

Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?

Lutz Kilian
University of Michigan
CEPR

October 22, 2005

Abstract: Since the oil crises of the 1970s there has been strong interest in the question of how oil production shortfalls caused by wars and other exogenous political events in OPEC countries affect oil prices, U.S. real GDP growth and U.S. CPI inflation. This study focuses on the modern OPEC period since 1973. The results differ along a number of dimensions from the conventional wisdom. First, it is shown that under reasonable assumptions the timing, magnitude and even the sign of exogenous oil supply shocks may differ greatly from current state-of-the-art estimates. Second, the common view that the case for the exogeneity of at least the major oil price shocks is strong is supported by the data for the 1980/81 and 1990/91 oil price shocks, but not for other oil price shocks. Notably, statistical measures of the net oil price increase relative to the recent past do not represent the exogenous component of oil prices. In fact, only a small fraction of the observed oil price increases during crisis periods can be attributed to exogenous oil production disruptions. Third, compared to previous indirect estimates of the effects of exogenous supply disruptions on real GDP growth that treated major oil price increases as exogenous, the direct estimates obtained in this paper suggest a sharp drop after five quarters rather than an immediate and sustained reduction in economic growth for a year. They also suggest a spike in CPI inflation three quarters after the exogenous oil supply shock rather than a sustained increase in inflation, as is sometimes conjectured. Finally, the results of this paper put into perspective the importance of exogenous oil production shortfalls in the Middle East. It is shown that exogenous oil supply shocks made remarkably little difference overall for the evolution of U.S. real GDP growth and CPI inflation since the 1970s, although they did matter for some historical episodes.

Key Words: Oil shock; war; counterfactual; oil supply; exogeneity; weak instruments.
JEL: E32, C32.

Acknowledgements: This paper has benefited from discussions with Robert Barsky, Dan Goldstein, Yuriy Gorodnichenko, Jim Hamilton, Atsushi Inoue, Matthew Shapiro, Jim Stock and Linda Tesar. Correspondence to: Lutz Kilian, 238 Lorch Hall, Department of Economics, 611 Tappan Street, University of Michigan, Ann Arbor, MI 48109-1220. Email: lkilian@umich.edu.

1. Introduction

Since the oil crises of the 1970s there has been strong interest in the question of how oil production shortfalls caused by wars and other exogenous political events in OPEC countries affect oil prices, U.S. real GDP growth and U.S. CPI inflation (see, e.g., Barsky and Kilian 2002, 2004; Bernanke, Gertler and Watson 1997; Hamilton 1983, 1996, 2003; Hoover and Perez 1994). This paper focuses on three key questions: How large are the exogenous fluctuations in the production of oil? To what extent do exogenous oil supply shocks explain changes in the price of oil? What are the dynamic effects of exogenous oil production shortfalls on U.S. real GDP growth and CPI inflation?

I introduce a new methodology for quantifying the shortfall of OPEC oil production caused by exogenous political events such as wars or civil disturbances. This methodology utilizes monthly production data for all OPEC countries and for aggregate non-OPEC oil production that are available from the U.S. Department of Energy since January 1973. It is based on the observation that any attempt to identify the timing and magnitude of these exogenous production shortfalls requires explicit assumptions about the counterfactual path of oil production in the absence of the exogenous event. The strategy is to generate the counterfactual production level for the country in question by extrapolating its pre-war production level based on the average growth rate of production in other countries that are subject to the same global macroeconomic conditions and economic incentives, but are not involved in the exogenous event. Which countries belong into this benchmark group must be decided on a case-by-case basis drawing on historical accounts and industry sources. This approach allows the construction of a monthly time series of exogenous oil production shortfalls since 1973. The change over time in this exogenous production shortfall series (expressed as a percent share of world oil production) provides a natural measure of the exogenous oil supply shock.

The proposed method of quantifying exogenous production shortfalls has five distinct advantages compared to the conventional approach based on quantitative dummy variables as discussed in Hamilton (2003): (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. This point is important as my analysis shows that wars in the Middle East, for example, may actually cause higher oil production, when the parties involved resort to oil exports to finance the war.

Measures of exogenous oil production shortfalls are particularly useful for estimating the dynamic effects of exogenous fluctuations in oil production on macroeconomic aggregates such as real GDP growth. One possibility is to use measures of the exogenous oil production shortfall as instruments in regressions that project macro aggregates on the price of oil, as suggested by Hamilton (2003). I demonstrate that such instrumental variable (IV) estimates tend to suffer from a weak instrument problem, calling into question the reliability of the empirical results. I propose an alternative, simpler approach to estimating the dynamic effects of exogenous oil supply shocks that does not require IV estimation, but only involves ordinary least squares (OLS) projections and standard methods of inference. The basic idea is that if we are simply interested in measuring the extent to which exogenous oil supply shocks cause lower real GDP growth, it suffices that we project real GDP growth on a constant, the current and lagged values of the exogenous oil supply shocks and lagged values of real GDP growth. This approach provides a direct estimate of the dynamic effects of exogenous oil supply shocks on macroeconomic aggregates.

Using this approach and the new exogenous oil supply shock measure, I find statistically significant evidence of a sharp drop in real GDP growth five quarters after an exogenous oil supply shock and of a spike in CPI inflation three quarters after the shock. I also study the question of how these macroeconomic aggregates would have evolved in the absence of exogenous oil production shocks. I show that the effects of exogenous oil supply shocks on U.S. real GDP growth and CPI inflation were comparatively small on average, but that they did matter for particular historical episodes such as the 1990/91 Persian Gulf War. The empirical conclusions based on the proposed measure of exogenous oil supply shocks are shown to be robust to many changes in the counterfactual.

In addition, I investigate the predictive content of exogenous oil supply shocks for changes in the real price of oil. My analysis suggests that exogenous oil production shortfalls are of limited importance in explaining oil price changes during crisis periods. This result is robust to the choice of the exogenous oil supply shock measure. Of the episodes studied, only the 1980/81 oil price increases can be attributed to exogenous oil supply disruptions. Exogenous oil

supply disruptions were an important contributing factor for the oil price increases in the third quarter of 1979, but do not help explain the subsequent oil price increases during that episode. Exogenous oil supply disruptions explain only a small fraction of the oil price increases during the 1973/74, 1990/91 and 2002/03 episodes. This finding is suggestive of an important role for other explanatory variables such as shifts in the demand for oil or shifts in the uncertainty about future oil supplies not driven by actual production shortfalls. The latter explanation is shown to explain the bulk of the oil price fluctuations during the 1990/91 Persian Gulf War episode.

The results in this paper differ along a number of dimensions from the conventional wisdom in the literature. First, I find that the timing, magnitude and even the sign of exogenous oil supply shocks may differ greatly from current state-of-the-art estimates. Second, the results in this paper differ markedly from the common view that “the case for exogeneity of at least the major oil price shocks is strong” (Bernanke, Gertler and Watson 1997, p. 93). The evidence suggests that this view is supported by the data for the 1980/81 and 1990/91 oil price shocks, but not for other oil price shocks. I also show that statistical measures of the net oil price increase relative to the recent past do not represent the exogenous component of oil prices. In fact, only a small fraction of the observed oil price increases during crisis periods can be attributed to exogenous oil production disruptions. Third, compared to previous indirect estimates of the effects of exogenous supply disruptions on real GDP growth that treated major oil price increases as exogenous, my direct estimates suggest a sharp drop after five quarters rather than an immediate and sustained reduction in economic growth for a year. I also find a spike in CPI inflation three quarters after the exogenous oil supply shock rather than a sustained increase in inflation, as is sometimes conjectured. Finally, my results put into perspective the importance of exogenous oil production shortfalls in the Middle East. A counterfactual historical analysis suggests that exogenous oil supply shocks made remarkably little difference overall for the evolution of U.S. real GDP growth and CPI inflation since the 1970s, although they did matter for some historical episodes.

The remainder of the paper is organized as follows. Section 2 reviews the exogenous oil supply shock dates since 1973. Section 3 proposes explicit counterfactuals for each of these events that allow the construction of a monthly time series of the exogenous OPEC oil production shocks. Section 4 uses these data to assess the impact of exogenous variation in the supply of oil on U.S. CPI inflation and real GDP growth. This section presents dynamic

responses to exogenous oil supply shocks as well as counterfactual simulations of the path of U.S. macroeconomic aggregates following exogenous oil supply shocks. Section 5 explores the tenuous link between exogenous oil supply shocks and oil price shocks. I show that standard measures of “oil price shocks” based on nonlinear transformations of oil price data do not identify the exogenous component of oil price changes, as is sometimes claimed. This section also highlights the importance of measuring oil price shocks in real as opposed to nominal terms. Section 6 contrasts the approach to constructing explicit measures of exogenous oil supply shocks introduced in this paper with the commonly used quantitative dummy variable approach. In section 7, I compare the explanatory power of these two measures of exogenous oil supply shocks for the change in the real price of oil. Finally, section 8 contrasts this paper’s approach to constructing impulse responses to the approach of using exogenous oil supply shocks as instruments for oil price changes in regressions that relate real GDP growth to the price of oil. Using recently developed statistical tests for weak instruments I show that these IV methods tend to be unreliable, regardless of the measure of exogenous oil supply shocks used. Section 9 concludes.

2. A Review of Exogenous Oil Shock Dates since 1973

The analysis of oil supply shocks in this paper focuses on the modern OPEC period that began in late 1973. It is based on data for the period of January 1973 through September 2004. One reason is that the proposed methodology relies heavily on monthly oil production data compiled by the Department of Energy that are only available back to January 1973. A second reason is that the year 1973 marks a watershed in the institutional structure of the oil market. It coincided with a radical shift in power from the major oil companies, who only a few years earlier had refused to even recognize the existence of OPEC, to the oil producing countries in the Middle East. By early 1974, OPEC members for the first time had taken into their own hands control of crude oil-producing operations (through majority interest in the producing ventures), control over prices and control over production volumes (see Seymour (1980), p. 125). Thus, there is reason to believe that the transmission of exogenous oil supply shocks to oil prices changed in late 1973.

We are interested in identifying shocks to OPEC oil production that are exogenous with respect to U.S. macroeconomic aggregates. Natural candidates for such shocks are the oil production disruptions caused by the Iranian revolution of 1978/79, by the Iran-Iraq war of 1980-

1988, by the Gulf war of 1990/91 and by the Iraq war of 2003. A good case can be made that the production cutback caused by the civil unrest in Venezuela in 2002/03 represents another such shock, as shown in section 3.¹ In contrast, there is no evidence that would suggest the occurrence of a major exogenous oil supply shock in Nigeria.²

Other studies of exogenous oil supply shocks after January 1973 have typically included in addition the Arab-Israeli war of October 1973 (which was followed by the Arab oil embargo from October 1973-March 1974). Should this 1973/74 episode be treated as an exogenous shock to oil production along with events such as the Iranian revolution or the Iran-Iraq war? The answer is not clear. As I will discuss below, no OPEC oil facilities were attacked during the October war, and there is no evidence of OPEC production shortfalls caused by military action. While the subsequent “Arab oil embargo” undoubtedly coincided with a major fall in Arab OPEC oil output (the extent of which I will quantify below), the question of whether this drop can be regarded as exogenous with respect to U.S. macroeconomic conditions remains a topic of debate (see Barsky and Kilian (2002) and Hamilton (2003) for further discussion). This is not the place to review the arguments for and against the exogeneity of the 1973/74 oil embargo. Instead I will present baseline results that include this event among the exogenous events, consistent with earlier studies, as well as additional results that exclude it. It turns out that my qualitative results are not affected by whether the production cutbacks associated with the 1973/74 Arab oil embargo are treated as exogenous or not.

3. Constructing a Counterfactual

The proposed methodology utilizes monthly production data for all OPEC countries and for aggregate non-OPEC oil production that are available from the U.S. Department of Energy since January 1973. Consider an exogenous event such as a war (or civil disturbance) that lowers oil production in an OPEC country. The central idea is to generate the counterfactual production level for this country by extrapolating its pre-war production level based on the average growth

¹ Based on Hamilton’s (2003) methodology, the last two events were each associated with a 3.5 % reduction in world oil supply. Although individually those shocks are only about half the size of the earlier shocks, they occurred within a few months of one another. By this metric, their joint effect should be of the same magnitude as that of the Iran-Iraq war, for example, and only slightly smaller than the Persian Gulf War, the Iranian revolution or the October 1973 war and subsequent oil embargo.

² Clearly, there are more exogenous political events in the region such as the Yemeni civil war from May to July 1994, the Egyptian-Libyan War of July 1977 or Israel’s invasion of Lebanon in June 1982. What distinguishes the events we focus on from the latter examples is that they caused a physical disruption of oil production in OPEC countries.

rate of production in other countries that are subject to the same global macroeconomic conditions and economic incentives, but not directly affected by the exogenous event in question. Which countries belong into this benchmark group must be decided on a case-by-case basis drawing on historical accounts and industry sources.

This method allows the construction of a monthly time series of exogenous oil production shortfalls that takes full account of the timing of the exogenous production shortfall and of variations over time in its magnitude. In sections 3.1-3.6., I illustrate the implementation of this approach for the exogenous events listed in section 2, starting with the 1973/74 event. In section 3.7., I construct a measure of the aggregate exogenous production shortfall across all OPEC countries. The change over time in this exogenous production shortfall series (expressed as a percent share of world oil production) provides a natural measure of the exogenous oil supply shock.

3.1. Counterfactual for the 1973 October War and the 1973/74 Arab Oil Embargo

The 1973/74 oil shock episode involves three conceptually distinct components: (1) the unplanned production shortfall caused by military action during the October 1973 Arab-Israeli war; (2) the oil export embargo of Arab oil producers targeted against selected OECD countries that were viewed as pro-Israel; and (3) the deliberate production cutbacks implemented by some Arab oil producers toward the end of 1973. Of these three distinct components only the last one matters for our purposes.

