A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries

Lutz Kilian
University of Michigan and CEPR

March 31, 2007

Abstract: A comparison of the effects of exogenous shocks to global crude oil production on seven major industrialized economies suggests a fair degree of similarity in the real growth responses. An exogenous oil supply disruption typically causes a temporary reduction in real GDP growth that is concentrated in the second year after the shock. Inflation responses are more varied. The median CPI inflation response peaks after three to four quarters. Exogenous oil supply disruptions need not generate sustained inflation or stagflation. Typical responses include a fall in the real wage, higher short-term interest rates and a depreciating currency with respect to the dollar. Despite many qualitative similarities, there is strong statistical evidence that the responses to exogenous oil supply disruptions differ across G7 countries. For suitable subsets of countries, homogeneity cannot be ruled out. A counterfactual historical exercise suggests that the evolution of CPI inflation in the G7 countries would have been similar overall to the actual path even in the absence of exogenous shocks to oil production, consistent with a monetary explanation of the inflation of the 1970s. There is no evidence that the 1973/74 and 2002/03 oil supply shocks had a substantial impact on real growth in any G7 country, whereas the 1978/79, 1980, and 1990/91 shocks contributed to lower growth in at least some G7 countries.

Key words: Oil supply; growth; inflation; exogeneity; dynamic effects; counterfactual.

JEL: E31, E32, Q43

Acknowledgements: I thank the editor and two anonymous referees for many helpful comments. I am indebted to Gary Saxonhouse for providing the Japanese real GDP series used in this study and to Thomas Knetsch for providing the French real GDP series. This paper also has benefited from discussions with Elena Pesavento, Ana María Herrera and Jim Hamilton.

1. Introduction

There is a large empirical literature on the effects of exogenous variation in the supply of crude oil on macroeconomic and industry aggregates in oil-importing countries.[1] Much of that literature has focused on the United States. The experience of other countries has remained understudied, notwithstanding some notable early exceptions such as Bohi (1989) who informally compared the inflation and real output experiences of Japan, Germany, the United Kingdom and the United States during the oil supply shocks of the 1970s. More recent examples of international comparisons of the effects of exogenous oil supply shocks include Mork et al. (1994) and Jiménez-Rodríguez and Sánchez (2005). In this paper I focus on the set of G7 countries comprised of the United States, Canada, Japan, France, Germany, Italy and the United Kingdom. The objective of the paper is to assess systematically the differences as well as similarities in the response of the G7 economies to exogenous oil supply shocks. For example, are there common patterns in the response of macroeconomic aggregates to exogenous oil supply disruptions? Which G7 economies have proved most resilient to exogenous oil supply shocks and which have been affected the most? Are these differences associated with different interest rate, real wage or exchange rate responses to these shocks? Are there systematic differences between G7 countries that produce oil and those that do not? Are there systematic differences between European and non-European G7 countries?

The analysis in this paper departs from the existing literature along several dimensions. First, the paper exploits recent methodological advances in measuring oil supply shocks that are exogenous with respect to global macroeconomic conditions. Specifically, I utilize a direct measure of these exogenous oil supply shocks proposed by Kilian (2006) rather than measures based on oil prices. Second, my analysis is based on a recently proposed approach to quantifying the dynamic effects of exogenous oil supply shocks that avoids some of the conceptual and econometric difficulties with earlier analyses. Third, whereas much of the existing literature has focused on the effects of exogenous oil supply shocks on real output, I put equal emphasis on the responses of output and inflation, with special emphasis on quantifying possible stagflationary effects of oil supply shocks. My analysis also distinguishes between consumer prices and the

[1] See, for example, Balke, Brown and Yücel (2002), Barsky and Kilian (2002, 2004), Bernanke, Gertler and Watson (1997), Bohi (1989, 1991), Burbidge and Harrison (1984), Bruno and Sachs (1982), Darby (1982), Davis and Haltiwanger (2001), Gisser and Goodwin (1986), Hamilton (1996, 2003), Hamilton and Herrera (2004), Hooker (1996, 2002), Hoover and Perez (1994), Lee and Ni (2002), Lee, Ni and Ratti (1995), Mork (1989), Mork, Olsen and Mysen (1994), Rasche and Tatom (1981), Shapiro and Watson (1988), Tatom (1988).

implicit GDP deflator. In addition, I study the effects on other key macroeconomic variables such as short-term interest rates, real wages and exchange rates, in an effort to shed light on the transmission of exogenous oil supply shocks. Finally, my data set extends to 2004.III and thus includes additional exogenous oil supply disruptions not covered by previous studies.

The regression analysis in this paper is designed to shed light on the dynamic effects of exogenous oil supply shocks on output, inflation and other macroeconomic aggregates in the G7 countries. In addition to conducting impulse response analysis, I use counterfactual simulations to assess how output and inflation would have evolved in these countries in the absence of exogenous oil supply shocks. The answer to these questions is relevant both for our understanding of the historical record and for the design of policy responses to future exogenous oil supply shocks. Specifically, I am interested in answers to the following questions: Do exogenous oil supply shocks generate sustained inflation? How does real GDP growth respond? How long does it take for the responses to set in? Are there long-run effects on the level of output and prices? Are the responses similar across countries? Are exogenous oil production shortfalls stagflationary? Are differences in these responses associated with differences in the responses of real wages, short-term interest rates and nominal dollar exchange rates? To what extent can the poor macroeconomic performance of G7 countries during specific historical episodes be attributed to exogenous shocks to oil production?[2]

My starting point is a linear regression model of the relationship between exogenous oil supply fluctuations and macroeconomic aggregates. Based on quarterly data for 1971.I-2004.III, I find a fair degree of similarity in the qualitative features of the estimated real GDP growth responses. The evidence suggests that an exogenous oil supply disruption causes a temporary reduction in real GDP growth that is concentrated in quarters 4 through 8 after the shock. CPI inflation responses appear more varied. The median CPI inflation response peaks three to four quarters after the shock. Exogenous oil supply disruptions do not necessarily generate sustained CPI inflation. As measured by real GDP and CPI data, with the exception of Germany, there is no statistically significant evidence that exogenous oil supply disruptions historically have caused stagflation. This result contradicts the popular notion that exogenous oil production shortfalls are by necessity stagflationary.

[2] See Barsky and Kilian (2004) for a review of the channels by which oil supply shocks may affect real GDP and inflation.

Despite qualitative similarities, there is strong statistical evidence that the responses to exogenous oil supply disruptions differ across G7 countries. For suitable subsets of G7 countries homogeneity cannot be ruled out. As measured by cumulative inflation and real growth responses, some countries fared well when faced with exogenous oil supply disruptions, whereas others did not. Only for Germany is there a statistically significant increase in the response of the CPI level after three years, and only for the U.K. and for Canada is there a statistically significant reduction in the level of real GDP at the 5% level. In most countries exogenous oil supply disruptions cause at least a temporary decline in real wages, a depreciation of the local currency against the dollar and a rise in short-term interest rates. There is no evidence that differences in these responses alone can explain the differences across countries in the inflation and output responses, although the results for Japan point to an important role for the interest rate response.

A counterfactual historical exercise reveals that the evolution of CPI inflation (and to a lesser extent of real GDP growth) in the G7 countries would have been quite similar overall to the actual path even in the absence of exogenous shocks to oil production. There are some differences, however, when it comes to interpreting specific historical episodes. For example, in all G7 countries the 1990/91 oil supply shock caused by the Persian Gulf War contributed to somewhat reduced real growth, albeit with a considerable delay. The oil shocks of 1978/79 and 1980 also left a mark in the data of some G7 countries. In contrast, the 1973/74 oil supply shock had hardly any impact on G7 real growth. Similarly, the effect of the 2002/03 oil supply shocks was negligible for all G7 countries.

The remainder of the paper is organized as follows: Section 2 reviews recent methodological advances in assessing the effects of exogenous oil supply shocks and describes and motivates the regression approach taken in this paper. In section 3, I describe the data used in the regression analysis, including the measure of exogenous oil supply shocks, and formally test some of the restrictions imposed on the baseline model. Section 4 contains the inflation and real GDP response estimates obtained from the baseline regression model. I also quantify the stagflationary effects of exogenous oil supply disruptions. In section 5, I explore an alternative less restrictive regression approach, compare the estimated responses to the baseline results, and provide statistical evidence in favor of the baseline model. Section 6 assesses the cumulative effects of exogenous oil supply shock on each country's real growth and CPI inflation rate for each of the major exogenous oil supply shock episodes since 1971. Section 7 extends the analysis to include responses for real wages, short-term interest rates, nominal dollar exchange rates, and the implicit GDP deflator. Section 8 assesses the similarities and differences of the response patterns in the G7 countries. In Section 9, I estimate the response of the real price of oil to exogenous oil supply shocks. Section 10 addresses some limitations of the analysis in this paper.

2. Alternative Approaches to Identifying the Effects of Exogenous Oil Supply Disruptions

2.1. Oil-Price Based Measures of Exogenous Oil Supply Disruptions

A common feature of all methodologies designed to learn about the dynamic effects of exogenous oil supply shocks is that they involve a projection of macroeconomic aggregates on some measure of the exogenous oil supply shock. Early studies sometimes treated the price of oil (or positive changes in that price) as the measure of the exogenous oil supply disruption. It is widely understood today that at least since 1973 the price of oil has been endogenous to global macroeconomic conditions and cannot be treated as exogenous (see Rotemberg and Woodford 1996; Barsky and Kilian 2002, 2004; Hamilton 2003).

Some studies have noted that at least the major oil price fluctuations in the 1970s and 1980s were arguably driven by exogenous political events in the Middle East (see, e.g., Shapiro and Watson 1988). This insight was subsequently formalized by Hamilton (1996, 2003), who proposed a statistical measure of the net oil price increase relative to the recent past designed to capture those major oil price increases presumably caused by exogenous political events. That measure also produces a time series very similar to fitted values from more sophisticated nonlinear models of the oil price (see, e.g., Lee, Ni and Ratti 1995, Hamilton 2003).

Such measures are problematic, however. Although three of the largest oil price increases since the early 1970s occurred near periods of large exogenous shocks to oil production, not all exogenous oil supply shocks have been associated with net oil price increases. For example, the 2002/03 twin shocks associated with civil unrest in Venezuela and the Iraq War were not associated with a net oil price increase in real terms. In addition, it is known that net oil price increases may arise even in the absence of exogenous political events in the Middle East as a result of strong demand for industrial commodities at large. Thus, net oil price increases are not an appropriate measure of exogenous oil supply shocks (see Kilian 2006).

The same reasoning casts doubt on the view that the innovations to the oil price series (or for that matter the innovations to the net oil price increase) represent exogenous oil supply shocks. In particular, there is reason to believe that oil price innovations in vector autoregressive models may also reflect shifts in the demand for crude oil, as suggested by sharp increases in industrial commodity prices during times of major global economic expansions.

2.2. Oil-Production Based Measures of Exogenous Oil Supply Disruptions

A different strand of the literature has instead used as identifying information the observable changes in the production levels of oil-producing countries that are subject to exogenous political shocks. For example, Hamilton (2003) proposes to use the drop in observable oil production following an exogenous event as a measure of the magnitude of the exogenous shock to the supply of oil. Typical examples of the events Hamilton considers exogenous are the Arab oil embargo of 1973/74, the Iranian Revolution of 1978/79, the Iran-Iraq War of 1980-88 and the Persian Gulf War of 1990/91. In each case, Hamilton focuses on the oil-producing countries directly involved in the event in question. Given the starting date of the event, he uses the level of oil production in the month prior to this date as a benchmark. He then compares that level to the level of production at some subsequent date. The difference in physical production levels over the period in question is expressed as a share of the average world oil production in the year, in which the exogenous event started (see Hamilton 2003, p. 390). The resulting "production shortfall" is treated as a measure of the magnitude of the shock that occurred in the first quarter of the exogenous event. This approach is in essence a quantitative version of the dummy variable approach used by Hoover and Perez (1994) to model oil supply shocks.

An alternative measure of exogenous oil supply shocks proposed by Kilian (2006) is based on crude oil production data for OPEC countries and non-OPEC countries that are available from the U.S. Department of Energy. This measure is based on the observation that any attempt to identify the timing and magnitude of exogenous production shortfalls requires explicit assumptions about the counterfactual path of crude oil production in the absence of the exogenous event. Using suitable assumptions about the counterfactual path of oil production based on the evolution of oil production in other oil-producing countries, the exogenous production shortfall is constructed as the difference between the actual path of crude oil production and the counterfactual path. The change over time in this exogenous production shortfall series (aggregated across OPEC countries and expressed as a percent share of world oil production) provides a natural measure of the exogenous oil supply shock.

This alternative method of quantifying exogenous production shortfalls has five distinct advantages compared to the conventional approach based on quantitative dummy variables: (1) It does not impose the assumption that the level of oil production would never have changed in the absence of the exogenous political event. (2) It allows the response of oil production to the exogenous event to be immediate or delayed. (3) It allows the response to be long-lasting. (4) It allows the response to be time-varying. (5) It allows the response to an exogenous political event to be negative or positive, possibly changing sign over time.

The analysis in this paper will utilize the baseline time series of exogenous oil supply shocks developed in Kilian (2006). A more detailed description of the derivation of this series is provided in section 3.1. For a comparison of the implications of this measure to the quantitative dummy approach the reader is referred to this earlier paper. The main findings reported in the empirical sections below are robust to the alterations in the baseline counterfactual discussed in Kilian (2006).

2.3 Inference based on Production Based Measures of Exogenous Oil Supply Disruptions

Given a production-based measure of the exogenous oil supply shock, there are two alternative approaches to quantifying the dynamic effects of exogenous fluctuations in oil production on macroeconomic aggregates. The first approach due to Hamilton (2003) is to use lags of the exogenous oil supply shock as identifying instruments in regressions that relate the macroeconomic aggregate of interest to past oil price changes and past values of the macroeconomic aggregates. While this approach is quite appealing, Kilian (2006) documents that such instrumental variable regressions in practice tend to suffer from a weak instrument problem that renders point estimates unreliable and standard inference misleading (see Stock, Wright and Yogo 2002). In response to this problem, Kilian (2006) proposed an alternative regression approach that will form the basis of the analysis in this paper.

This alternative approach follows the convention in the literature of treating changes to oil production induced by political events such as wars or revolutions in the Middle East as exogenous with respect to the macroeconomic aggregates in G7 economies. Specifically, it treats the oil supply shock series as strictly exogenous in the sense that there is no feedback from current or lagged values of the dependent variable to the exogenous variable. Let x_t denote the date t observation of the exogenous oil supply shock series, Δy_t the corresponding percent growth rate in real GDP and π_t the percent change in the consumer price index. The objects of interest are the impulse responses $\partial \Delta y_{t+i} / \partial x_t$ and $\partial \pi_{t+i} / \partial x_t$, $i = 1, 2, 3...$ For each country, the first-order effect of a given increase in x_t on Δy_{t+i} and π_{t+i}, respectively, may be computed based on the fitted value of the linear ordinary least squares (OLS) regressions:

(1) $$\Delta y_t = \alpha + \sum_{i=1}^{4} \beta_i \Delta y_{t-i} + \sum_{j=0}^{8} \gamma_j x_{t-j} + u_t$$

and

(2) $$\pi_t = \delta + \sum_{i=1}^{4} \lambda_i \pi_{t-i} + \sum_{j=0}^{8} \eta_j x_{t-j} + v_t,$$

where the error terms u_t and v_t are serially uncorrelated, given the inclusion of four lags of the dependent variable and eight lags of the exogenous oil supply shock.[3] Provided that the exogenous oil supply shock regressors are not correlated with any omitted exogenous variables, the implied impulse responses will measure the causal effects of the exogenous variations in oil supply. Level responses for real GDP and the level of consumer prices may be obtained by cumulating the estimated impulse responses. Confidence intervals for these impulse responses may be constructed by bootstrap methods or by drawing from the asymptotic normal distribution of the slope parameters and simulating the standard errors of the impulse response estimators.

The estimated responses provide a measure of the expected response of macroeconomic aggregates to exogenous oil production shortfalls based on historical data. They represent consistent estimates of the causal effects of a unit change in the exogenous oil supply shock measure. It is important to keep in mind that these impulse response estimates are reduced-form in that they will capture the average effect of country-specific endogenous policy responses and other propagation mechanisms that prevailed at the time of the exogenous oil supply shocks. In this sense, the estimated responses do not reflect inevitable facts, but empirical regularities. They are designed to quantify the historical tendency of the G7 economies to perform poorly as a consequence of exogenous oil supply disruptions.

