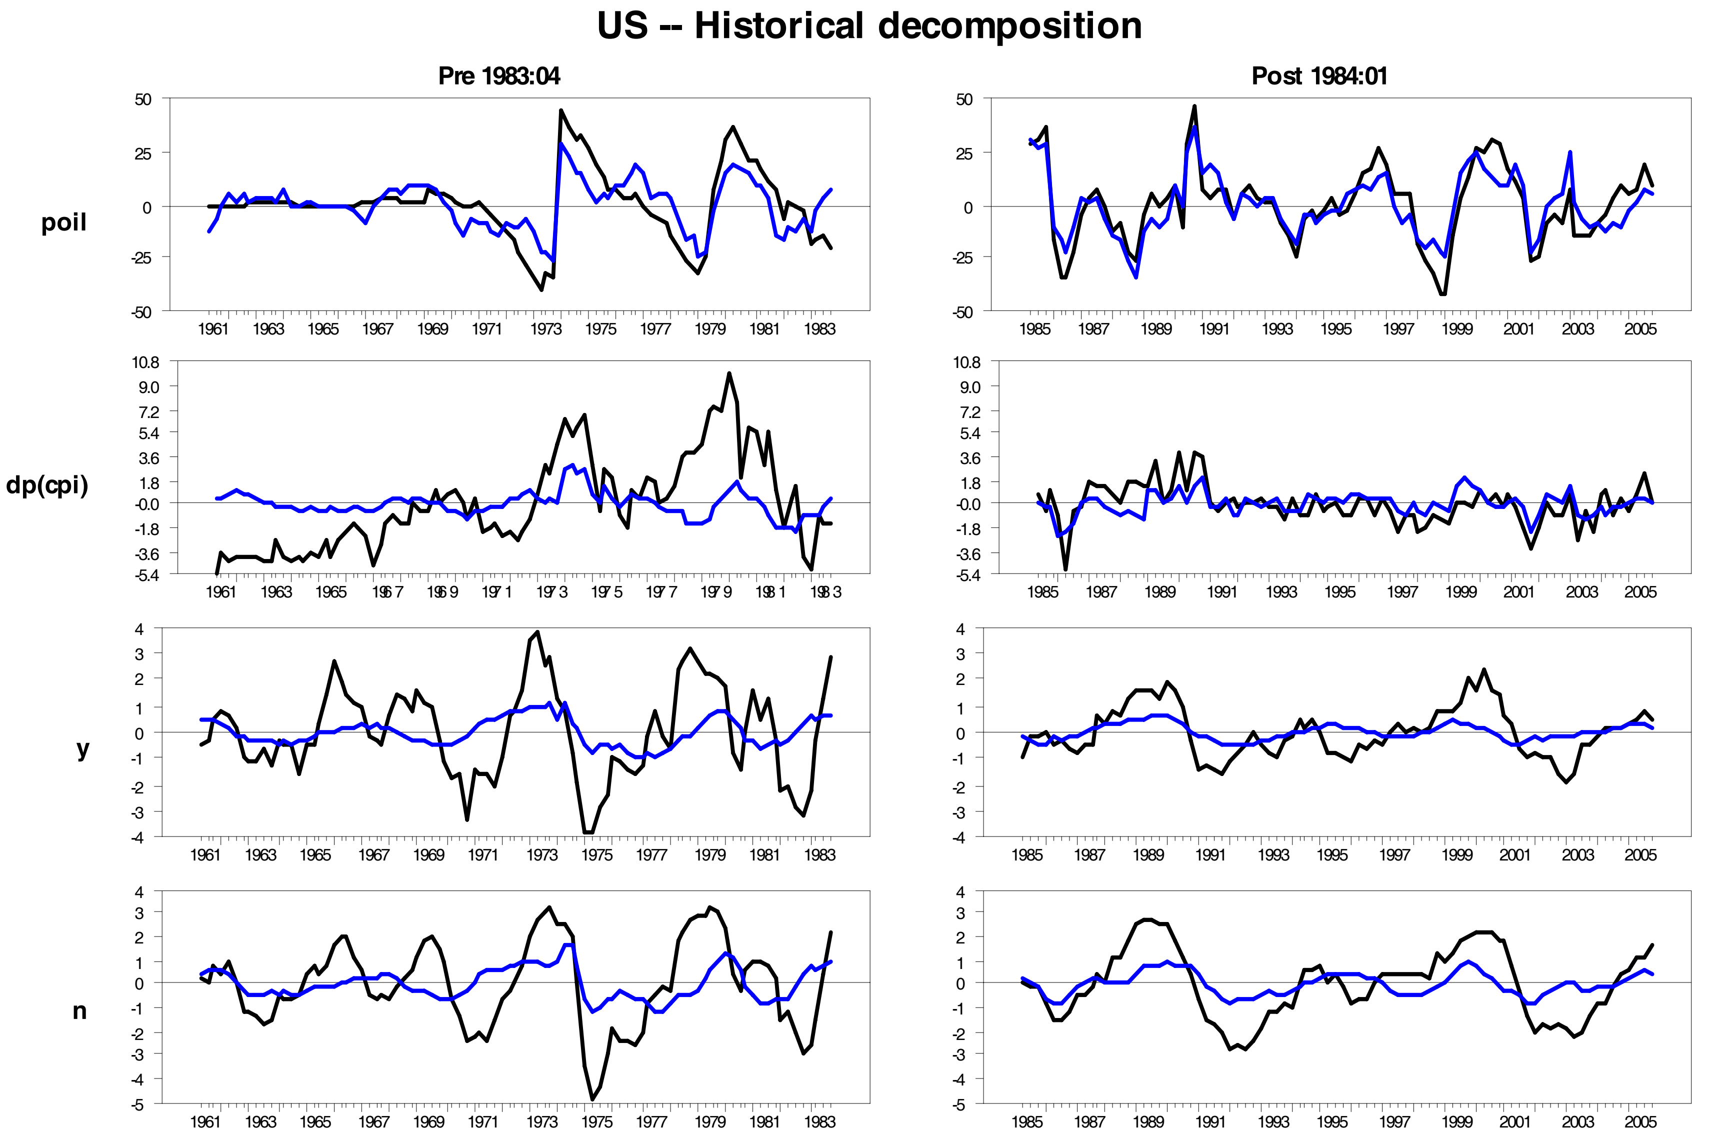
US -- Historical decomposition
Pre 1983:04
Post 1984:01
poil
dp(cpi)
y
n

Figure 11

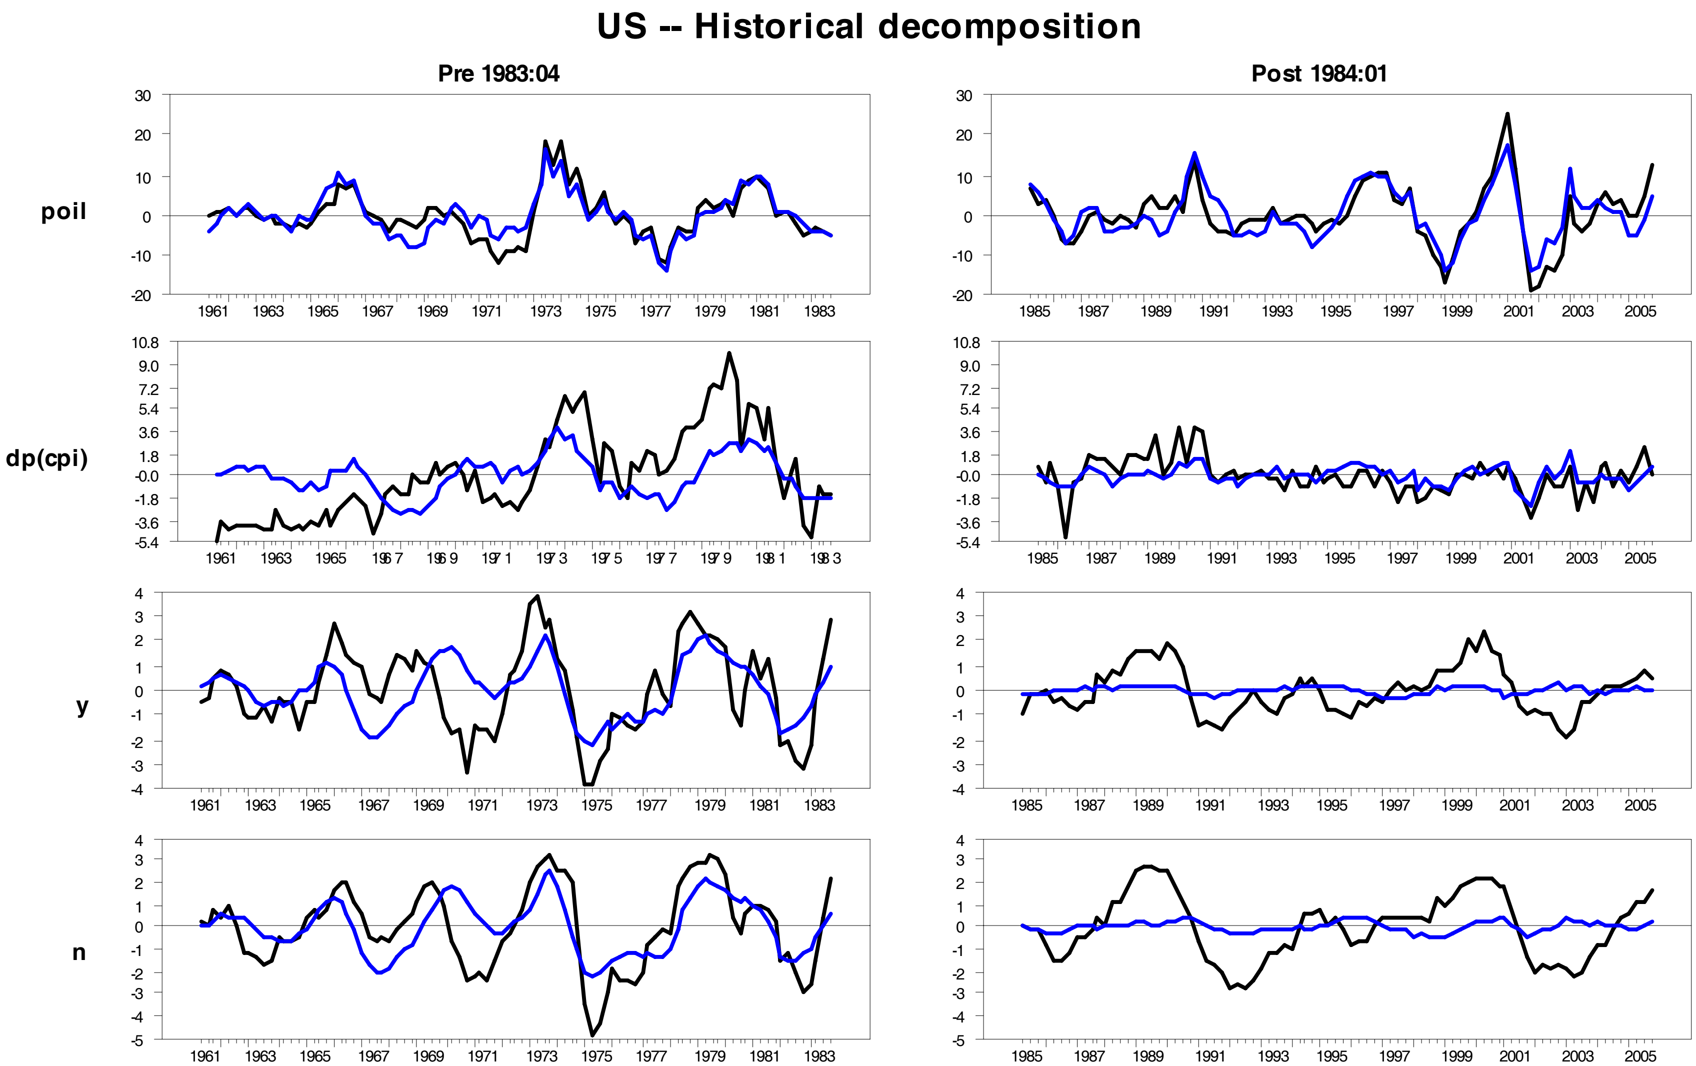
US -- Historical decomposition
Pre 1983:04
Post 1984:01
poil
dp(cpi)
y
n