Although some of the decline in OPEC oil production after September 1973 occasionally is attributed to the destruction of oil facilities during that war, there is little evidence for that view. The weekly Oil and Gas Journal on October 15, 22 and 29 reported in detail about war-related damage to oil facilities in the region, notably in Syria, Lebanon, Israel and Egypt. No damage in Iraq or any other OPEC country was reported. Thus, we can rule out the notion that, by destroying oil facilities in OPEC countries, the war caused an unplanned production shortfall.³ Similarly, the politically motivated embargo against specific countries was not by itself the cause of a reduction in oil supplies. Although shipments of oil to some oil-importing

³ Although there is no evidence of a reduction in OPEC oil output directly caused by the war, there was a reduction in Iraqi oil exports to the Mediterranean when the Tripoli oil terminal in Lebanon was attacked in October 1973. The Oil and Gas Journal conjectured that the pre-war export volume on that route may have been as high as 0.500 million barrels per day. The reduction in Iraqi production in October is 0.320 million barrels per day relative to September. It is unclear to what extent that production cut was related to the Tripoli attack.

countries were initially conditioned on specific diplomatic demands, it was quickly understood by all sides that such targeting would not and could not be effective because oil can be re-sold or simply diverted (see Terzian 1985, p. 178-180 for examples). For that reason I focus on the production cutbacks that took place during the Arab oil embargo between October 1973 and March 1974 as a result of deliberate policy decisions. In assessing the effect of this policy, the key question is how large this temporary production shortfall actually was. The construction of a counterfactual involves, first, an assessment of what normal production levels would have been prior to the oil shock. Second, it involves an assumption about how production would have developed in the absence of the war and embargo.

It is well known that oil production from Arab OPEC countries fell between September and November of 1973, whereas oil production in the rest of the world did not. It may seem natural at first to attribute the entire differential to the effects of the embargo. That approach would be misleading, however, because Middle Eastern OPEC oil producers were subject to different economic incentives than producers in the rest of the world in the period leading up to October 1973. The difference in incentives arose as a result of the 1971 Tehran/Tripoli agreements between the oil companies and Middle Eastern OPEC oil producers. These five-year agreements in short provided a moderate improvement in the financial terms that host governments received from oil companies for each barrel of oil extracted by the oil companies in exchange for assurances that these governments would allow the oil companies to extract as much oil as they saw fit on those terms (see Seymour 1980, p. 80).⁴ The latter option principally affected Saudi Arabia as the country with the largest spare capacity. As reported by the Oil and Gas Journal of November 12, 1973, projections of future Saudi oil production under the leadership of U.S. oil companies up to October 1973 envisioned ever expanding oil production “based on the premise that this is how much oil the Middle East would need to produce to balance the world’s oil demand.”

When global demand for oil accelerated in 1972/73, reflecting a worldwide economic boom (see Seymour 1980, p. 100), some Middle Eastern countries were operating close to capacity already and unable to increase oil output; whereas others, notably Saudi Arabia, had the capacity to increase oil output and did increase output prior to October 1973, albeit reluctantly.

⁴ Wilkins (1976, p. 168) cites the 1974 testimony of a representative of Exxon before the U.S. senate that the oil companies in Tehran “... made every effort to minimize posted price increase and to assure some security of supply.”

This reluctance can be attributed to the fact that the posted price agreed upon in 1971 might have been reasonable at the time, but was quickly eroded in real terms as a result of dollar devaluations and rising U.S. inflation. This development caused increasing opposition to the Tehran/Tripoli agreements that intensified in March of 1973 and culminated in the repudiation of the agreements by the oil producing countries in early October of 1973. According to the Oil and Gas Journal article, the Saudi refusal to supply virtually limitless quantities of oil after September 1973, though perhaps prompted by the Mideast war and embargo, had “little to do with politics”, but with the fact that Saudi Arabia in October 1973 had become independent of foreign oil companies and for the first time was able to make a “rational decision” about her production levels. In other words, with the repudiation of the Tehran/Tripoli agreements Saudi production levels should have reverted to normal levels, consistent with the level of production in non-OPEC countries, even in the absence of an embargo. Thus, a substantial fraction of the observed decline in Arab oil output in late 1973, notably in Saudi Arabia and Kuwait, was simply a reversal of an unusual increase in Arab oil production (relative to non-OPEC oil producers) earlier that year in response to pressure from the oil companies.

Put differently, as noted by Barsky and Kilian (2002), even in absence of an embargo one would have expected a reduction in oil production in the OPEC countries to the extent that these countries prior to October 1973 were forced by long-term contracts to supply higher quantities of oil to the oil companies than the countries would have wished for at the low price agreed upon with the oil companies prior to the shift of the demand curve for oil. This conjecture may be verified empirically. Monthly oil production data for the OPEC countries are available starting in January 1973. A direct implication of the argument in Barsky and Kilian (2002) is that – in the absence of an exogenous oil production cut - one would expect to see reductions in oil production in late 1973 in those countries that expanded output in early 1973, but not in the other countries. Only to the extent that there are additional cuts beyond the production level that prevailed before the expansion, can these cuts be associated with the embargo.

A natural benchmark against which to judge the validity of this explanation is the oil production of non-OPEC countries over the same time period. The upper panel of Figure 1 shows the number of extra barrels each country produced per day since March 1973 relative to the production level one would have expected, had that country’s production grown at the same rate as non-OPEC oil production. March 1973 is a natural starting date, since at this point the

pressure to abandon the existing oil market regime intensified, suggesting that OPEC countries were forced to produce more than they would have wanted to at the prevailing price of oil. I also explored earlier starting dates; the results are qualitatively similar, but the production shortfall associated with the embargo is smaller than under the March scenario. Figure 1 shows that Saudi Arabia and to a lesser extent Kuwait produced disproportionately more oil after March 1973 than other countries. Their output peaked in mid-1973. As conjectured earlier, all of the reduction in Saudi and Kuwaiti oil production in October 1973 (and some of the decline in November) can be understood as a return to normal production levels by international standards. Only the fall in oil production below the counterfactual in November and December of 1973 is *prima facie* evidence of the effects of the Arab oil embargo.

In contrast, the other Arab members of OPEC combined never experienced an unusual increase in oil production in early 1973.⁵ Thus, the decline in their oil output relative to international levels that started in November 1973 can be attributed in its entirety to the oil embargo. The data in Figure 1 also show that the Saudi and Kuwaiti embargo was effectively over by January 1974 (when the price of oil stabilized at its new, much higher level). Among other Arab OPEC countries the cutback persisted much longer, yet at diminishing levels. In both cases, the bulk of the output reduction is concentrated in November and December. Finally, Figure 1 shows that Iraq never was a major contributor to the production shortfall associated with the oil embargo, consistent with its public rhetoric (see Seymour 1980, p. 119).

Since the production levels in Figure 1 are normalized relative to non-OPEC output, they already account for the increase in output stimulated by the initial increase in oil prices in early October. Since the increase in prices announced in early October 1973 had been determined prior to the war and was motivated on economic grounds, the case can be made that it would have been announced in any case (see, e.g., Seymour 1980, p. 98, p. 107, p. 113).⁶ This interpretation is also consistent with the participation of Iran, a non-Arab country that was not a participant in the war, in the early October oil price increase and with the fact that the price

⁵ This result also holds when considering each of these countries separately.

⁶ Consistent with this widely held view, the *Oil and Gas Journal* (October 15, 1973, p. 44) summarized the purpose of the October 8 OPEC talks as follows: “The [oil] producing countries are seeking revision of the 1971 Tehran Agreement to take account of inflation and of increased product prices in world markets. The Tehran pact [which was initially supposed to be binding until 1976] provides for only 2.5% annual inflation”. For a detailed discussion of this point see Penrose (1976).

increase was adopted by other non-Arab OPEC countries such as Venezuela and Nigeria.⁷

The lower panel of Figure 1 shows the total production shortfall associated with the October War and the Arab oil embargo. The totals shown represent the contributions of all Arab OPEC countries, i.e., Saudi Arabia, Kuwait, the United Arab Emirates, Qatar, Libya, Algeria and Iraq. There is no evidence of a production shortfall in October 1973. At its peak in November and December the embargo involved an output reduction of about 2.67 million barrels per day. In January and February, the shortfall drops to 0.80 and 0.57 million barrels per day. Thereafter, the shortfall is zero.

3.2. Counterfactual for Iran

The beginning of the Iranian revolution can be dated in October 1978, when striking workers paralyzed oil installations across the country. Although the abdication of the Shah and the transition to the Khomeini government was complete in January 1979, the effects of the revolution on Iranian oil production continue even today. How would Iranian production have developed in the absence of the Iranian revolution? Presumably OPEC would have operated much as it did in the period January 1974-September 1978, with overall OPEC production being determined in accordance with oil market conditions and macroeconomic developments, as described by Barsky and Kilian (2002), and with fairly stable market shares for each member country.

Any unusual departure of a country's OPEC market share from historical levels thus would be an indication that this country changed its production levels in response to the Iranian revolution. Table 1 shows production shares of selected OPEC countries before and after the Iranian revolution. The shares in the first two columns are computed with respect to total OPEC production minus Iranian production. With the exception of Saudi Arabia and Iraq, the production shares remain remarkably constant. The latter two countries greatly increased their production share (as well the number of barrels produced per day) after the Iranian revolution. As the last two columns of Table 1 show, after eliminating Iraq and Saudi Arabia from the total, OPEC production shares are roughly constant over time.

⁷ To the extent that subsequent oil price increases were driven by the embargo and/or the war (as opposed to strong demand for oil), and to the extent that these price increases stimulated additional non-OPEC oil production over the period in question, the relative fall in Arab oil output after October 1973 may be inflated. We ignore this possibility given the existence of capacity constraints in most non-OPEC oil-producing countries that made short-term increases in oil output difficult (see Seymour 1980, p. 100). The same constraints also applied to the remaining OPEC countries.

This finding is not surprising. I will show in sections 3.3 and 3.5 that Iraq and Saudi Arabia as well experienced exogenous production disturbances as a result of the Iranian revolution. The growth of OPEC production excluding Iran, Iraq and Saudi Arabia thus provides a natural benchmark against which to compare the actual growth in Iranian oil production immediately after the revolution. This benchmark allows us to construct a counterfactual path for Iranian oil production from October 1978 until August 1980.

The outbreak of the Iran-Iraq war in September 1980 does not require any modification in this benchmark since we already have excluded all countries whose production was directly affected by this second exogenous shock. The relevant counterfactual continues to be the behavior of total OPEC production minus the combined production of Iran, Iraq and Saudi Arabia until August 1990, at which point the benchmark becomes total OPEC minus Iran, Iraq, Saudi Arabia and Kuwait, as Kuwait experiences an exogenous production shortfall of its own and must be dropped from the benchmark. Finally, following the civil unrest of December 2002, I also drop Venezuela from the benchmark. The upper panel of Figure 2 shows the complete counterfactual path together with the actual Iranian production levels. The vertical difference between these paths may be viewed as a measure of the exogenous shortfall of Iranian oil production. This distance is plotted in the second panel.

In constructing this counterfactual I ignore the possibility that the OPEC countries in the benchmark group in response to higher oil prices triggered by the Iranian revolution may have expanded their production of oil if not immediately, then over time. To the extent that they did, I will overestimate the extent of the exogenous production shortfall. Of course to the extent that the observed increase in the price of oil after the Iranian revolution was due to strong demand for oil driven by a worldwide economic expansion (see Barsky and Kilian (2002) for further discussion), any response of oil production to higher oil prices must be appropriately treated as part of the benchmark.

Moreover, the notion that surely the benchmark countries listed above must have increased their oil production after the Iranian revolution (and again after the outbreak of the Iran-Iraq war) is at odds with the fact that actual oil production in these countries fell sharply starting in 1979, presumably reflecting the cooling of the world economy (see upper panel of Figure 2). It also seems at odds with the strong degree of co-movement for the remainder of the sample between oil production in the benchmark countries and in Iran at all but the highest

frequencies (see upper panel of Figure 2). It is of course possible that production would have fallen even more sharply in the absence of these exogenous events. In this sense the counterfactual proposed here may be biased in the direction of overestimating the extent of the exogenous production shortfall.

Finally, I make the implicit assumption that the global recession of the early 1980s that was mainly responsible for the oil production cutbacks in the benchmark countries at the time was driven by a global monetary contraction rather than the delayed effects of the Iranian revolution and the outbreak of the Iran-Iraq war. As the historical decompositions in section 4.3 will show this assumption may be questioned for some periods in the sample, but still provides a fairly good approximation for a first pass at the data.

It is useful to reflect for a moment upon the differences between this series and conventional quantitative dummy measures (of the type more fully discussed in section 6) that implicitly assume that Iranian oil production would have continued at September 1978 levels. On impact, the counterfactual in Figure 2 suggests that conventional measures underestimate the production shortfall because Iranian production would in all likelihood have increased slightly in 1979 if there had not been a revolution. Starting in 1980, however, OPEC production fell for reasons presumably not related to the Iranian revolution, implying that conventional oil shock measures grossly overestimate the shortfall in production caused by the Iranian revolution.

It is difficult (and for our purposes unnecessary) to construct a separate counterfactual for the outbreak of the Iran-Iraq war. Nevertheless, Figure 2 suggests that – apart from a very short-lived fall in production in 1980/81 – this war on balance had a positive effect on the supply of Iranian oil (as evidenced by the narrowing of the two bands in Figure 2). The necessity to finance military expenditures and easy access to the Gulf waters is likely to have spurred efforts to increase oil production in Iran in the 1980s.

3.3. Counterfactual for Iraq

The counterfactual for Iraq begins in September 1978. Although Iraq was not directly involved in the Iranian revolution, this exogenous event by releasing previous constraints set a chain of events in motion that culminated in the Iran-Iraq war of September 1980. Given the close relationship between the Shah and the United States and the military strength of Iran under the Shah, the idea of Iraq invading Iran prior to October 1978 would have been unthinkable. In fact, relations between the two countries were cordial to the point that Iraqi authorities imposed

increasingly severe restrictions on the movements of Iranian dissidents in Iraq such as the fundamentalist Iranian Shia leader Ayatollah Khomeini, prompting Khomeini to leave Iraq for France in September of 1978 (see Terzian 1985, p. 278). The fall of the Shah and Khomeini's seizure of power led to a break with the United States and a deterioration of Iranian military capabilities. In addition, the secular and socialist Baathist regime in Iraq felt threatened by the fundamentalist Shia movement in Iran, given its own majority Shia community. Thus, the Iranian revolution provided both opportunity and added motive for Iraq to attack Iran.