This direct method of estimating the effects of exogenous oil supply shocks may be motivated as follows. Consider a dynamic simultaneous equation model for x_t and z_t, where x_t

[3] This lag structure corresponds to the assumptions used by Hamilton (2003) and Kilian (2005).

denotes the exogenous variable and z_t denotes the macroeconomic aggregate of interest. It is understood that higher-dimensional vector processes including additional endogenous variables may always be marginalized and expressed in terms of a bivariate vector process for $[x_t \; z_t]'$. Suppose that this process can be represented as a linear vector autoregression:

$$x_t = \sum_i a_{1i} x_{t-i} + \sum_j b_{1j} z_{t-j} + e_{1t}$$

$$z_t = \sum_i a_{2i} x_{t-i} + \sum_j b_{2j} z_{t-j} + e_{2t} .$$

with serially uncorrelated reduced-form innovations e_{1t} and e_{2t}. The intercept terms have been suppressed for simplicity. Consider the corresponding structural form of this process:

$$x_t = \theta z_t + \sum_i \alpha_{1i} x_{t-i} + \sum_j \beta_{1j} z_{t-j} + \varepsilon_{1t} \tag{3}$$

$$z_t = \gamma x_t + \sum_i \alpha_{2i} x_{t-i} + \sum_j \beta_{2j} z_{t-j} + \varepsilon_{2t} . \tag{4}$$

where the structural innovations ε_{1t} and ε_{2t} are contemporaneously and serially uncorrelated. We say that x_t is pre-determined if $\theta = 0$. If $\beta_{1j} = 0 \, \forall j$ in addition, then x_t is strictly exogenous with respect to z_t (see, e.g., Cooley and LeRoy 1985).[4] Under strict exogeneity, there is no current or lagged feedback from z_t to x_t, and we can consistently estimate the effect of a change in x_t on z_t based on equation (4) alone.[5]

Equations (1) and (2) above are specific examples of equation (4). Compared to working with structural VAR models, the single-equation regression approach taken in this paper has two main advantages. First, equations (1) and (2) can be constructed without taking a stand on the structural VAR representation of $[x_t, \Delta y_t, \pi_t]'$, making the results invariant to alternative identifying assumptions. Second, equations (1) and (2) are relatively parsimonious, facilitating

[4] Pre-determinedness and strict exogeneity in this framework correspond to the notion of weak and strong exogeneity, respectively, in the parlance of Engle, Hendry and Richard (1983).

[5] The strict exogeneity of x_t also ensures that any higher-dimensional process can be reduced to the structural representation (3) and (4), since equation (3) is invariant to marginalization under exogeneity, and ε_{1t} retains its original structural interpretation. Of course, the innovation ε_{2t} after marginalization will be a composite of the structural innovations in the underlying higher-dimensional VAR model and will not in general have the same structural interpretation as the corresponding innovation in the original model.

their estimation on short samples.

The regression models (1) and (2) will be used as the baseline model for this paper. Statistical evidence in support of this model can be found in section 3.3. In section 5, an alternative, less restrictive regression model will be explored and compared with the baseline model.

3. Data

3.1. Exogenous oil supply shocks

Table 1 lists the major political events that for the purpose of this study will be considered exogenous with respect to inflation and real output growth in the G7 countries. Barsky and Kilian (2002) observe that there is reason to be skeptical about the exogeneity of the 1973/74 oil embargo decision (see Hamilton (2003) for a different view). Here I will follow the convention in the literature in treating this event as exogenous.

The measure of exogenous oil supply shocks used in the empirical analysis is based on monthly data on crude oil production by country provided by the U.S. Department of Energy, starting in January 1973. The approach can be summarized as follows: Any attempt to identify the timing and magnitude of an exogenous production shortfall requires explicit assumptions about the counterfactual path of oil production in the absence of the exogenous event in question. For a given exogenous event, say a war in some OPEC country, the first step is to identify a group of oil producing countries that was subject to the same global macroeconomic conditions and same economic incentives as the war-stricken country, but whose production was not affected by the war. Which countries belong in the benchmark group must be decided on a case-by-case basis drawing on historical accounts and industry sources. The counterfactual production level for the country affected by the war is generated by extrapolating its pre-war production level based on the average growth rate of production in the benchmark countries. Table 2 summarizes the assumptions about the benchmark countries used in constructing the counterfactual.[6]

[6] In constructing the counterfactual, I distinguish between changes in crude oil production that operate directly in response to the exogenous event in question (such as the temporary production increases in Saudi Arabia during the Persian Gulf War or the Iraqi production increase following the Iranian revolution) on the one hand and endogenous changes in crude oil production in response to higher oil prices on the other hand. The latter are treated as part of the propagation mechanism (and hence are captured by the impulse response estimates), whereas the former are explicitly incorporated into the counterfactual.

In the second step the exogenous production shortfall for each OPEC country is obtained by constructing the difference between actual crude oil production in that country and the counterfactual level of crude oil production at each point in time, resulting in one time series for each of the nine OPEC countries in Table 2 representing that country's exogenous shortfall of crude oil production. The aggregate measure of the exogenous crude oil production shortfall for all of OPEC is obtained by summing these time series. Finally, the exogenous OPEC oil supply shock is constructed by normalizing that aggregate shortfall series as a percent of world crude oil production and taking first differences. This approach amounts to treating the exogenous OPEC production shortfall series as a random walk and is consistent with the lack of serial correlation in the differences.

Figure 1 shows the exogenous OPEC oil supply shock series derived on the basis of Table 2, suitably aggregated to quarterly frequency. Although the Department of Energy crude oil production data only begin in January of 1973, it is possible to extend the exogenous oil supply shock series back to January 1971, as shown in Figure 1. We know from historical records that there were no exogenous events in 1971.I-1972.IV that could have triggered an exogenous oil production shortfall, so the shortfall must have been zero.[7]

The major oil dates have been imposed in Figure 1 as vertical lines. As expected, the most important spikes in the series occur near those dates. Unlike existing quantitative dummy measures of exogenous oil supply shocks, the series in Figure 1 contains negative as well as positive shocks. This feature is essential in that many historical production shortfalls triggered by exogenous events have been at least partially reversed. Thus the corresponding change in the exogenous production shortfall should be characterized by an initial negative spike followed by one or more subsequent positive spikes. Only if the exogenous OPEC production shortfall were permanent would there be no positive realizations of the oil supply shock measure. As Figure 1 shows, virtually all exogenous shortfalls were followed by at least a partial reversal. Positive values of the oil shock measure also reflect the fact that wars in the Middle East may actually cause higher oil production over time, when the parties involved resort to oil exports to finance

[7] While earlier exogenous events in OPEC countries could have had effects that extended into the 1971-1972 period in principle, that outcome does not seem plausible, as the only candidate would be the 1967 Six-Day War and the associated Arab oil embargo. Both events were short-lived and their effects on oil production were long over before 1971.

the war. Another key difference is that this new measure allows for repeated exogenous oil supply shocks of the same sign. A good example is the period following the Persian Gulf War.

The high volatility of the shock measure between 1980 and 1988 reflects the varying success with which Iran and Iraq were able to increase their crude oil exports to finance the Iran-Iraq War. In contrast, the volatility of the shock measure in 1998-2002 is mainly due to the uneven enforcement of U.N. sanctions on Iraq that resulted in highly volatile Iraqi crude oil production after 1998, whereas Iraqi production had been virtually flat in the first years after the Persian Gulf War. For further discussion of the derivation of the series shown in Figure 1 the reader is referred to Kilian (2006).[8]

3.2. Real GDP and CPI Inflation

The regression analysis is based on seasonally adjusted quarterly real GDP growth and CPI inflation rates for 1971.I-2004.III. The starting date of 1971.I ensures that the regression models can be used to construct counterfactuals starting as early as 1973.I (before the first major oil supply shock). Since the value of the exogenous oil supply shock is zero prior to 1973.IV, the possibility of a structural change in the transmission of exogenous oil supply shocks induced by the rise of OPEC in late 1973 has no effect on the regression analysis.

The CPI data are from the IFS with the exception of the U.S. CPI which was obtained from the Bureau of Labor Statistics. The German inflation series was obtained by combining West German CPI inflation data for 1971.I-1991.I with CPI inflation rates for the unified Germany starting in 1991.II. The GDP data for Canada, Germany, and Italy are from the IFS website. The French GDP data are from the Banque of France, as detailed in Knetsch (2005). Missing observations for the last one or two quarters of Italian, German and French data were replaced by the quarter-on-quarter seasonally adjusted real GDP growth rates provided on the *Eurostat* website. The Japanese real rates of growth have been computed based on the consistent quarterly real GDP series (in 1990 prices) provided by the Economic and Social Research Institute, Cabinet Office, Japan. That series ends in 2000.IV. The growth rates for the subsequent quarters are based on seasonally adjusted IFS data in 1995 prices. For the United Kingdom, I used the quarter-on-quarter real GDP growth series provided on the website of *U.K. National Statistics*, which unlike the IFS data is seasonally adjusted. The seasonally adjusted U.S. real GDP data were obtained from the Bureau of Economic Analysis. All rates have been annualized.

[8] The data are publicly available at http://www-personal.umich.edu/~lkilian/oilshock.txt.

3.3. Testing the Specification of the Baseline Model

A central proposition of this paper is that x_t is strictly exogenous in the sense described in section 2. Strict exogeneity involves two conditions: (1) Pre-determinedness of the oil supply shock series and (2) Granger-noncausality from macroeconomic aggregates to the oil supply shock series. While the second condition is testable, the first one is not.

Evidence of Granger causality from macroeconomic aggregates to the oil supply shock series would contradict the presumed exogeneity of that series. It is this Granger non-causality implication of strict exogeneity that is tested in Hamilton (1983), when he makes the case for the exogeneity of nominal oil prices in the period up to 1972. Table 3 reports asymptotic *p*-values for the same type of Granger-noncausality tests applied to x_t. The first column shows the test results for lagged real GDP growth. Table 3 shows that with the exception of Germany the null of no Granger causality cannot be rejected at the 5 percent level for any country. The remaining *p*-values range from 12 percent to 71 percent. The *p*-value for Germany is 1 percent. This apparent evidence of predictability is highly implausible not only a priori, but also in that one would not expect predictability to hold for one country alone. Indeed, the highly significant rejection for Germany appears to be driven by the coincidence of German re-unification and the aftermath of the Persian Gulf War, despite the inclusion of suitable dummy variables. After restricting the German sample to the pre-unification period, the significance at the 5 percent level vanishes. The results for lagged CPI inflation are shown in the second column. All *p*-values range from 13 percent to 86 percent. Again after limiting the German sample to the pre-unification period, the German *p*-value rises sharply. These test results suggest that the data are consistent with the absence of feedback from inflation and real growth to the oil shock series, as required by strict exogeneity.

Equation (4) in general allows for a contemporaneous effect from x_t on z_t. Table 4 formally tests the null that $\gamma = 0$ in equation (4). There is no evidence of a contemporaneous link from the exogenous oil supply shock to CPI inflation. No test result is significant at the 5 percent level. The *p*-values range from 10 percent to 63 percent, depending on the country. Similarly, for real GDP growth *p*-values generally range from 27 percent to 100 percent, suggesting that the contemporaneous regressor may be safely omitted from equations (1) and (2). The glaring exception in the latter case is the U.K. Again it seems implausible that a contemporaneous link would exist only for the U.K. and not for any other country. In the baseline results reported below I therefore impose $\gamma = 0$ for all regressions.[9]

4. Response Estimates by Country in the Baseline Model

This section contains a description of the baseline estimation results for each G7 country. All results are based on OLS estimates of models (1) and (2).

4.1 Real Growth and Inflation Responses in the Baseline Model

Given estimates of models (1) and (2) it is straightforward to assess the impact of an exogenous 1% reduction in global oil production by simulation. The dynamic responses to an exogenous 1 percent permanent reduction in crude oil production are shown in Figure 2 along with the corresponding responses of the levels of real GDP and of the CPI.

United States

The first column of Figure 2 shows a sharp reduction in U.S. real GDP growth between five and seven quarters after the shock that is significant at the 10% level. After cumulating the growth rate effects, I find a significant reduction of the level of real U.S.GDP in the second and third year following the shock at the 10% level. There also is a sharp increase in consumer price inflation three quarters after the shock (significant at the 10% level), but no significant effect on the level of consumer prices.

Italy

Italy as well experienced a significant temporary reduction in real GDP growth at the 10 percent significance level (between six and nine quarters after the shock), and again there is a significant reduction of the level of real GDP (between eight and twelve quarters after the shock) at the 10 percent significance level. At the same time, there is a significant spike in consumer price inflation four quarters after the shock, but no significant increase in the price level. These results are qualitatively similar to those for the United States.

France

The third column shows a significant reduction in French real GDP growth at the 10 percent level between seven and ten quarters after the exogenous oil supply disruption. The effect on the

[9] Not imposing that restriction would result in an implausibly large effect of the exogenous oil supply shock on U.K. real GDP growth on impact. At longer horizons the estimated responses are not very different. To save space these alternative estimates are not reported in the paper.

level of real GDP is significant at the 10 percent level starting eight quarters after the shock. Exogenous oil supply disruptions also cause a sustained and significant increase in consumer price inflation that extends from the second to the fifth quarter following the shock, before inflation reverts to zero. This hump in the inflation response implies a significant increase in the level of consumer prices at horizons of three to nine quarters.

Germany

The German data pose special problems due to German re-unification in 1990.IV. It is not obvious how to account for the effects of this event on German inflation and real growth, since it coincided with the effects of the Persian Gulf War. I model the transition dynamics following re-unification using dummies representing level shifts and slope shifts. Similar results were obtained under a number of alternative dummy specifications.

The baseline regressions for German real GDP growth account for a one-time increase in growth in 1991.I due to the re-definition of the German GDP series and allow for a change in the average growth rate after 1990.IV. The implied response of real growth is very similar to that obtained for data up to 1990.II only (excluding the re-unification period). In contrast, there is no evidence of a mean shift in inflation due to unification, but there is evidence of a temporary increase in inflation after 1990.IV. My baseline regression estimates include one level dummy each for 1991.I-IV intended to account for the inflationary pressures immediately following re-unification. The shape of the resulting impulse response estimate is very similar to that of the response estimated on data up to 1990.II only, but the magnitude of the responses is somewhat lower than based on the pre-1990.II data. Adding dummies for the 1993.I VAT increase and the1994.I mineral oil tax increase, which arguably were related to the cost of financing unification and which coincided with sharp spikes in inflation, does not affect the results appreciably.

The baseline results for Germany are shown in the fourth column of Figure 2. There is evidence of a temporary reduction in real GDP growth between the fourth and eight quarter, most of which is significant at the 10% level. Additional significant reductions in real growth at horizons ten and eleven are too small in magnitude to have much of an effect. The implied output response is significantly negative at the 10 percent level at horizons 7 through 12.

A clear anomaly of the German case is the implication of Figure 4 that an exogenous oil supply disruption causes a statistically significant *increase* in real GDP growth after one quarter (and a significant increase in real GDP for the first three quarters). It can be shown that this result is not driven by the fact that the 1990/91 oil crisis coincided with temporarily high German real growth driven by German re-unification. In fact, the same anomaly arises when I estimate the responses on German data up to 1990.II only (excluding the 1990.III oil crisis). Thus, this implausible result is likely to reflect a spurious sample correlation between economic expansions in Germany and exogenous oil supply disruptions. One would expect this correlation to vanish as the sample size increases. Apart from this anomaly, the German real growth and real GDP responses are qualitatively similar to those for Italy, France and the United States.

German CPI inflation shows a significant peak at horizons 2 through 4, similar to France, Italy and the United States, but that peak is followed by sustained increases in inflation from horizon 6 through 12 with additional peaks after 6 and 10 quarters. The level response shows a significant increase in consumer prices for all but the first quarter. This pattern is also unique in the sample of G7 countries. Again it can be shown that this result is robust to excluding the post 1990.II data.

United Kingdom

The analysis of the U.K. data is complicated by a possible structural break in the relationship between oil supply shocks and macroeconomic aggregates in the late 1970s, when U.K. domestic oil production gained momentum. Oil production data from the Department of Energy show that U.K. crude oil production, having been negligible until early 1975, reached 10,000 barrels a day in 1975.III, 500,000 barrels a day in 1977.I and finally 1 million barrels a day in 1978.II. The estimated inflation and output responses based on post 1975.III and post-1977.I data, however, are similar to the full-sample estimates, suggesting that the possibility of a structural break in the U.K. regression model can be ignored. I therefore focus on the full-sample results in Figure 2, although I will report additional results based on a possible 1978.II break in section 8.1.