Figure 12a

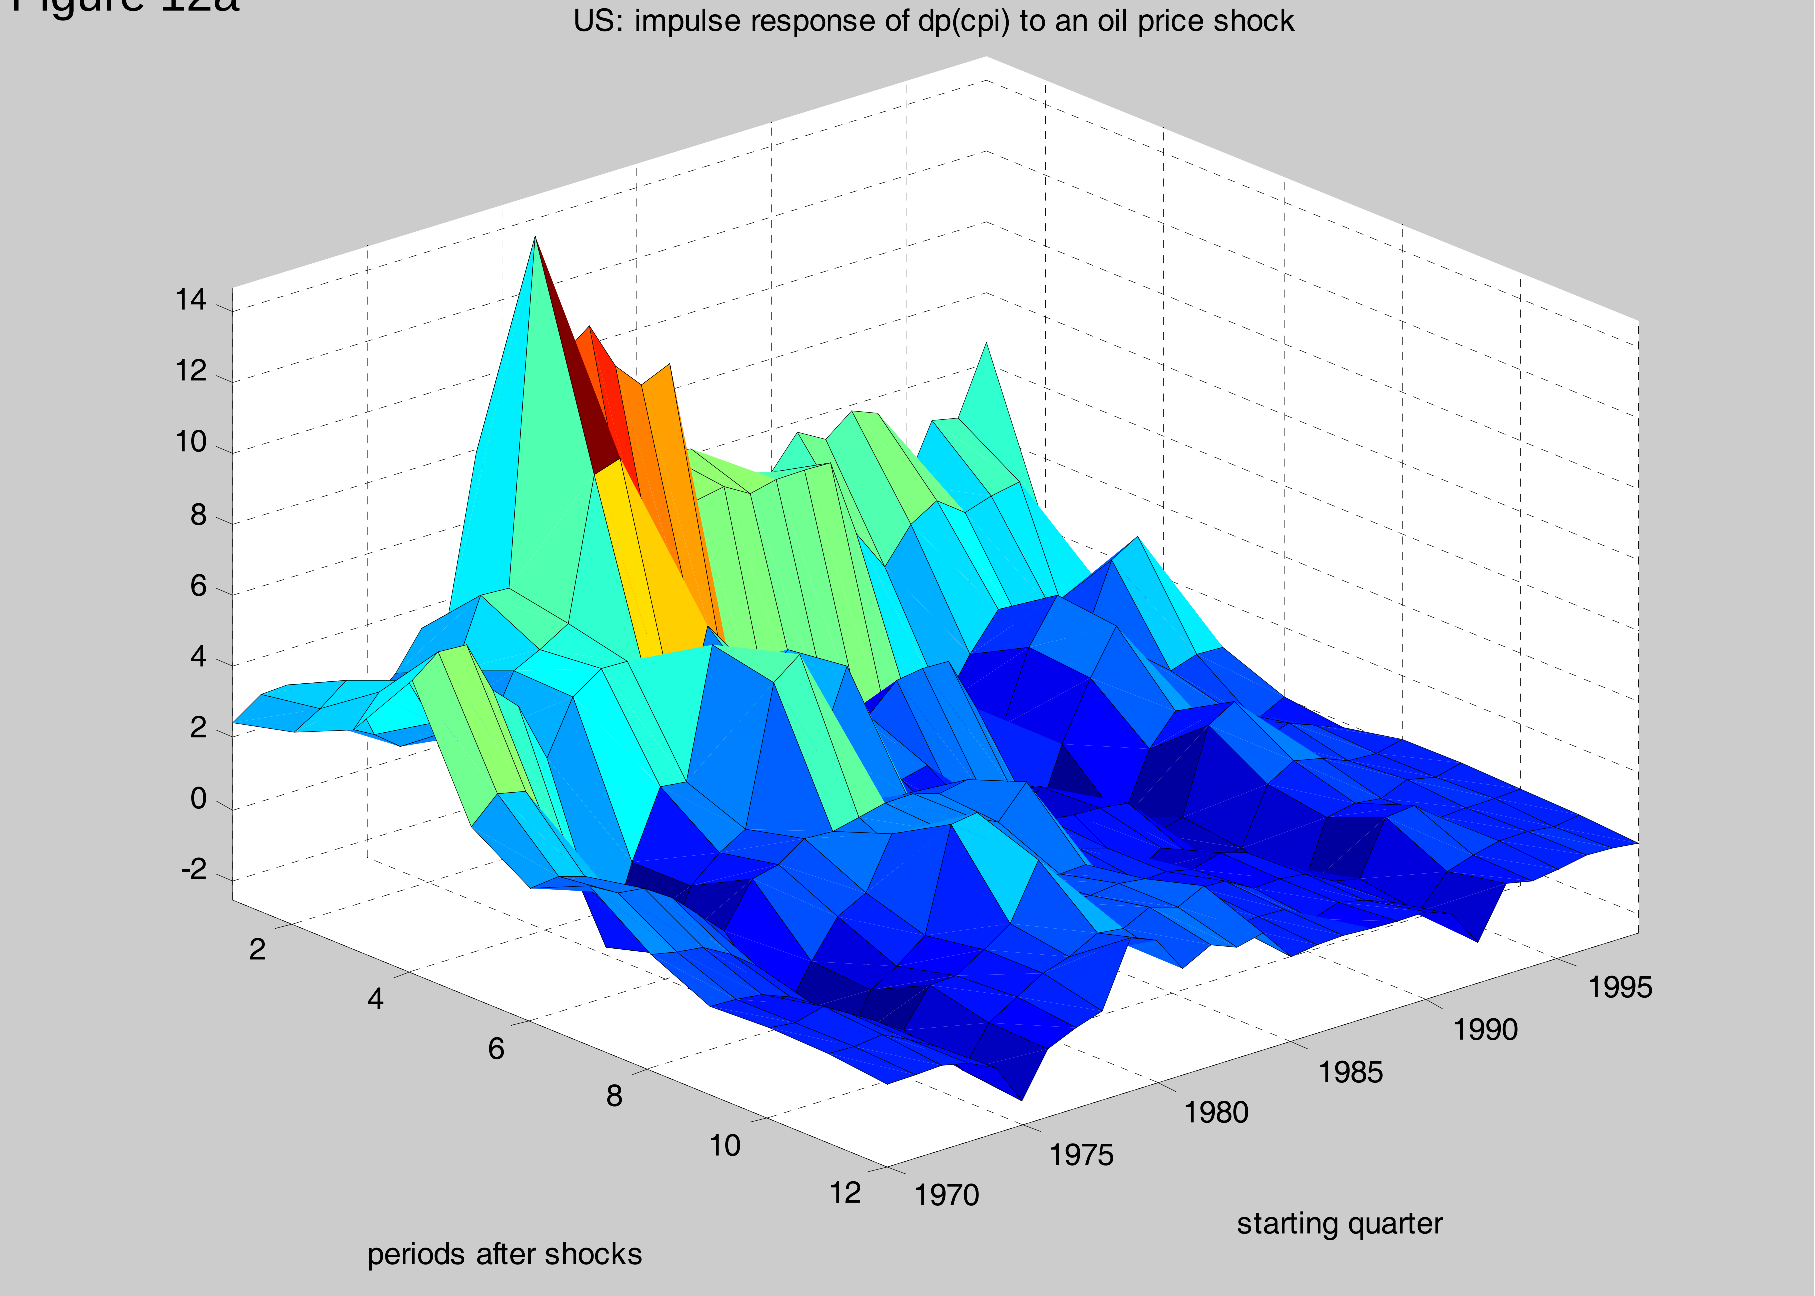

US: impulse response of dp(cpi) to an oil price shock
14
12
10
8
6
4
2
0
-2
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 12b

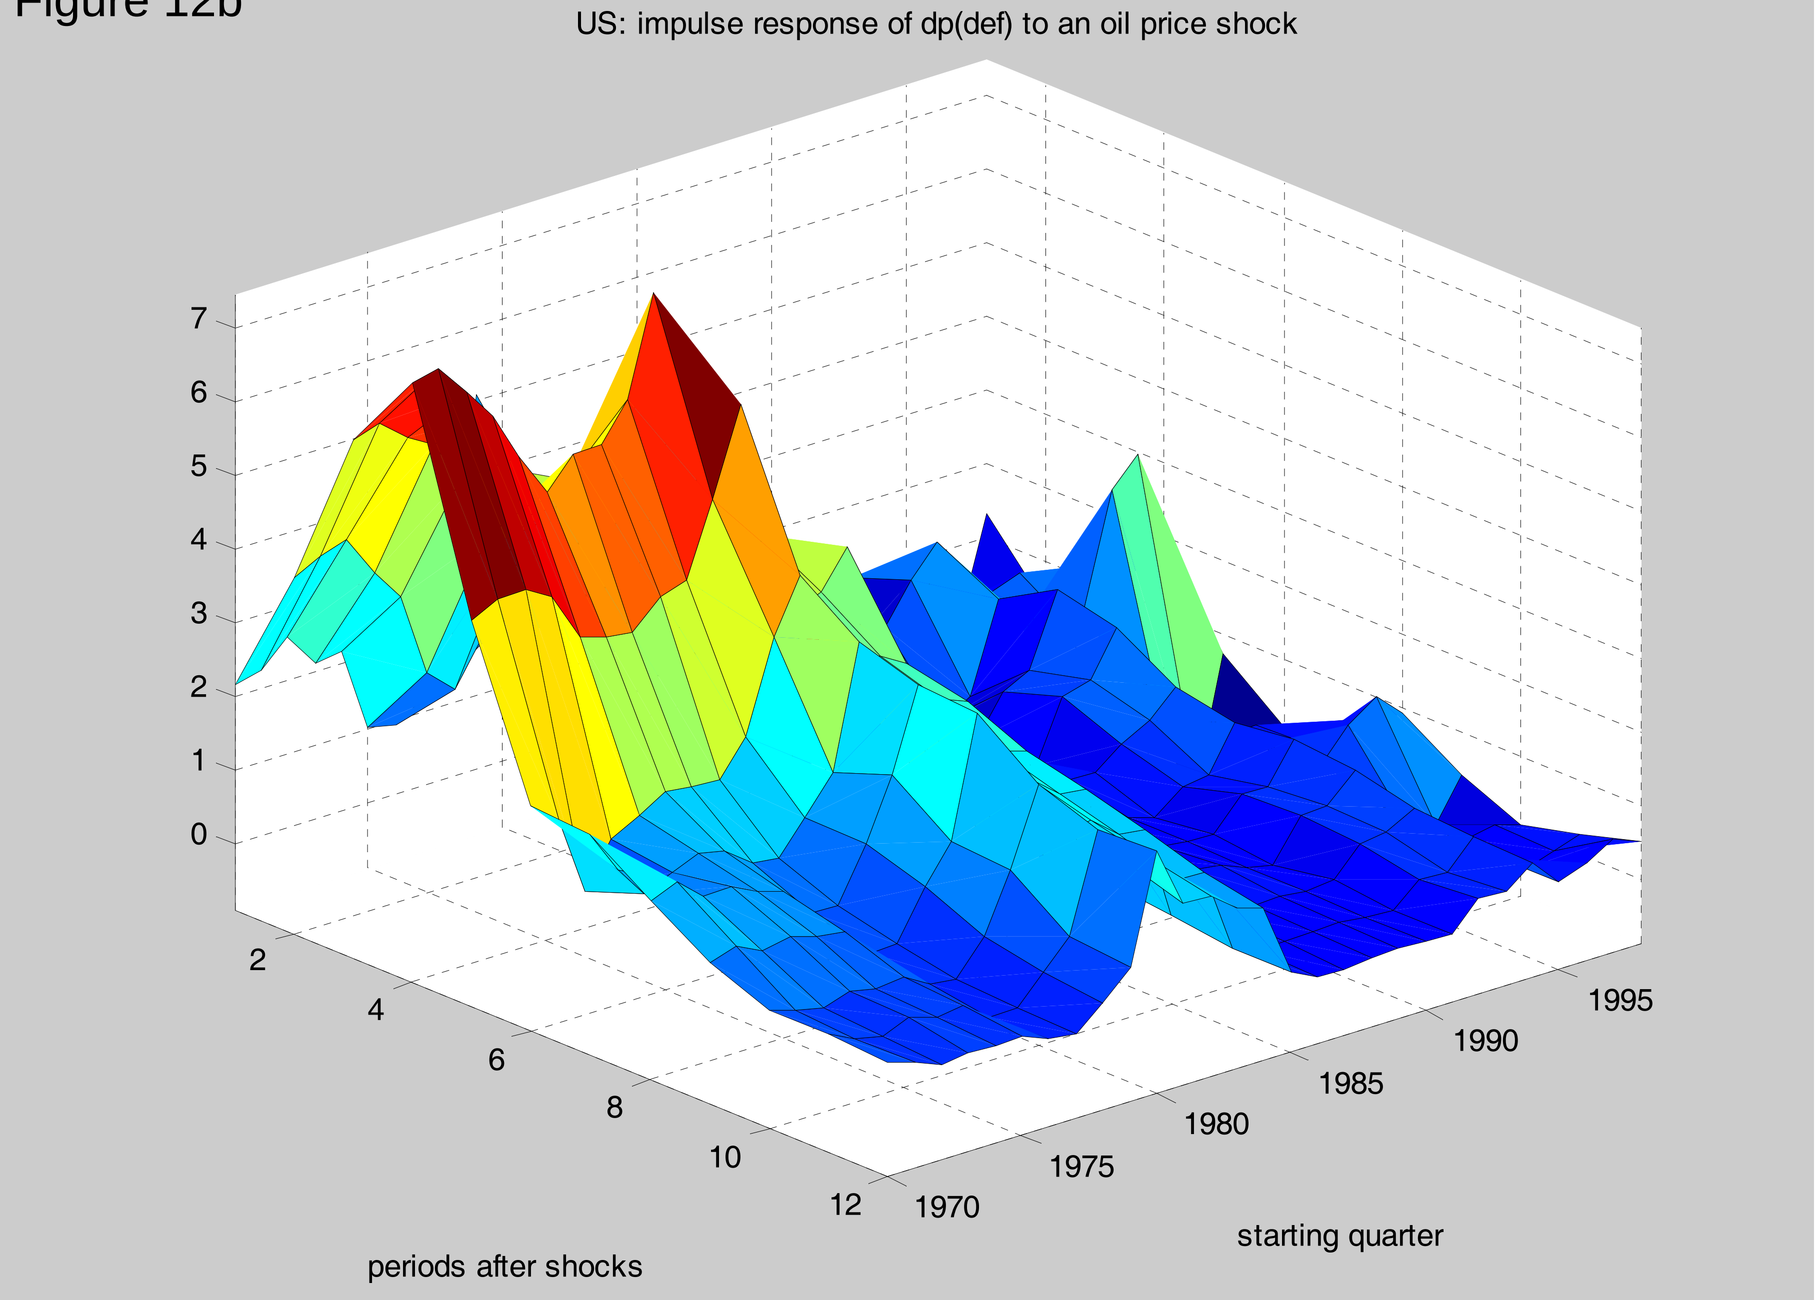
US: impulse response of dp(def) to an oil price shock
0
1
2
3
4
5
6
7
2
4
6
8
10
12
1970
1975
1980
1985
1990
1995
periods after shocks
starting quarter

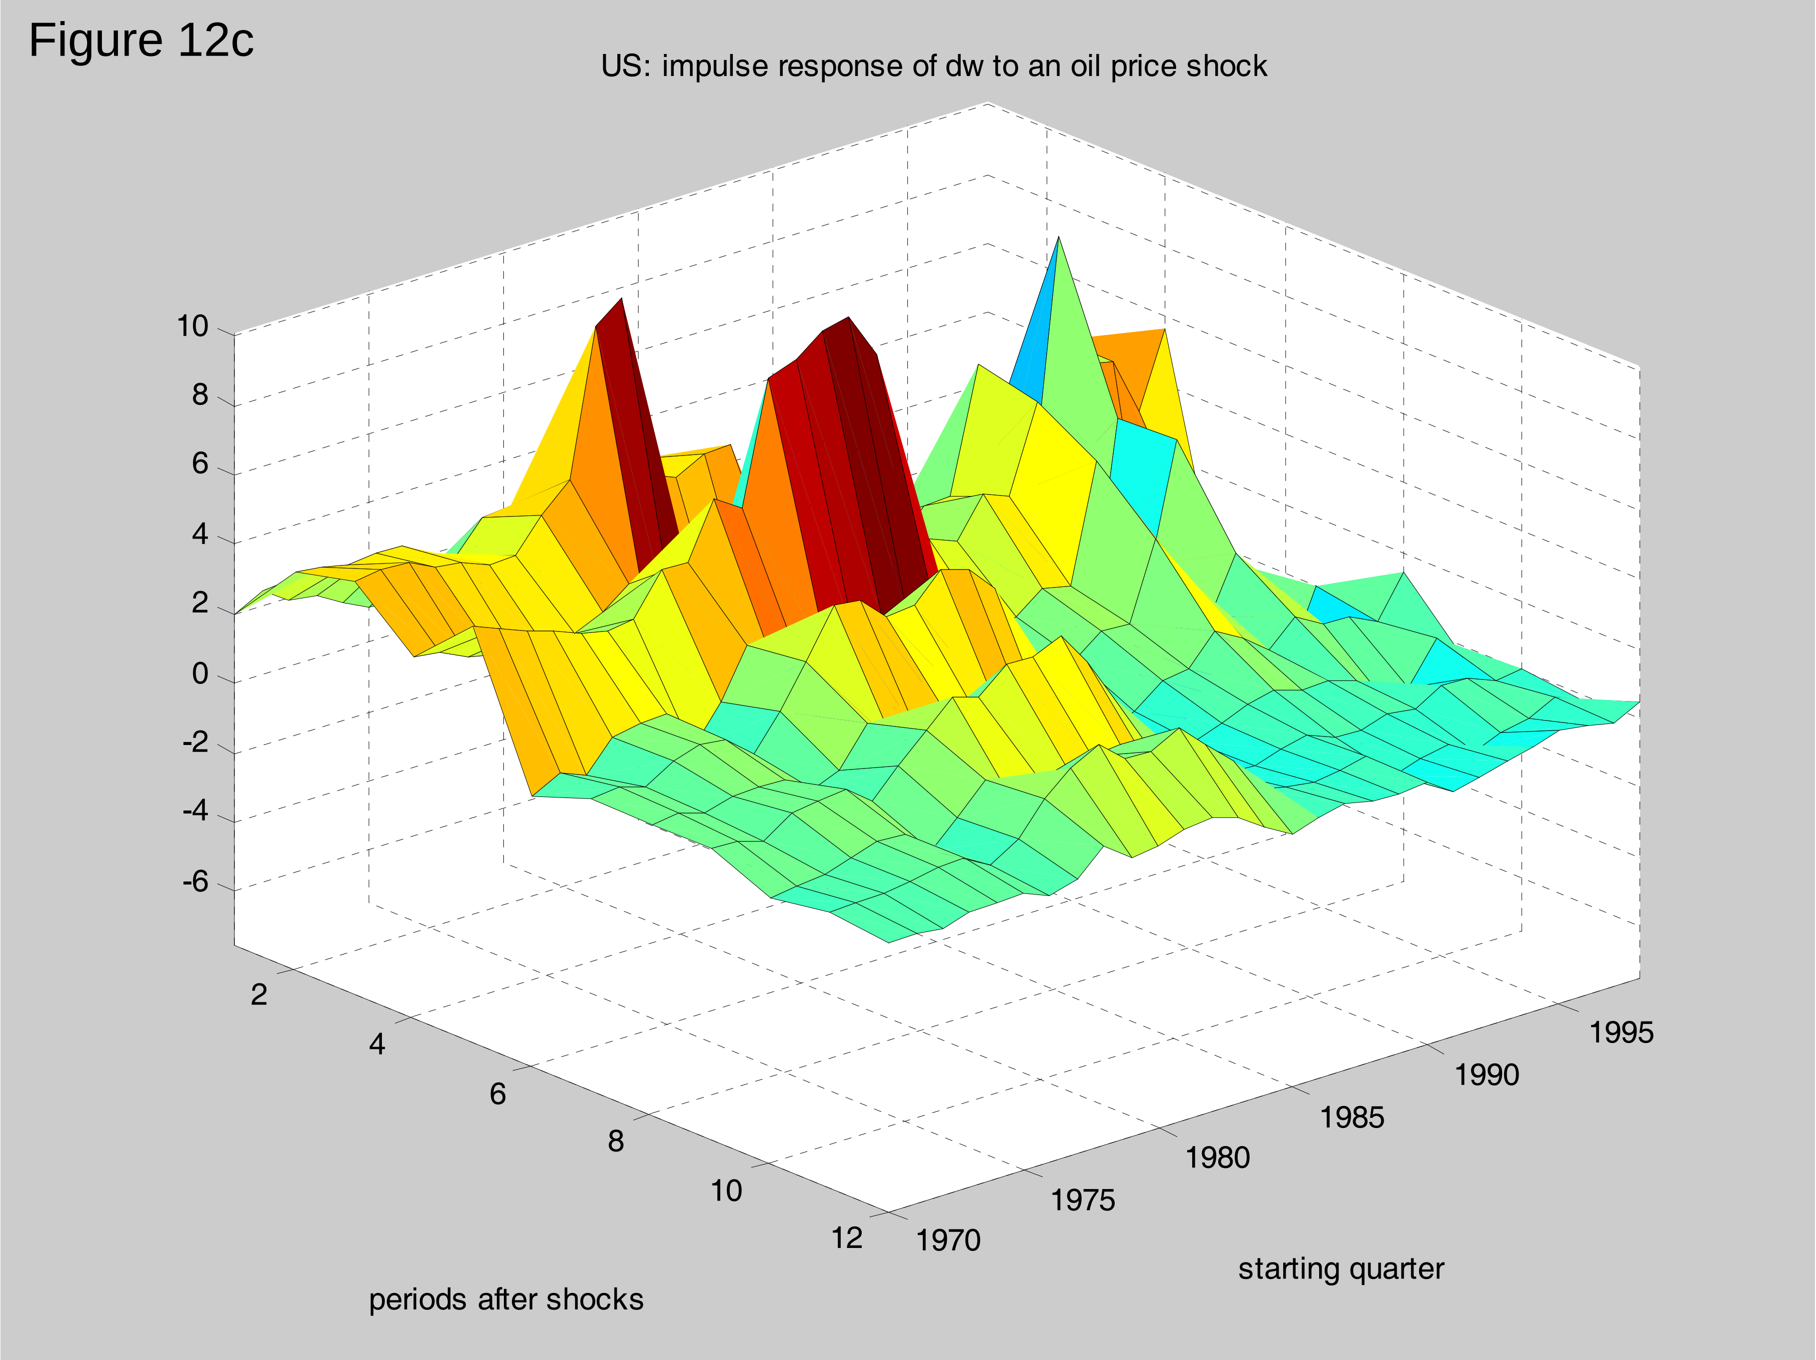
Figure 12c
US: impulse response of dw to an oil price shock
10
8
6
4
2
0
-2
-4
-6
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 12d