The change in Iraqi policy towards its neighbor clearly would not have happened, had the Iranian revolution not occurred. The Iraqi preparation for the coming war started immediately after the fall of the Shah. Iraq began stockpiling substantial stocks of arms and spare parts as well as foreign exchange reserves amounting to \$35 billion right before the war. Just three weeks after the war broke out Iraqi authorities boasted that they would be able to maintain the war effort for a year without exporting any oil at all (see Terzian 1985, p. 282, for details).⁸ The ability to build such reserves was crucial for Iraq since geography made oil exports much more vulnerable for Iraq than for Iran in case of a war. The only way Iraq could accumulate such reserves and pay for arms and spare parts was to export oil at an unprecedented rate. In this sense, the Iranian revolution caused an increase in Iraqi oil production. It is this increase that is reflected in the increased production share of Iraq after September 1978 in Table 1 and that motivated the exclusion of Iraq from the benchmark, when discussing the Iranian counterfactual. This reasoning further suggests that we view the increase in Iraqi oil production relative to other OPEC countries between October 1978 and August 1980 as an exogenously driven oil supply shock rather than a supply response to market conditions. Put differently, the preparations for the Iran-Iraq war were arguably an endogenous response to the Iranian revolution, yet clearly exogenous with respect to global macroeconomic conditions.⁹

⁸ There are no publicly available data documenting this build-up of foreign exchange reserves. Since Iraq stopped reporting data on its foreign exchange reserves to the IMF in 1978, the series ends in December 1977.

⁹ The shift in Iraqi policy in 1979 was not mirrored by similar preparations for war in Iran, which for the time being was pre-occupied with the revolution. In fact, when Iraq launched a large-scale military assault in September 1980, the Iranian leadership was taken by surprise, despite earlier warnings of Iraqi intentions. Part of the reason was that the Iranian army at this point had been rendered ineffective by the revolution; part that the Iranian regime was preoccupied with an internal power struggle and part that the Iranian government considered news of the Iraqi war preparations a deliberate piece of Soviet misinformation, a reasonable response given the Soviet invasion of neighboring Afghanistan (see Terzian 1985, p. 278-281).

A suitable counterfactual for Iraq can be constructed by comparing actual Iraqi oil output to the production level that would have prevailed, had Iraqi production after September 1978 grown at the same rate as that of total OPEC minus the combined production of Iraq, Iran and Saudi Arabia. The rationale for excluding Iran and Iraq from the benchmark is self-evident. I exclude Saudi Arabia given the evidence, discussed in section 3.5, that Saudi Arabia itself was subject to three temporary exogenous production shocks in 1978/79, 1990/91 and 2002/03. Moreover, as discussed in section 3.5., Saudi production decisions starting in mid-1979 were heavily influenced by price developments within OPEC that in turn may have been driven in part by exogenous events. This observation suggests that it is safer to exclude Saudi oil production data from the baseline benchmark, although I will consider including Saudi production data as part of my sensitivity analysis. Finally, further adjustments to the counterfactual were made for Kuwait starting in August 1990 and for Venezuela starting in December 2002, when those countries experienced their own exogenous production shortfalls.

The upper panel of Figure 3 plots the actual and counterfactual production levels. There is clear evidence of the disproportionate increase in Iraqi production after the Iranian revolution. Accounting for this temporary increase in output is essential for assessing the shortfall following the outbreak of the Iran-Iraq war in September 1980. A substantial proportion of the actual fall in Iraqi output in 1980 is simply a reversal of the earlier positive shock, so conventional measures of oil shocks are bound to overestimate the exogenous fall in output. As in the case of Iran, oil output recovered in the early 1980s (as Iraq was able to make use of newly built pipelines beyond the reach of Iran) and exceeded the counterfactual even before the war had ended in 1988. Again, the enormous financial cost of the war is likely to have been the cause. Even after the war had ended, this build-up continued, presumably in preparation for the invasion of Kuwait.

Figure 3 shows that conventional oil production shortfall measures again overestimates the production shortfall that occurred in August 1990. At the same time, Figure 3 shows that Iraqi production remained at very low levels for years following the defeat in the Gulf war, owing no doubt to the sanctions and export restrictions imposed by the U.N. This effect as well will be completely ignored by conventional measures of the oil production shock. By 2000 Iraqi oil output had recovered to near-normal levels on average, as the effectiveness of sanctions had been eroded. Iraqi oil output collapsed again in March 2003 with the outbreak of the Iraq war.

3.4. Counterfactual for Kuwait

Kuwait was not affected directly by any exogenous event until the invasion of Kuwait in August 1990. A natural counterfactual for the remaining period can be based on the oil production of total OPEC minus the combined production of Kuwait, Iran, Iraq and Saudi Arabia, with Venezuela added to this list after December 2002, when that country experienced its own exogenous production shortfall. The rationale for the exclusion of Saudi Arabia is discussed in section 3.5. Figure 4 shows that by 1994 most of the production shortfall had been recovered, although production remained slightly below what one would have expected based on the benchmark.

3.5. Counterfactual for Saudi Arabia

Saudi Arabia differs from other OPEC oil producers in that it had the ability to act as a swing producer. Over the period in question it also served on several occasions as a supplier of last resort. This poses special challenges for the construction of the counterfactual.

The unusual increase in Saudi oil production starting in late 1978 has already been mentioned. There were two distinct motivations for this increase. Initially, the Saudis increased production to offset the incipient shortfall caused by the Iranian revolution. It was understood that this increase would be temporary and lapse as soon as the Iranian crisis was over. Indeed, the Saudis sharply curtailed their production in April of 1979.¹⁰ Since this initial spike in Saudi oil production in 1978/79 was a discretionary act and a direct response to the Iranian revolution (rather than to higher oil prices triggered by the Iranian revolution), the case can be made that it must be treated as exogenous with respect to the U.S. economy. The Saudi counterfactual for the Iranian revolution involves simply comparing its production during October 1978-April 1979 to that of other OPEC countries excluding Iran and Iraq. Since the response of Saudi production to the revolution – like the Arab oil embargo – is treated as strictly temporary, there is no need to model the counterfactual beyond April 1979.

¹⁰ Skeet (1988, p. 163) suggests that the April decision reflected the belief that the balance of supply and demand in the global oil market had reverted to an acceptable degree of normality given the Iranian production increase in March (also see Seymour 1980, p. 183). Global oil production data lend support to the Saudi view. Figure 5 shows global production levels during the Iranian revolution. During the first two months of the Iranian revolution global production actually increased. After the peak in November, there was a temporary fall from December 1978 until March 1979. From April 1979 on global production levels exceeded the September 1978 level of 62.477 million barrels.

In contrast, when oil prices accelerated again after April 1979, the Saudis sharply increased production for a different reason. The latter sustained increase was intended to force other OPEC members to bring oil prices down to levels deemed acceptable by the Saudis.¹¹ This response was clearly not exogenous and cannot be treated as an exogenous disturbance of oil production. Another (smaller and less pronounced) increase in Saudi output coincided with the outbreak of the Iran-Iraq war, but there is no historical evidence that this increase was motivated by anything else than the power struggle within OPEC. For that reason I do not include this event either.

There are, however, two more exogenous events that prompted an increase in Saudi production. After the invasion of Kuwait in August 1990, Saudi Arabia temporarily increased its production with the aim of offsetting the shortfall created by the war. The obvious counterfactual consists of the production growth of total OPEC-Saudi Arabia-Iraq-Iran-Kuwait. Similarly, the 2003 Iraq war prompted a pre-emptive increase in Saudi production. The Iraq war differs from all earlier episodes in that the war was anticipated and without a well-defined end. Nevertheless, it is clear that Saudi Arabia temporarily increased its oil production to offset jitters in the oil market. A counterfactual may be constructed by focusing on the period from July 2002 when the war option was increasingly discussed in public until October 2003, when the focus had shifted from open war to the pacification of Iraq.

Figure 6 shows that all three episodes that have been identified based on historical evidence were associated with sharp positive spikes in Saudi oil output relative to the relevant OPEC benchmarks for each event, lending credence to the view that Saudi Arabia's response was caused by the exogenous events in questions. It may be possible to identify smaller war-related peaks of a similar nature in the production data of other countries, but the magnitudes will be so much smaller that I abstract from this possibility. Although there are compelling reasons to think of the Saudi supply response on those three occasions as exogenous, it is interesting to pursue the alternative view that this response might have been endogenous after all. This possibility will be investigated in sections 4 and 8. As it turns out, the qualitative results are not affected.

¹¹ In April 1981, in an interview with NBC, the Saudi oil minister Yamani declared that if there was an oil surplus it was because his country had 'engineered it' to force other OPEC members to lower prices (see Terzian 1985, p. 287).

3.6. Counterfactual for Venezuela

The Venezuelan oil shock of December 2002 is not usually included in the list of exogenous oil shocks, but, as Figure 7 shows, is a natural candidate for an exogenously induced oil shock. The civil unrest in Venezuela was followed by a sharp and well-defined reduction in crude oil output. The counterfactual is based on total OPEC production minus the combined production of Venezuela, Iraq, Iran, Kuwait and Saudi Arabia, following the reasoning used in the analysis of the previous counterfactuals.

3.7. Constructing a Historical Series of Exogenous OPEC Oil Supply Disturbances

Figure 8 plots the sum of all OPEC oil production shortfalls discussed so far (including the 1973/74 event which may or may not have been exogenous). The upper panel displays levels and the lower panel plots shares in world oil output. Normalizing by world output makes little difference. A natural measure of the exogenous OPEC oil supply shock is the change in the normalized production shortfall over time. For the remainder of the paper, I will treat the change in the exogenous production shortfall shown in Figure 9 as the baseline measure of exogenous oil supply shocks (suitably aggregated to quarterly frequency where appropriate). Figure 9 shows that, as expected, all oil dates with the exception of the 2003 Iraq war are associated with large exogenous swings in oil production. The production shocks range from almost +4 % to almost -7 % of world crude oil production at monthly frequency. Exogenous changes in crude oil production driven by political events in OPEC countries account for about 6 % of the variability in world crude oil production changes, according to the baseline counterfactual.

As I will show in sections 4 and 8, my main empirical findings based on this baseline counterfactual are remarkably robust to alterations in this benchmark. These counterfactuals undoubtedly could be refined further. They ignore, for example, some more subtle considerations such as long-term declines in the productive capacity of countries, the damage done to oil fields by excessive production, the possible existence of capacity constraints in OPEC countries and in non-OPEC countries, distributional bottlenecks in the oil market, or the fact that crude oil is not a homogeneous commodity. It seems unlikely, however, that any of these factors would make a large difference in the current context.

4. The Dynamic Effects of Exogenous Variations in the Production of Oil on U.S.

Macroeconomic Aggregates

An important question from a policy point of view is how exogenous oil production disruptions, as defined in section 3, affect CPI inflation and real GDP. Since GDP data are available only at quarterly frequency, for this section, I rely on quarterly analogues of the data shown in Figure 9. A more detailed discussion of this quarterly series can be found in section 6.2, where I compare this series to the conventional approach of using quarterly quantitative dummies.

A central point of this paper is that it suffices that we regress real GDP growth on a constant, eight lagged values of the exogenous oil supply shocks and four lagged values of real GDP growth by OLS, if we are interested in whether exogenous oil supply shocks cause lower real GDP growth.¹² Provided that the exogenous regressor in question is independent of all other potential exogenous regressors (a plausible assumption in the present context), we can think of the exogenous variations in oil supply as a “treatment”, the effect of which can be measured by the dynamic response of real GDP growth to an exogenous shock in oil production. The latter response by construction will quantify the “tendency of the U.S. economy to perform poorly in the wake of ... historical conflicts [in the Middle East]” (Hamilton 2003, p. 364). Inference for this response is standard and OLS estimates are consistent. This type of regression approach provides the basis for the answer to the following two questions.

4.1. How Does U.S. Real GDP Growth Respond to Exogenous Oil Supply Shocks?

Point estimates of the dynamic effect of exogenous oil supply shocks on real U.S. GDP growth are shown in the upper panel of Figure 10. The estimates show a sharply negative growth rate five quarters after the oil supply shock, before the response reverts back to zero. This negative growth according to the 68% confidence band to some extent persists until the seventh quarter; according to the 95% confidence band it lasts only for one quarter.

Qualitatively similar results were obtained for a number of alternative specifications of the counterfactual. Specifically, one alternative counterfactual differed from the baseline counterfactual in that it ignores the contribution of the Saudi production response to the Iranian revolution, the Gulf war and the Iraq war in constructing the exogenous movements in oil output. Another counterfactual differed from the baseline in that it treats the production shortfall caused by the oil embargo of 1973/74 as endogenous. The third alternative counterfactual differed from

¹² Regressions that also included the contemporaneous value of the exogenous oil supply shock gave similar results.

the baseline in that Saudi Arabia with the exception of 1973/74 is included in the benchmark for the OPEC countries affected by exogenous shocks.

4.2. How Does U.S. CPI Inflation Respond to Exogenous Oil Supply Shocks?

The same OLS regression approach may also be used to study the response of CPI inflation. Point estimates of the dynamic effects of exogenous oil supply shocks on U.S. CPI inflation (in the lower panel of Figure 10) suggest a sharp spike in inflation three quarters after the shock. Otherwise the response is flat. The spike is significant at the 68% confidence level, but not the 95% confidence level. At all other horizons the response of CPI inflation is insignificant at the 68% level. These results are qualitatively unchanged for the alternative specifications of the counterfactual, although the point estimate of the spike in CPI inflation tends to be lower. The evidence that exogenous oil supply shocks over the post-1973 sample period have not caused sustained inflation is consistent with the theoretical arguments and related empirical evidence in Barsky and Kilian (2002, 2004).

4.3 How Would U.S. CPI Inflation and Real GDP Growth Have Evolved in the Absence of Exogenous Oil Supply Shocks?

A question of obvious interest is how the U.S. economy would have evolved since the 1970s in the absence of exogenous oil supply shocks. The tentative answer is that exogenous oil supply shocks made little difference overall, although they did matter for specific episodes.¹³ Figure 11 shows the path of CPI inflation and real GDP growth after subtracting from the actual value of these series the fitted value obtained from the exogenous oil supply shocks alone in the OLS regression. These estimates naturally are subject to considerable sampling uncertainty, so the resulting counterfactuals are only suggestive. All estimates are based on the baseline counterfactual constructed in this paper. These results are robust to the alterations in the counterfactual discussed earlier.¹⁴

Figure 11 suggests that real GDP growth would have been similar in the absence of exogenous oil supply shocks with occasional exceptions. The main effect of these exogenous shocks was lower growth in 1975, 1980, 1982, and 1991-93. There also is evidence that growth would have been slower in 1982-84, 1993-95 and 1999-2001 without the (mostly positive) oil

¹³ A similar conclusion was reached by Shapiro and Watson (1988) using different methods and data.