The fifth column of Figure 2 indicates a significant reduction in U.K. real GDP growth two, four, five, seven and eight quarters after the shock. Although the growth response appears somewhat erratic, it is broadly consistent with earlier results for other countries. There is a significant and persistent reduction in the level of real GDP starting two quarters after the shock. CPI inflation spikes two quarters after the shock, followed by a secondary spike six quarters out. Unlike in the German data, the two spikes are separated by periods of falling consumer prices, however, preventing a significant change to the price level.

Canada

The Canadian real growth response in the sixth column of Figure 2 shows the familiar pattern of a temporary reduction in growth at horizons of 4 to 9 quarters, accompanied by a significant reduction in real GDP in the second and third year after the shock. CPI inflation shows a significant increase in the first two quarters, a pattern unique to Canada, and another significant peak after six quarters. With the exception of the first three quarters, the increase in the price level is not significant at the 10% level.

Japan

The last column of Figure 2 shows that apart from a drop six quarters after the shock, the response of Japanese real growth to exogenous oil supply disruptions is not significant nor is there a significant effect on the level of real GDP. The inflation response is continuously positive for the first seven quarters, but only the first quarter response is significantly positive at the 10% level. Given the apparent imprecision of the estimate, it is unclear whether there are multiple peaks or just one hump as in the case of France. For the remaining quarters the response is mostly negative. The implied increase in the level of consumer prices is marginally significant at the 10 percent level for the first six quarters.

4.2. Quantifying the Stagflationary Effects in the Baseline Model

A popular notion is that exogenous oil supply shocks can be held responsible for triggering stagflation, defined as the coincidence of rising price levels and falling output. For example, *The Economist* (November 27, 1999), in referring to the 1973 and 1979/80 oil supply shocks, writes that "both … shocks resulted in double-digit inflation and global recession". Figure 3 provides a formal assessment of this proposition based on the conditional co-movement measure proposed by Den Haan (2000). In Figure 3, this measure is applied to the responses of CPI inflation and real GDP growth to an exogenous oil supply disruption. Following Den Haan and Summer (2004, p. 1340), the plot shows conditional covariances rather than conditional correlations. This normalization facilitates a comparison of the statistic across horizons. The conditional covariance at horizon *h* is constructed as

$$C(h) = \Delta y_h^{imp} \ \pi_h^{imp}$$

where z_h^{imp} denotes the response of variable z_t at horizon *h* to a 1 percent exogenous oil supply disruption (see Den Haan 2000, p. 8). Stagflation in the form of rising prices and falling output will cause this measure to be negative. It is natural to conduct a one-sided test of the null of zero conditional covariance against the stagflationary alternative. Figure 3 plots 80 percent and 90 percent bootstrap confidence intervals along with the point estimates. The coverage rates are chosen such that the rejection probability in the upper tail corresponds to 10 percent and 5 percent, respectively. Notwithstanding some negative point estimates, there is large sampling uncertainty. Only for Germany is there some evidence of stagflation arising in response to an exogenous oil supply disruption. Thus, the evidence in Figure 3 confirms that exogenous oil supply disruptions do not necessarily cause stagflation. The stagflationary responses for Germany are concentrated in the second year after the shock. While these estimates are significant at the 10 percent level, much of that significance vanishes at the 5 percent level. Only the estimates at horizon 6 and beyond horizon 10 remain marginally significant.

5. An Alternative Regression Approach

All results so far were based on the baseline model represented by equations (1) and (2). An alternative model specification that allows consistent OLS estimation of the dynamic effects of exogenous oil supply disruptions is:

(5) $$\Delta y_t = \alpha + \sum_{i=1}^{12} \phi_i x_{t-i} + u_t$$

(6) $$\pi_t = \delta + \sum_{i=1}^{12} \psi_i x_{t-i} + v_t$$

where u_t and v_t are potentially serially correlated errors. That model shares with the baseline model the assumption that x_t is predetermined. It differs from the baseline model in that it does not impose the Granger non-causality restriction. Since that restriction is not rejected by the data, as shown earlier, this generalization is of limited value. More importantly, however, equation (5) also relaxes the assumption that the data are well approximated by a linear VAR representation. Thus, it is of some interest to compare the response estimates from the alternative regression models (5) and (6) to the estimates from the more tightly parameterized baseline models (1) and (2). In the alternative regression model the impulse response coefficients at horizon *h* correspond to ϕ_h and ψ_h, respectively. Thus, the number of lags is determined by the maximum horizon of the impulse response function, which is set to 12 quarters to maintain consistency with Figure 2. As in the baseline model there is no instantaneous feedback from x_t to inflation and real growth.

In conducting inference on these response estimates, the presence of serial correlation in the error term requires the use of suitable robust standard errors and bootstrap methods.

5.1. Response Estimates in the Alternative Model

Figure 4 shows the response estimates obtained from the alternative model in equations (5) and (6). Notwithstanding some minor differences, the real GDP growth responses in the first panel are remarkably similar to those in Figure 2. Similarly, the real GDP level responses in the second panel are qualitatively similar to Figure 2. The differences are largely limited to changes in the significance of the real GDP responses at longer horizons, some results being more and some less significant than in Figure 2. Even the CPI inflation responses in the third panel are qualitatively similar to those in Figure 2, if somewhat higher in magnitude. The largest discrepancies relate to the point estimates of the price level responses, which tend to be larger than in Figure 2 and often more significant, albeit at most at the 10 percent level. At the 5 percent significance level, the results remain fully consistent with those in Figure 2 in that only the German CPI response is significant. Figure 5 presents the corresponding conditional covariance measure based on the alternative regression model. While there are some differences, the results are overall similar to Figure 3. None of the estimates are significant at the 5 percent significance level. At the 10 percent level, for some horizons the German, Canadian and U.K. responses are marginally significant.

5.2. Model Selection Results

The evidence in Figure 4 suggests that for real GDP growth and real GDP levels it makes little difference whether we use the baseline model or the alternative model; for CPI inflation (and more importantly for the consumer price level) there are some differences, in contrast, although none that could not be explained by sampling uncertainty. This makes it important to compare formally the fit of these two models.

As shown by Sin and White (1996) suitable information criteria will choose the best approximating model among a set of candidate models with probability one asymptotically. Given that both of our models contain exactly the same number of parameters, their ranking by the prediction mean squared error is identical to the ranking according to any information criterion. Table 5 shows that for real GDP growth neither model fits the data better for all countries. Since the empirical results for real GDP growth are very similar in any case, this means that for real GDP growth we can focus on the baseline model without loss of generality.

For CPI inflation, in sharp contrast, for all seven countries the baseline model fits the data better than the alternative model. The improvement in fit can be substantial. In some cases, the prediction mean-squared error of the alternative model is more than four times as large as for the baseline model. Thus, to the extent that there are any differences in the results for CPI inflation, the formal model selection criteria in Table 5 establish conclusively that the baseline regression which imposes additional structure on the data has more statistical support. In the remainder of this paper, I thus will focus on the results from the baseline regression.

6. Counterfactual Simulations in the Baseline Model

There is a tendency to think of exogenous oil supply shocks as one-time adverse oil supply shocks. This approach is implicit in the impulse response analysis provided so far. Such response estimates also postulate that the initial shock is never reversed. Effectively, this amounts to postulating a permanent shortfall of crude oil production. This assumption is not realistic. As noted earlier, historically, exogenous production shortfalls have tended to be temporary. In that case, by construction, negative shocks to oil production are followed by positive shocks. In assessing the overall impact of the exogenous oil supply shocks associated with a given event such as the Iranian Revolution or the Persian Gulf War, it therefore is essential to go beyond dynamic multipliers and to conduct a counterfactual historical simulation that takes account of the cumulative effect of the full sequence of exogenous oil supply shocks associated with a given historical event.

Such historical decompositions can be constructed from the baseline model by simulating counterfactual data from the estimates of equations (1) and (2) under the assumption that $x_t = 0 \; \forall t.$ The cumulative effect of the sequence of realizations for x_t is measured by the difference between this counterfactual time series and the historical observations for the left-hand side variable. Such historical decompositions allow a comprehensive assessment of the effects of specific historical events via their effect on oil supply.

6.1. Comparing the Effects of Exogenous Oil Supply Shocks Across Countries and Time

Tables 6a and 6b characterize the aftermath of each of the exogenous events shown in Table 1. Each column refers to a different exogenous event, starting with the 1973/74 oil shock and ending with the oil crises of 2002/2003. The first line for each country and episode shows average real GDP growth and average CPI inflation rates. All rates have been normalized relative to their respective averages for 1973.I-2004.III such that a negative value corresponds to below average rates. The second line shows the average effect on real GDP growth and CPI inflation caused by exogenous oil supply shocks over the same subperiods. All rates have been annualized. The corresponding time series plots are shown in Figures 6a and b.

6.1.1. Real GDP Growth

There is a presumption that periods following exogenous oil supply shocks tend to be characterized by unusually low real growth. Table 6a shows that this is true for the aftermath of the 1973/74 shock and – with the exception of Japan – for the 1980 shock. It also is true for the 1990 shock, abstracting from Germany which was subject to the effects of German re-unification at the time. On the other hand, after the 1978/79 and 2002/03 shocks three of seven countries were able to maintain average or above average real growth rates. In 1978/79 these countries were Italy, Germany and Japan; in 2002/03 they were the United States, the U.K. and Japan.

There also is a presumption that the observed unusually low real growth may be attributed to the effects of the preceding exogenous oil supply shocks. As expected, Table 6a shows that the estimated effect of exogenous oil supply shocks on real growth is negative in each instance. There are, however, important differences across episodes and countries that become readily apparent in Figure 6a. For example, in all countries that experienced distinctly below average growth after the 1973/74 shock, the effects of the exogenous oil supply shock appear negligible.

For the 1978/79 shock there is some evidence in Figure 6a that the exogenous oil supply shock contributed to a small extent to fluctuations in real growth in the United States and the United Kingdom, but for most countries the contribution is again negligible. Stronger evidence is obtained for the 1980 oil shock. Especially for the United States, the oil shock contributed to the decline in real growth in 1982, although the effect is mainly limited to one quarter in 1982. Similarly, this shock contributed to lower Canadian real growth in late 1981 and in 1982 and to lower U.K. growth in 1981 and early 1982. The evidence is strongest for Germany, where the effect of the oil shock seems to account for most of the growth fluctuations during 1980.IV-1983.I. There also is some apparent effect in the Italian and French data, but hardly any for Japan.

Exogenous oil supply shocks also seem to matter in the period after the 1990 shock, although this effect does not explain the immediate drop in real growth during late 1990 and early 1991 observed in all countries but Germany (for the reasons alluded to above) and Japan. Rather, the cumulative effect of exogenous oil supply shocks on real growth makes itself felt mainly in 1992 and 1993. Finally, following the 2002/03 shocks, there is hardly any evidence of a reduction in real growth being caused by exogenous oil supply shocks.

6.1.2. CPI Inflation

There is a similar presumption that periods following exogenous oil supply shocks are periods of abnormally high inflation. Table 6b suggests otherwise. While the shocks of 1973/74 and 1978/79 and – with the exception of Japan – that of 1980 are indeed followed by unusually high CPI inflation rates, disregarding the unique case of Germany in the early 1990s, the more recent shocks are followed by below average inflation rates. Table 6b also shows that the average contribution of exogenous oil supply shocks to the inflation rate tends to be small or negligible. This evidence is particularly striking for the first two shock episodes. These observations cast serious doubt on the view that exogenous oil supply shocks are responsible for the sustained inflation of the 1970s and early 1980s.

Figure 6b shows that after 1973/74 the increase in CPI inflation that can be attributed to exogenous oil supply cuts is negligible compared to existing levels of CPI inflation (which leaves open the possibility, of course, that higher oil prices due to endogenous factors such as high global demand for oil may have played a more important role). This is true for all G7 countries. As noted by Bohi (1989), the striking differences in economic performance across the United States, Japan, and Germany, in particular, when faced with the same exogenous shock, are suggestive of an important role for domestic economic policies. Qualitatively similar results hold for the 1978/79 and 1990 shocks. After the 1980 shock there is some effect on CPI inflation, notably in Germany and in the U.K., but again exogenous oil supply cuts are not the primary cause of higher inflation. Finally, the effects of the 2002/03 shock on inflation are negligible. These results underscore the earlier point that the unusual inflation experience of the 1970s and early 1980s was not caused by exogenous oil production cuts, but is likely to have been the consequence of countries' macroeconomic policy choices, which would help explain both the wide variation in inflation experiences across countries over the same subperiod and the wide variation across different episodes in the same country (as in the case of Japan, for example).

7. Further Evidence on the Transmission of Exogenous Oil Supply Disruptions

While evidence on inflation and output responses is central to assessing the effects of exogenous oil supply disruptions, there are other variables that one would expect to play a key role in the transmission of these shocks. In this section, I focus on the responses of the real wage (defined as hourly earnings deflated by the CPI), short-term interest rates and nominal dollar exchange rates. For a precise definition of these variables the reader is referred to the comparative tables provided in the OECD Main Economic Indicators. All results are based on the baseline model.

7. 1. Responses of Real Wages, Interest Rates and Exchange Rates by Country

A potentially important channel of the transmission of exogenous oil supply shocks is based on the notion of real wage rigidities (see Bruno and Sachs 1982, 1985). A downwardly rigid real wage (for example as a result of pressure from labor unions) would amplify the effects of exogenous oil supply shocks on real GDP. Although this explanation is popular, its empirical support is scant. The available evidence for the United States suggests that more likely than not real wages fell following the oil shocks of the 1970s, although that evidence is based on different methodologies from the approach used in this paper (see, e.g., Rotemberg and Woodford 1996). Even less is known about the response of real wages to exogenous oil supply shocks in other G7 countries, although the real wage rigidity hypothesis when first introduced seemed particularly suitable to explaining the poor economic performance of European economies following the oil supply shocks of the 1970s and early 1980s.

The first panel of Figure 7 provides a comprehensive comparison of the real wage responses by country. It shows that for five of the G7 countries exogenous oil supply disruptions cause a decline in real wages. These include the United States, France, Germany, the U.K. and Japan. The decline is usually significant at least at some horizons. For the United States the rejections are only marginally significant at the 10 percent level. For other countries such as Germany, Japan or the U.K. the rejections are more decisive. For the remaining G7 countries, Italy and Canada, point estimates show a slight increase in the real wage, but, especially for Italy, the confidence intervals are unusually wide, and the positive responses are highly insignificant for both countries. Overall, the point estimates suggest clear evidence of falling real wages for most countries in response to exogenous oil supply disruptions, although the confidence bands in most cases are also consistent with a return to the original real wage after three years. Only for Japan is the decline in real wages statistically significant at the 10 percent level even after three years. Thus, despite evidence for at least a temporary decline of real wages in many countries, we are unable to rule out the real wage rigidity hypothesis at conventional significance levels.

A second potential channel of transmission is an increase in short-term interest rates in response to exogenous oil supply disruptions, as would be consistent with a monetary tightening in anticipation of future inflation (see, e.g., Bernanke, Gertler and Watson 1997). The second panel of Figure 7 shows that the oil-importing European countries tend to experience a sustained and significant rise in short-term interest rates in response to exogenous oil supply disruptions, whereas Japanese interest rates essentially do not respond. In contrast, in the United States and Canada exogenous oil supply shocks are followed by a temporary rise in interest rates. Taking account of sampling uncertainty the differences between the U.S. and Canada on the one hand and continental Europe on the other are not as pronounced, however. The latter intervals are also consistent with a temporary interest rate increase in Italy, France and Germany. The response for the U.K. is somewhat erratic, but shows no evidence of an increase in short-term interest rates.

Finally, the last panel of Figure 7 investigates the responses of the six nominal dollar exchange rates. The currencies of non-oil producing G7 countries significantly depreciated against the dollar in response to exogenous oil supply disruptions. In contrast, the depreciation of the British pound is smaller and typically insignificant. The Canadian dollar even appreciates slightly in the short run in response to an exogenous oil supply disruption, the latter result being marginally significant at the 10 percent level, although it depreciates insignificantly in the long run. While these estimates seem to suggest that the dollar exchange rates of non-oil producing G7 countries depreciate more in response to exogenous oil supply disruptions than the currencies of oil producers such as the U.K. and Canada, formal statistical tests in section 8 do not reject the null of equal responses for these countries.