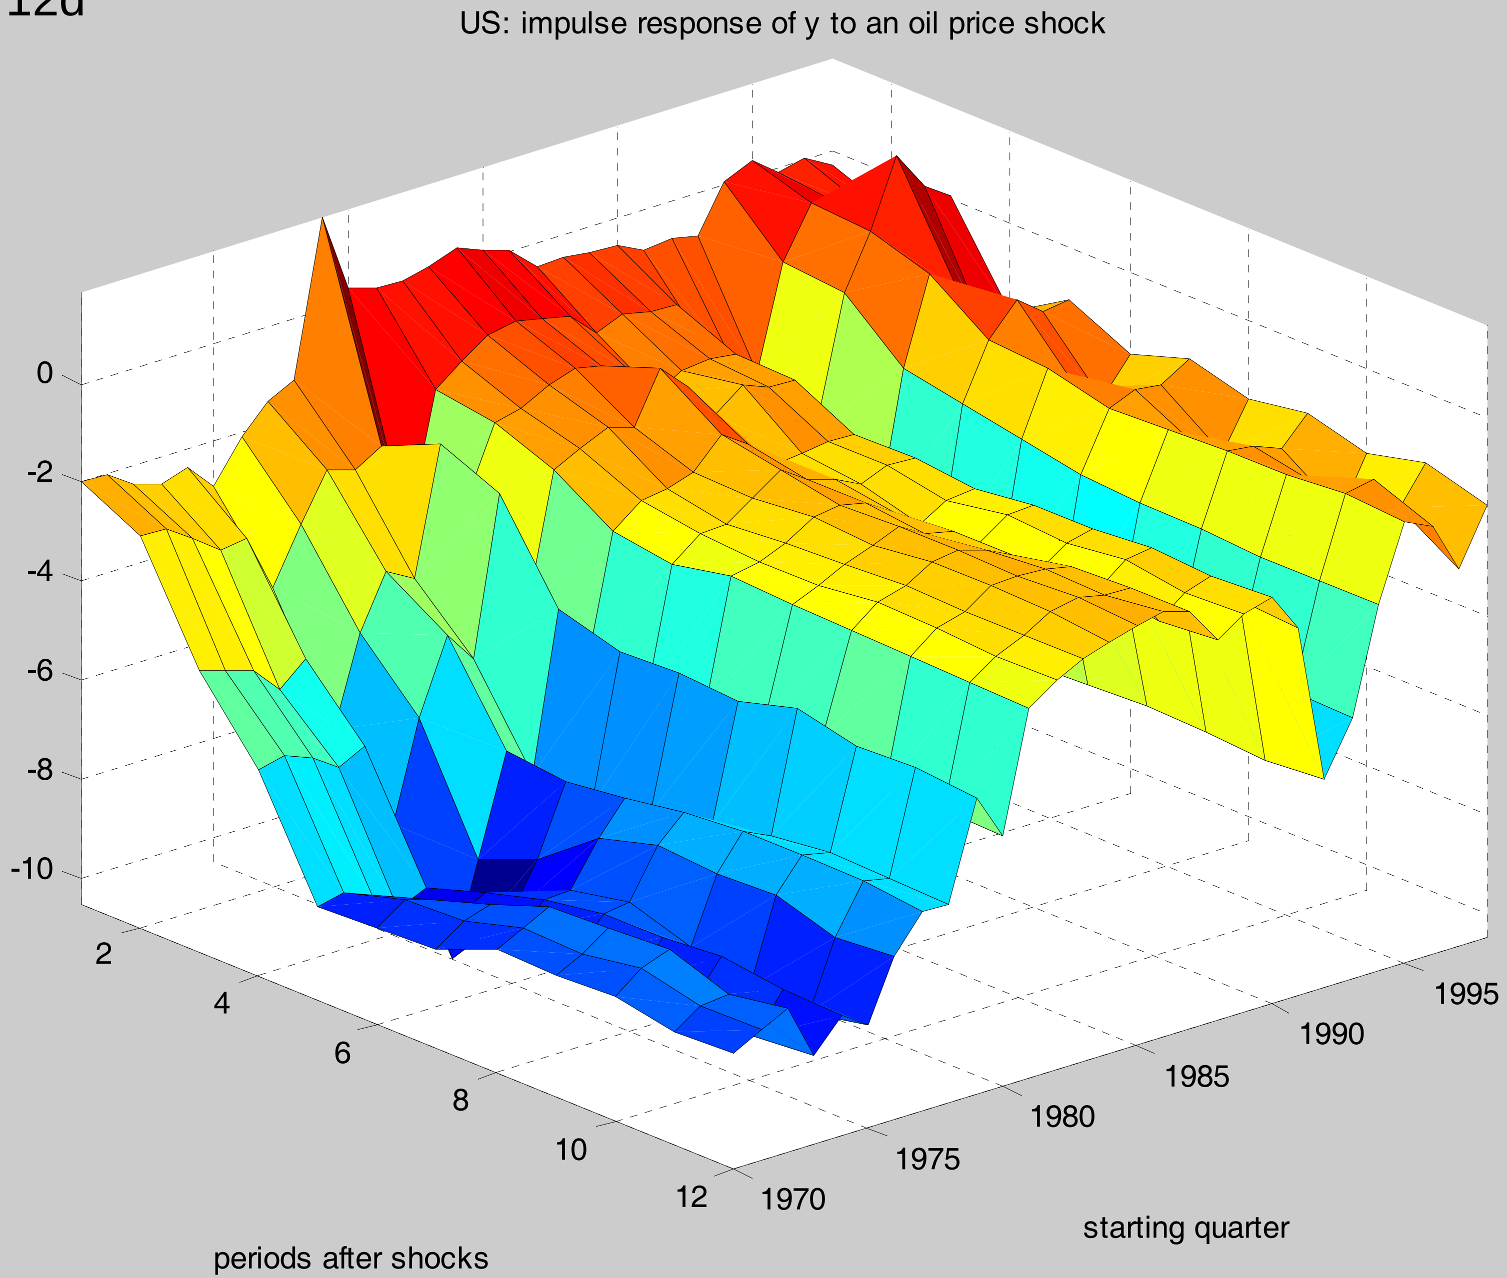
US: impulse response of y to an oil price shock
0
-2
-4
-6
-8
-10
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 12e

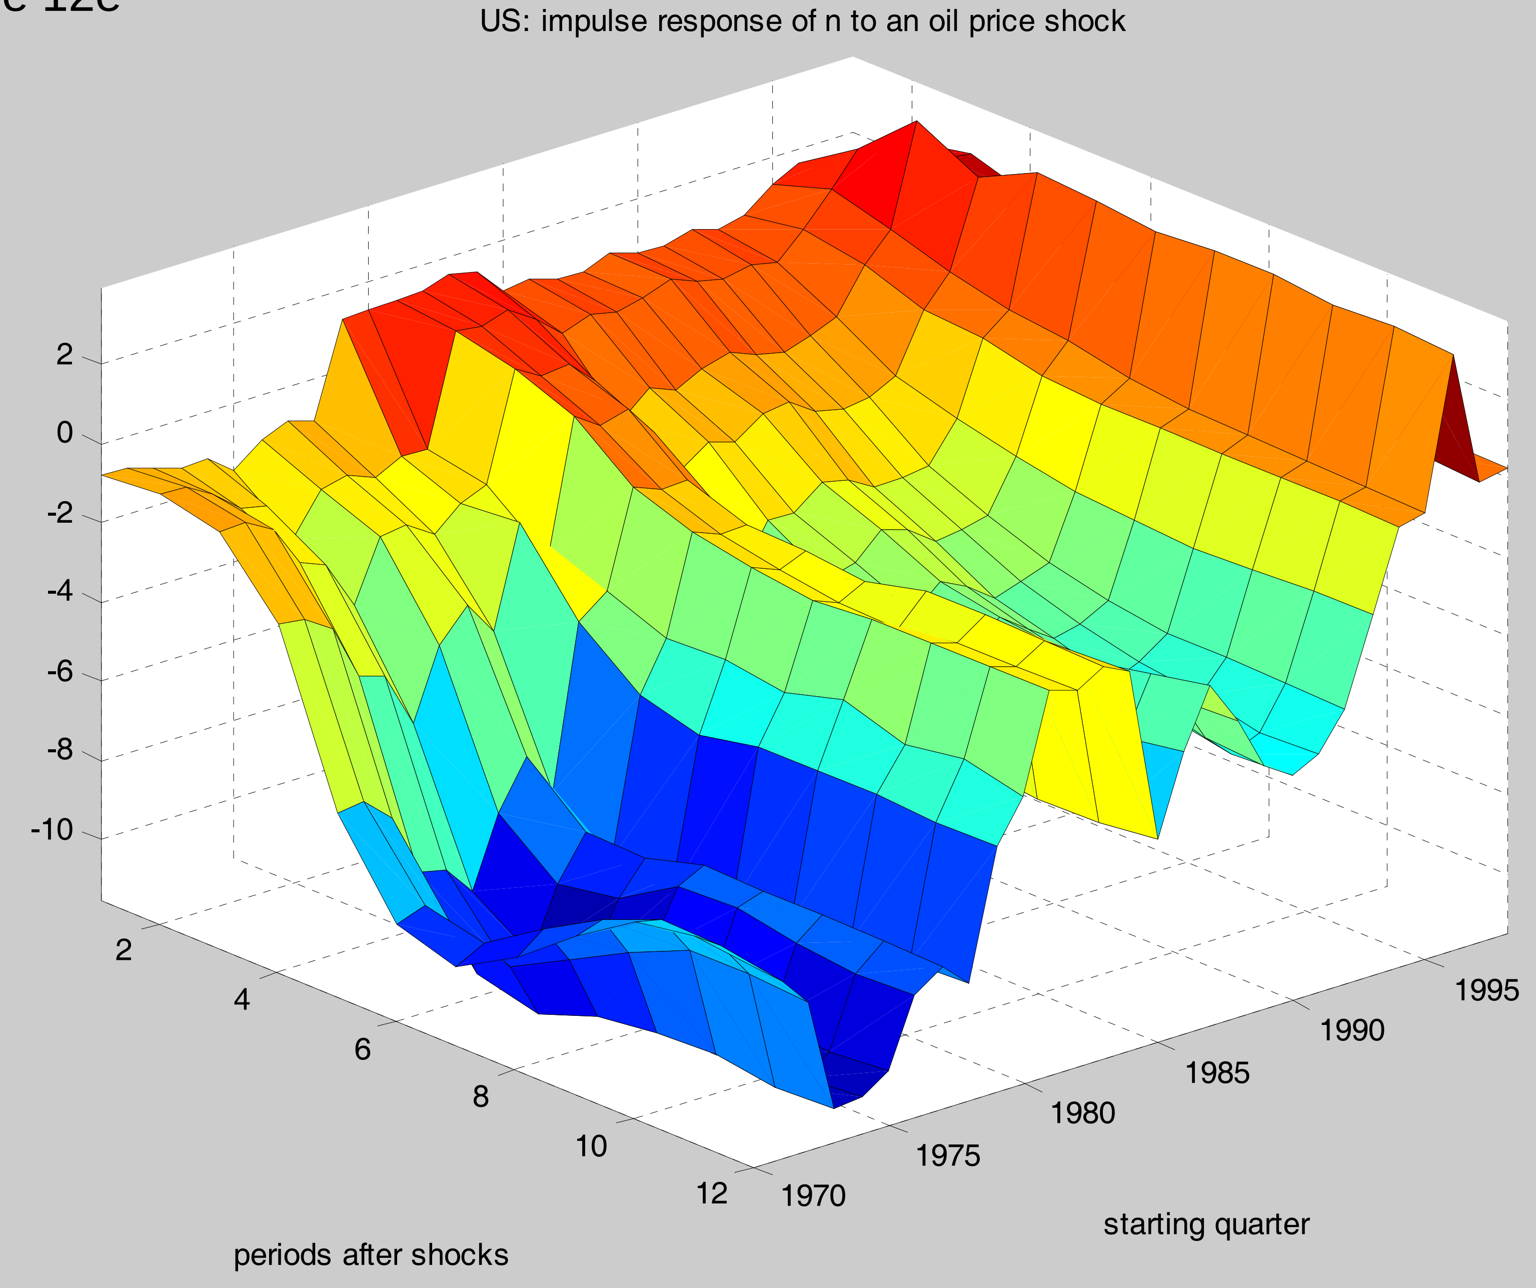
US: impulse response of n to an oil price shock
2
0
-2
-4
-6
-8
-10
2
4
6
8
10
12
1970
1975
1980
1985
1990
1995
periods after shocks
starting quarter