¹⁴ The counterfactual in Figure 11 is based on data since 1971.I to allow the inclusion of the 1973/74 episode in the counterfactual. The values of the exogenous oil supply shock series for 1971.I-1972.IV are set to zero for this purpose, since there were no exogenous oil supply shocks in 1971-1972.

supply shocks at that time. Real growth in 2003 and 2004 appears slightly lower due to exogenous oil shocks. There also is evidence that CPI inflation would have been lower in 1979, 1981 and 1991, but higher in 1982 and 1992-94 without exogenous oil shocks. Overall, CPI inflation is less sensitive to exogenous oil supply shocks than real GDP growth.

5. The Tenuous Link from Exogenous Oil Supply Shocks to Oil Price Shocks

Early work on the effect of exogenous oil supply shocks on U.S. real GDP was facilitated by the unique institutional arrangements in the oil market in the post-war period prior to 1973. Hamilton (1983) made the case that nominal oil prices over this period were effectively exogenous. None of these arguments apply to the oil market since 1973, however, and there is widespread recognition that today oil prices must be considered endogenous with respect to global macroeconomic conditions. Recently the case has been made that, nevertheless, nonlinear transformation of the price of oil designed to capture “oil price shocks” effectively identify the exogenous component of the price of oil. From our point of view these measures are of interest, because they seem to provide a simple alternative to the approach proposed in this paper. As the analysis below will demonstrate this is not the case.

5.1. Nominal Oil Price Shocks

Oil price shocks refer to unusually high oil prices relative to recent experience. Building on work by Mork (1989), Lee, Ni and Ratti (1995) and Hamilton (1996), Hamilton (2003) proposed a formal definition of oil price shocks based on the net increase in the nominal price of oil relative to the maximum of the price of oil over the previous three years. This definition has been used in studying the responses of U.S. sectoral and macroeconomic aggregates to oil price shocks (see, e.g., Bernanke, Gertler and Watson 1997; Davis and Haltiwanger 2001; Lee and Ni 2002, Hamilton 2004).

Figure 12 applies this definition to the refiner acquisition cost of imported crude oil as defined by the Department of Energy (and extrapolated backwards until January 1972 based on the wholesale price index of crude oil). The upper panel of Figure 12 shows that by this definition major nominal oil price shocks occurred in 1973/74, 1979-1980, 1981, 1990/91, 2000/01, 2002/03 and 2004. Additional minor nominal oil price shocks occurred in 1975/76, 1989/90 and 1996/97. Hamilton (2003) suggested that the predictive power of these net oil price increases for macroeconomic aggregates such as U.S. real GDP growth can be attributed to their

ability to filter out influences on oil prices that do not come from exogenous political or military events abroad.¹⁵ This interpretation is contradicted by the fact that not every oil price shock in Figure 12 can be associated with an exogenous political event in the Middle East (indicated by vertical lines). This is true not just for the relatively minor shocks, but also for some of the more important spikes such as those in 2000/01 and in 2004. Given that oil price shocks are not necessarily preceded by exogenous Middle Eastern events, it is not clear that the observed increase in oil prices after those exogenous events was actually caused by those events.

5.2. Real Oil Price Shocks

My analysis has focused on nominal oil prices so far, following the convention in the literature. From an economic point of view it is, of course, the real price of oil that matters in defining an oil price shock. The distinction makes a difference in defining the oil price shock. The middle panel of Figure 12 is based on the same oil price data as the upper panel, but deflated by the U.S. CPI. The first observation is that nominal oil price shocks are not necessarily real oil price shocks. Second, there is no real oil price shock at the time of the Venezuelan crisis in late 2002 or of the 2003 Iraq war, demonstrating that exogenous oil supply shocks need not be followed by real oil price shocks. Third, as in the nominal case, there are oil price shocks that are not immediately preceded by exogenous events, notably in 1996/97, 1999/2000, and 2004. Fourth, there are dramatic changes in the relative importance of oil price shocks. For example, the 1974 event is much more important than the 1991 event in real terms, whereas focusing on nominal data causes a reversal of this relationship. Fifth, the timing of oil price shocks may be affected by the distinction between nominal and real prices. For example, the oil price shock after the Iranian revolution in real terms occurs only in 1979, not in late 1978 as the nominal data would suggest. We will focus on the real price of oil in the remainder of the paper.

5.3. Did OPEC Cause Price Shocks in Other Industrial Commodities?

As discussed by Barsky and Kilian (2002, 2004), sharp increases in the real price of oil may occur even in the absence of exogenous events, for example, when increases in the demand for oil occur in the presence of capacity constraints in the oil market, which would explain the existence of oil price shocks at dates not related to exogenous events in the Middle East. By the

¹⁵ For example, Hamilton (2003, p. 395) writes that nonlinear transformations of this type "... filter out many of the endogenous factors that have historically contributed to changes in oil prices" and "seem in practice to be doing something rather similar to isolating the exogenous component of oil price changes" (p. 391).

same token one cannot simply attribute to exogenous events in the Middle East oil price shocks that occurred after those events because at least part of the observed oil price shock may have occurred even in the absence of that event.

To illustrate this point the last panel of Figure 12 displays the real price shocks computed using Hamilton's net price increase method for the linearly detrended CRB series of real non-oil industrial commodity prices.¹⁶ That panel demonstrates that the 1973/74 and 1978/1979 oil dates were associated with substantial non-oil industrial commodity price shocks. An uninformed reader may be forgiven for concluding that OPEC must have caused those commodity price shocks too, but it is well known of course that these price shocks had nothing to do with exogenous events in the Middle East, but were driven by global demand pressures that affected all industrial commodities, including oil. This example illustrates by analogy the danger of spuriously associating exogenous events with oil price shocks. In addition, the magnitude of the price shock in industrial commodities in the early 1970s suggests that the bulk of the 1973/74 oil price shock may have been due to rising demand for oil rather than the Arab oil embargo. I will return to this point below when I assess the explanatory power of exogenous oil production shortfalls for real oil prices. The delayed increase of the oil price relative to other industrial commodity prices can be explained by the institutional constraints of the oil market prior to 1974 that were discussed in section 3. Also note that for the oil dates of 1980 and 1990/91 there is no evidence of aggregate demand pressures in industrial commodity markets.

6. Quantifying Exogenous Oil Production Shortfalls: A Comparison with the Status Quo

The evidence in Figure 12 suggests that the only way of estimating the exogenous component of oil prices is by regressing changes in the real price of oil on measures of the exogenous oil supply shock. Before addressing this question, I will briefly digress and compare the measure of exogenous oil supply shocks proposed in this paper to an alternative approach that has been used in the literature. This will allow me to highlight both the differences and commonalities that arise when using these measures.

¹⁶ The data were downloaded from the Commodity Research Bureau (CRB) at <http://www.crbtrader.com/crbindex>. The industrial commodities include burlap, copper scrap, cotton, hides, lead scrap, print cloth, rosin, rubber, steel scrap, tallow, tin, wool tops, and zinc. Qualitatively similar results are obtained based on the CRB spot price index for all non-oil commodities. The detrending is required because of the long-term downward trend in non-oil commodity prices.

6.1. Conceptual Problems with the Quantitative Dummy Approach

The state-of-the-art approach to constructing series of exogenous oil supply disturbances is described in Hamilton (2003). Hamilton 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 exogenous events Hamilton considers are the Arab oil embargo of 1973/74, the Iranian revolution of 1978/79, the Iran-Iraq war of 1980-88 and the 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 a 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). This approach is in essence a quantitative version of the dummy variable approach used by Hoover and Perez (1994) to model oil shocks.

Compared with the methodology proposed in this paper, this quantitative dummy approach has several disadvantages. First, implicit in the quantitative dummy approach is the assumption that oil production would have remained at its level in the month prior to the exogenous event, if that event had never occurred. This assumption is implausible. Not only do the data compiled by the Department of Energy suggest that countries' oil production levels in general change considerably over time, even in the absence of exogenous shocks, but one would expect OPEC countries in particular to adjust their levels of production in response to changes in global demand for oil, which in turn are affected by global economic growth, interest rates, inflation rates and exchange rates (see, e.g., Barsky and Kilian 2002, 2004). Thus, there is no reason to believe that Iraqi and Iranian oil production in 1980, for example, would have remained at their August 1980 levels for the next three months (or for that matter the next eight years).

Second, it seems intuitive that a shock to OPEC oil production that persists for years will be much more serious than a short-lived disruption of production. Quantitative dummy measures do not allow such distinctions. A temporary production shortfall lasting one month is treated the same as a persistent loss of OPEC productive capacity, as long as the magnitude of the initial fall is the same. Nor does the quantitative dummy measure make a distinction between a gradual

decline in production and a sudden drop or between an immediate and a delayed drop in production. It also fails to account for the positive effect of a partial or complete reversal of the exogenous supply shortfall over time. Such exogenous dynamics of oil supply (that are distinct from the endogenous supply responses triggered by an exogenous oil shock) must not be ignored in assessing the dynamic effect of exogenous events on the price of oil. The analysis in section 3 documented the quantitative importance of such dynamics. In short, the use of quantitative dummy measures will distort the statistical relationship between estimated production shortfalls and oil prices, which plays a central role in our understanding of the effects of oil supply shocks.

Third, the dating conventions used by Hamilton (2003) to quantify the production shortfall caused by an exogenous event are questionable. Hamilton's methodology involves a comparison of production levels at two different dates. The choice of the initial date is usually uncontroversial because most exogenous events are easily dated.¹⁷ The choice of the subsequent comparison date on the other hand is not. In some cases, Hamilton measures the shortfall one month after the event began; in others seven months later. Hamilton does not discuss how he arrived at these comparison dates, but it can be shown that if one were to search systematically for the date that yields the largest drop in oil production following the exogenous event, for each of the oil shocks studied by Hamilton one would arrive at the comparison date that he reported. Having identified the comparison date, Hamilton assigns the production shortfall computed for this comparison date in its entirety to the month in which the exogenous event started, even if the actual drop in production only occurred many months later.¹⁸ It is unclear what the rationale for that procedure is.

It may seem that these dating problems could be avoided by simply imposing the beginning and end dates of the exogenous event, as provided by historians, and treating the discrepancy between some benchmark and actual production levels as the shortfall. This is not

¹⁷ The main exception is Hamilton's dating of the outbreak of the Iran-Iraq war in August 1980. Although Teheran reported sporadic outbreaks of violence along its border with Iraq in August 1980, large-scale armed conflict only commenced in September (see Terzian 1985, p. 280). This difference in dating is inconsequential for Iraq, whose oil production remained unchanged from July to August, but it does make a difference for Iran. Iranian oil production fell somewhat in August, from a peak of 1.699 million barrels per day in June and July 1980 back to 1.472 million barrels per day in August, the same level it produced in May 1980 under peacetime conditions. This suggests that Hamilton overestimates the 1980 fall in Iranian oil output by 0.227 million barrels per day. A similar problem occurs with Hamilton's (2003) dating of the October 1973 war as a November 1973 shock. Unlike the dating of the outbreak of the Iran-Iraq war, however, that distinction is inconsequential when the data are measured at quarterly frequency.

¹⁸ For example, for the Persian Gulf War episode the decline in production is computed over the period from August 1990 until February 1991, but assigned to the third quarter of 1990.

necessarily the case, unless the effect of the exogenous event on oil production is purely temporary. There is reason to believe that some events, such as the Gulf war, had long-term effects on oil production in Iraq that persisted far beyond the formal conclusion of the war in 1991. In the case of the Iranian revolution of 1978/79 there is no natural end date at all; some effect on oil production is likely to remain even today.

This discussion reinforces the point that simple before-after comparisons will be not be sufficient if we want to quantify the shortfall in oil production caused by exogenous events. Rather we must construct an explicit counterfactual. In other words, the relevant benchmark is the level of oil production in the countries affected that would have prevailed in the absence of the exogenous event. Unless there is a compelling reason to believe that the shortfall is temporary, we have to allow for the counterfactual to differ from the actual production levels indefinitely. The approach used in this paper, unlike the quantitative dummy measure, allows for that possibility.

It is clear that there may be more than one counterfactual that one could construct, and the implied measures of the exogenous shortfall in oil production may differ. On the other hand, *any* analysis of the shortfall will by construction involve a counterfactual. Hamilton's methodology is no exception. Thus, the virtue of the approach proposed in this paper is to make the assumptions that go into the counterfactual explicit and to foster a critical discourse about these assumptions. It is worth stressing that there is nothing in the methodology outlined here that would a priori bias the results in favor of assigning a smaller (or for that matter a larger) role to exogenous oil shocks or to their effects on oil prices and GDP growth. There is no doubt that the baseline assumptions outlined in this paper will be improved and refined over time. Nevertheless, I maintain that these assumptions are more realistic than the implicit assumptions used by Hamilton (2003). While some of the details may be open to disagreement, as we have seen, many of the qualitative difference between Hamilton's oil shock series and the series constructed in this paper are likely to arise for a wide range of assumptions.

The approach to constructing the counterfactual that I follow in this paper is not without its own shortcomings, however, because it involves some degree of judgment. This point is best illustrated by the example of the Iranian revolution. In essence, I argue that Iran would have lowered its oil production after 1979 even in the absence of a revolution because slower economic growth in OECD countries brought down the demand for oil. Implicit in this argument

is the view that the economic downturn of the early 1980s was caused primarily by a shift in monetary policy under Paul Volcker. An alternative view would be that the Iranian revolution, through its effect on the real price of oil, caused that economic downturn. In this sense, the particular counterfactual for Iran used in this paper corresponds to a particular view of history, not the only possible one. In addition, the approach in this paper will tend to result in less biased, but potentially noisier measures of the exogenous oil supply shock. These facts make it important to contrast the implications of my baseline counterfactual to those of alternative counterfactuals including the quantitative dummy approach. The remainder of the paper is devoted to this comparison.