A comparison of Figures 2 and 7 suggests that differences in these responses alone cannot explain the differences in inflation and output responses across countries in Figure 2, although it is noteworthy that Japan, which shows essentially no short-term interest rate response in Figure 7 and a response of real wages and nominal exchange rates similar to other countries, also experiences a much smaller real GDP growth response than other countries.

7. 2. Responses of Deflator Inflation by Country

While the CPI is commonly used as a measure of the price level, it is not the most natural choice for measuring the price level in relation to real GDP. This subsection therefore presents additional results based on the implicit GDP deflator, as available in the OECD Main Economic Indicators. The first panel of Figure 8 shows the responses of deflator inflation to an exogenous oil supply disruption. In contrast to the result for CPI inflation in Figure 3, the response of U.S. deflator inflation in Figure 8 shows no spike after three quarters. More generally, with the exception of the estimate for Italy, deflator inflation responses are more muted and less significant than the corresponding responses of CPI inflation. The same applies to the deflator responses in the second panel. This finding is consistent with the view that exogenous oil supply shocks, to the extent that they work through oil prices, would be more directly reflected in gross output price measures such as the CPI than in the price of value added (see Barsky and Kilian 2004). For completeness the third panel of Figure 8 shows the stagflationary effects of a 1 percent exogenous oil supply disruption when stagflation is defined in terms of the implicit GDP deflator rather than CPI. The results shown are qualitatively similar to the results in Figure 3 in that only for Germany is there marginally significant evidence of stagflation.

8. Are All G7 Responses Alike?

Although the sample available for this study is small, there is a fair degree of similarity in the real growth responses across G7 countries (see Figure 2). In all G7 countries but Japan there is significant evidence that an exogenous oil supply disruption causes a decline in real growth. This effect with the exception of the U.K. is delayed until at least the fourth quarter after the shock. With the exception of Germany and France there are no significant reductions in real growth after the ninth quarter. Moreover, only in one of seven countries is there significant evidence that exogenous oil supply disruptions cause a reduction of real GDP within the first year following the shock. In contrast, in five of the seven countries exogenous oil supply disruptions cause a significant reduction in real GDP in the second and third year following the shock.

Turning to CPI inflation, there is some heterogeneity among G7 countries in that in the United States and Italy exogenous oil supply disruptions tend to cause sharp spikes in CPI inflation, in France the response is hump-shaped, and in Canada, the U.K., Germany (and to a much lesser extent Japan) the responses are characterized by repeated spikes. There also is evidence of a significant increase in the German CPI level three years after the shock. For the United States and the U.K. the price level three years after the oil supply disruption is no higher than before the shock. For all other countries the price level rises as a consequence of the disruption, but not significantly so.

8.1. Median Responses in the Baseline Model

One way of summarizing these data is to focus on the median peak of CPI inflation responses and median trough of the responses of real GDP in the G7 countries. This avoids imposing the strong assumption that all G7 responses are the same, an assumption that will be formally tested in section 8.3. The upper panel of Table 7 suggests that exogenous oil supply disruptions cause a trough in G7 real growth 7 quarters after the shock. The median magnitude of the trough caused by a 10% reduction in world oil supplies is -1.7%. Real GDP reaches a trough about 11 quarters after the shock. The median cumulative effect on real GDP is -5.9%. The lower panel shows that the median inflation peak comes three quarters after the exogenous oil supply disruption with a median magnitude of 1.25%. The CPI level peaks after 7 quarters with a median cumulative effect of about 4%.

I also report separate results for the four European G7 countries to address the question of whether there are important differences in the European experience compared to the United States, Canada and Japan. The timing and magnitude of the growth and output responses is very similar. For the inflation responses as well the differences are small. For price level responses, in contrast, the data seem to suggest a median of 9.5 quarters lag for the peak response for the European G7 countries, compared with only 7 quarters for the G7. Upon inspection, this result is driven by the French CPI response which exhibits a primary (statistically insignificant) peak of magnitude 4.09 after 12 quarters, as well as a secondary, but statistically significant peak of magnitude 4.07 after 6 quarters. If we focus on the secondary peak instead of the primary peak, the median delay shrinks from 9.5 to 6.5, very close to the G7 median peak time, with virtually no change in the median cumulative effect. These results suggest that there is no systematic difference in the qualitative features of the responses.

As mentioned earlier, the U.K. responses, while still similar to the baseline case, change somewhat when the sample is reduced to the post 1978.III period, effectively eliminating the 1973/74 oil shock episode from the sample. I therefore present additional results in which these alternative U.K. results have been substituted for the baseline results. Table 7 shows that this substitution somewhat lowers the estimated median impact of exogenous oil supply disruptions on real growth and inflation, but otherwise the median results are quite similar.

8.2. Which Countries are Affected the Most by Exogenous Oil Supply Shocks?

While the evidence of a temporary reduction in real GDP growth is consistent across all G7 countries, the inflationary consequences of an oil supply shock are less clear-cut. An important lesson from the impulse analysis is that it is possible for countries to fare comparatively well when faced with exogenous oil supply shocks. For example, in the United States exogenous oil supply disruptions tend to cause only a brief spike in CPI inflation after 3 quarters rather than a sustained increase in CPI inflation. Qualitatively similar results hold for Italy with a spike after 4 quarters and to a lesser extent for France, which experiences a hump in CPI inflation from the second through the fifth quarter. This evidence contrasts with the pattern of repeated spikes in CPI inflation responses found in German data, and to a lesser extent Canadian, Japanese and U.K data. Only for Germany is there some evidence of sustained inflation being caused by exogenous oil supply shocks. This pattern is also reflected in the estimated cumulative effect on the price level after three years. Whereas the cumulative effect for the United States and the United Kingdom is negative (although not significantly so), Italy (2.69), Japan (2.63), Canada (4.38) and France (4.09) experienced moderate price level increases after three years. None of those increases is statistically significant. The only G7 country with a significant long-run increase in consumer prices is Germany (8.03).

Similarly, the estimated cumulative output responses after three years suggest that countries such as Italy (-3.80), France (-3.45) and Japan (-2.82) suffer less severe reductions in real GDP growth due to exogenous oil supply disruptions than the U.K. (-7.87), Canada (-7.60), Germany (-5.93), or the United States (-6.21). This pattern is further evidence that policy responses and institutional characteristics may matter in dealing with exogenous oil supply shocks.

8.3. How similar are the Response Patterns across G7 Countries?

While there are clear patterns in the response estimates, it is not clear whether the apparent differences in impulse responses are also statistically significant. It is possible to test formally the equality of these response estimates. I start with a joint test of the equality of the impulse response coefficients across all G7 countries. Table 8 shows the results for each macroeconomic aggregate. The first row shows that for all six variables considered in this paper, the hypothesis of equal responses in all G7 countries is strongly rejected with *p*-values between 1 and 4 percent. Of particular interest is the question of how much of that result is due to the impact of one of the three largest G7 countries: the United States, Japan and Germany. This question may be assessed by observing the change in *p*-values, as these countries are dropped from the sample one at a time. There is some evidence in Table 8 that for real GDP both Germany and to a lesser extent the United States contribute to the lack of homogeneity. Likewise Japan appears to be partially responsible for the lack of homogeneity in the short-term interest rates responses. Yet overall the evidence of non-homogeneity is quite resilient.

Next consider the subset of European G7 countries. Here none of the equality tests reject at the 10 percent level, suggesting a greater degree of homogeneity than for the set of all G7 countries. Excluding Germany from this group does not raise the *p*-value for all variables. In fact, the equality of the real wage responses now is rejected at the 10 percent level. One might suspect that the inclusion of the U.K., as the only European oil producer, may have contaminated these results, but focusing on the continental European G7 countries does not systematically raise the *p*-values, whether Germany is excluded or not.

Alternatively, one may break down the G7 economies along the lines of oil-producers (United States, Canada, U.K.) and non-oil producers. Both groups are relatively homogeneous in that in neither case the homogeneity null can be rejected at the 5 percent level, but in both cases there is at least one rejection at the 10 percent level. The most homogeneous results are obtained for the set of non-oil producing G7 countries excluding Germany, for which the lowest *p*-value is 18 percent.

A final question of interest is whether the responses for some macroeconomic aggregates are systematically more homogeneous than those for others. Focusing on the results for various subsets of G7 countries, by far the most homogeneous responses are obtained for the nominal dollar exchange rates (with *p*-values between 0.78 and 1.00), followed by the implicit deflator (with *p*-values between 0.18 and 0.99) and the short-term interest rate (with *p*-values between 0.11 and 0.81). The least homogeneous responses are found for the real wage (with *p*-values between 0.04 and 0.56).

While the results in Table 9 are instructive, it is important to keep in mind that such tests may be unduly restrictive in that impulse responses can be qualitatively similar across countries, yet not numerically equal. In particular, a rejection of the equality null would still be consistent with the responses being qualitatively similar as measured by statistics such as the timing or magnitude of response peaks or troughs. In this sense, the results in section 8.3 are complementary to those presented in section 8.2.

9. How Does the Real Price of Crude Oil Respond to Exogenous Oil Supply Shocks?

While the regression approach used in this paper does not rely on oil price data, an important check of the plausibility of the identification of exogenous oil supply shocks is the response of the real price of crude oil to exogenous oil supply shocks. Figure 9 presents two alternative estimates, one based on the baseline model that imposes exogeneity and the other based on the alternative distributed lag model introduced in section 5. As expected, both response estimates are consistent with a persistent increase in the real price of crude oil. The baseline model implies a significant increase of the real price of oil by near 2% on impact in response to a 1% reduction of oil supplies. After seven quarters, the response peaks at 3%. In contrast, the alternative model implies a response as high as 5%, but the estimates from the alternative model are very imprecisely estimated and not significantly different from zero.

The implied elasticity of the price with respect to exogenous supply disruptions after two years is near 3 in the baseline model and about 5 in the alternative model. Thus, a permanent exogenous reduction of oil supplies by 5% all else equal would approximately raise the real price of oil by 15% or 25%, respectively, depending on the regression model. These estimates already incorporate the offsetting effects on the real price of crude oil of endogenous crude oil production increases elsewhere in the world in response to higher oil prices.

10. Conclusion

Building on recent work in Kilian (2006), this study estimated and compared the effects of exogenous shocks to global oil production on inflation and real output in major industrialized countries. The analysis in this paper is subject to two important caveats. One is that I have abstracted from possible nonlinearities in the relationship between exogenous oil supply shocks and macroeconomic aggregates. This approach is consistent with the analysis of the quantitative dummy measure of exogenous oil supply shocks in Hamilton (2003), but differs sharply from the large empirical literature that measures exogenous oil supply shocks based on nonlinear transformations of the price of oil. It is widely believed that there is no linear causal relationship between oil prices and macroeconomic aggregates. Thus, the purpose of using nonlinear transformations of oil prices is to filter out the exogenous component of oil prices such that standard linear regression analysis can be applied on the filtered series (see Hamilton 2003).

Since my approach has been to identify the exogenous oil supply shock directly, it is no longer obvious what role nonlinearities have to play, if any. Unlike linear oil price regressions the results reported in this paper look quite ‘reasonable’ without resorting to nonlinear functional forms.

While nonlinearities are not obviously required in modeling the effects of exogenous oil supply shocks, this does not mean that they could not exist. Clearly, equations (1) and (2) will not capture any higher-order effects of exogenous oil supply shocks and only provide a linear approximation to the true effects. Under certain conditions, the true effects may be larger (or smaller) than my estimates suggest. Specifically, there are reasons to suspect that the effects of exogenous oil supply disruptions may be amplified by a number of factors including the strength of the world business cycle, the presence of capacity constraints in crude oil production, high cost shares of energy in production, and high oil price volatility. A detailed investigation of the specific forms of nonlinearity that might affect the relationship of exogenous oil supply disruptions, while interesting in its own right, is left for future research.

A second caveat is that this study has followed Hamilton (2003) in focusing on observable changes in crude oil production. The approach of Kilian (2006) is designed to quantify exogenous oil *supply* shocks, but it does not capture all effects of exogenous events in OPEC countries. The effects of events such as the Persian Gulf War or the Iranian Revolution may also include exogenous oil *demand* shocks that arise from sudden shifts in expectations about the future availability of oil supplies. For example, the imminent military threat to Saudi oil fields in August of 1990 (which lasted until late 1990) caused a large temporary increase in precautionary demand for oil, which was reflected in sharp swings in the price of oil in 1990/91 (see Kilian (2006) for a more detailed analysis of this episode). Those fears about future oil supplies never materialized. As Saudi oil supplies never were disrupted, the production-based measure of the exogenous oil supply shocks appropriately does not incorporate this effect. This does not mean that exogenous shocks to the demand for oil (as reflected in the real price of oil) were not important for the G7 economies. Thus, an important topic for future research will be the measurement of exogenous expectation shifts in the crude oil market. This topic is addressed in Kilian (2007).

References

Balke, N.S., Brown, S. and M.K. Yücel (2002), "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" *Energy Journal*, 23, 27-52.

Barsky, R.B., and L. Kilian (2002), "Do We Really Know that Oil Caused the Great Stagflation? A Monetary Alternative," in: *NBER Macroeconomics Annual 2001*, B.S. Bernanke and K. Rogoff (eds.), MIT Press: Cambridge, MA, 137-183.

Barsky, R.B., and L. Kilian (2004), "Oil and the Macroeconomy Since the 1970s," *Journal of Economic Perspectives*, 18(4), 115-134.

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

Bohi, D.R. (1989), *Energy Price Shocks and Macroeconomic Performance*. Washington, D.C.: Resources for the Future.

Bohi, D.R. (1991), "On the Macroeconomic Effects of Energy Price Shocks" *Resources and Energy*, 13, 145-162.

Bruno, M., and J. Sachs (1982), "Input Price Shocks and the Slowdown in Economic Growth: the Case of U.K. Manufacturing," *Review of Economic Studies*, 49, 679-705.

Bruno, M., and J. Sachs (1985), *Economics of Worldwide Stagflation*, Cambridge, MA: Harvard University Press.

Burbidge, J., and A. Harrison (1984), "Testing for the Effect of Oil Price Rises Using Vector Autoregressions," *International Economic Review*, 25, 459-484.

Cooley, T.F., and S. LeRoy (1985), "Atheoretical Macroeconometrics: A Critique," *Journal of Monetary Economics*, 16, 283-308.

Darby, M.R. (1982), "The Price of Oil and World Inflation and Recession," *American Economic Review*, 72, 738-751.

Davis, S.J., and J. Haltiwanger (2001), "Sectoral Job Creation and Destruction Responses to Oil Price Changes," *Journal of Monetary Economics,* 48, 465-512.

Den Haan, W.J. (2000), "The Comovement Between Output and Prices," *Journal of Monetary Economics*, 46, 3-30.

Den Haan, W.J., and S.W. Summer (2004), "The Comovement between Real Activity and Prices in the G7," *European Economic Review*, 48, 1333-1347.

The Economist (1999), "Oil's Pleasant Surprise," November 27, 16.

Engle, R.F., Hendry, D.F., and J.-F. Richard (1983), "Exogeneity," *Econometrica*, 51, 277-304.

Gisser, M., and T.H. Goodwin (1986), "Crude Oil and the Macroeconomy: Tests of Some Popular Notions," *Journal of Money, Credit and Banking*, 18, 95-103.

Gonçalves, S., and L. Kilian (2004), "Bootstrapping Autoregressions with Conditional Heteroskedasticity of Unknown Form," *Journal of Econometrics*, 123, 89-120.

Hamilton, J.D. (1996), "This Is What Happened to the Oil-Macroeconomy Relationship," *Journal of Monetary Economics*, 38(2), 215-220.

Hamilton, J.D. (2003), "What is an Oil Shock?" *Journal of Econometrics*, 113, 363-398.

Hamilton, J.D., and A.M. Herrera (2004), “Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy,” *Journal of Money, Credit and Banking*, 36, 265-286.

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

Hooker, M.A. (2002), “Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications versus Changes in Regime,” *Journal of Money, Credit and Banking*, 34, 540-561.

Hoover, K.D., and S.J. Perez (1994), “Post Hoc Ergo Propter Hoc Once More: An Evaluation of ‘Does Monetary Policy Matter?’ in the Spirit of James Tobin,” *Journal of Monetary Economics*, 34, 89-99.