Figure 13. Policy Tradeoffs with Flexible Wages

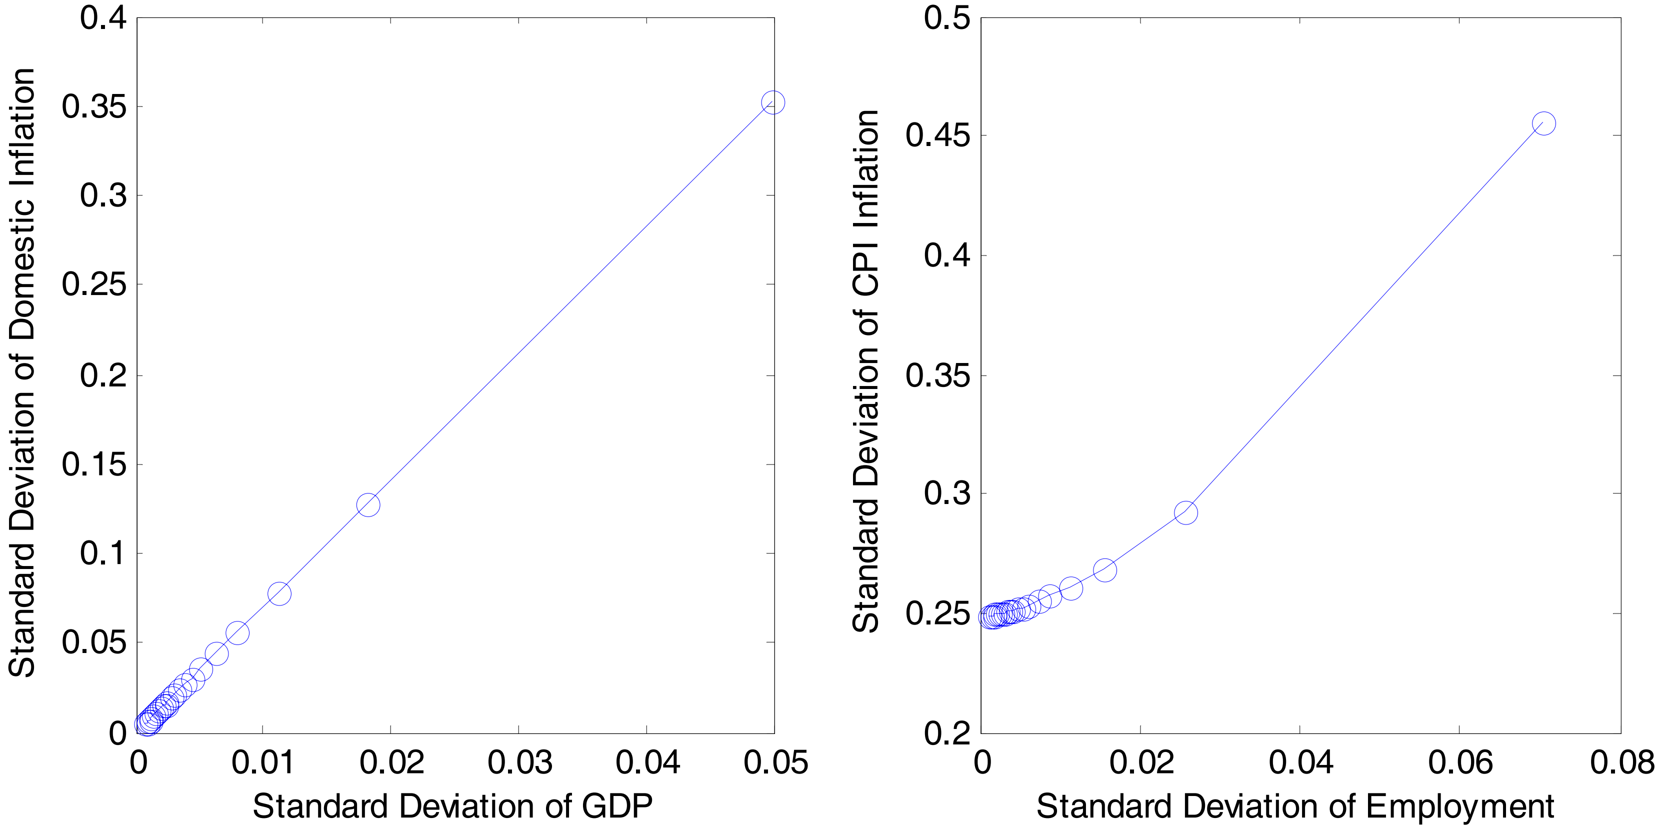

Figure 14. Real Wage Rigidities and Policy Tradeoffs

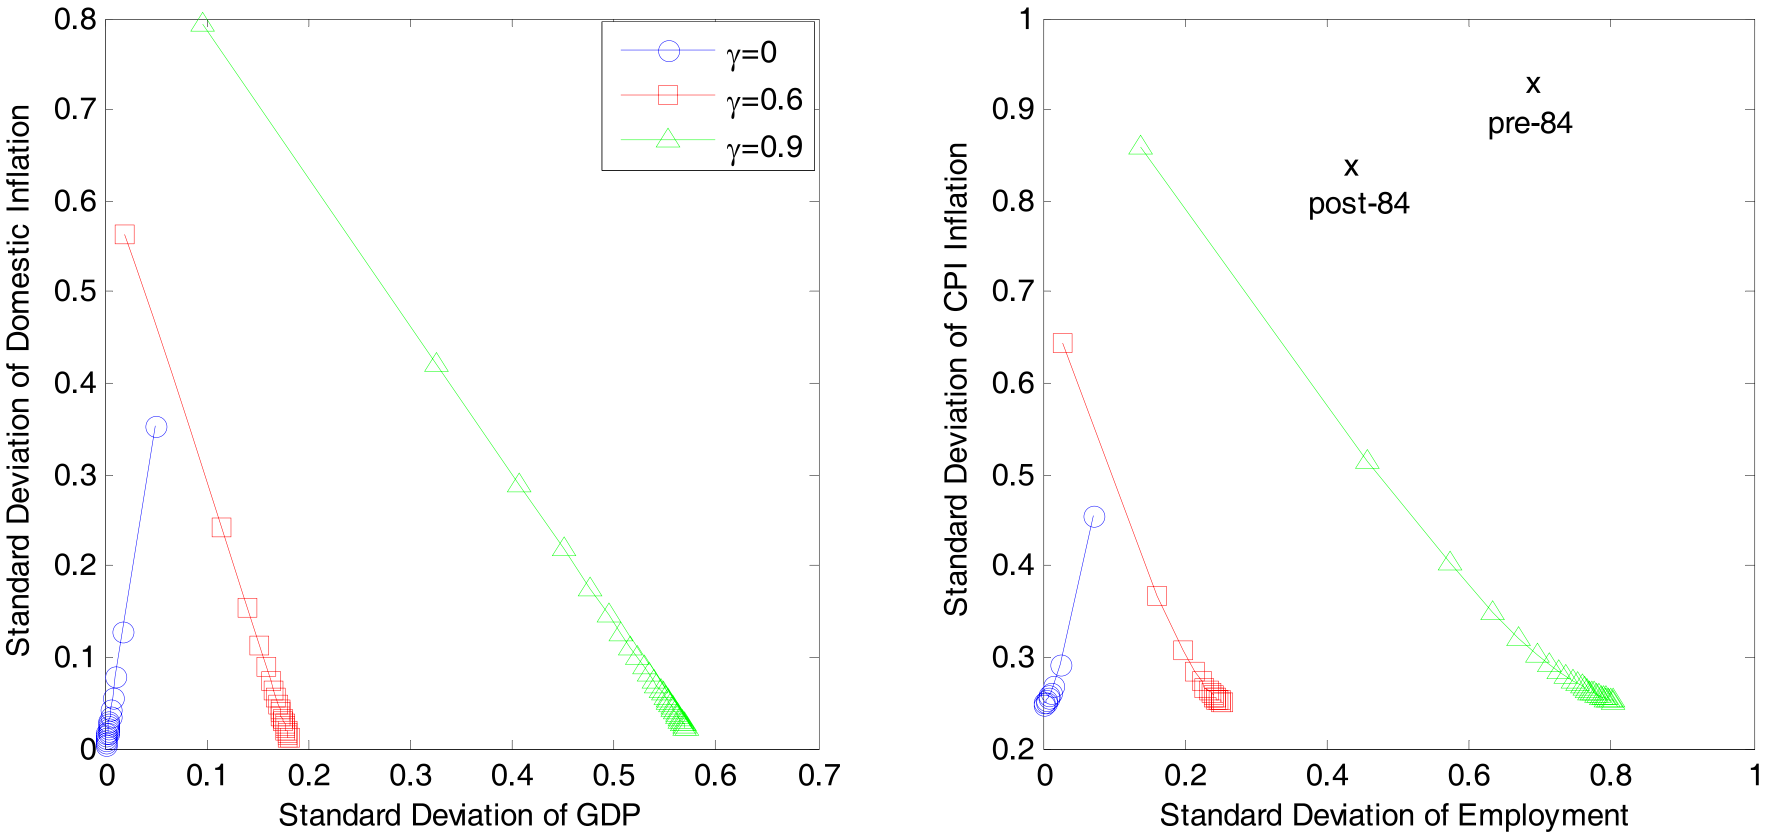

Figure 15a

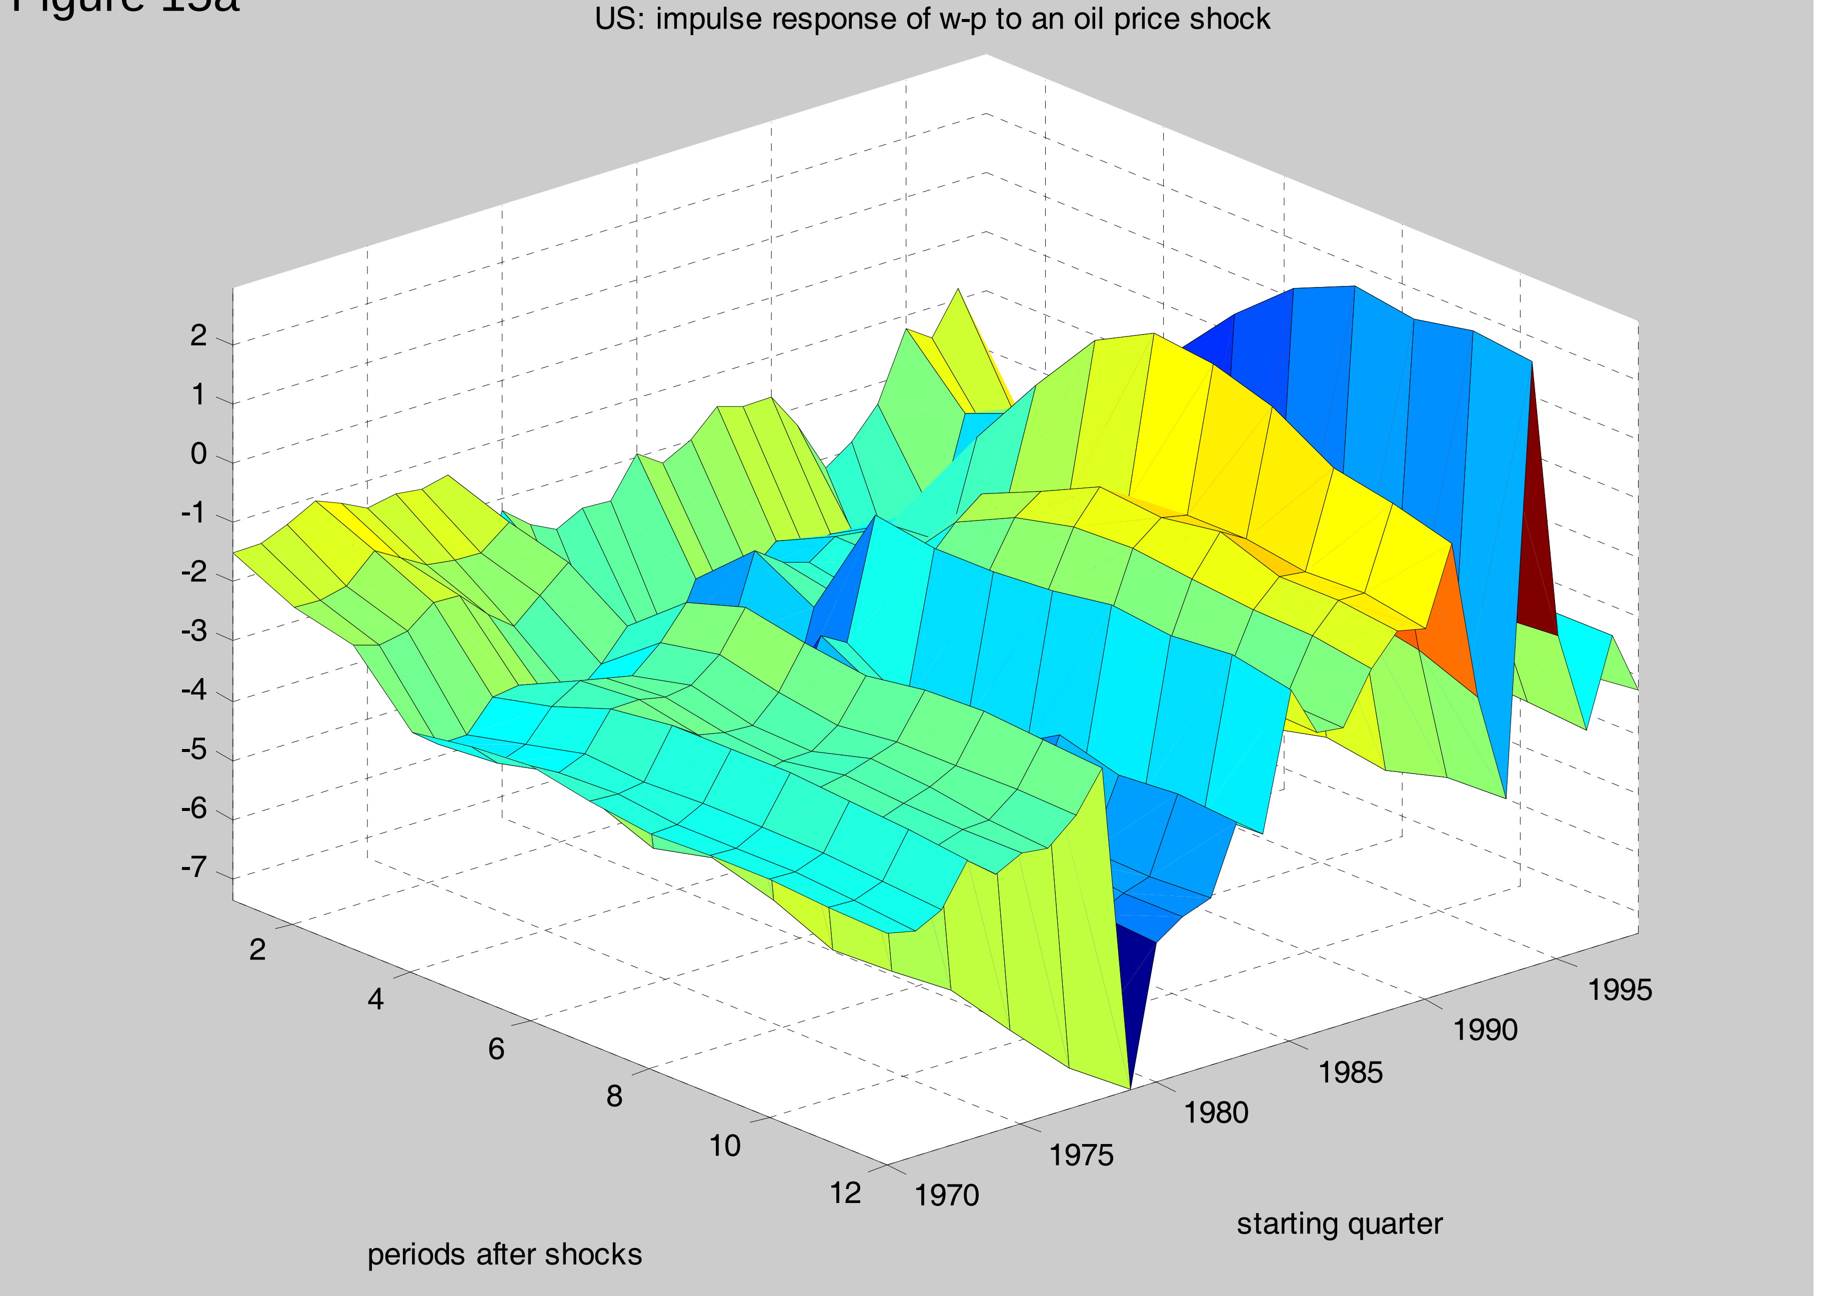
US: impulse response of w-p to an oil price shock
2
1
0
-1
-2
-3
-4
-5
-6
-7
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 15b