6.2. Empirical Comparison

Having compared the quantitative dummy approach and the approach taken in this paper from a conceptual point of view, we now compare them from an empirical point of view. Figure 13 is based on quarterly averages of the share series in Figure 8. This allows a direct comparison with the corresponding series based on Hamilton (2003) with suitable updates. The procedure used in constructing the Hamilton series follows exactly the description in Hamilton (2003).

From the upper panel of Figure 13, it is apparent that the two series look quite different. The quantitative dummy series captures only a small fraction of the variation in the exogenous production shortfall series proposed in this paper. Part of the reason is that the counterfactual measures the combined effect of all exogenous events up to that point, whereas the quantitative dummy series only measures the effect of the most recent event. In this sense, the Hamilton (2003) measure is more akin to a measure of the change in the production shortfall than of the shortfall itself. In other words, his intention is to capture the “initial shock” in the country where the event takes place (p. 391). He discounts all other exogenous variation in oil output triggered by that event. Leaving aside the inherent ambiguity in defining the “initial” impact, this fact suggests that a more appropriate comparison would be with the quarterly values of the exogenous oil supply shock series defined in section 3.7., as shown in the lower panel of Figure 13. Indeed, the two series in the lower panel line up much more closely, although important qualitative and quantitative differences remain. For example, the quantitative dummy series always stays below zero, even during times when production recovered for exogenous reasons such as in the course of the Iran-Iraq war or toward the end of the Arab oil embargo of 1973/74. Moreover, the quantitative dummy series indicates shocks of potentially very different

magnitudes. The quantitative dummy measure in all cases indicates a larger shock; in some cases more than twice as large as the new measure. The only exogenous event for which both measures indicate a shock of similar magnitude is the outbreak of the Iran-Iraq war in 1980.

7. How Well Do Measures of Exogenous Oil Supply Shocks Explain Real Oil Price Increases During Crisis Periods?

Having shown that oil price shocks as measured by the net oil price increase are not necessarily related to exogenous oil supply shocks, we now turn to the central issue of measuring the extent to which percent changes in the real price of oil can be attributed to current and lagged values of the exogenous oil supply shock. For this purpose I project the quarterly percent change in the real price of oil on an intercept, the current value and four lags of the exogenous oil supply shock. For a similar approach applied to the change in nominal oil prices see Hamilton (2003).

7.1. Actual and Fitted Real Oil Price Changes

Figure 14 plots the fitted values for the two alternative oil supply shock measures against the actual percent change in real oil prices. First consider the quantitative dummy measure of exogenous oil supply shocks. This measure predicts only about 20 % of the change in oil prices that occurred in 1973/74, suggesting that the remainder is not related to exogenous oil production shocks. The new measure proposed in this paper assigns an even less important role to the production shortfalls associated with the Arab oil embargo. Either way we conclude that exogenous changes in oil production do not seem to have played a major role in the observed price increases during that episode. This finding is consistent with the earlier evidence that oil prices during this episode behaved similarly to other industrial commodity prices that were not subject to exogenous supply shocks.

For 1979/80, the quantitative dummy series not only suggests that the oil price increase should have occurred well before it actually did, but it also predicts a sharp spike, whereas the actual oil price increase in 1979/80 was more gradual. It is hard to imagine how this effect could have been offset by other factors, causing the actual data to show no spike at all at that point. The new measure, in contrast, indicates a much smaller effect of exogenous oil supply shocks, the timing of which is consistent with the timing of the actual oil price increase. For 1980/81 once again the quantitative dummy measure predicts a peak earlier than the actual data. Both exogenous oil supply shock measures over-predict the magnitude of the 1980/81 oil price

increase, but the quantitative dummy measure much more so than the measure proposed in this paper. The fact that this oil supply shock had less of an effect than predicted may be attributed to the global economic slowdown at the time.

For 1990/91, the quantitative dummy measure provides a better fit with the sharp increase in actual oil prices, but even that measure predicts only about one third of the actual oil price increase in 1990/91. In contrast, the new measure suggests that production shortfalls played a minor role in explaining the oil price increases of 1990/91. Moreover, according to the new measure, the production shortfalls affected the price only with a delay and caused sustained oil price increases rather than a spike. Either way, a substantial part of the 1990/91 oil price shock was not driven by the exogenous production shortfall. Finally, for 2002/03, neither measure of the exogenous oil supply shock does a good job at predicting the observed pattern of oil price changes. The quantitative dummy measure predicts higher oil price increases than occurred in 2002/03.

Apart from illustrating the timing problems encountered with the quantitative dummy approach, the regression evidence in Figure 14 underscores that neither measure of exogenous oil production disturbances does a good job at explaining the magnitudes of the major oil price increases since 1973. For example, even the quantitative dummy measure explains only one fifth of the 1973/74 spike and less than a third of the 1990/91 spike in oil prices. What then could possibly explain the additional sharp increase in the price of oil during these and other episodes? Shifts in the demand for oil are likely to have played a crucial role. Increases in demand, of course, tend to be smooth and would not in general be expected to generate a sudden and potentially large increase in the real price of oil - except in the presence of capacity constraints. There is strong evidence of demand for oil hitting capacity constraints in 1973, as evidenced, for example, by the rapid and sustained increase in tanker freight rates long before the exogenous event in question (also see Seymour 1980, p. 100, for related discussion). There is similar evidence for 2004/05, when strong demand due to the business cycle coincided with additional demand for oil from Asian countries that had not traditionally competed for this resource, and to a somewhat weaker extent for 1979. On the other hand, there is no evidence that capacity constraints played a role in 1980/81 and 1990/91.

An alternative and complementary explanation is that concerns about the future availability of oil supplies, especially in the face of high demand for oil, could trigger a sudden

and potentially large increase in precautionary demand for oil. There is evidence of such panic buying both in late 1973 and in 1979, for example. The latter phenomenon reflects fears of yet unrealized oil supply disruptions rather than concerns about already realized exogenous production shortfalls. While typically linked to exogenous events, shifts in precautionary demand may be highly idiosyncratic and differ from one episode to the next. They will be captured by the regression evidence underlying Figure 14 only to the extent that they are linearly related to the actual exogenous oil supply disruptions that took place since 1973.

7.2. Beyond Exogenous Production Shortfalls: The Role of Uncertainty about Future Oil Supplies

Figure 14 suggests only a minor role for exogenous oil supply shocks in explaining fluctuations in the real price of oil. There is reason to believe that the underlying regressions are missing something important, however. The 1990/91 episode in particular is instructive because for that period there is no reason to believe that global demand for oil increased enough to generate a sharp increase in oil prices; yet exogenous oil production shortfalls do not explain this episode well.¹⁹ This is true for the new measure of exogenous oil supply shocks as well as the quantitative dummy measure. Even the latter measure explains less than a third of the observed oil price increase in 1990. The only plausible alternative explanation are increased fears of future supply disruptions, not simply in the countries directly affected by the Persian Gulf war, but in other countries such as Saudi Arabia as well. The discrepancy between the smaller and more sustained increase in oil prices predicted by the new measure, in particular, and the observed sharp spike suggests the existence of a large uncertainty premium.

If a positive spike in oil prices is driven by increased uncertainty about future oil supplies, it should be followed by a negative spike when the uncertainty is resolved. Indeed there is evidence in Figure 14 of such a reversal after the end of the 1991 Persian Gulf War. This evidence suggests that a shift in uncertainty rather than the physical supply disruption caused by these wars may explain the bulk of the observed sharp and immediate spikes in oil prices that is

¹⁹ The poor predictive performance of the exogenous oil supply shock measures during 1990/91 cannot be explained by estimation error. If we were to select coefficients that ensure that this price spike can be explained by the exogenous oil supply shock alone, then by construction we would end up with an even worse fit during most other exogenous supply shock episodes. Thus, there must be other factors at play that are not related to oil production shortfalls.

left unexplained by exogenous oil production disturbances.²⁰ The uncertainty-based explanation seems plausible in that the 1990/91 episode was fundamentally different from any other episode in our data in that the invasion of Kuwait for the first (and only) time in history created an imminent military threat to the oil supplies of Saudi Arabia. The U.S. responded to this threat immediately by moving troops to Saudi Arabia, but only in November of 1990 enough troops had been assembled to neutralize the threat to Saudi oil fields (see United States Central Command 1991). This unprecedented shift in uncertainty may be captured by adding a dummy variable to the regressions underlying Figure 14 that takes on a value of 1 for the third quarter of 1990, when the imminent threat to Saudi Arabia first arose, and of -1 for the fourth quarter, when it was eliminated.

Including the current value and one lag of this dummy variable in the regression on the quantitative dummy measure as well as in the regression on the new measure, greatly improves the fit of *both* models in question for the 1990/91 episode, as shown in Figure 15. There now is little to choose between the two measures of exogenous oil supply shocks for that period, suggesting that the earlier regressions confounded the effect of exogenous oil production shortfalls in Iraq and Kuwait in 1990/91 with the effect of the imminent threat to Saudi oil supplies. Similarly, for the 1980 oil supply shock, the two measures now give virtually identical results. The timing problem of the quantitative dummy measure in this episode has vanished. Both exogenous supply shock measures over-predict the effect of the oil supply shock on the price of oil in 1981, consistent with the view that a global economic slowdown largely offset the adverse effects of the 1980 oil supply shock. For the 1979/80 episode, the new exogenous oil supply shock measure now predicts a much larger increase in the price of oil, the timing of which is consistent with the actual data. The new supply shock series accounts for three quarters of the increase in the price of oil in the third quarter of 1979, but not for any of the subsequent increases. The quantitative dummy series implies an even larger increase in the price of oil, but the timing of the predicted increase is premature. Similarly, for the 1973/74 episode both exogenous supply shock series now predict a much larger oil price increase. Nevertheless, they

²⁰ For the earlier episodes the evidence on the role of uncertainty is less clear-cut. Exogenous events such as major wars or revolutions in the Middle East should have a fairly immediate effect on the price of oil, if they cause a major increase in uncertainty about future supplies of oil. This was not the case after the Iranian revolution in October 1978 or the outbreak of the Iran-Iraq war in September 1980, for example. It took more than half a year before prices started taking off after the Iranian revolution and the process continued into 1980. Moreover, to the extent that these events increased uncertainty about future oil supplies, that uncertainty dissolved gradually, as these events dragged on, making it difficult to separate the uncertainty effect from other factors.

only account for a third or a sixth of the actual oil price increase in 1973/74, respectively, reaffirming the earlier result that much of the observed oil price increase for this episode cannot be attributed to exogenous oil supply shocks. For the 2002/03 there is no improvement in fit. Exogenous oil supply shocks do little to explain the observed oil price changes.

To sum up, the analysis so far suggests three conclusions: (1) In general, exogenous oil production shortfalls are of limited importance in explaining oil price changes during crisis periods. This result is robust to the choice of the exogenous oil supply shock measure. (2) The predictions of the new exogenous supply shock measure are easier to reconcile with the actual oil price data than the predictions of the quantitative dummy measure in that their timing in 1979 is more consistent with the actual data. (3) Of the episodes studied, only the 1980/81 oil price increases can be attributed to exogenous oil supply disruptions. Exogenous oil supply disruptions were an important contributing factor for the oil price increases in the third quarter of 1979, but do not help explain the subsequent oil price increases during that episode. Exogenous oil supply disruptions explain only a small fraction of the oil price increases during the 1973/74, 1990/91 and 2002/03 episodes.²¹

8. Implications for the IV Approach

The choice of the counterfactual also is important for constructing instrumental variable estimates of the response of real GDP growth to lagged oil price changes. This section will illustrate some of the important empirical differences that may arise due to the choice of counterfactual. I also will highlight some of the deficiencies of the IV approach to estimating the dynamic effects of oil supply shocks compared with the approach proposed in section 4. The discussion builds on Hamilton (2003) who proposed the use of lags of the exogenous oil supply shock as instruments in regressions that relate real GDP growth to past oil price changes and past real GDP growth. Table 2 shows IV estimates of the type proposed by Hamilton based on alternative assumptions about the counterfactual and for alternative sample periods.²² As in

²¹ We also repeated this analysis using monthly data with twelve lags of the exogenous oil supply shock series of Figure 9 and two lags of the dummy variable, dated as August 1990 and December 1990. Using monthly data allows us to capture more precisely the dynamics of the price adjustment. The results were remarkably similar to the quarterly results in Figure 15. The main difference is that the counterfactual provides a somewhat better fit for the 1980/81 episode.

²² We follow Hamilton (2003) in focusing on the F -test of the contribution of lagged oil price changes to real GDP growth. Similar test results hold when the F -test is replaced by the Wald test based on the conventional asymptotic approximation.

Hamilton (2003), the estimation method is two-stage least-squares. In all cases the instruments include a constant, four lags of real GDP growth and eight lags of the exogenous oil supply shock series. The precise definitions of these variables may differ, as discussed below.

8.1. Baseline results

Given the lack of monthly production data prior to 1973, one cannot compare IV estimates based on the explicit counterfactual proposed in this paper to the estimates reported by Hamilton (2003) for the period 1947.II-2001.III. Nevertheless, I begin by replicating and extending Hamilton's analysis. In interpreting the results in Table 2, it is useful to keep in mind that the quantitative dummy measure used as instruments in the regressions underlying columns (1)-(4) includes not only the exogenous events after 1973, as discussed in sections 3 and 4, but also all additional exogenous events prior to 1973 considered by Hamilton (2003).

Column (1) replicates Hamilton's findings on his original data set. A key finding is that the F -test of the null that lagged oil price changes do not affect GDP growth is clearly rejected with a p -value of 2 %. If we replace Hamilton's nominal oil price series by the real price of oil (obtained by deflating nominal prices by the U.S. CPI), as shown in column (2), the p -value rises to 5 %, but otherwise the result remains intact. Columns (3) and (4) repeat this analysis on a data set that has been extended to 2004.III, using the same conventions as Hamilton (2003). The qualitative results are not affected.

In contrast, columns (5) and (6) focus on the period since 1973.I. This amounts to excluding the Suez crisis of 1956 from the set of exogenous events. Unlike in Hamilton (2003), this omission raises the p -value to 15 % when nominal oil prices are used and to 22 %, when real prices are used. The difference is that Hamilton (2003) in his sensitivity analysis includes all observations prior to 1973 in the sample, even after dropping the Suez crisis. Columns (5) and (6) show that, even taking Hamilton's exogenous oil supply shock measure at face value, his results are no longer significant if we focus on the modern OPEC period since the early 1970s.