Jiménez-Rodríguez, R., and M. Sánchez (2005), “Oil Price Shocks and Real GDP Growth: Empirical Evidence for some OECD Countries,” *Applied Economics*, 37, 201-228.

Kilian, L. (2006), “Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?” Forthcoming: *Review of Economics and Statistics*.

Kilian, L. (2007), “Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market,” mimeo, Department of Economics, University of Michigan.

Knetsch, T. (2005), “Short-Run and Long-Run Comovement of GDP and some Expenditure Aggregates in Germany, France and Italy,” mimeo, Economics Department, Deutsche Bundesbank.

Mork, K.A. (1989), “Oil and the Macroeconomy when Prices Go Up and Down: An Extension of Hamilton’s Results,” *Journal of Political Economy*, 91, 740-744.

Mork, K.A., Olsen,O., and H.T. Mysen (1994), “Macroeconomic Responses to Oil Price Increases and Decreases in Seven OECD Countries,” *Energy Journal*, 15, 15-38.

Lee, K., and S. Ni (2002), “On the Dynamic Effects of Oil Price Shocks: A Study Using Industry Level Data,” *Journal of Monetary Economics*, 49(4), 823-852.

Lee, K., Ni, S. and R. Ratti (1995), “Oil Shocks and the Macroeconomy: The Role of Price Variability,” *Energy Journal*, 16, 39-56.

Rasche, R.H., and J.A. Tatom (1981), “Energy Price Shocks, Aggregate Supply and Monetary Policy: the Theory and International Evidence,” *Carnegie-Rochester Conference Series on Public Policy*, 14, 125-142.

Rotemberg, J., and M. Woodford (1996), “Imperfect Competition and the Effects of Energy Price Increases on Economic Activity,” *Journal of Money, Credit, and Banking*; 28(4), Part 1, 550-77

Tatom, J.A. (1988), “Are the Macroeconomic Effects of Oil Price Changes Symmetric?” *Carnegie-Rochester Conference Series on Public Policy*, 28, 325-368.

Shapiro, M.D., and M.W. Watson (1988), “Sources of Business Cycle Fluctuations,” in: *NBER Macroeconomics Annual 1988*, S. Fischer (ed.), MIT Press: Cambridge, MA, 111-148.

Sin, C.-Y., and H. White (1996) “Information Criteria for Selecting Possibly Misspecified Parametric Models,” *Journal of Econometrics*, 71, 207–225.

Stock, J.H., J.H. Wright, and M. Yogo (2002), “A Survey of Weak Instruments and Weak Identification in Generalized Method of Moments,” *Journal of Business and Economic*

Statistics, 20(4), 518-529.

Table 1: Key Oil Dates

Date	Political Event
October 1973	Yom-Kippur War
October 1973	Arab Oil Embargo
October 1978	Iranian Revolution
September 1980	Iran-Iraq War
August 1990	Persian Gulf War
December 2002	Civil Unrest in Venezuela
March 2003	Iraq War

Table 2: Production Benchmarks Used in Constructing the Exogenous Crude Oil Production Shortfalls

	OPEC Countries Subject to Exogenous Production Shocks in 1971.1-2004.9								
	Algeria	Iran	Iraq	Kuwait	Libya	Quatar	Saudi Arabia	UAE	Venezuela
1971.1-1973.10									
1973.11-1974.3*	*a*		*a*	*a*	*a*	*a*	*a*	*a*	
1978.10-1979.4		*b*	*b*				*b*		
1979.5-1990.7		*b*	*b*						
1990.8-1991.3		*c*	*c*	*c*			*c*		
1991.4-2002.6		*c*	*c*	*c*					
2002.7-2002.11		*c*	*c*	*c*			*c*		
2002.12-2003.10		*d*	*d*	*d*			*d*		*d*
2003.11-2004.9		*d*	*d*	*d*					*d*

NOTES: Based on monthly data for crude oil production levels available from the U.S. Department of Energy. See Kilian (2005) for details. Benchmarks: [a] Non-OPEC Oil Producers. [b] OPEC excluding Iran, Iraq, and Saudi Arabia. [c] OPEC excluding Iran, Iraq, Saudi Arabia and Kuwait. [d] OPEC excluding Iran, Iraq, Saudi Arabia, Kuwait and Venezuela.

* Level of shortfall for 1973.10-1974.3 derived from counterfactual starting in 1973.3, reflecting the contractual constraints on Arab OPEC oil producers prior to 1973.10.

Table 3: *p*-Values of Granger Non-Causality Tests

	$H_0 : \beta_{1j} = 0 \forall j$	
	Lagged Real GDP Growth	Lagged CPI Inflation
United States	0.711	0.126
Italy	0.197	0.340
France	0.120	0.677
Germany	**0.005**	0.167
U.K.	0.253	0.667
Canada	0.260	0.858
Japan	0.096	0.180
Germany[a]	0.056	0.492

NOTES: Significant rejections at 5% level shown in boldface. [a] Pre-unification sub-sample.

Table 4: *p*-Values of Tests for the Exclusion of the Contemporaneous Effect of the Exogenous Oil Supply Shock

	$H_0 : \gamma = 0$	
	Real GDP Growth	CPI Inflation
United States	0.266	0.562
Italy	0.953	0.205
France	0.813	0.626
Germany	0.999	0.097
U.K.	**0.002**	0.577
Canada	0.688	0.481
Japan	0.507	0.522

NOTES: Significant rejections at 5% level shown in boldface. Under the maintained assumption of exogeneity $\beta_{1j} = 0 \forall j$ and $\theta = 0$ in equation (3).

Table 5: Comparison of the Fit of the Baseline and Alternative Models

	Prediction Mean Squared Error of Fitted Model			
	Real GDP Growth		CPI Inflation	
	Baseline	Alternative	Baseline	Alternative
USA	**9.44**	10.11	**4.67**	11.44
Italy	8.34	**7.69**	**12.22**	41.83
France	**4.17**	4.56	**4.29**	19.73
Germany	**12.24**	13.04	**4.09**	5.93
U.K.	14.03	**11.66**	**18.90**	42.94
Canada	**6.40**	8.00	**6.35**	15.10
Japan	14.23	**13.61**	**11.97**	23.78

NOTES: The prediction mean square error ranking coincides with the ranking by any consistent information criterion since the number of regressors is the same for both models. Boldface indicates the models selected by the criterion.

Table 6a: Real GDP Growth Rates Relative to Long-Run Average and Average Estimated Effect of Exogenous Oil Supply Shocks

		Episodes of Exogenous Oil Supply Shocks				
		1973.IV-1975.II	1978.IV-1980.III	1980.IV-1983.I	1990.III-1993.III	2002.IV-2004.III
United States	Growth	-3.84	-2.64	-1.87	-1.30	0.68
	Effect	-0.28	-0.77	-1.07	-1.44	-0.79
Italy	Growth	-2.01	2.10	-1.66	-1.96	-1.50
	Effect	-0.26	-0.35	-1.00	-0.98	-0.32
France	Growth	-1.06	-0.24	-0.37	-1.72	-0.96
	Effect	-0.08	-0.19	-0.89	-0.83	-0.20
Germany	Growth	-3.38	0.15	-2.01	2.33	-1.70
	Effect	-0.24	-0.24	-1.54	-1.36	-0.50
U.K.	Growth	-3.50	-2.45	-1.14	-2.02	0.50
	Effect	-0.18	-1.07	-1.05	-1.64	-0.93
Canada	Growth	-0.24	-0.41	-2.56	-2.71	-0.50
	Effect	-0.36	-0.75	-1.53	-1.79	-0.81
Japan	Growth	-1.75	1.00	0.17	-1.19	0.15
	Effect	-0.15	-0.35	-0.45	-0.56	-0.24

NOTES: The effect of exogenous oil supply shocks is computed by simulating the predicted value of regression (1) for $x_t = 0 \; \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. This effect is averaged for each sub-period and compared to the deviation of real GDP growth from its average value for 1971.I-2004.III, averaged over the same sub-period. All rates are annualized.

Table 6b: CPI Inflation Rates Relative to Long-Run Average and Average Estimated Effect of Exogenous Oil Supply Shocks

		Episodes of Exogenous Oil Supply Shocks				
		1973.IV-1975.II	1978.IV-1980.III	1980.IV-1983.I	1990.III-1993.III	2002.IV-2004.III
United States	Inflation	4.92	6.91	1.54	-1.33	-2.33
	Effect	-0.02	0.08	-0.64	-0.16	0.17
Italy	Inflation	9.84	8.92	8.07	-2.96	-5.67
	Effect	0.01	0.39	0.22	0.52	0.45
France	Inflation	7.25	6.42	5.92	-2.58	-3.08
	Effect	0.15	0.65	0.40	0.82	0.59
Germany	Inflation	4.32	2.00	2.27	1.39	-1.48
	Effect	0.24	0.85	1.29	1.38	0.44
U.K.	Inflation	14.34	8.35	1.04	-3.20	-3.82
	Effect	0.24	0.42	-1.00	-0.27	0.43
Canada	Inflation	5.32	4.82	4.84	-2.02	-2.90
	Effect	0.17	0.59	0.64	0.71	0.12
Japan	Inflation	12.88	2.66	-0.51	-0.94	-3.24
	Effect	0.15	0.75	-0.35	0.48	0.53

NOTES: The effect of exogenous oil supply shocks is computed by simulating the predicted value of regression (2) for $x_t = 0 \ \forall t$ and subtracting the predicted value from the realized value of quarterly CPI inflation. This effect is averaged for each sub-period and compared to the deviation of CPI inflation from its average value for 1971.I-2004.III, averaged over the same sub-period. All rates are annualized.

Table 7: Median Peaks and Troughs of the Responses to a 1% Permanent Reduction in World Oil Supplies in the Baseline Model

		Baseline		Same as baseline, except U.K. sample starts in 1978.III	
Trough		Real GDP Growth	Real GDP	Real GDP Growth	Real GDP
Median Timing (quarters)	G7	7	11	7	11
	European G7[a]	7	10	7	10.5
Median Magnitude (%)[b]	G7	-0.69	-2.37	-0.61	-2.37
	European G7[a]	-0.68	-2.03	-0.61	-2.03

		Baseline		Same as baseline, except U.K. sample starts in 1978.III	
Peak		CPI Inflation	CPI Level	CPI Inflation	CPI Level
Median Timing (quarters)	G7	3	7	3	7
	European G7[a]	3.5	9.5	4	9.5
Median Magnitude (%)[b]	G7	0.50	1.64	0.41	1.64
	European G7[a]	0.52	1.41	0.46	1.41

NOTES: [a] The European G7 countries are France, Germany, Italy and the U.K.
[b] All percent rates are annualized quarterly rates.

Table 8: Asymptotic *p*-values of Wald Tests for Equality of Impulse Responses

	Real GDP	CPI	Implicit GDP Deflator	Real Wage	Short-term interest rate	Local currency/ US$
G7	**0.038**	**0.025**	**0.011**	**0.035**	**0.007**	-
G7 excluding Germany	0.287	**0.026**	**0.004**	**0.179**	**0.002**	-
G7 excluding Japan	**0.050**	**0.021**	**0.022**	**0.017**	0.145	-
G7 excluding U.S.	0.160	**0.031**	**0.022**	0.063	**0.002**	0.876
European G7	0.164	0.676	0.904	0.104	0.316	0.948
European G7 excluding Germany	0.628	0.872	0.830	0.067	0.114	0.782
Continental European G7	0.093	0.941	0.990	0.081	0.813	1.000
Continental European G7 Excluding Germany	0.335	0.921	0.946	**0.040**	0.621	0.982
Oil producing G7	0.797	0.094	0.308	0.564	0.624	-
Oil producing G7 excluding U.S.	0.925	0.181	0.177	0.242	0.244	0.867
Non-oil producing G7	0.099	0.972	0.977	0.238	0.265	0.998
Non-oil producing G7 excluding Germany	0.240	0.967	0.959	0.175	0.148	0.979
Non-oil producing G7 excluding Japan	0.093	0.941	0.990	0.081	0.813	1.000

NOTES: Significant rejections at 5% level shown in boldface. All tests are based on bootstrap estimates of the variance-covariance matrix of the impulse response estimators obtained using block bootstrap methods that preserve the correlations across countries.

Figure 1: Measure of Exogenous Oil Supply Shocks 1971.I-2004.III

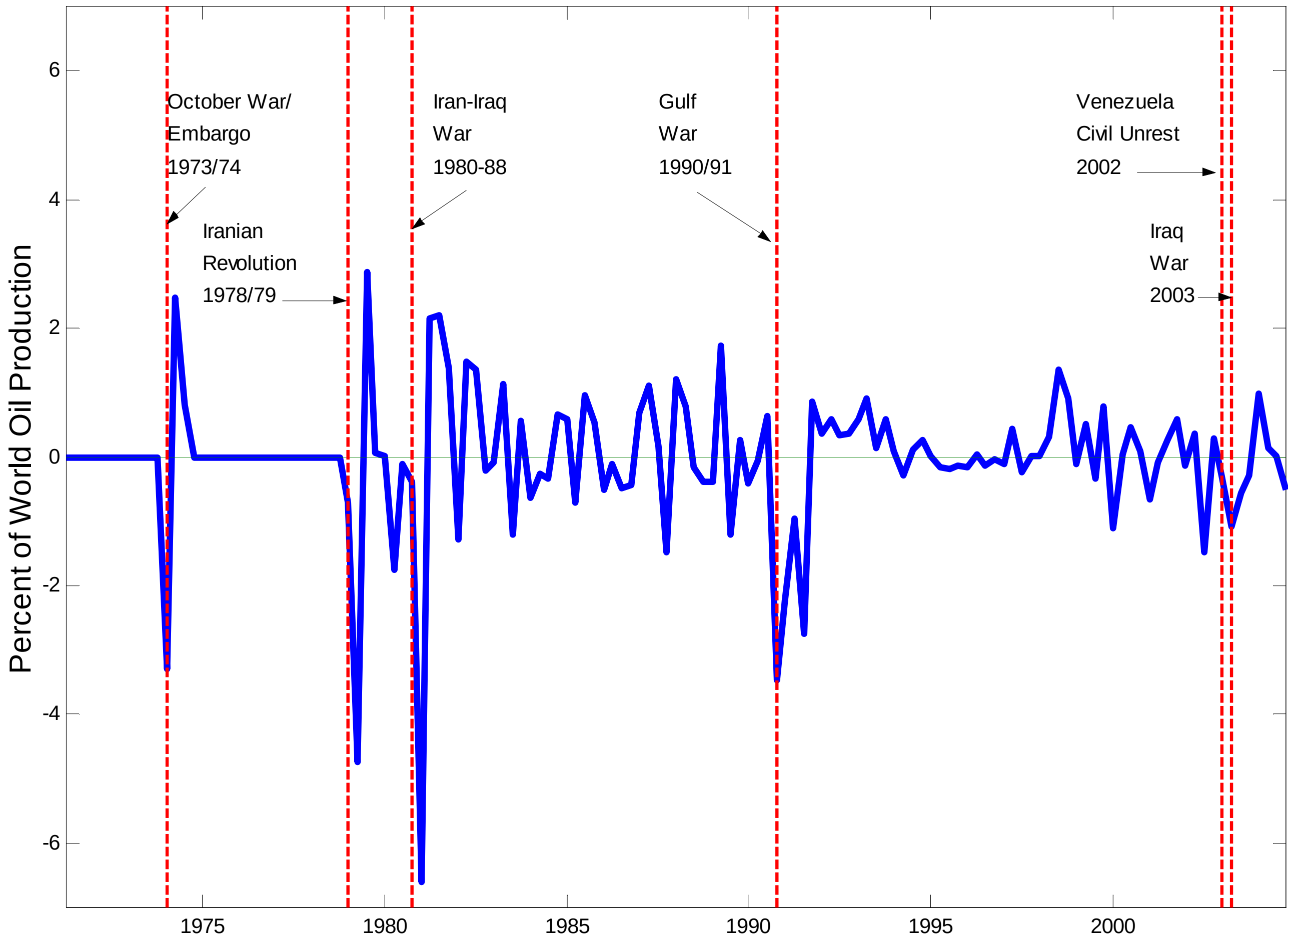

Source: Kilian (2005)

Figure 2: Dynamic Effects of a Permanent 1% World Oil Supply Disruption
OLS Point Estimates from Baseline Model with One- and Two-Standard Error Bands