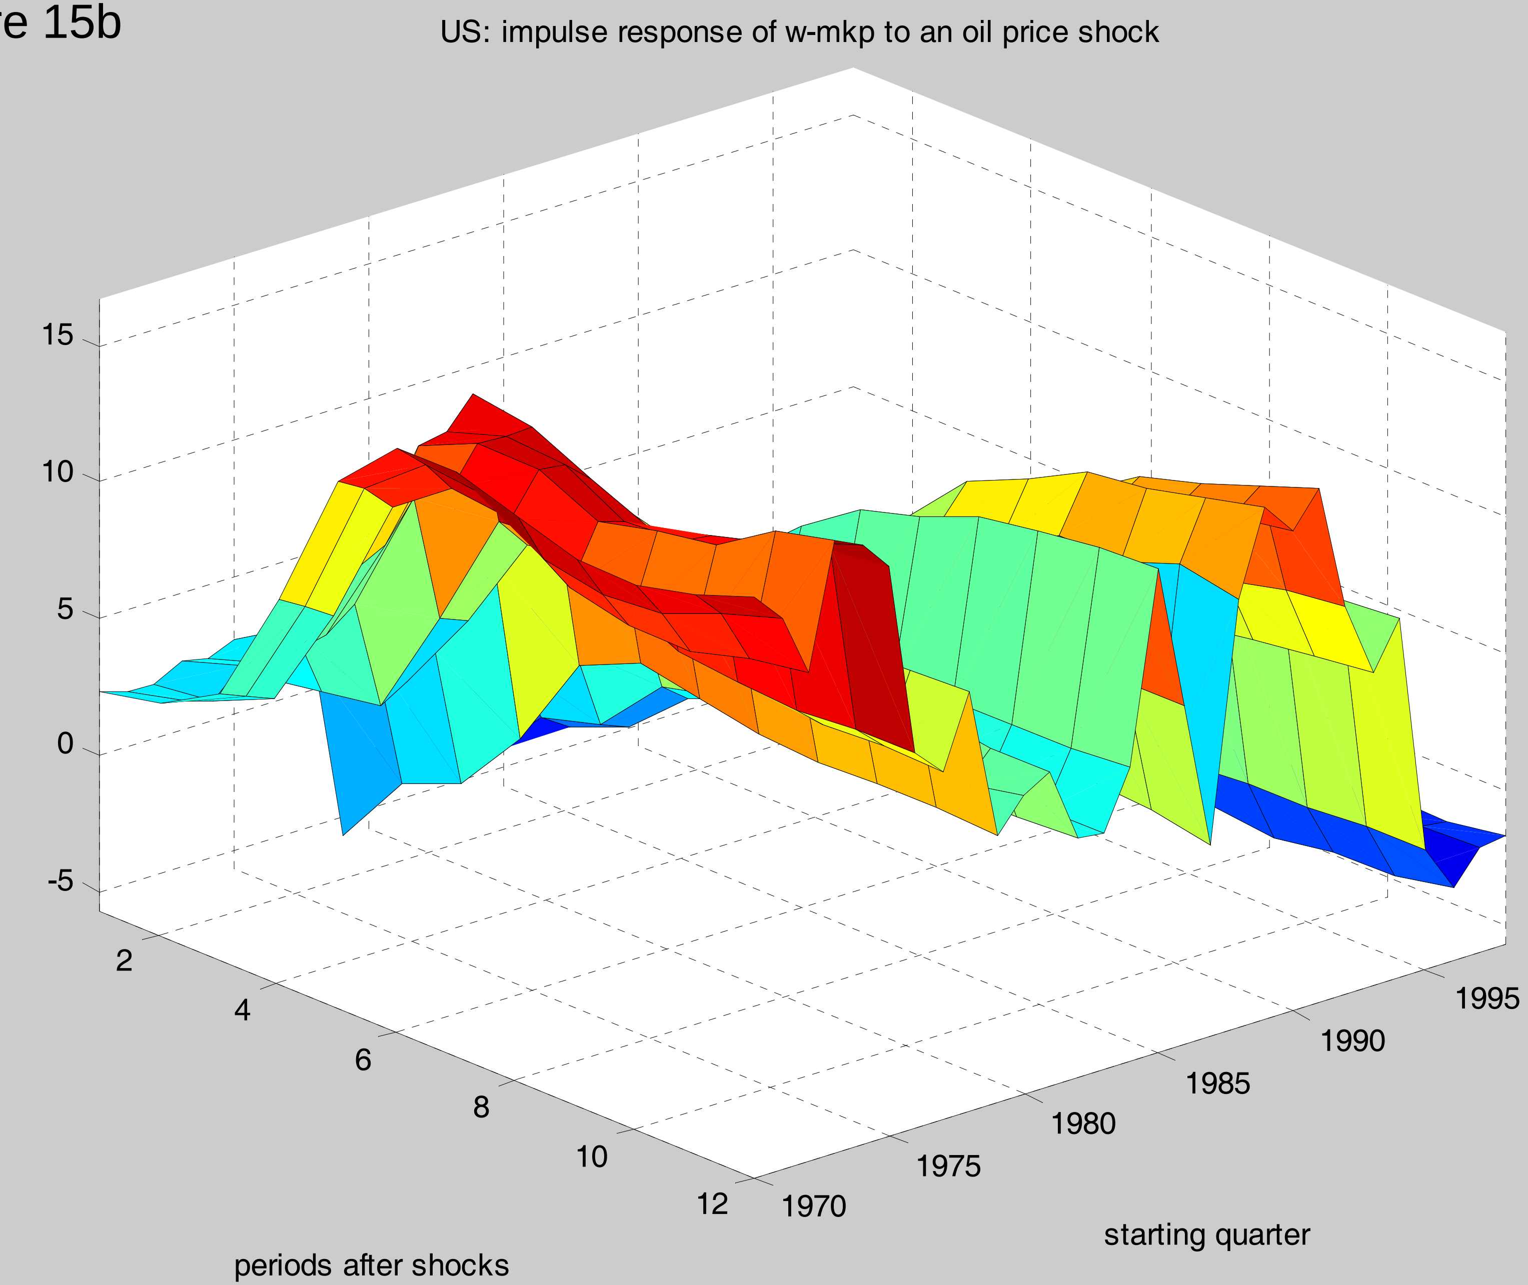
US: impulse response of w-mkp to an oil price shock
15
10
5
0
-5
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 15b

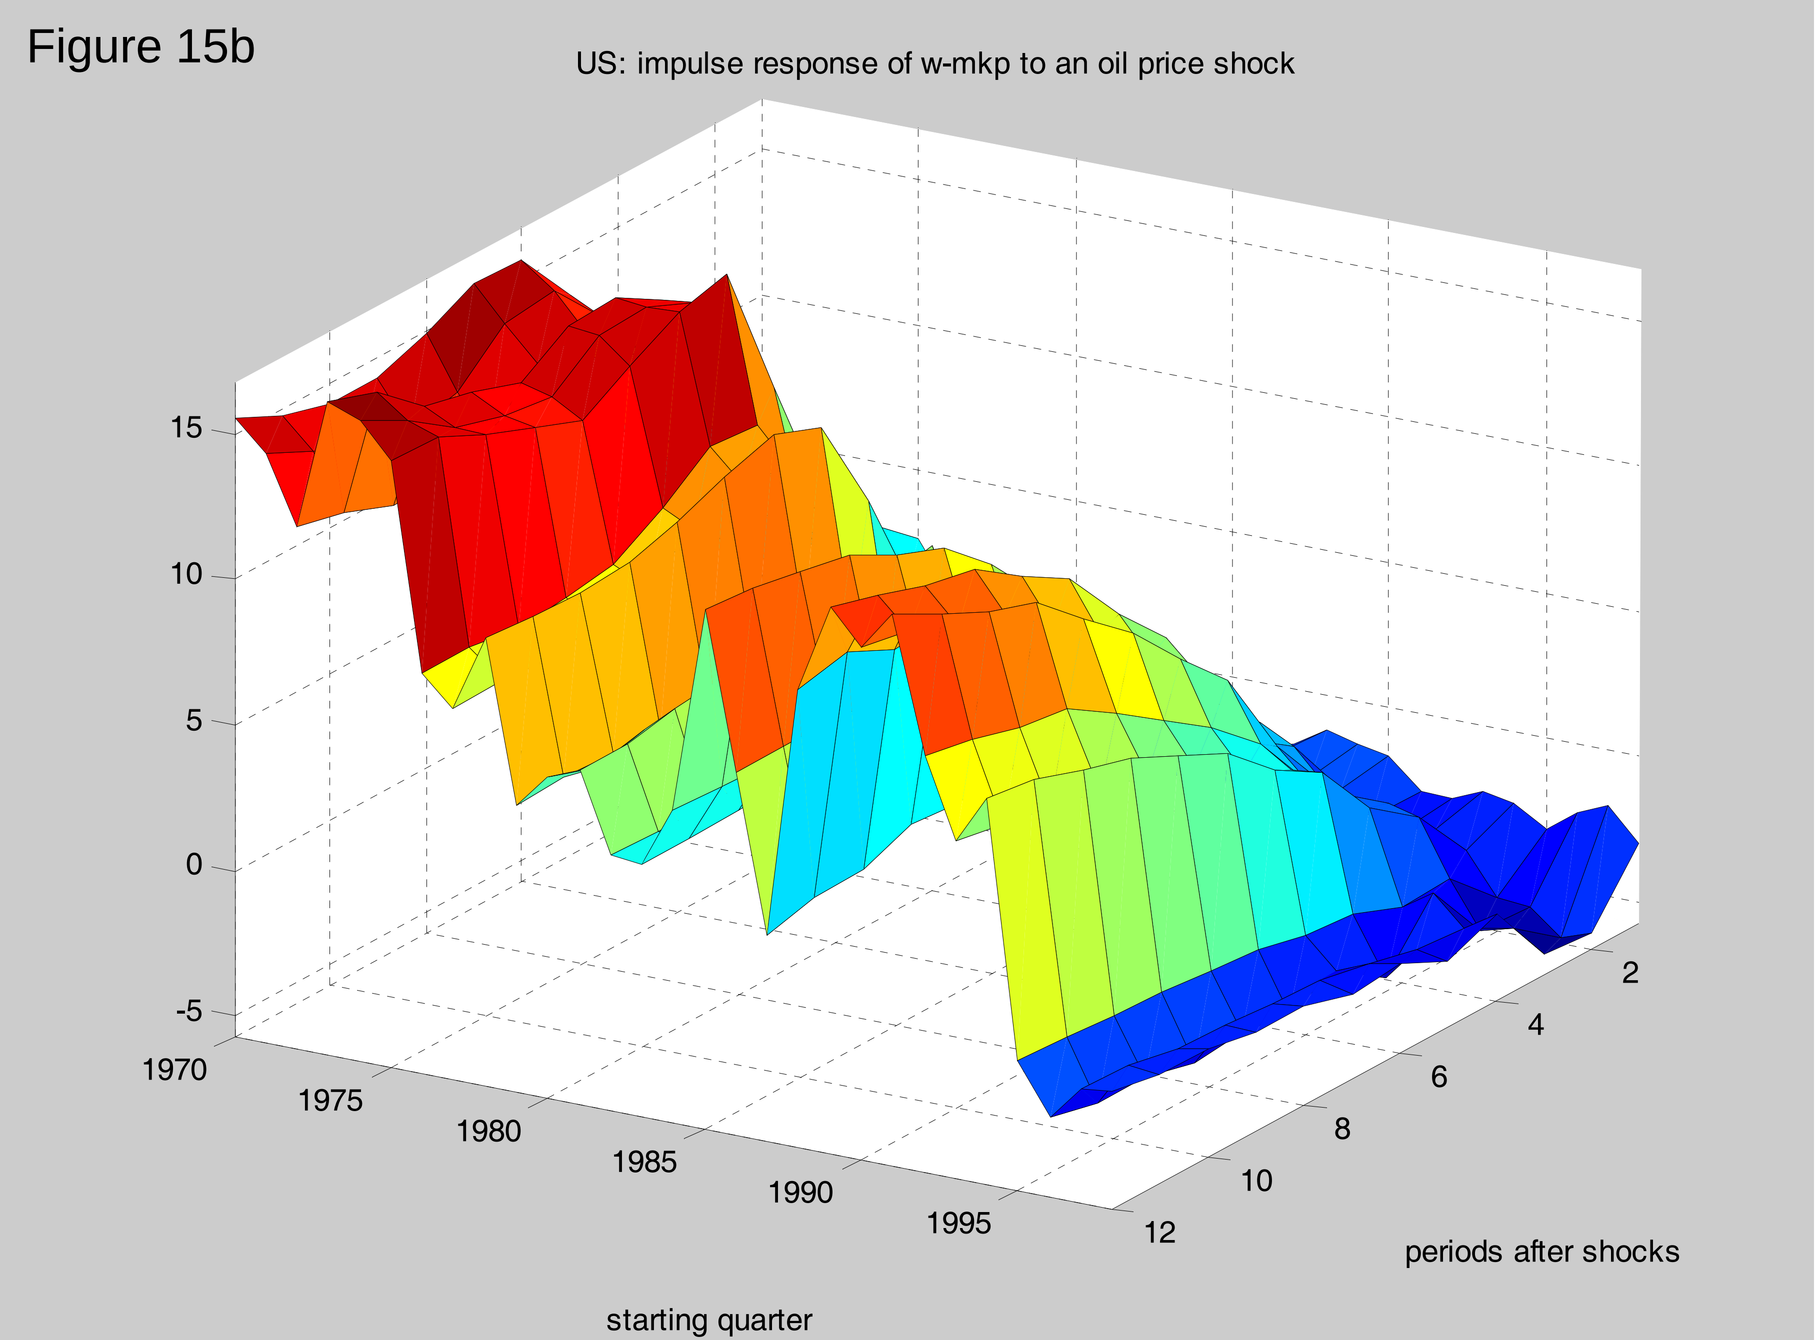

US: impulse response of w-mkp to an oil price shock
15
10
5
0
-5
1970
1975
1980
1985
1990
1995
2
4
6
8
10
12
periods after shocks
starting quarter

Figure 15c

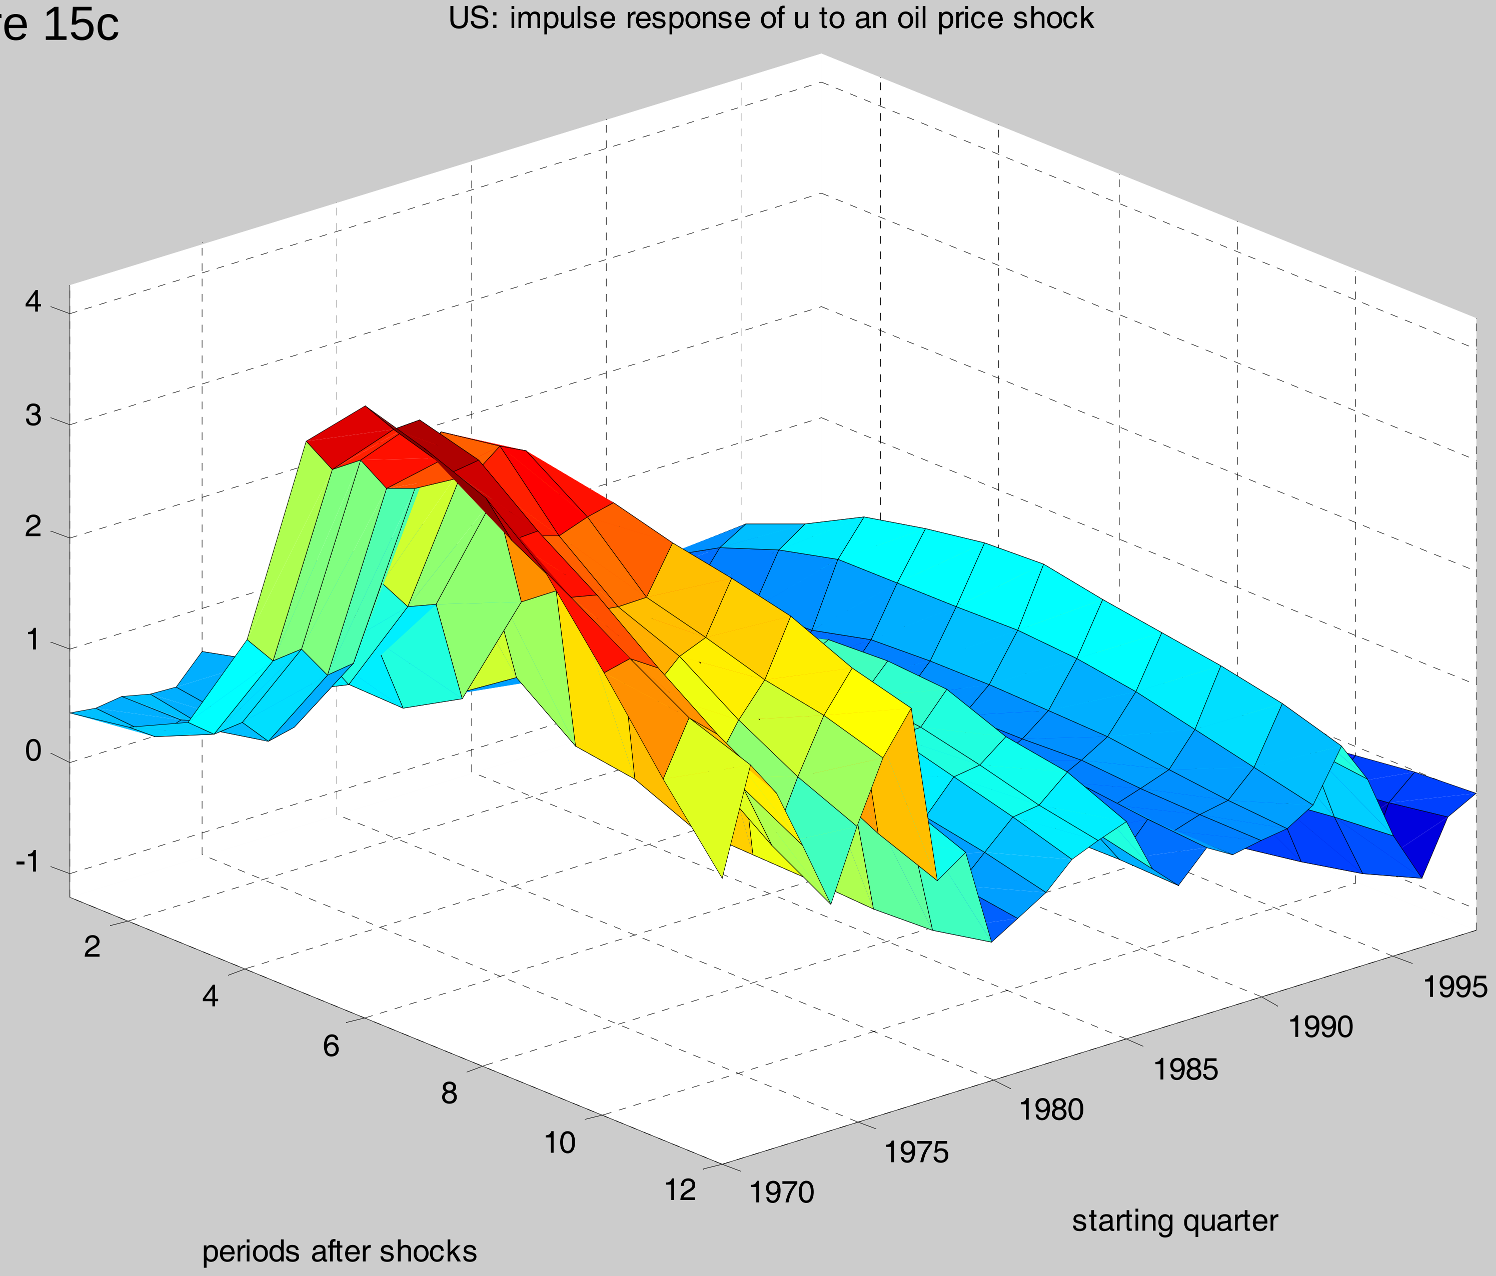
US: impulse response of u to an oil price shock
4
3
2
1
0
-1
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 15c