This sensitivity of the IV results is a reason for concern, since there are strong institutional reasons to expect the transmission mechanism from wars to oil prices to have changed since the 1940, 1950s and 1960s, as noted earlier. Additional reasons to be skeptical of the pre-1973 analysis in Hamilton (2003) have to do with the set of events considered. For example, it is unclear why Hamilton, having included the 1973/74 Arab-Israeli war and Arab oil embargo excludes the Arab-Israeli war in June of 1967 (which resulted in a closure of the Suez

Canal and which was followed by an Arab oil embargo from June 1967-August 1967). Terzian (1985, p. 108), using a methodology similar to Hamilton (2003), reports a drop of Arab oil production of 4.336 million barrels/day from May 1967 to June 1967. Hamilton (2003) also excludes several other events that Hamilton (1983) classifies as exogenous oil supply shocks such as the nationalization of Iranian oil in 1951 and its subsequent boycott by the major oil companies, causing Iranian production to drop in 1951-54 (see Terzian 1985, p. 14, Table 5.2).²³ These facts cast doubt on the reliability of the 1973-2004 regression evidence on the effects of exogenous oil supply shocks.

Hamilton's (2003) study uses as its oil price measure the producer price index for crude oil that is produced in the U.S. It is well known that domestic crude oil prices at times behaved differently from international crude oil prices. The latter prices are not available at monthly frequency for the pre-OPEC period, but they are available for the period when oil prices were endogenous. For that reason, columns (7) and (8) focus on the average U.S. import price of crude oil, as reported by the U.S. Department of Energy. The results are similar to columns (5) and (6), suggesting that the results are not driven by the choice of oil price index. For the remaining analysis I will rely on average U.S. import prices for crude oil.

Columns (8) and (9) differ from the previous two columns only in that I replace the quantitative dummy measure of the exogenous oil supply shock by the newly proposed exogenous oil supply shock measure, as shown in Figure 9. The difference is dramatic. The *p*-value jumps to 83 % for nominal prices and 86 % for real prices. Since the real price is clearly more relevant for the question at hand, we will focus on real oil prices only from now on. Column (9) provides strong evidence that Hamilton's (2003) results are highly dependent not

²³ Hamilton (1983) also cites the "accidental" rupture of the trans-Arabian pipeline (the repair of which was delayed by Syrian authorities for nine months) and Libyan oil production cutbacks in 1970 as exogenous oil supply shocks. The latter events cannot be regarded as exogenous. The events in Libya, for example, had less to do with Middle East politics, but with the fall in the real price of oil combined with fixed posted prices since 1960. Since the independent oil companies operating in Libya (unlike the oil companies operating in other countries) paid taxes only on the basis of real prices, the Libyan government received only 30 cents per barrel of oil exported, compared to 90 cents in other OPEC countries, and was naturally displeased (see Terzian 1985, p. 114). To the extent that the fall in real prices reflected rising U.S. inflation, it seems hard to contend that the Libyan government's reaction was exogenous to U.S. macroeconomic aggregates. Not surprisingly the Libyan move was tolerated by the major oil companies who were not thrilled by the "unfair" competition from independent companies operating in Libya. Similarly, the move by other countries to raise the price of oil (or at any rate their share in the profits) in 1970-71, which led to the "accidental" rupture of the pipeline while negotiations were being conducted, occurred explicitly in response to the strengthening demand in international oil markets.

just on the sample period and the use of nominal oil price changes, but even more so on the specific definition of the exogenous oil supply shock series.

8.2. Sensitivity Analysis

As stressed earlier, the specification of the counterfactual will by construction affect the definition of the exogenous production shortfall. I focused on a baseline counterfactual that is plausible, but some of the underlying assumptions may be debated. It turns out that reasonable departures from the baseline counterfactual do not change the results much. Columns (11)-(13) present results for some alternative counterfactuals. Column (11) differs from the baseline counterfactual in that I ignore the contribution of the Saudi production response to the Iranian revolution, the Gulf war and the Iraq war in constructing the exogenous movements in oil output. Treating the Saudi response as endogenous lowers the p -value of the F -test somewhat, but only to 77 %. Column (12) returns to the point made earlier that it is not clear that the Arab oil embargo of 1973/74 can be treated as exogenous. The results in column (12) differ from the baseline in that we assume that the production shortfall caused by the oil embargo of 1973/74 was endogenous. In this case, the p -value jumps to 92 %. Finally, column (13) shows results that differ from the baseline in that Saudi Arabia has been included in the benchmark for except for the 1973/74 episode. This counterfactual implies exogenous oil supply shocks closer to Hamilton's series (although important differences remain), and results in a p -value of 55 %. Regardless of these variations in the counterfactual, none of the p -values comes even close to establishing significance at reasonable levels.

8.3. Impulse Response Analysis

The F -tests in Table 2 provide one way of judging whether oil supply shocks cause lower real GDP growth. A more informative approach is the construction of the response over time of real GDP growth to an exogenous oil price increase based on the IV regression. Confidence intervals for the pointwise responses may be constructed by simulating draws from the asymptotic distribution of the instrumental variable estimator in column (10) of Table 2. Figure 16 shows the dynamic response of real GDP growth to a 10% increase in oil prices, augmented by one- and two-standard deviation error bands. Unlike the estimate shown in Figure 10, the response of real GDP growth is erratic and very imprecisely estimated. The intervals constructed from column (10) cover the zero line at all horizons, contradicting the view that the data tell us that oil price

increases cause lower growth. Qualitatively similar results were obtained based on the alternative counterfactuals underlying columns (11), (12) and (13).²⁴ Hamilton's original IV analysis suggested that an oil price increase implies a persistent and statistically significant reduction of real output growth for about five quarters (see Hamilton 2003, Figure 14). This conclusion is not supported by the results in Figure 16.

8.4. Weak Instrument Problems

Regardless of the results, there is reason to be skeptical of the IV results in section 8.3. As is well known, weak instruments produce biased IV estimators and hypothesis tests with large size distortions (see Stock, Wright and Yogo (2002) for a review). One way of detecting a weak instrument problem is to compare the value of the first-stage F -statistic to the critical values compiled by Stock and Yogo (2003). As argued by Stock et al. (2002), the first-stage F -statistic must be large, typically exceeding 10, for two-stage least-squares inference to be reliable. In the presence of many variables to be instrumented, as in the IV regressions in Table 2, it is not enough to consider the individual F -statistics for each first-stage regression. Cragg and Donald (1993) have proposed an alternative statistic designed to handle multiple first-stage regressions. Critical values for a transformation of this statistic, denoted by g_{\min} , can be found in Stock and Yogo (2003), who propose a test of the null of weak instruments against strong instruments based on g_{\min} . Table 2 shows that all the estimated g_{\min} -statistics are far below the least conservative critical value of about 4. Even for columns (1) through (4) the weak instrument null cannot be rejected.

The evidence in Table 2 suggests that IV estimates of the effect of exogenous variation in the price of oil of the type proposed by Hamilton (2003) tend to be inherently unreliable due to the presence of a weak instrument problem. This problem renders the IV estimates imprecise and standard inference on the dynamic effects of exogenous variations in the price of oil invalid. It is not straightforward to conduct inference for the impulse response coefficients in the presence of weak instruments. In contrast, the alternative approach proposed and implemented in section 4 does not require IV estimation. Estimation and inference is standard, and the empirical results are sharp and intuitive.

²⁴ For comparison, using the IV estimates of column (8) based on Hamilton's exogenous supply shock series, one would rule out at the 68 % confidence level a zero response of real GDP growth at horizons 2 and 4, but not at the other horizons. None of the responses are significantly different from zero based on two-standard error bands. Again, this evidence underscores that the choice of the counterfactual matters.

Given the conceptual difficulties of comparing the indirect estimates of the dynamic effects of exogenous oil supply shocks in Figure 16 to the direct estimates in Figure 10, we conclude this section with an exercise that highlights the differences in the two measures of exogenous oil supply shocks, while controlling for the method of estimation. Figure 17 shows the impulse responses that are obtained based on least-squares projections on lags of the quantitative dummy measure of exogenous oil supply shocks. The results are directly comparable to Figure 10. Starting with the response of real GDP growth in the upper panel, the key difference is that the quantitative dummy measure suggests a sharp reduction in real growth two quarters after the shock that is significant at the 95% level. Moreover, the point estimates of the responses remain negative from the first until the seventh quarter, quarters 4, 5, and 6 being significantly different from zero at the 95% significance level. It is also worth noting that the maximum response in Figure 17 is only about -1.5 % compared to -2.2 % in Figure 10. The corresponding least-squares projections on lags of the quantitative dummy measure of exogenous oil supply shocks produce inflation responses that are not significant at the 68% level except for a negative response after 7 quarters (see lower panel of Figure 17). While the point estimate is positive for the first three quarters, the maximum increase is only about 0.4 % (compared to a maximum of almost 1.0 % in Figure 10). There is no evidence of a spike in CPI inflation at any horizon or of a sustained increase. These estimates appear considerably less intuitive than the results in the lower panel of Figure 10.

9. Conclusion

Since the oil crises of the 1970s there has been strong interest in the question of how oil production shortfalls caused by wars and other exogenous political events in OPEC countries affect oil prices, U.S. real GDP growth and U.S. CPI inflation. This paper focused on the modern OPEC period of the oil market since 1973. The analysis provided answers to the following questions:

- *How large are the exogenous fluctuations in the supply of oil?*

The conventional measure of the exogenous production shortfall discussed by Hamilton (2003) is implausible on a priori grounds. I proposed an alternative measure that involves the construction of an explicit counterfactual for each country involved in an exogenous event. Based on this counterfactual I provided time series of the exogenous production shortfall by OPEC country and

for all OPEC countries combined. The change over time in this series expressed as a share of world oil production may be viewed as a measure of the exogenous oil supply shocks. The resulting shocks range from about -7 % to +3 % of world crude oil production at quarterly frequency. Exogenous changes in crude oil production driven by political events in OPEC countries account for about 6 % of the variability in world crude oil production changes.

- *To what extent can oil price changes be explained by variations in exogenous oil supply?*

Exogenous oil production shortfalls are of limited importance in explaining oil price changes both overall and during crisis periods. This result is robust to the choice of the exogenous oil supply shock measure. Of the episodes studied, only the 1980/81 oil price increases can be attributed to exogenous oil supply disruptions. Exogenous oil supply disruptions were an important contributing factor for the oil price increases in the third quarter of 1979, but do not help to explain the subsequent oil price increases during 1979/80. Exogenous oil supply disruptions explain only a small fraction of the oil price increases during the 1973/74, 1990/91 and 2002/03 episodes. These results highlight the importance of alternative explanations such as shifts in the demand for oil driven by global macroeconomic conditions or the imminent military threat to the Saudi oil fields in 1990. This and other evidence suggests that measures of the net oil price increase relative to the recent past do not represent the exogenous component of the price of oil.

- *How reliable are IV estimates of the effect of exogenous oil supply shocks on the U.S. economy?*

IV estimates of the effect of exogenous variations in the price of oil as proposed by Hamilton (2003) tend to be inherently unreliable due to the presence of a weak instrument problem, which also renders standard inference on the dynamic effects of exogenous variations in the price of oil invalid. If all one is interested in is quantifying the dynamic effects of exogenous variations in the supply of oil on real GDP growth, a much simpler approach suffices that relies on OLS estimates of a regression of real GDP growth on a constant, lagged real GDP growth and lagged exogenous oil supply shocks. Inference for this regression is standard and OLS estimates are consistent. The latter approach provides the basis for the answer to the following two questions.

- *What are the effects of exogenous oil supply shocks on U.S. real GDP growth and CPI inflation?*

Under the baseline counterfactual, point estimates of the dynamic effect of exogenous oil supply

shocks on real U.S. GDP growth suggest a sharply negative growth rate five quarters after the oil supply shock, before the response reverts back to zero. This negative growth according to the 68% confidence band to some extent persists until the seventh quarter; according to the 95% confidence band it lasts only for one quarter. Point estimates of the response of U.S. CPI inflation suggest a sharp spike in inflation three quarters after the shock. Otherwise the response is flat. The spike is significant at the 68% confidence level, but not the 95% confidence level. At all other horizons the response of CPI inflation is insignificant at the 68% level. There is no evidence that exogenous oil supply shocks cause sustained inflation. Similar results for the inflation and output growth responses are obtained for a number of alternative specifications of the counterfactual. Historical counterfactuals suggest that exogenous oil supply shocks made little difference on average for the evolution of U.S. real economic growth and inflation since the 1970s, although they did matter for particular episodes, especially for U.S. real growth.

- *Why are the effects of exogenous oil supply shocks not larger?*

A partial answer is that the approach in this paper has been causal, albeit reduced-form, and hence incorporates the average effects of endogenous policy responses. Another part of the answer is that this paper has followed the existing literature on the effects of exogenous oil supply shocks in focusing on the role of physical supply interruptions. The analysis in this paper suggests that exogenous oil production shortfalls narrowly defined can account for only a comparatively small part of oil price movements. This does not necessarily mean that “oil shocks” are of lesser importance, however. Even if physical production does not move, expectations of future oil supply interruptions alone may have powerful effects, as is evident during the 1990/1991 Persian Gulf War. Such expectations are not directly observable, but a natural extension of the current line of work would be to identify observables other than oil production data that are likely to drive fears of future oil supply disruptions. This task will be taken up in a separate paper.

References

- 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.
- Cragg, J.G., and S.G. Donald (1993), "Testing Identifiability and Specification in Instrumental Variable Models," *Econometric Theory*, 9, 222-240.
- Davis, S.J., and J. Haltiwanger (2001), "Sectoral Job Creation and Destruction Responses to Oil Price Changes," *Journal of Monetary Economics*, 48(3), 465-512.
- Hamilton, J.D. (1983), "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91(2), 228-248.
- 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. (2004), "Historical Effects of Oil Shocks," mimeo, ftp://weber.ucsd.edu/pub/jhamilto/interpreting_oil_regressions.pdf
- 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.
- 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.
- 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., S. Ni, and R.A. Ratti (1995), "Oil Shocks and the Macroeconomy: the Role of Price Variability," *Energy Journal*, 16, 39-56.
- Oil and Gas Journal (1973), volume 71, issues 42-46.
- Penrose, E. (1976), "The Development of Crisis," in: R. Vernon (ed.), *The Oil Crisis*, Norton; New York, 39-57.
- Seymour, I. (1980), *OPEC. Instrument of Change*. MacMillan Press: London.
- 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.
- Skeet, I. (1988), *OPEC: Twenty-five Years of Prices and Politics*, Cambridge University Press: Cambridge.
- 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.