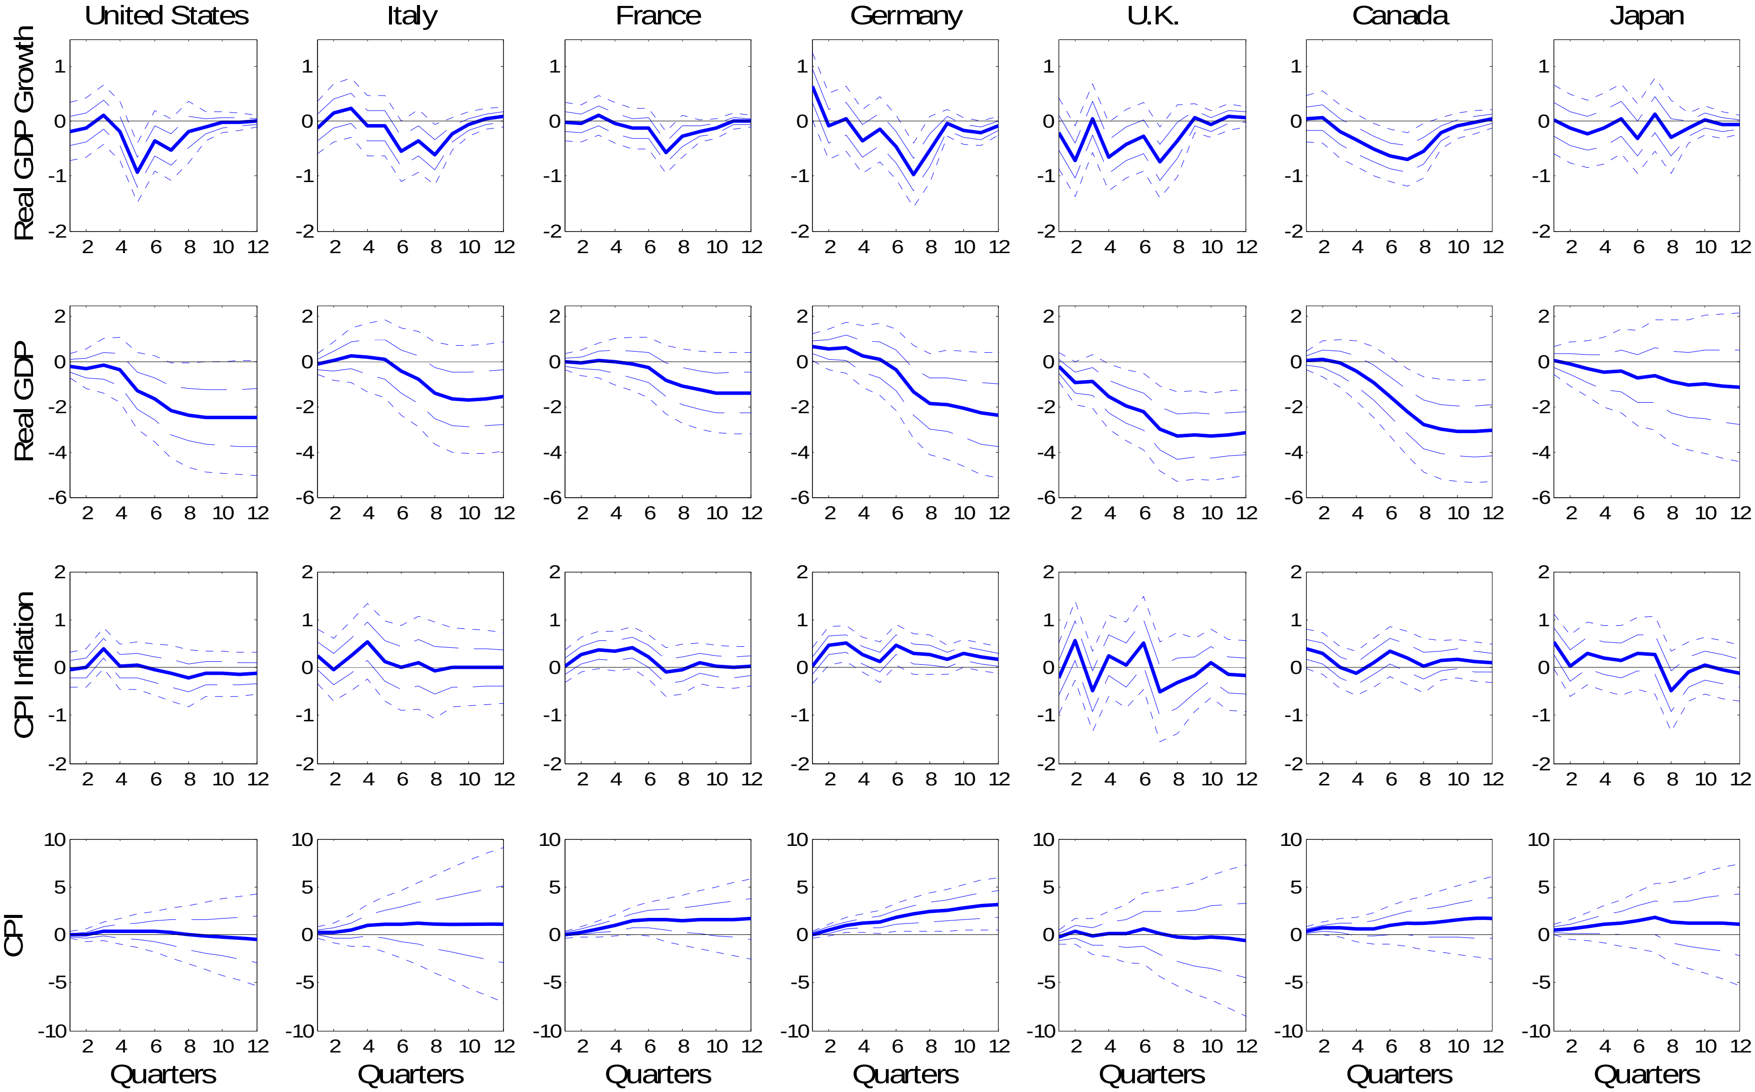

NOTES: The asymptotic confidence intervals were constructed based on Monte Carlo methods.

Figure 3: Stagflationary Effects of a Permanent 1% World Oil Supply Disruption in Baseline Model Conditional Covariance of Inflation and Real GDP Growth with 80% and 90% Confidence Bands

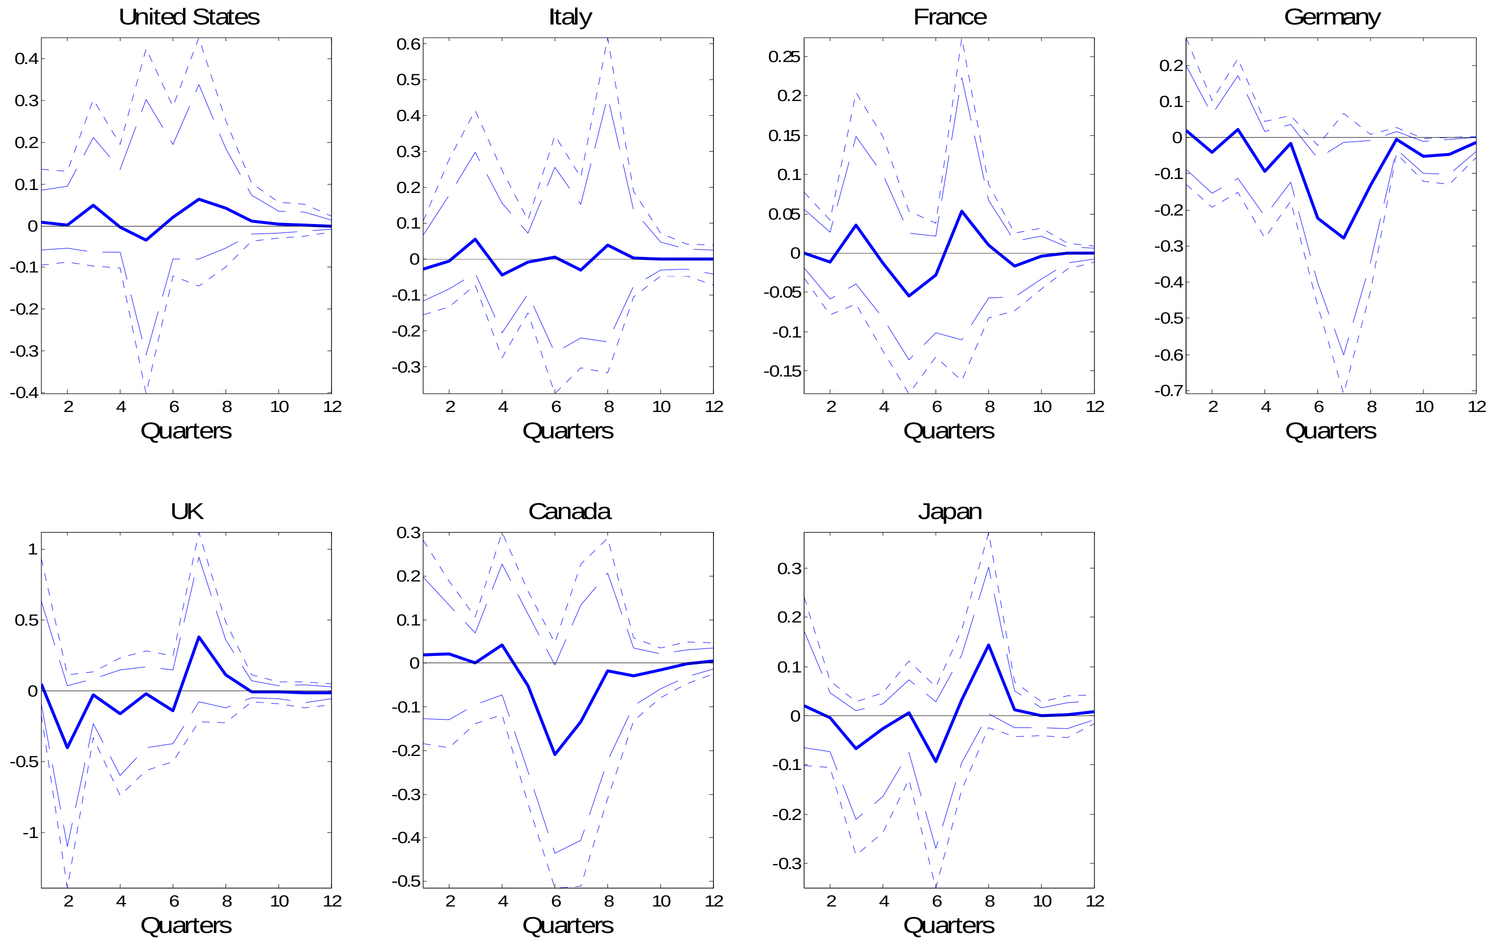

NOTES: The plot shows a statistical measure of the conditional co-movement between real GDP growth and inflation, as defined in Den Haan (2000) and Den Haan and Summer (2004), based on the impulse responses in Figure 2. Stagflation in the form of rising prices and falling output causes this measure to be negative. The confidence intervals were constructed using the fixed-design wild bootstrap method (see Gonçalves and Kilian 2004).

Figure 4: Dynamic Effects of a Permanent 1% World Oil Supply Disruption
OLS Point Estimates from Alternative Model with One- and Two-Standard Error Bands

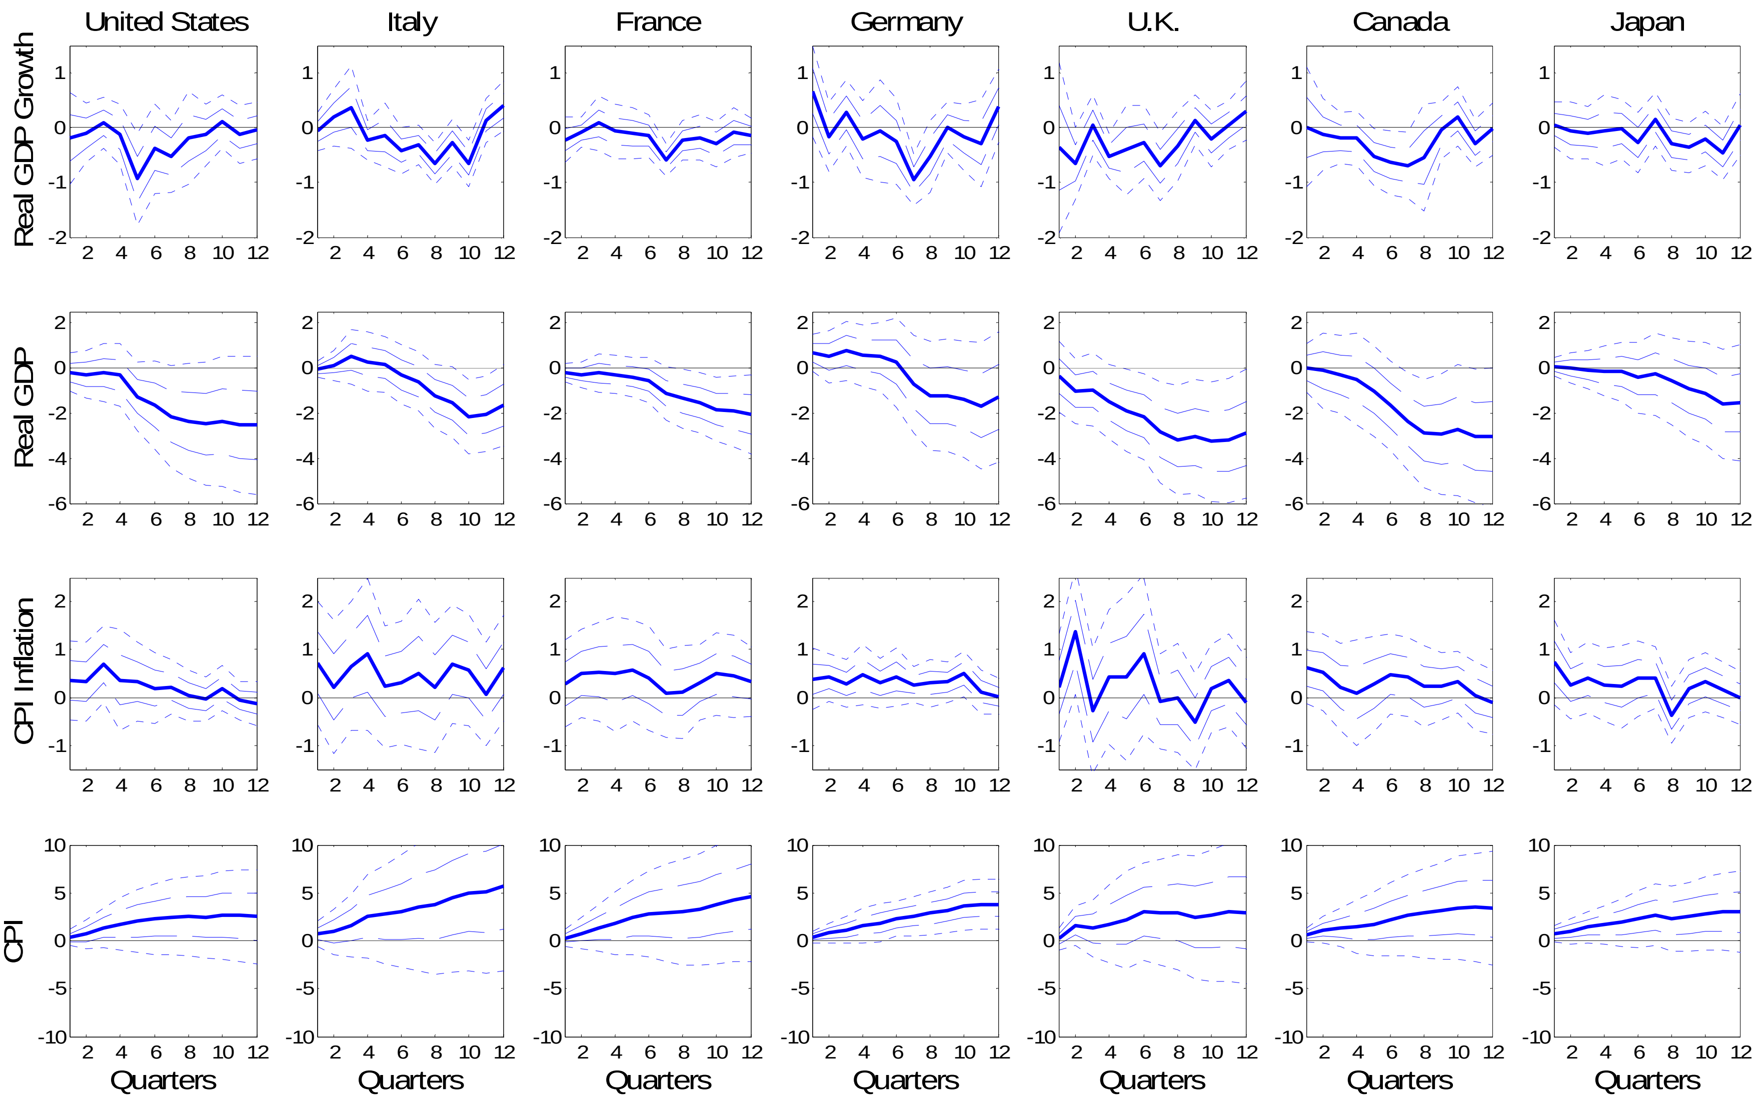

NOTES: The confidence intervals were constructed using the block bootstrap, taking account of possible serial correlation in the error term.