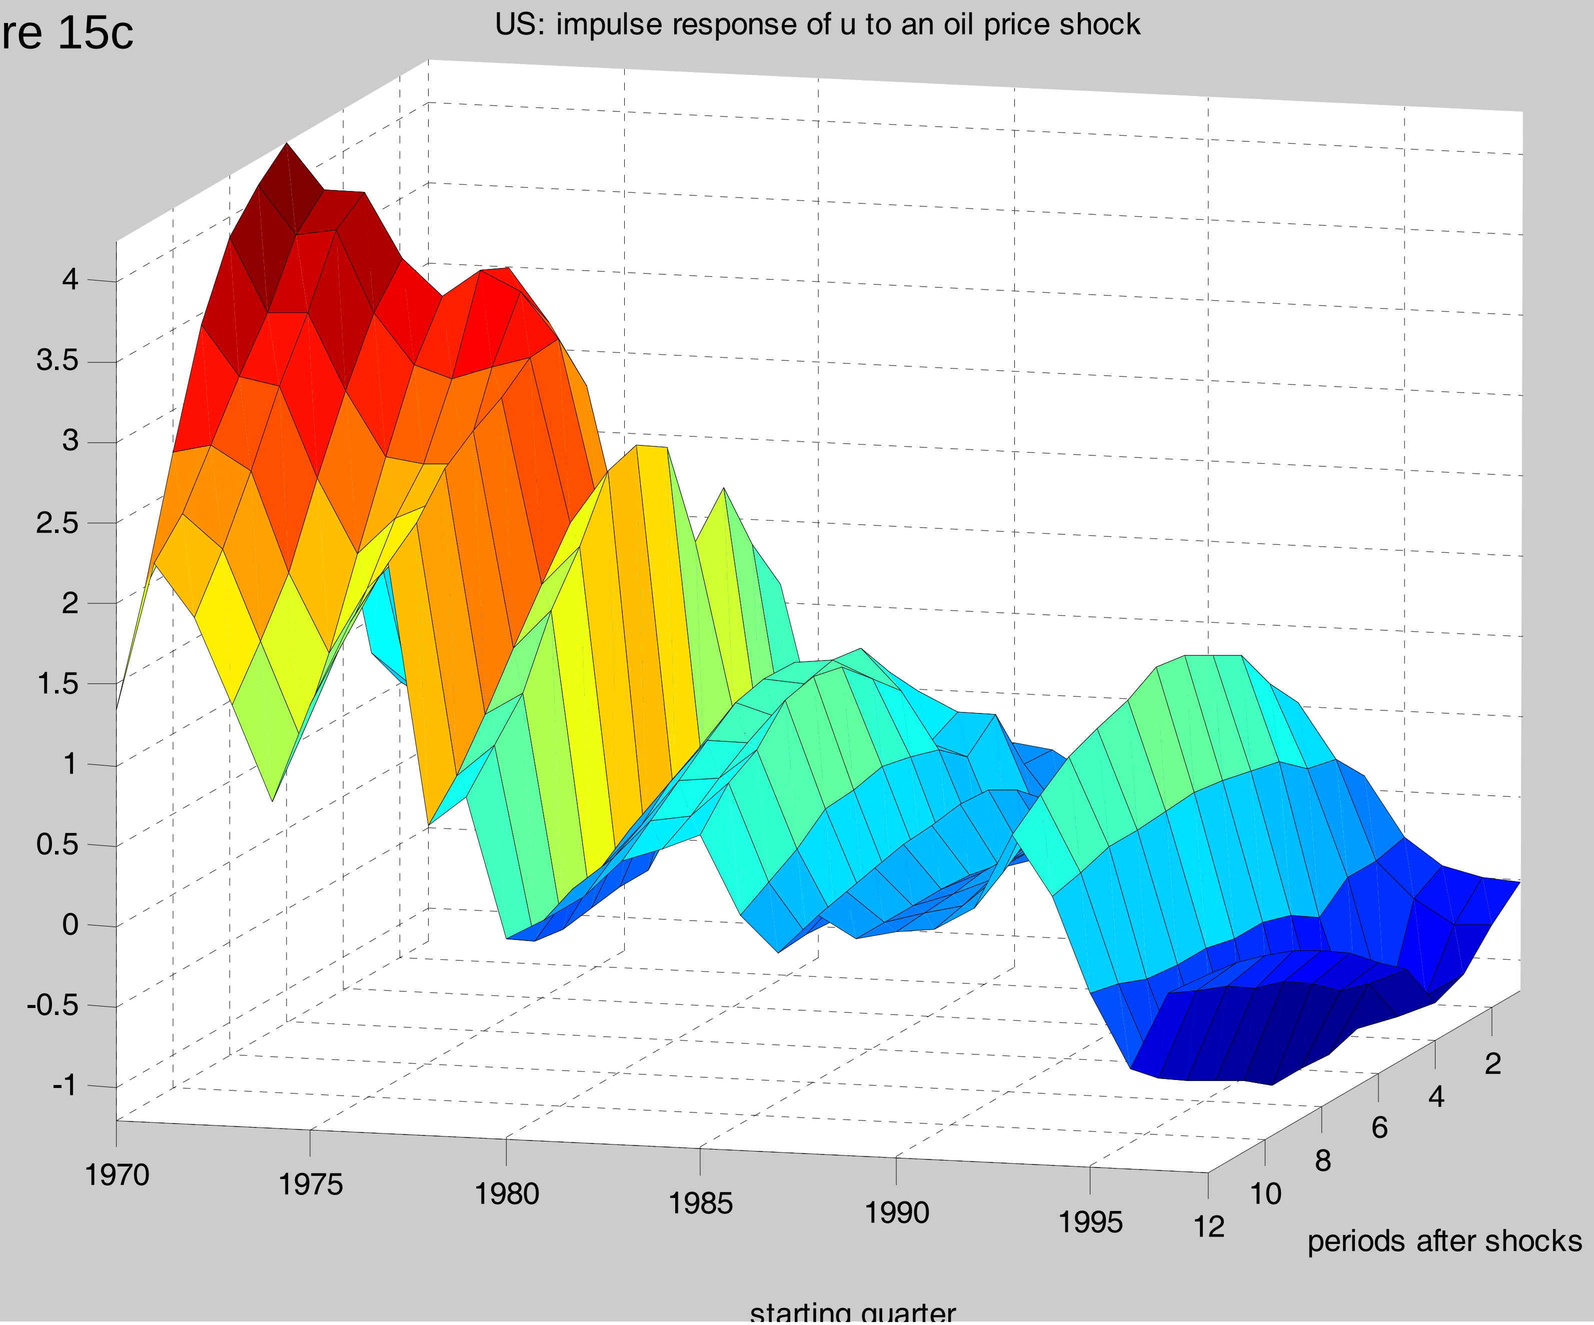

US: impulse response of u to an oil price shock
4
3.5
3
2.5
2
1.5
1
0.5
0
-0.5
-1
1970
1975
1980
1985
1990
1995
starting quarter
2
4
6
8
10
12
periods after shocks