- Stock, J.H., and M. Yogo (2003), "Testing for Weak Instruments in Linear IV Regression," manuscript, Department of Economics, Harvard University.
- Terzian, P. (1985), *OPEC: The Inside Story*, Zed Books: London.
- United States Central Command (1991), *Operation Desert Shield/Desert Storm*, 11 July 1991. Executive Summary. Top Secret, 31 pages (declassified).
(<http://www.gwu.edu/~nsarchiv/NSAEBB/NSAEBB39/document6.pdf>)
- Wilkins, M. (1976), "The Oil Companies in Perspective," in: R. Vernon (ed.), *The Oil Crisis*, Norton: New York, 159-178.

Table 1: Production Shares of Selected OPEC countries

OPEC Country	Percent share in oil production of Total OPEC-Iran		Percent share in oil production of Total OPEC-Iran-Iraq-Saudi Arabia	
	1974.1-1978.9	1978.10-1980.8	1974.1-1978.9	1978.10-1980.8
Algeria	4.5	4.5	8.1	8.8
Indonesia	6.3	5.9	11.3	11.5
Iraq	9.6	12.5	-	-
Kuwait	9.1	8.2	16.4	15.9
Libya	7.5	7.5	13.3	14.5
Nigeria	8.5	8.3	15.2	16.2
Quatar	2.0	1.9	3.6	3.6
Saudi Arabia	34.7	36.1	-	-
UAE	7.7	6.7	13.8	13.0
Venezuela	10.2	8.5	18.2	16.5

Source: Department of Energy.

Table 2: Instrumental Variable Regressions for Real GDP Growth

Regressand: Δgdp_t	Exogenous oil supply shocks measured as quantitative dummies								Exogenous oil supply shocks as proposed in this paper				
	1947.II-2001.III	1947.II-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	1973.I-2004.III	
Regressors:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
c	0.95	0.78	0.87	0.73	0.68	0.54	0.74	0.58	0.57	0.57	0.56	0.40	0.49
(std. dev)	(0.17)	(0.16)	(0.16)	(0.15)	(0.19)	(0.18)	(0.22)	(0.22)	(0.36)	(0.39)	(0.30)	(0.52)	(0.22)
Δgdp_{t-1}	0.20	0.22	0.20	0.21	0.20	0.20	0.16	0.17	0.17	0.21	0.23	0.15	0.19
(std. dev)	(0.09)	(0.10)	(0.10)	(0.10)	(0.14)	(0.14)	(0.15)	(0.16)	(0.27)	(0.27)	(0.20)	(0.34)	(0.16)
Δgdp_{t-2}	0.11	0.12	0.13	0.14	0.12	0.13	0.06	0.07	0.11	0.08	0.03	0.28	0.08
(std. dev)	(0.10)	(0.10)	(0.10)	(0.10)	(0.15)	(0.15)	(0.17)	(0.18)	(0.38)	(0.38)	(0.24)	(0.54)	(0.18)
Δgdp_{t-3}	-0.04	-0.04	-0.05	-0.05	0.02	0.03	0.03	0.05	0.11	0.09	0.06	0.18	0.06
(std. dev)	(0.10)	(0.10)	(0.10)	(0.10)	(0.14)	(0.15)	(0.15)	(0.16)	(0.33)	(0.35)	(0.21)	(0.41)	(0.16)
Δgdp_{t-4}	-0.17	-0.18	-0.15	-0.15	-0.07	-0.07	-0.06	-0.06	-0.09	-0.08	-0.01	-0.14	-0.00
(std. dev)	(0.09)	(0.10)	(0.09)	(0.10)	(0.14)	(0.15)	(0.14)	(0.16)	(0.27)	(0.29)	(0.19)	(0.37)	(0.15)
Δnp_{t-1}^{oil}	-0.028	-	-0.022	-	-0.018	-	-0.020	-	0.025	-	-	-	-
(std. dev)	(0.029)		(0.028)		(0.023)		(0.025)		(0.110)				
Δnp_{t-2}^{oil}	-0.052	-	-0.050	-	-0.040	-	-0.047	-	-0.093	-	-	-	-
(std. dev)	(0.027)		(0.026)		(0.021)		(0.024)		(0.134)				
Δnp_{t-3}^{oil}	-0.013	-	-0.002	-	-0.004	-	-0.008	-	0.055	-	-	-	-
(std. dev)	(0.026)		(0.026)		(0.022)		(0.027)		(0.104)				
Δnp_{t-4}^{oil}	-0.064	-	-0.066	-	-0.032	-	-0.036	-	-0.011	-	-	-	-
(std. dev)	(0.031)		(0.031)		(0.025)		(0.026)		(0.065)				
Δrp_{t-1}^{oil}	-	-0.032	-	-0.025	-	-0.021	-	-0.026	-	0.045	0.057	-0.041	-0.012
(std. dev)		(0.033)		(0.031)		(0.025)		(0.030)		(0.120)	(0.096)	(0.196)	(0.067)
Δrp_{t-2}^{oil}	-	-0.059	-	-0.056	-	-0.044	-	-0.055	-	-0.083	-0.063	-0.115	-0.059
(std. dev)		(0.031)		(0.030)		(0.025)		(0.031)		(0.151)	(0.055)	(0.153)	(0.035)
Δrp_{t-3}^{oil}	-	-0.019	-	-0.006	-	-0.008	-	-0.016	-	0.070	0.036	0.062	-0.023
(std. dev)		(0.030)		(0.030)		(0.027)		(0.035)		(0.122)	(0.079)	(0.133)	(0.054)
Δrp_{t-4}^{oil}	-	-0.070	-	-0.072	-	-0.035	-	-0.042	-	-0.015	-0.011	0.027	-0.008
(std. dev)		(0.034)		(0.034)		(0.028)		(0.031)		(0.068)	(0.052)	(0.092)	(0.036)
F test	2.948	2.426	2.827	2.388	1.738	1.456	1.868	1.443	0.370	0.328	0.448	0.225	0.773
(p-value)	(0.021)	(0.049)	(0.026)	(0.049)	(0.147)	(0.220)	(0.121)	(0.225)	(0.829)	(0.859)	(0.774)	(0.924)	(0.545)
g_{\min}	1.568	1.524	1.172	1.121	0.660	0.559	0.563	0.423	0.063	0.055	0.104	0.055	0.137

Notes: The instruments include a constant, four lags of real GDP growth and eight lags of the oil supply shock series. Columns (1)-(6) are based on the PPI for domestic crude oil as reported by the BLS and Hamilton (2003); columns (7)-(13) are based on the average price of imported crude oil as reported by the U.S. Department of Energy, extended backwards from 1974.I to 1973.I as in Barsky and Kilian (2004). The F-test refers to the null that oil prices changes have no effect on real GDP growth. Column (11) excludes the Saudi production response; column (12) drops the embargo; column (13) includes Saudi Arabia in the benchmark starting in 1974. The last line shows the g_{\min} -statistic of Stock and Yogo (2003).