Figure 5: Stagflationary Effects of a Permanent 1% World Oil Supply Disruption in Alternative Model Conditional Covariance of Inflation and Real GDP Growth with 80% and 90% Confidence Bands

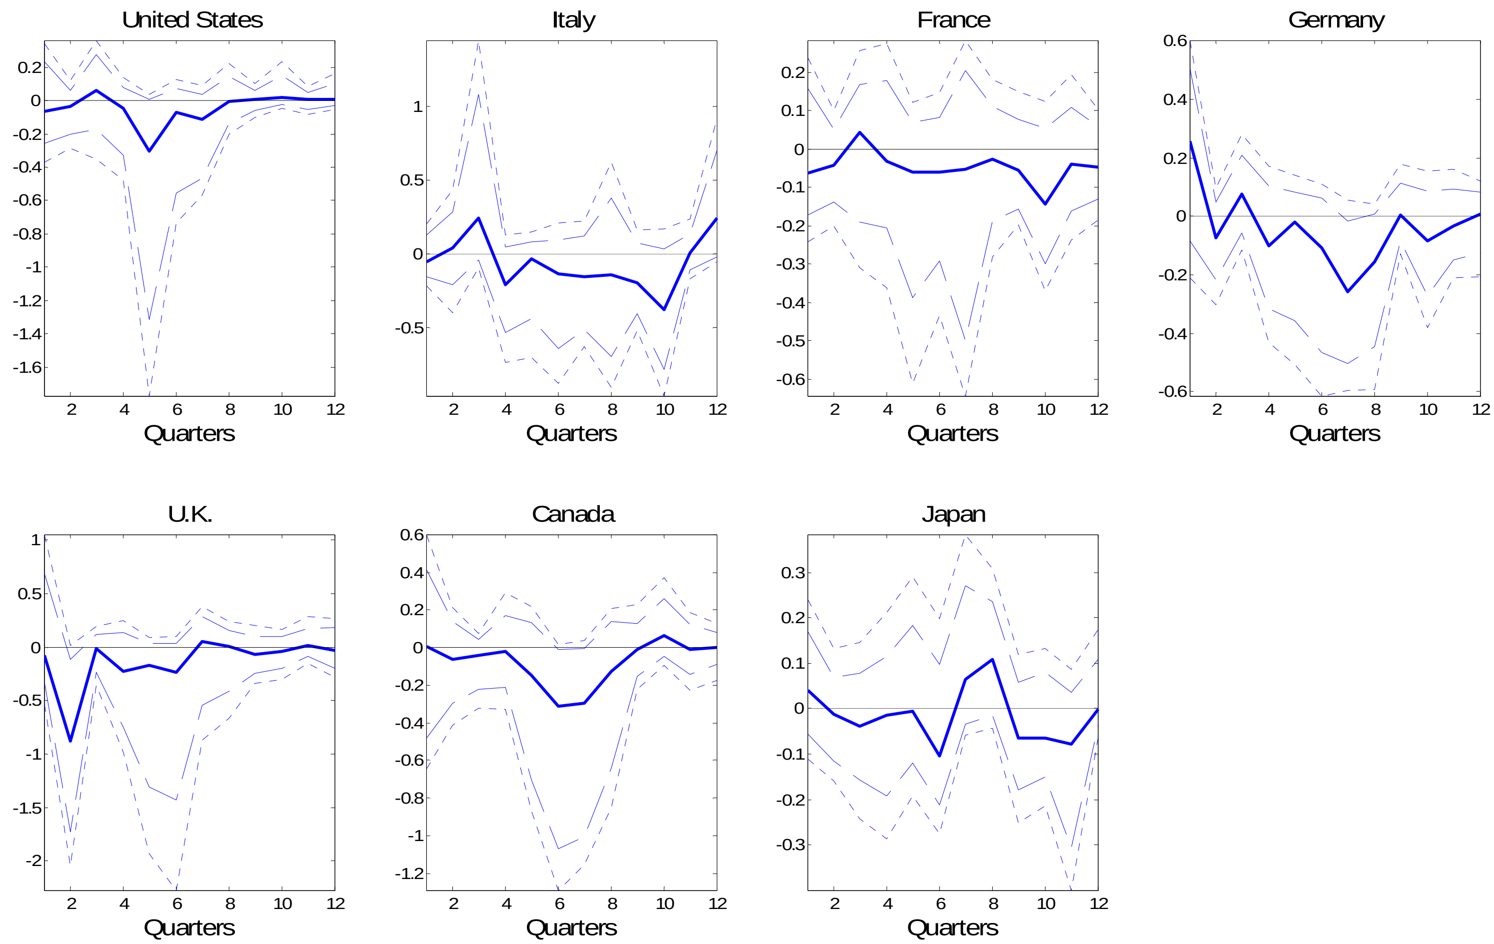

NOTES: The plot shows a statistical measure of the conditional co-movement between real GDP growth and CPI inflation, as defined in Den Haan (2000) and Den Haan and Summer (2004), based on the impulse responses in Figure 3. Stagflation in the form of rising prices and falling output causes this measure to be negative. The confidence intervals were constructed using the block bootstrap, taking account of possible serial correlation in the error term.

Figure 6a: Comparison of Real Growth Experience by Oil Shock Episode

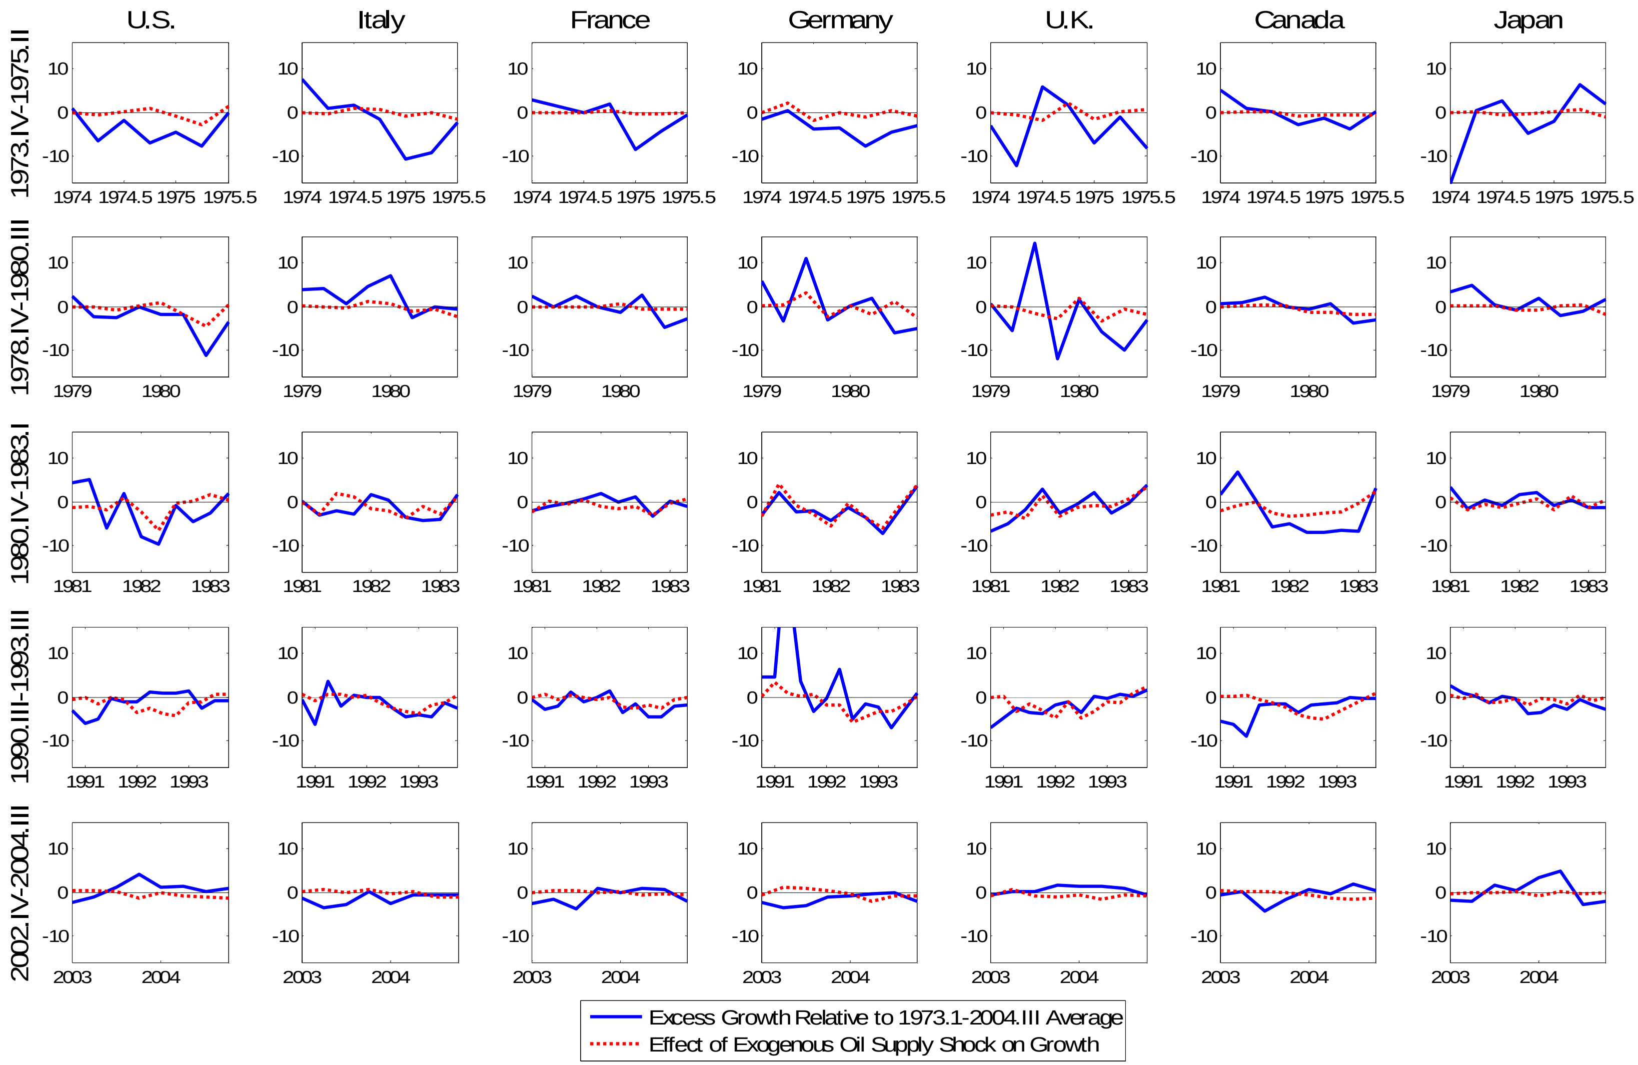

NOTES: The effect of exogenous oil supply shocks is computed by simulating the predicted value of regression (1) for $x_t = 0 \; \forall t$ and subtracting the predicted value from the realized value of quarterly real GDP growth. All rates are annualized.

Figure 6b: Comparison of CPI Inflation Experience by Oil Shock Episode

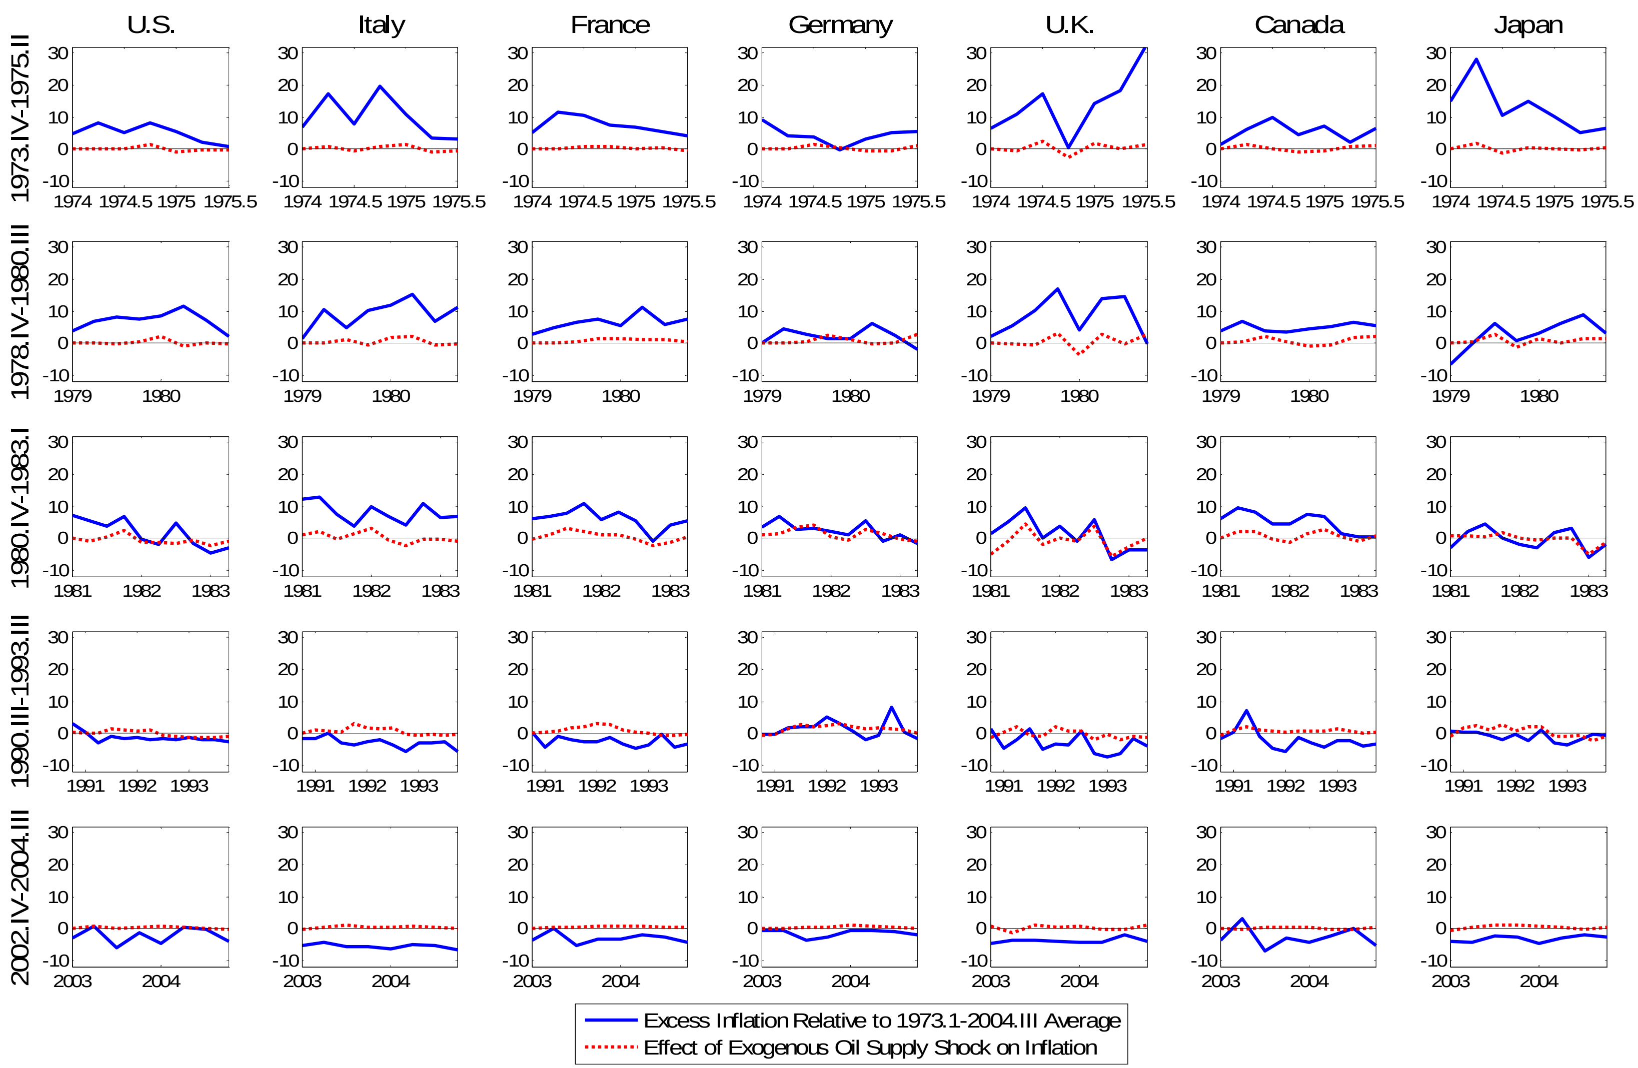

NOTES: The effect of exogenous oil supply shocks is computed by simulating the predicted value of regression (2) for $x_t = 0 \ \forall t$ and subtracting the predicted value from the realized value of the quarterly CPI inflation rate. All rates are annualized.