Figure 16

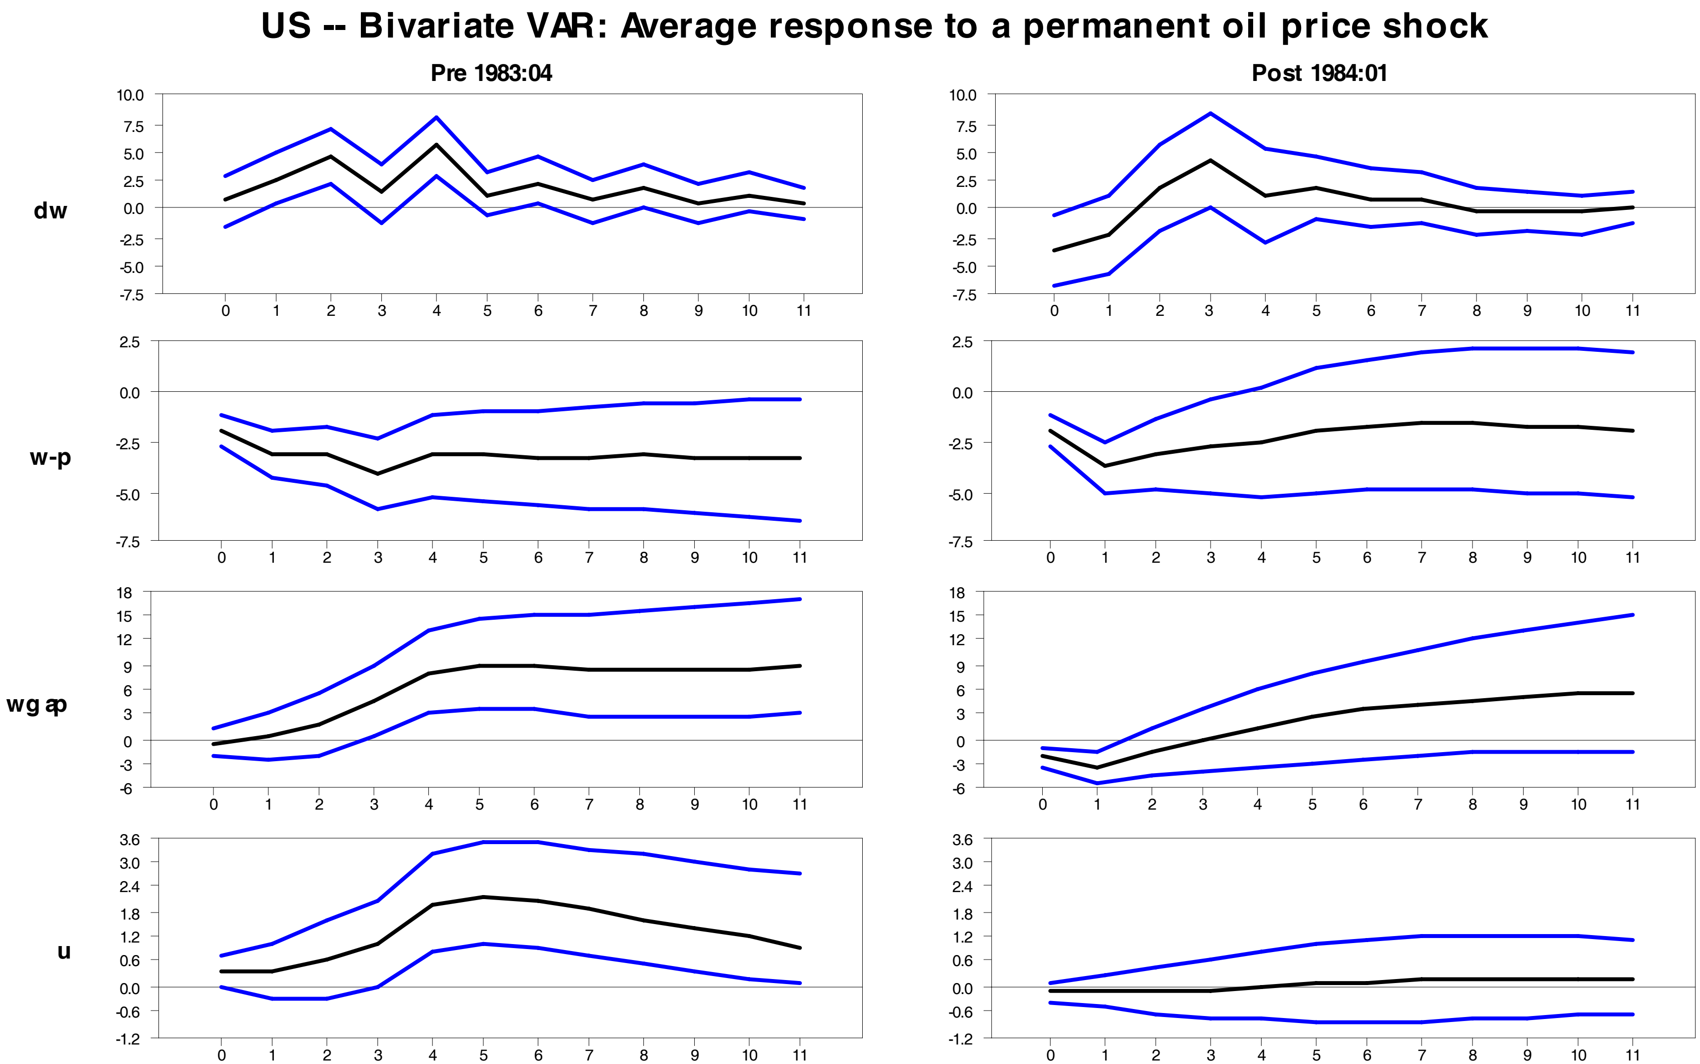

Bivariate VARs – Labor Market Variables

Figure 17: Credibility and Policy Tradeoffs

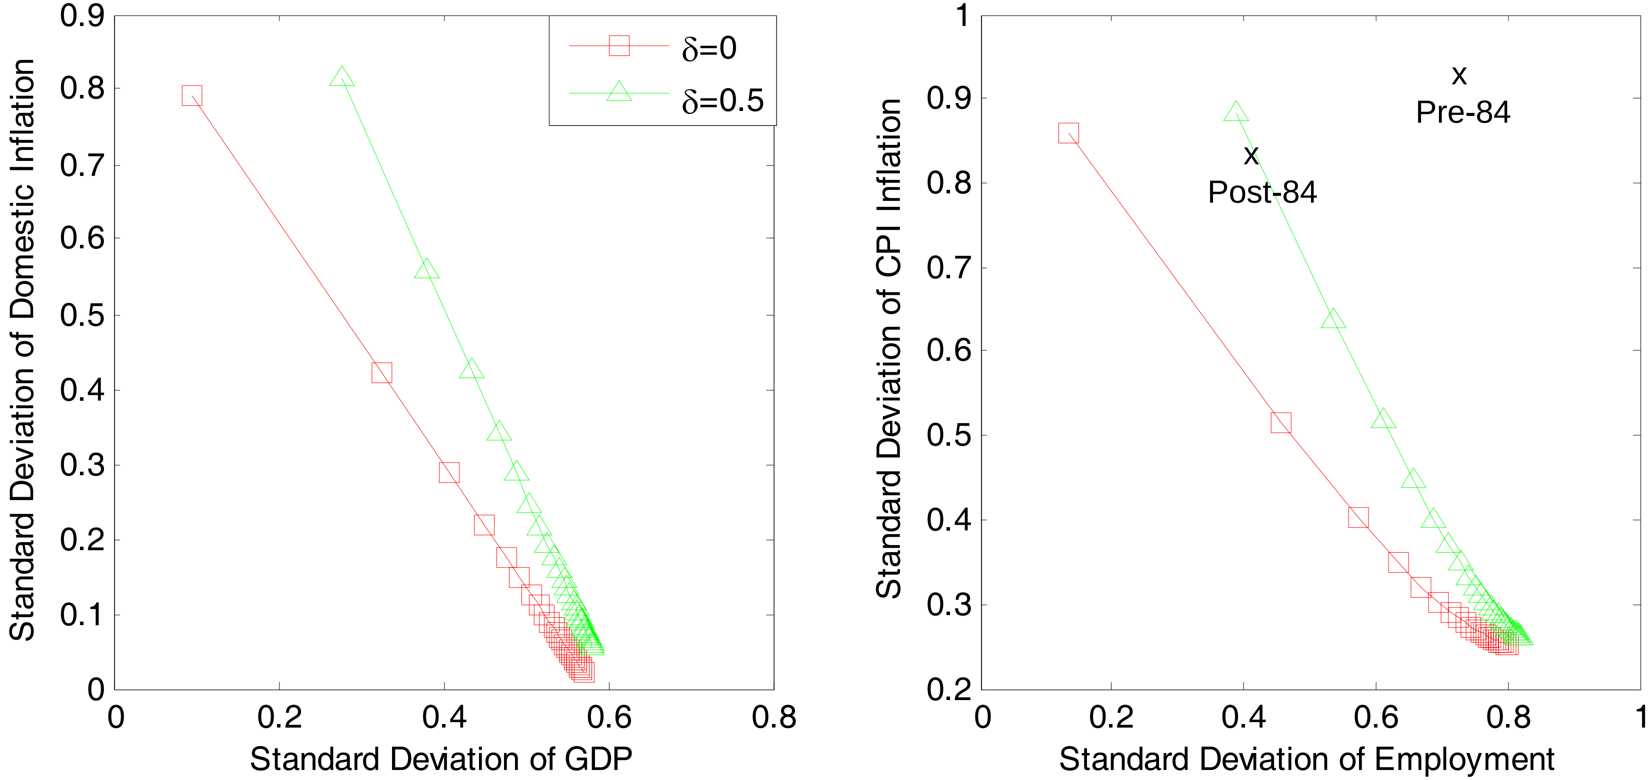

Figure 18a

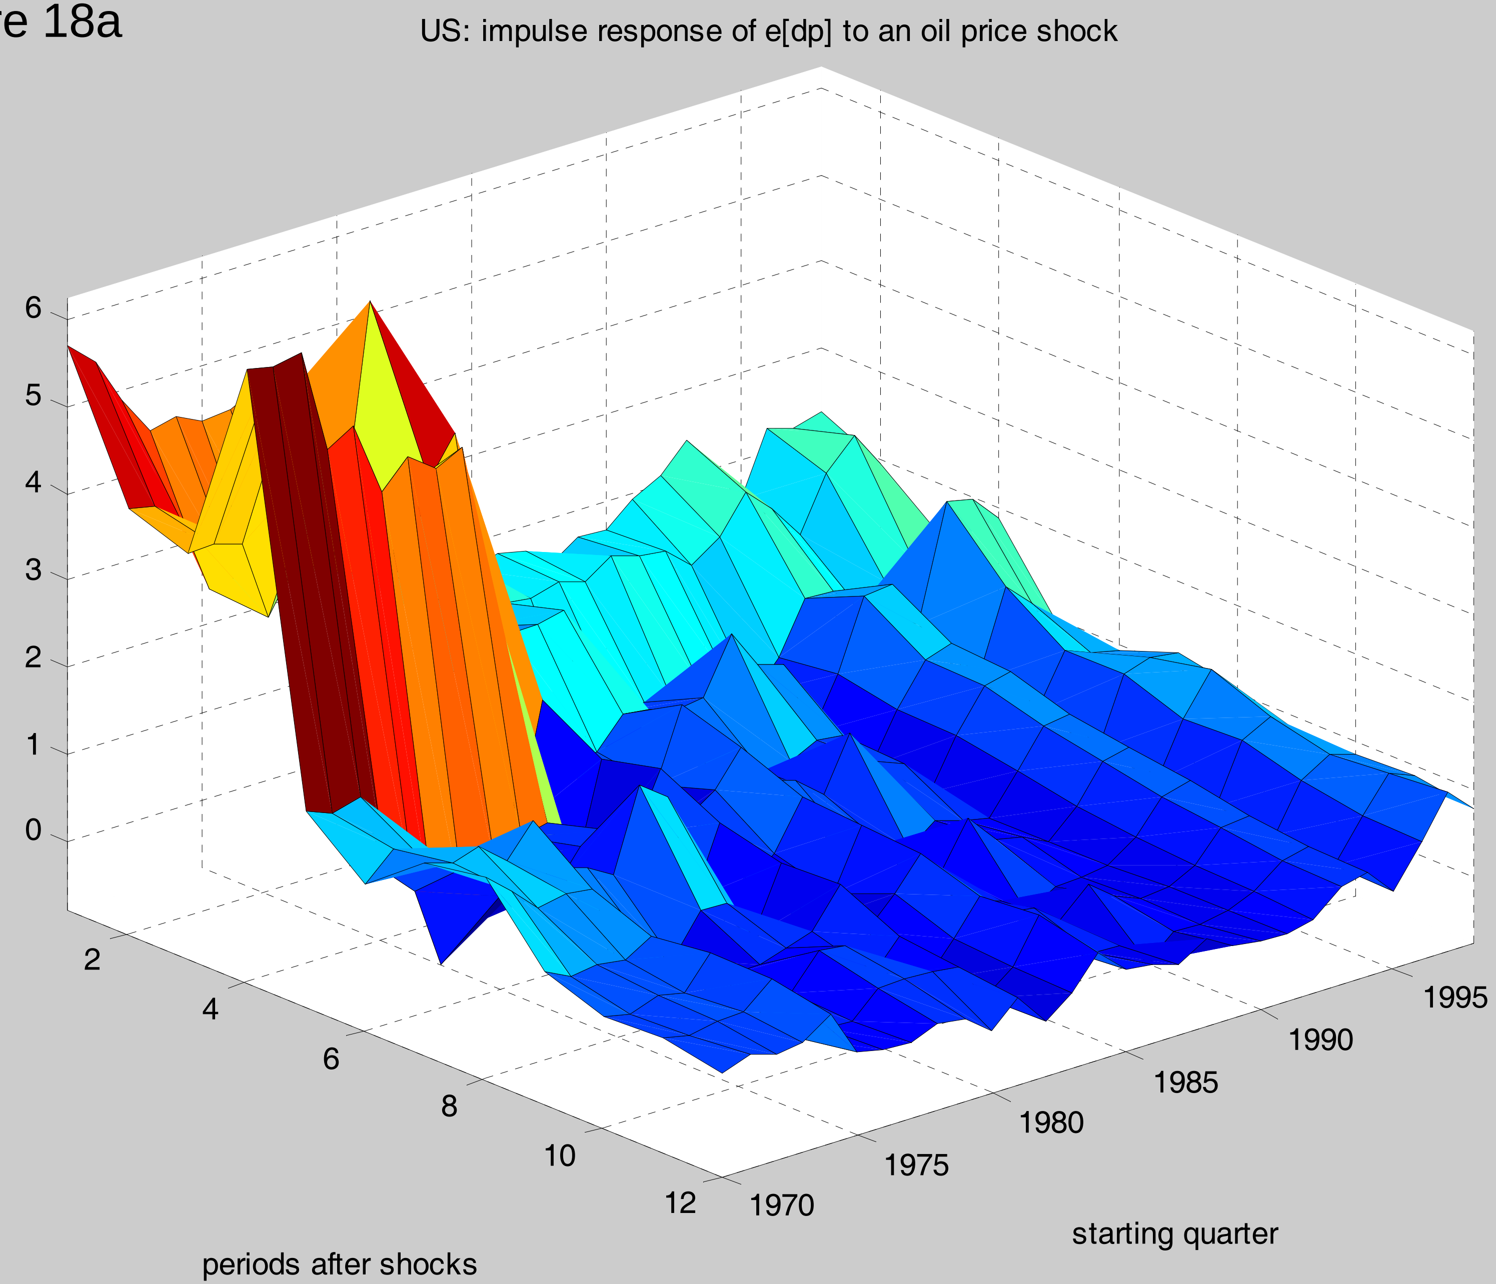
US: impulse response of e[dp] to an oil price shock
6
5
4
3
2
1
0
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 18a

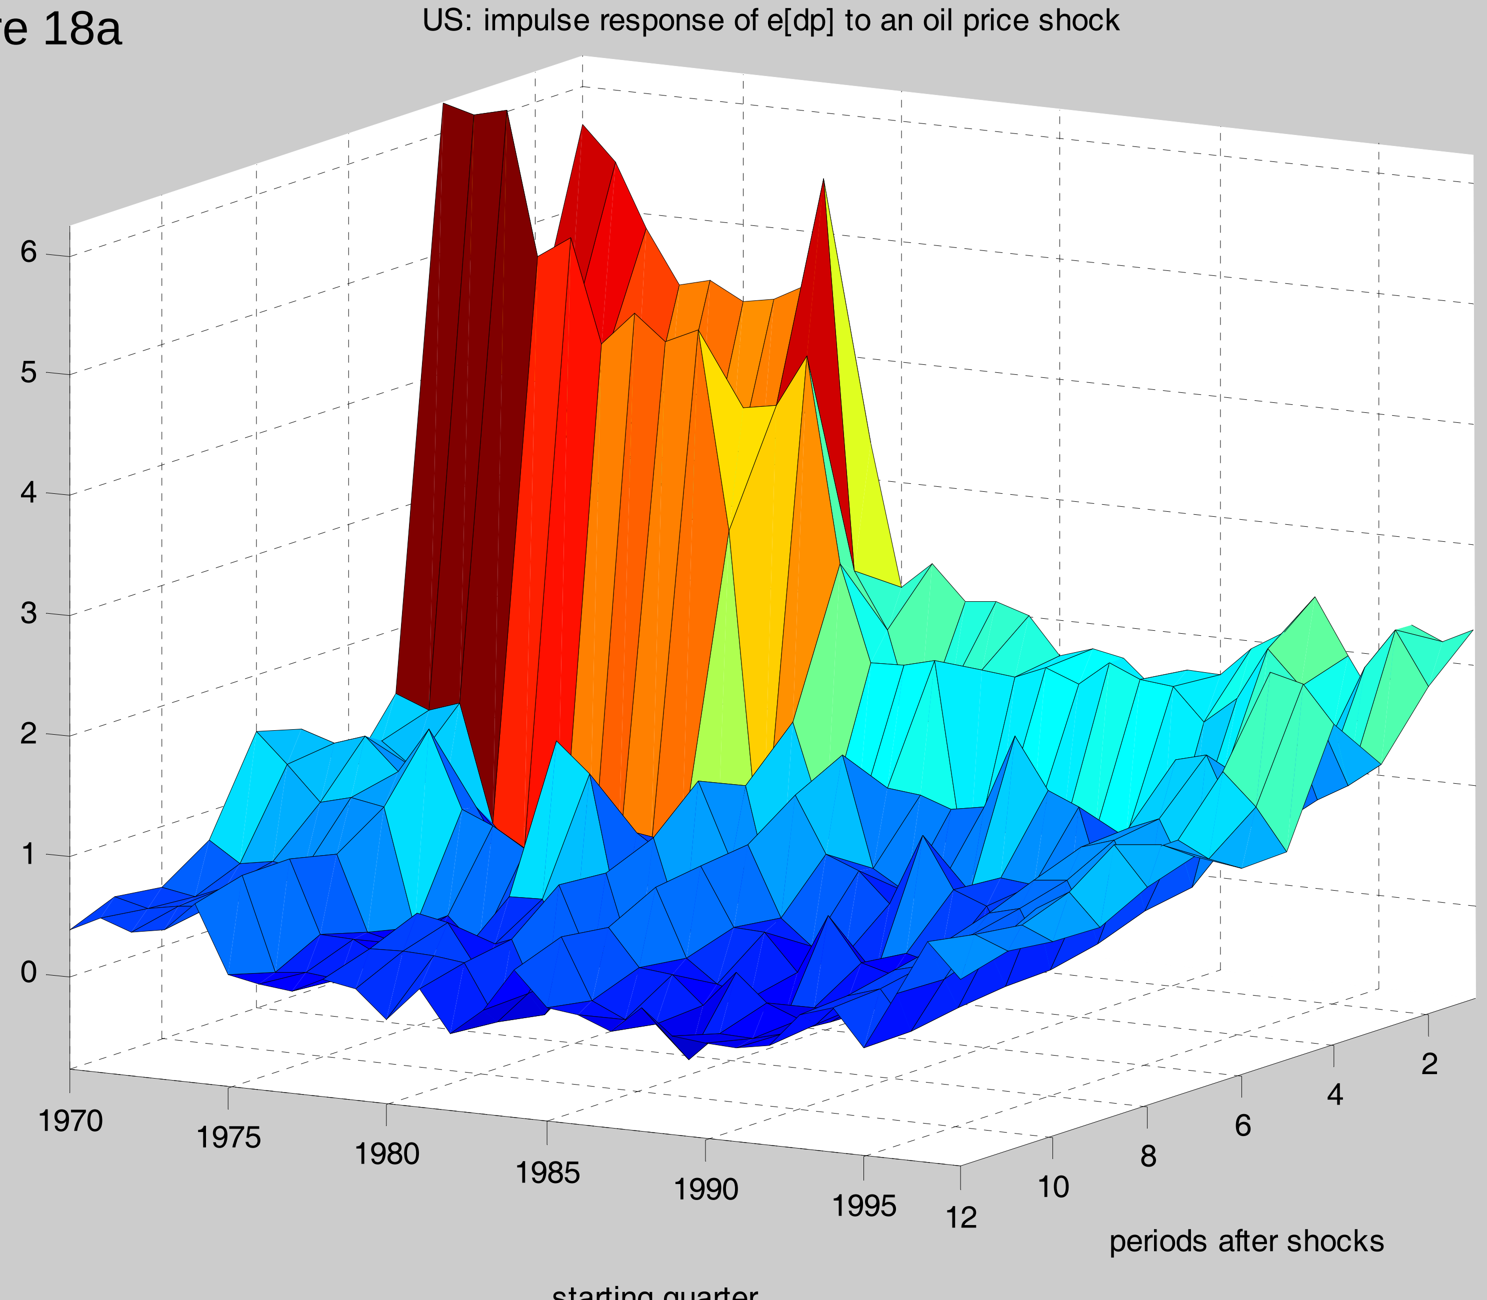
US: impulse response of e[dp] to an oil price shock
6
5
4
3
2
1
0
1970
1975
1980
1985
1990
1995
starting quarter
12
10
8
6
4
2
periods after shocks

Figure 18b

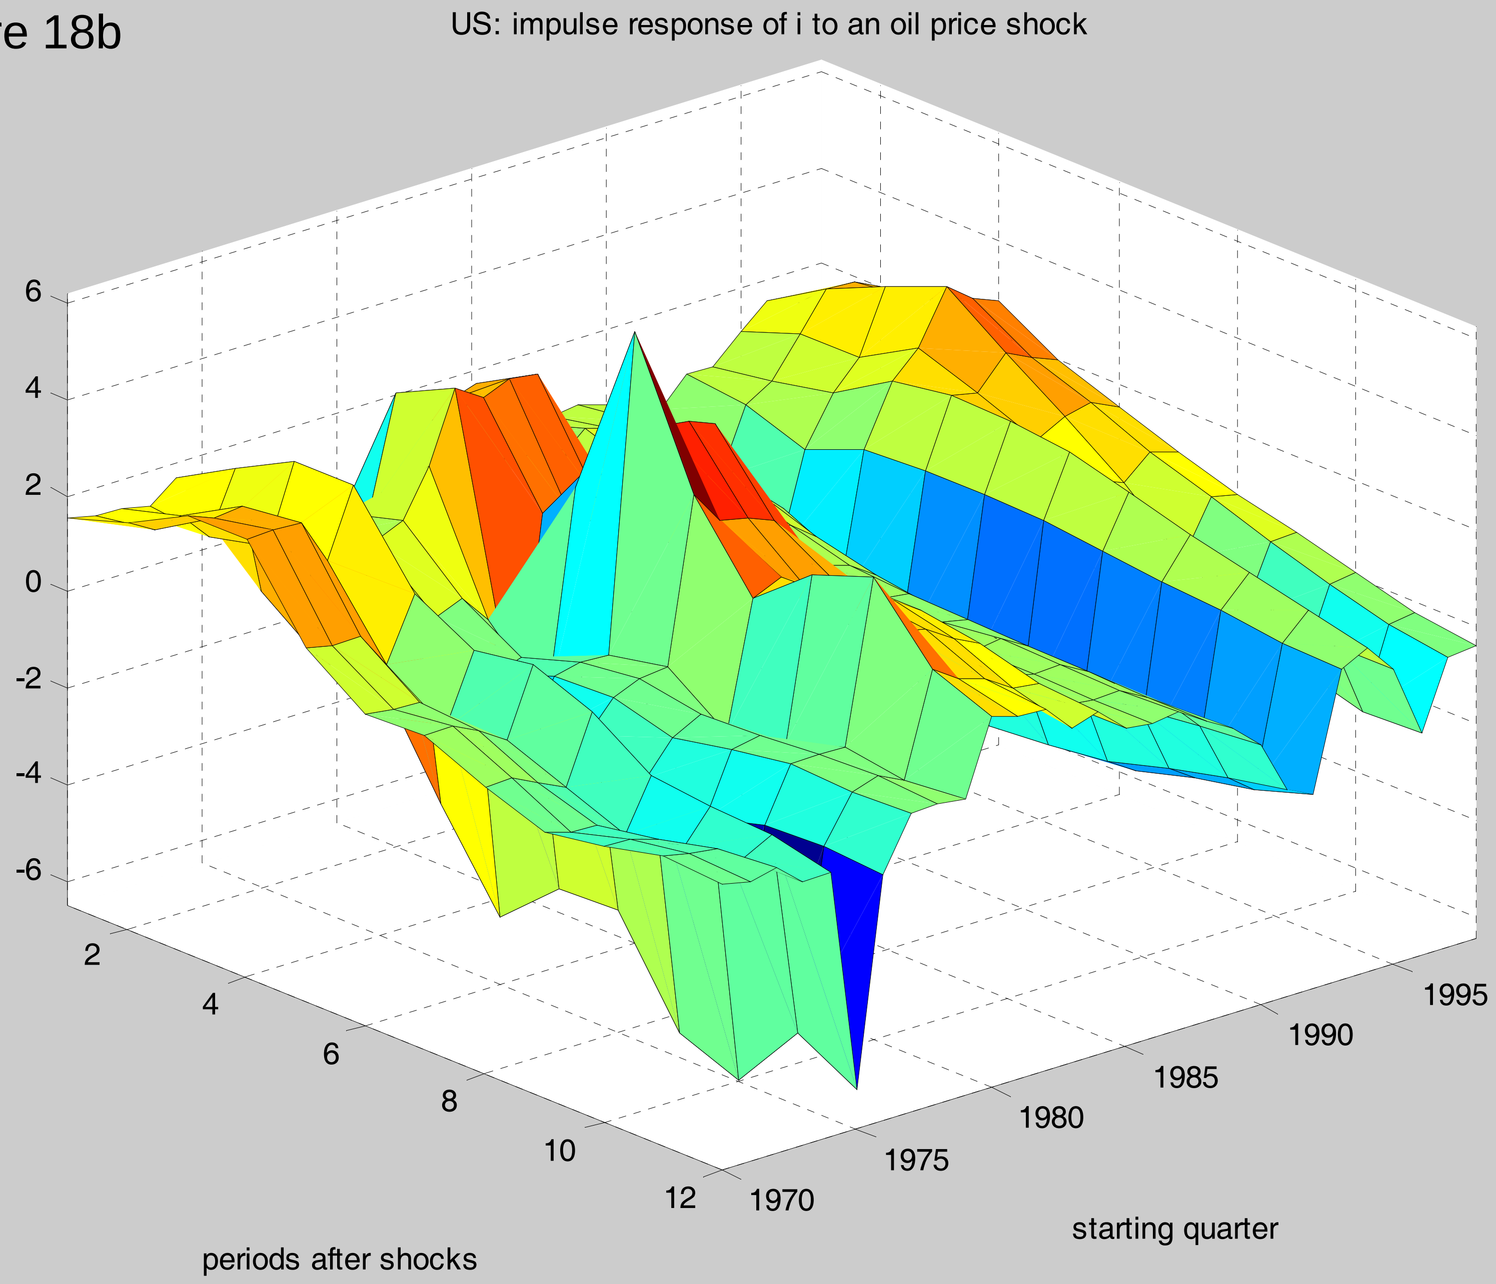
US: impulse response of i to an oil price shock
6
4
2
0
-2
-4
-6
2
4
6
8
10
12
periods after shocks
1970
1975
1980
1985
1990
1995
starting quarter

Figure 18c

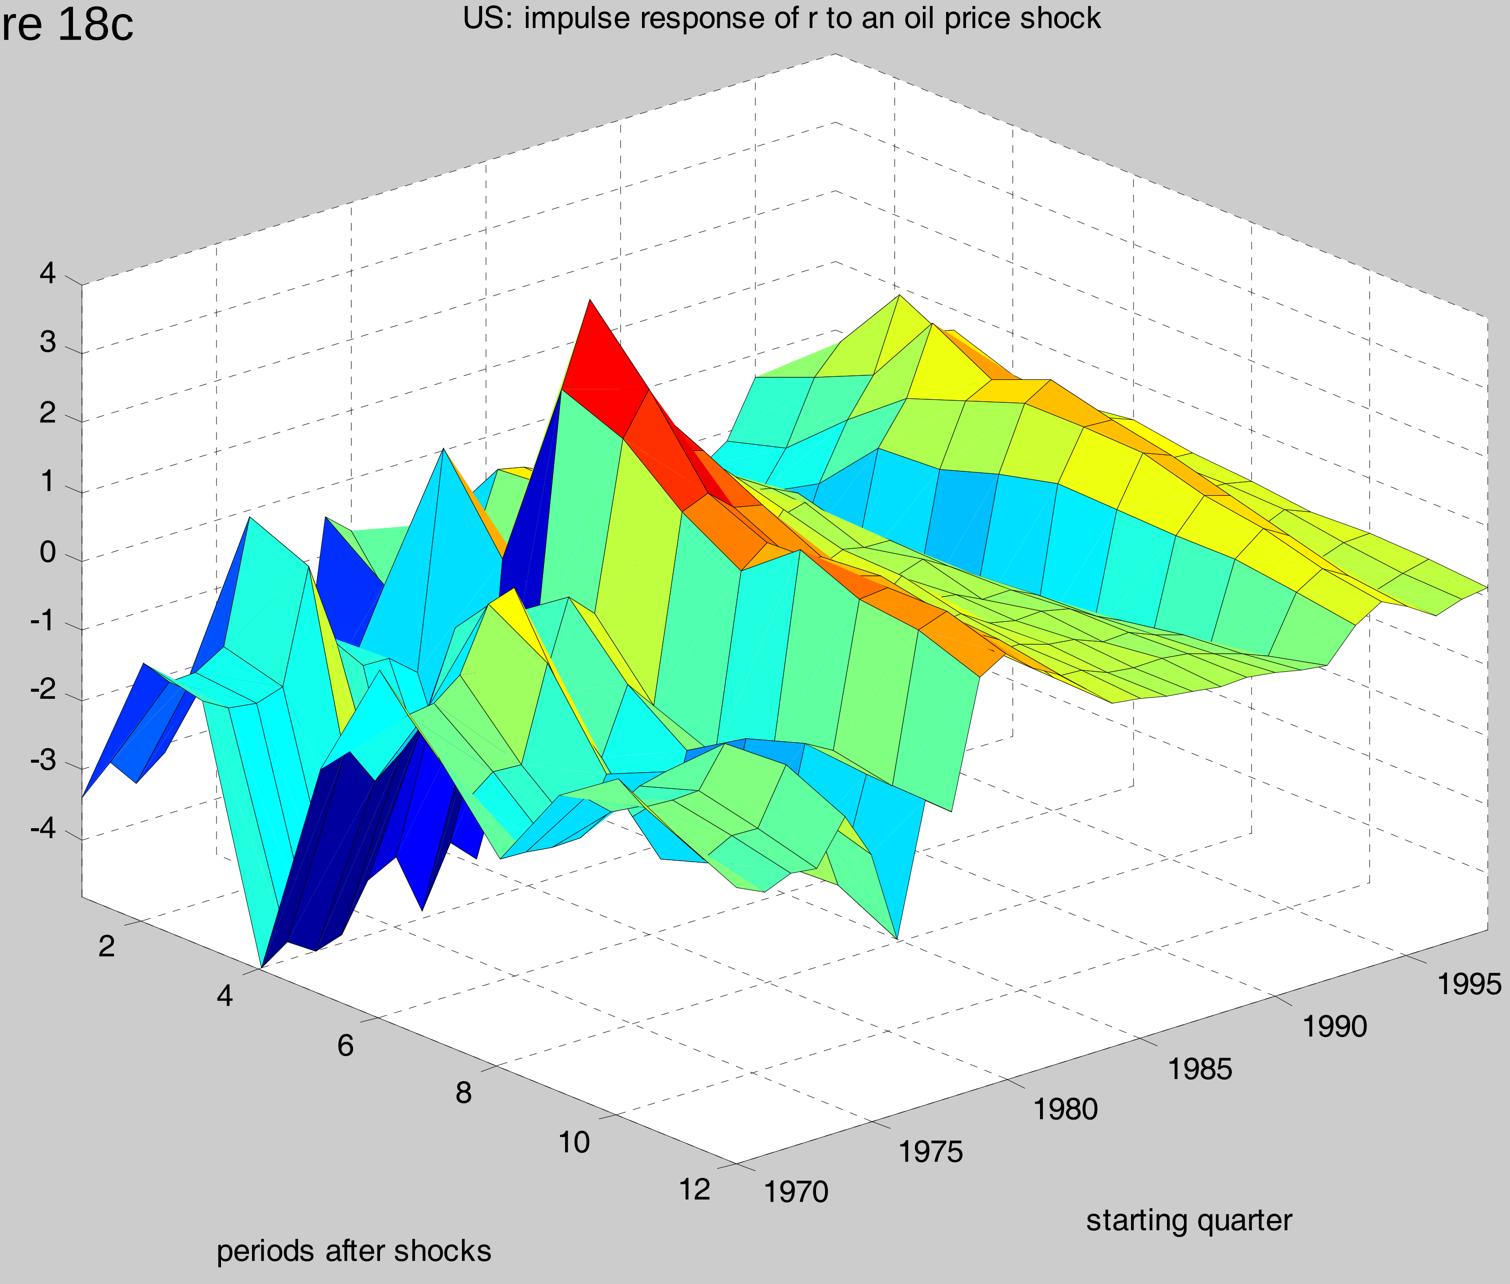
US: impulse response of r to an oil price shock
4
3
2
1
0
-1
-2
-3
-4
2
4
6
8
10
12
1970
1975
1980
1985
1990
1995
periods after shocks
starting quarter

Figure 19

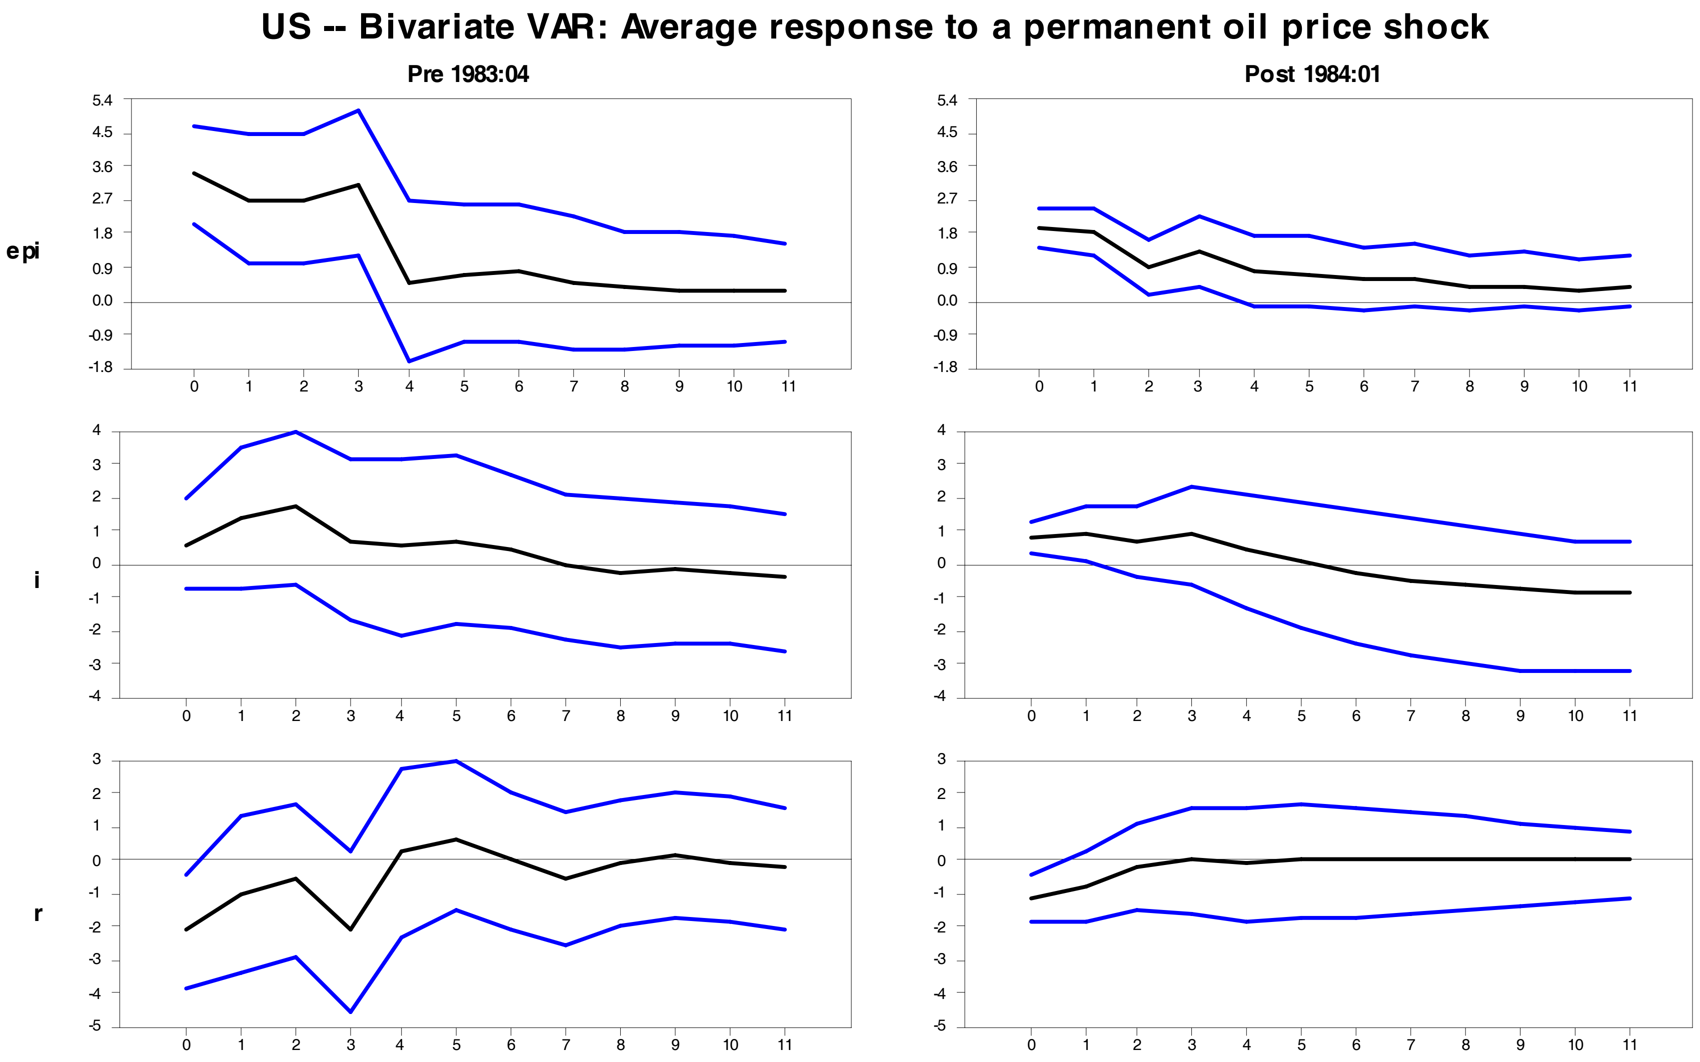

Bivariate VARs – Expected Inflation and Interest Rates

Figure 20: Changing Oil Shares and Policy Tradeoffs