Figure 7: Dynamic Effects of a Permanent 1% World Oil Supply Disruption
OLS Point Estimates from Baseline Model with One- and Two-Standard Error Bands

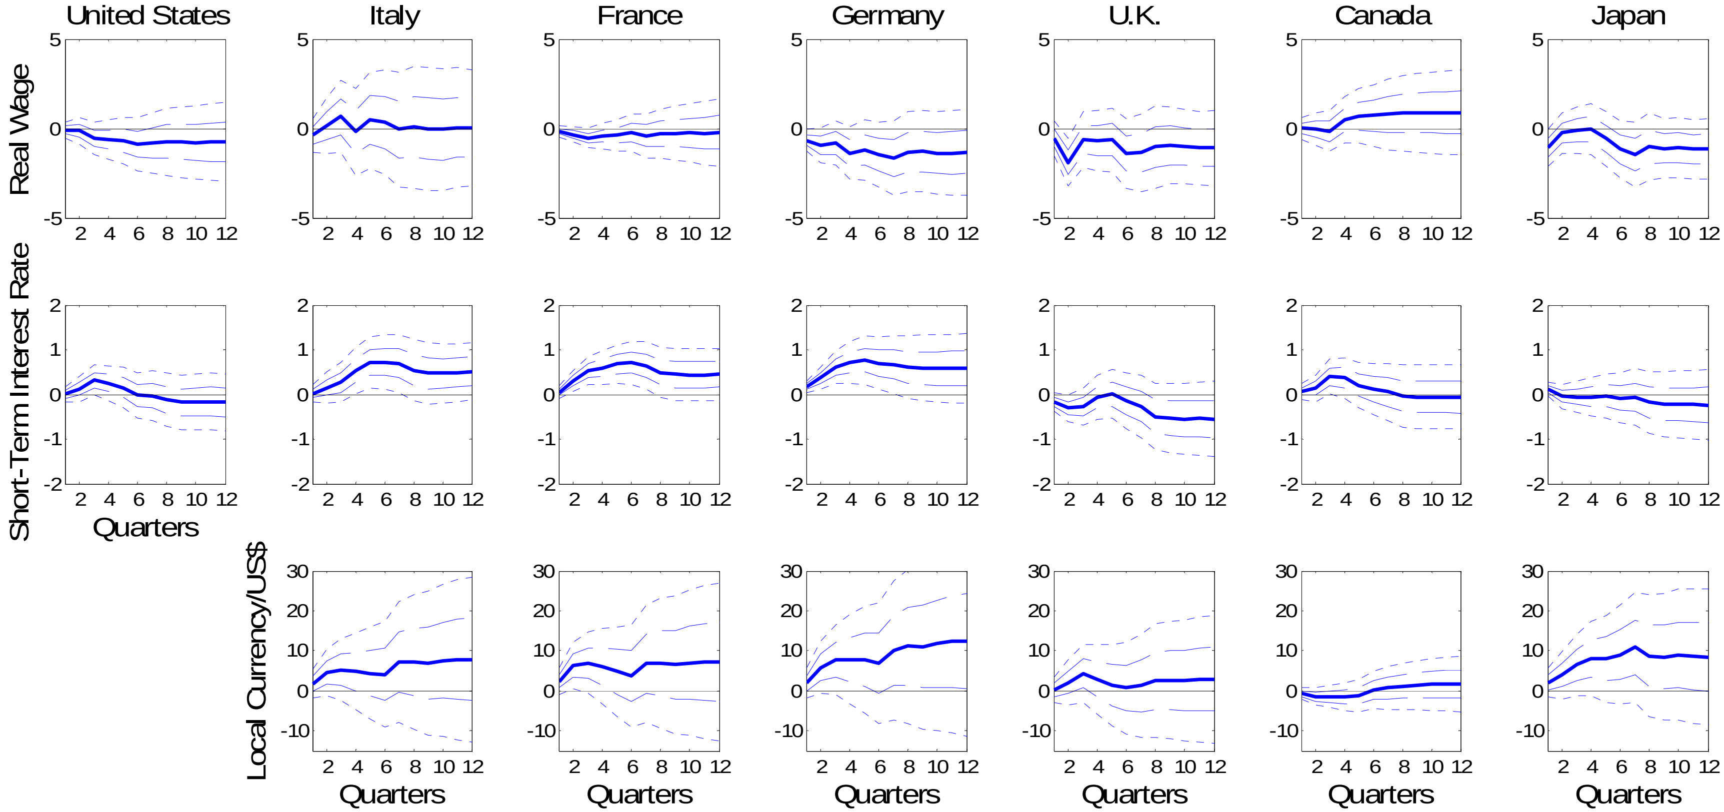

NOTES: All variables have been specified in first differences and the responses cumulated to obtain the level responses shown. The figures show asymptotic confidence intervals based on Monte Carlo methods.

Figure 8: Dynamic Effects of a Permanent 1% World Oil Supply Disruption OLS Point Estimates from Baseline Model Based on the Implicit GDP Deflator

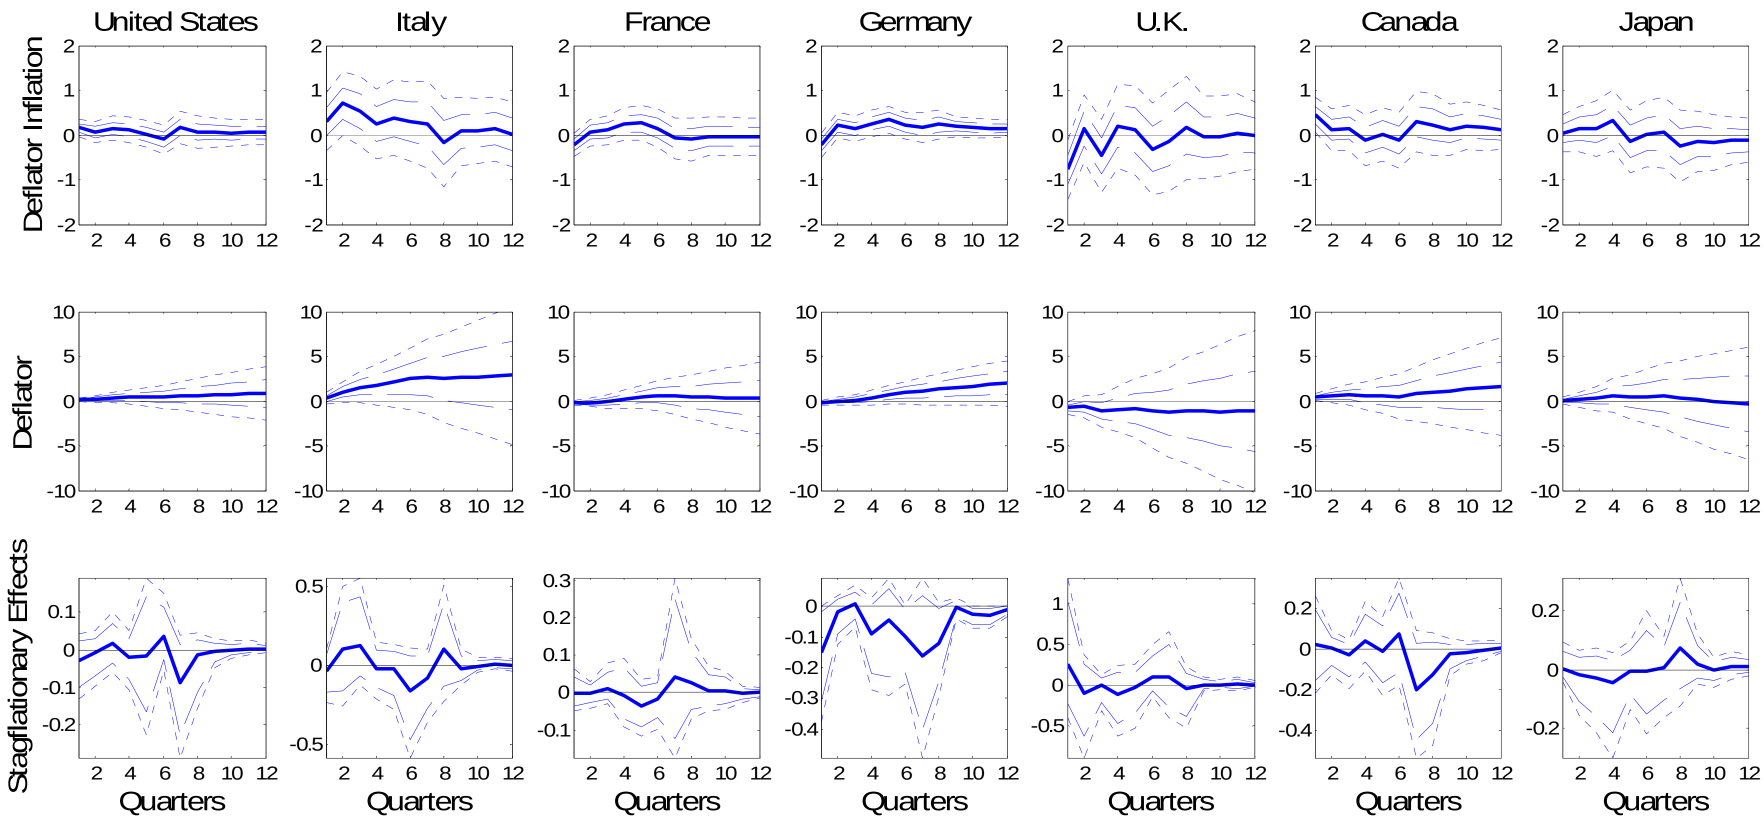

NOTES: The first two panels show impulse response estimates with one- and two-standard error bands. The third panel shows Den Haan's (2000) statistical measure of the conditional co-movement between real GDP growth and deflator inflation with 80% and 90% confidence intervals. Stagflation in the form of rising prices and falling output causes this measure to be negative. The confidence intervals have been constructed as in Figures 2 and 3.

Figure 9: Dynamic Effects of a Permanent 1% World Oil Supply Disruption on the Real Price of Oil
OLS Point Estimates with One- and Two-Standard Error Bands

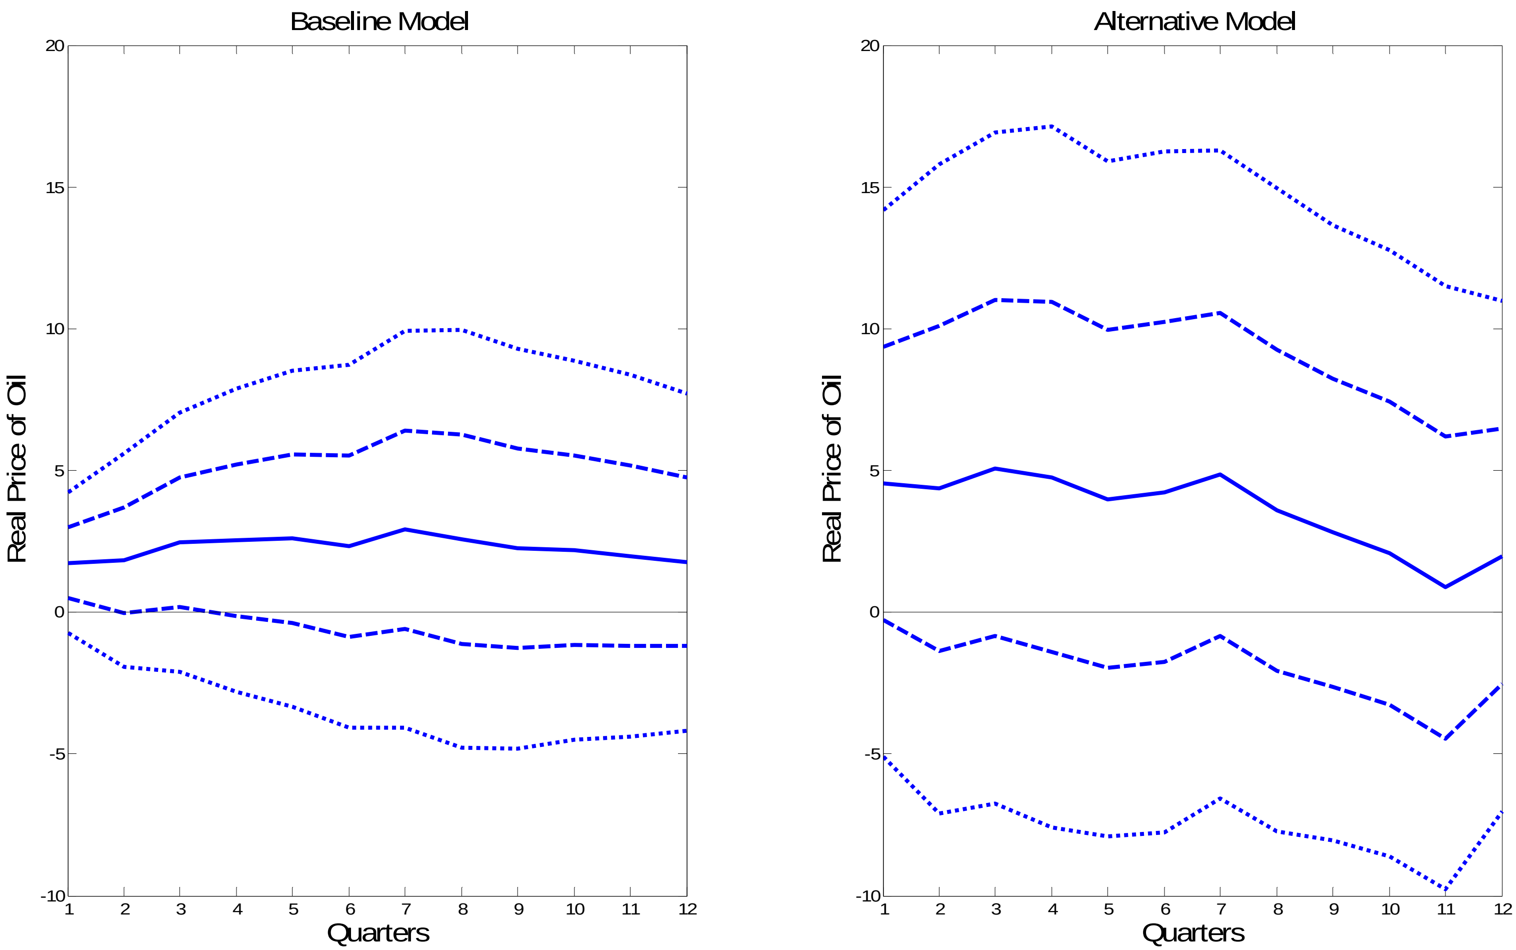

NOTES: The real price of oil has been specified in levels. Similar results would have been obtained if the real price of oil had been specified in differences and the impulse responses had been cumulated. The confidence intervals have been constructed using appropriate bootstrap methods as in Figures 2 and 4.