Figure 1: The October 1973 War and 1973/74 Embargo





Figure 2: Iran

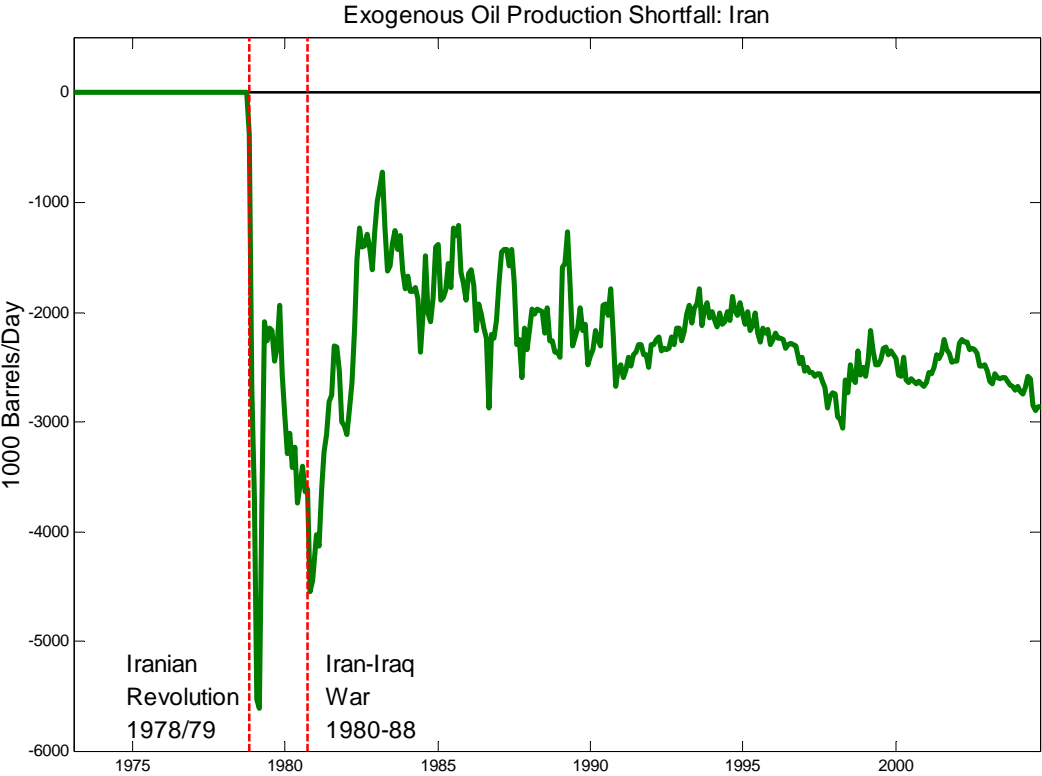



Figure 3: Iraq


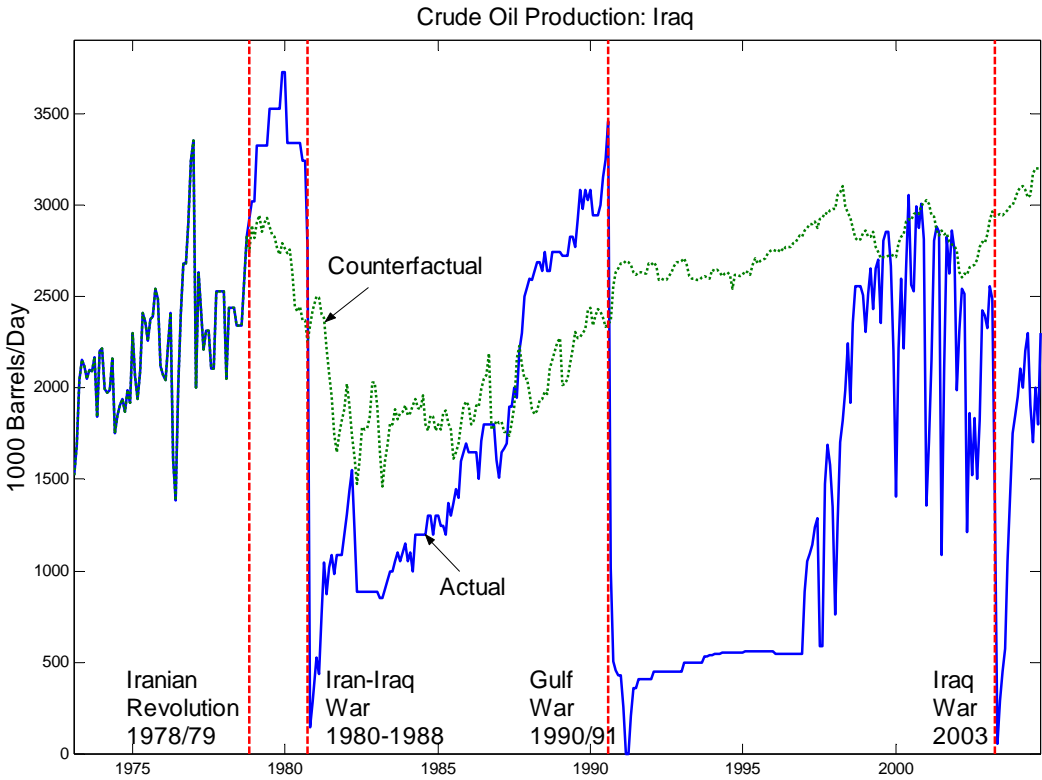


Figure 4: Kuwait

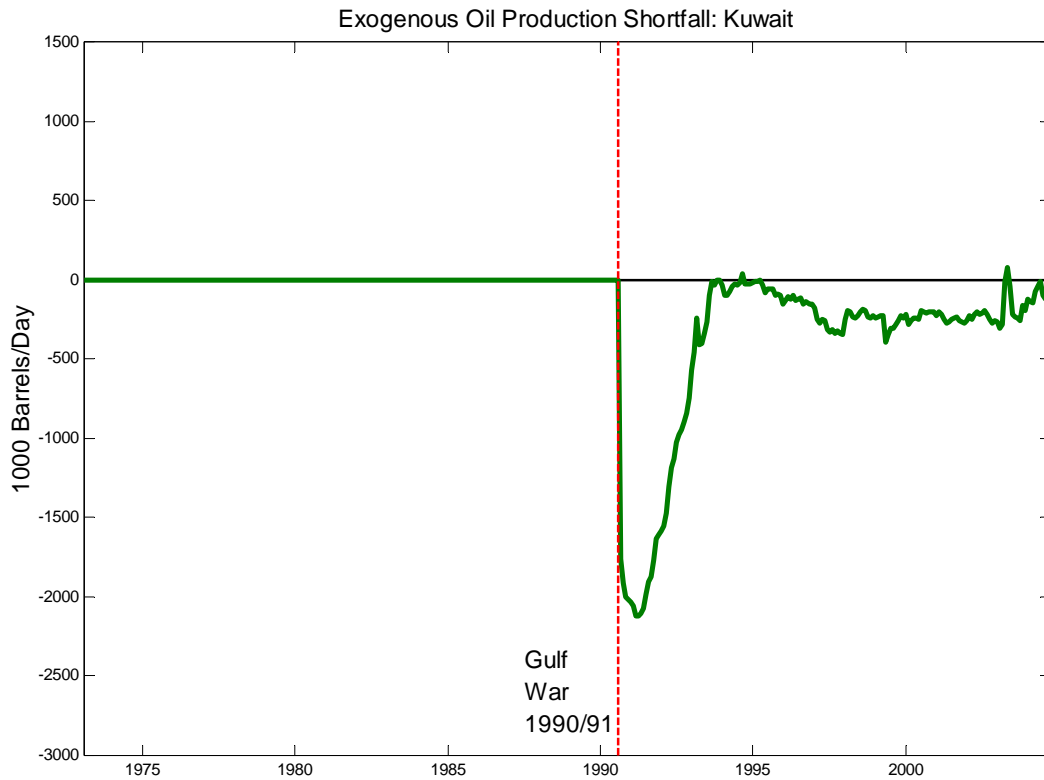
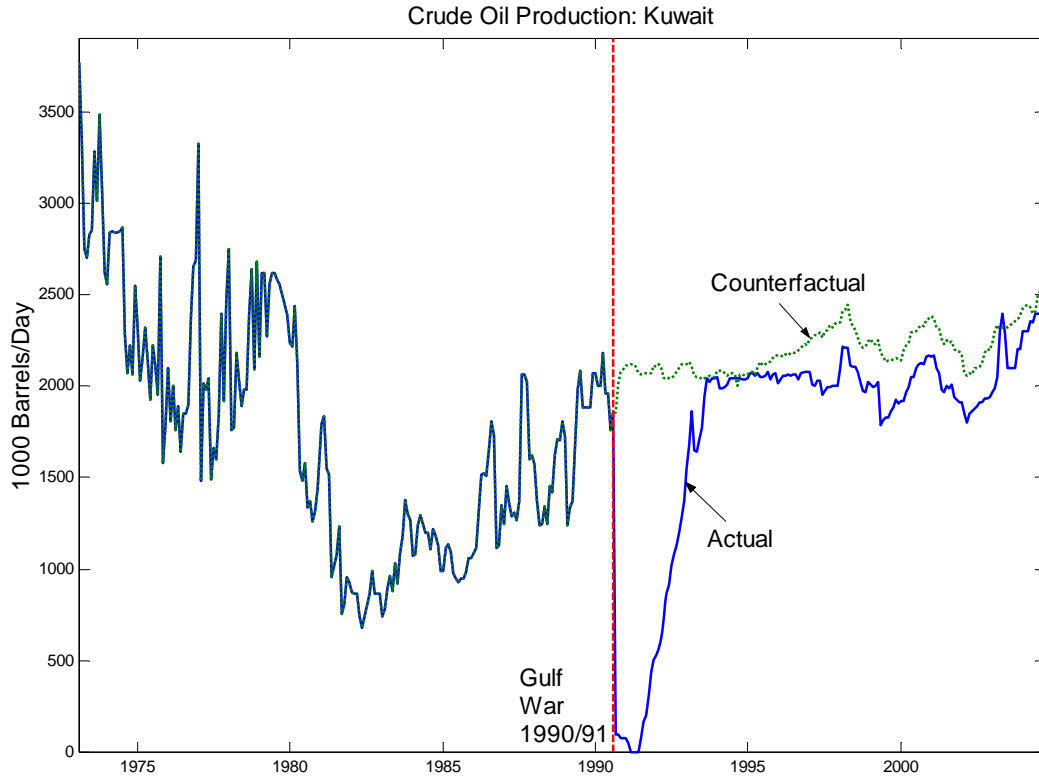


Figure 5: World Oil Supply during Iranian Revolution




Figure 6: Saudi Arabia

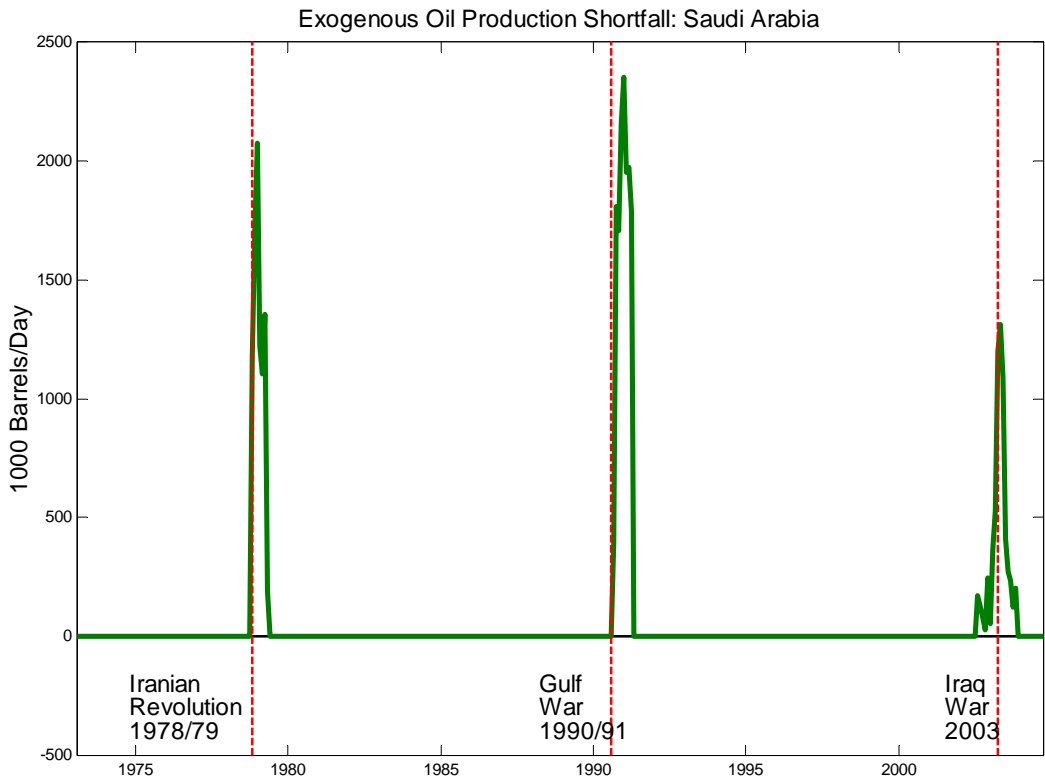
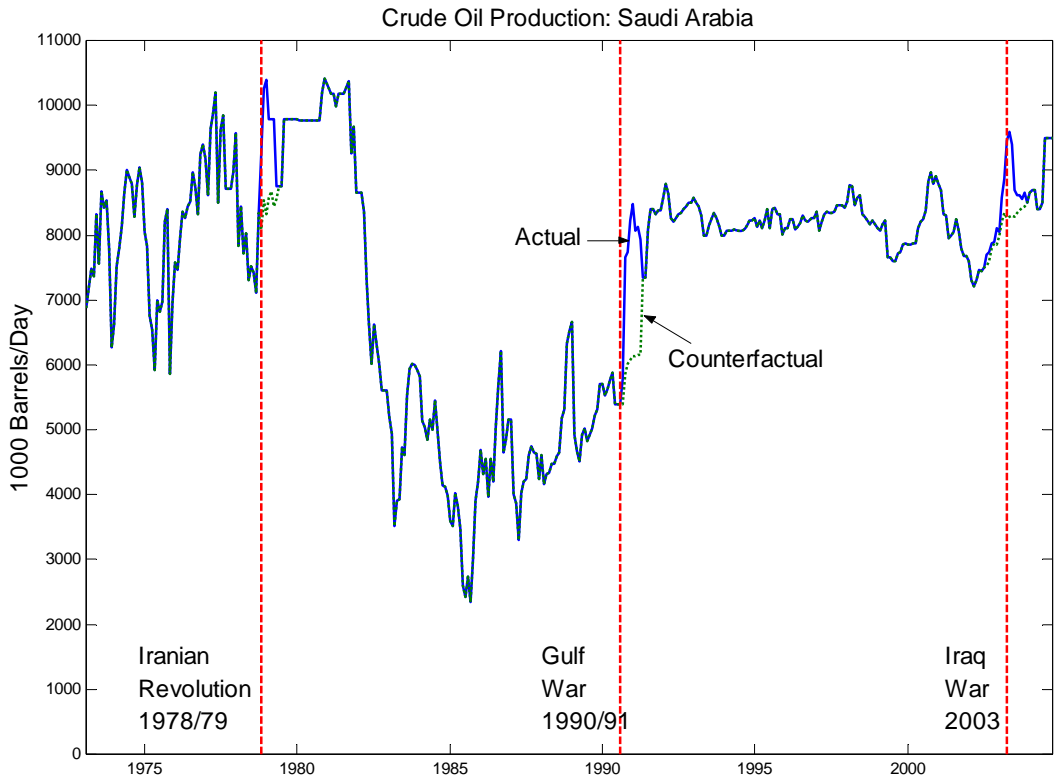


Figure 7: Venezuela


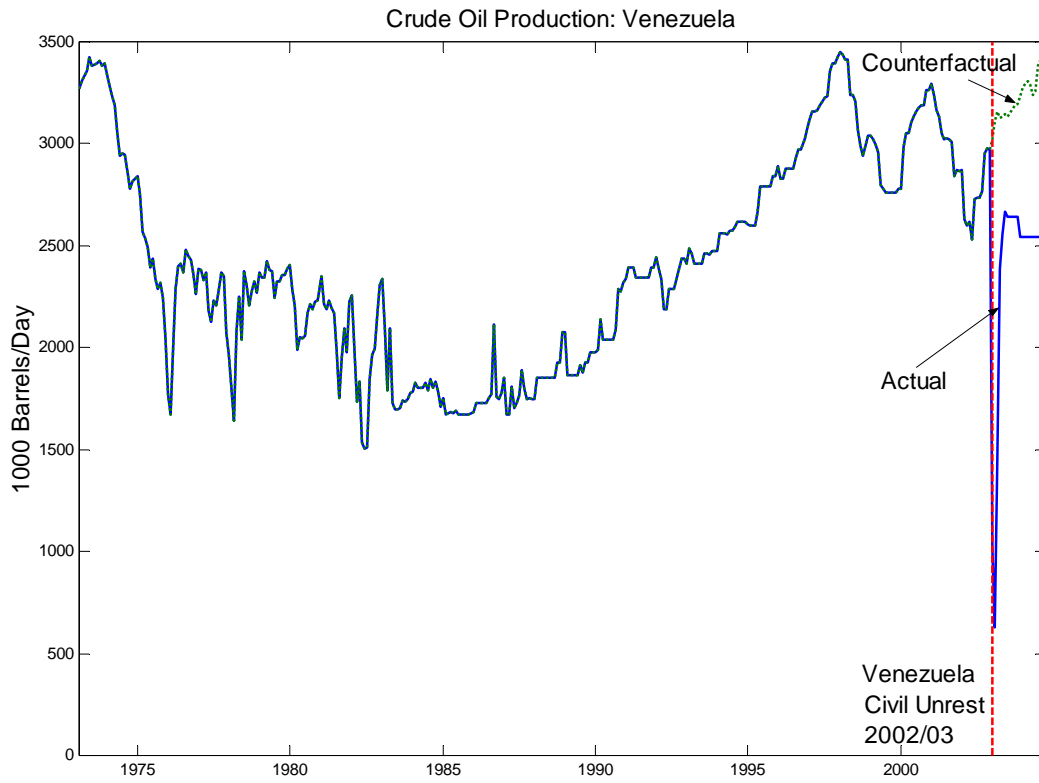


Figure 8: Exogenous Oil Production Shortfalls: OPEC

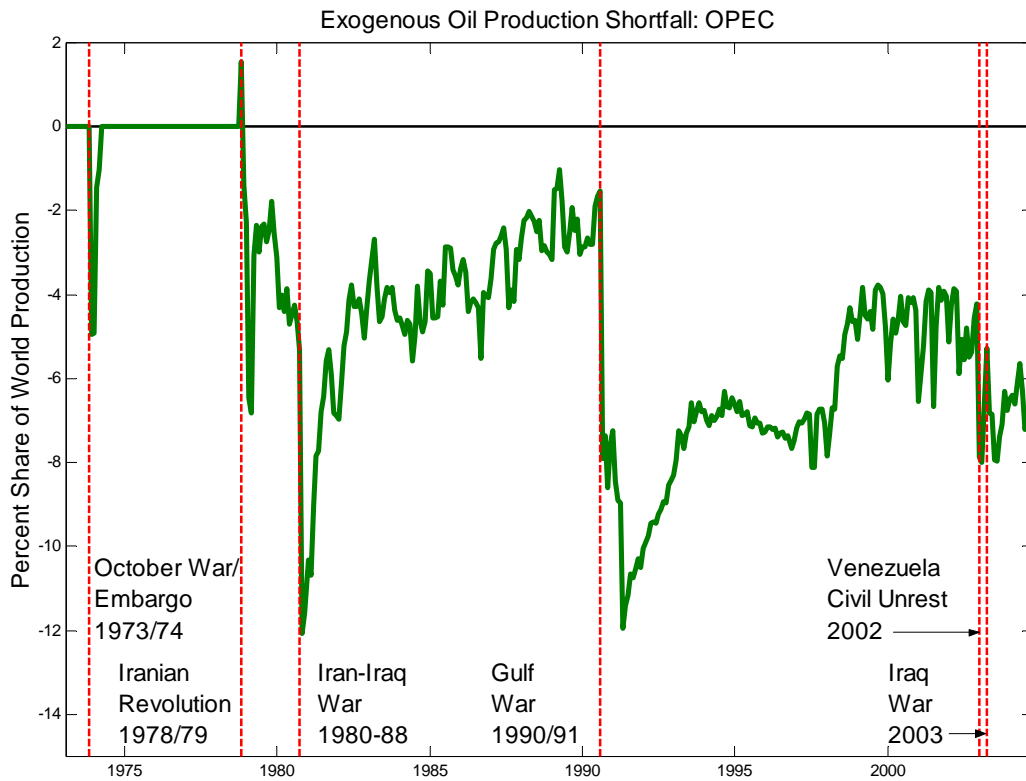

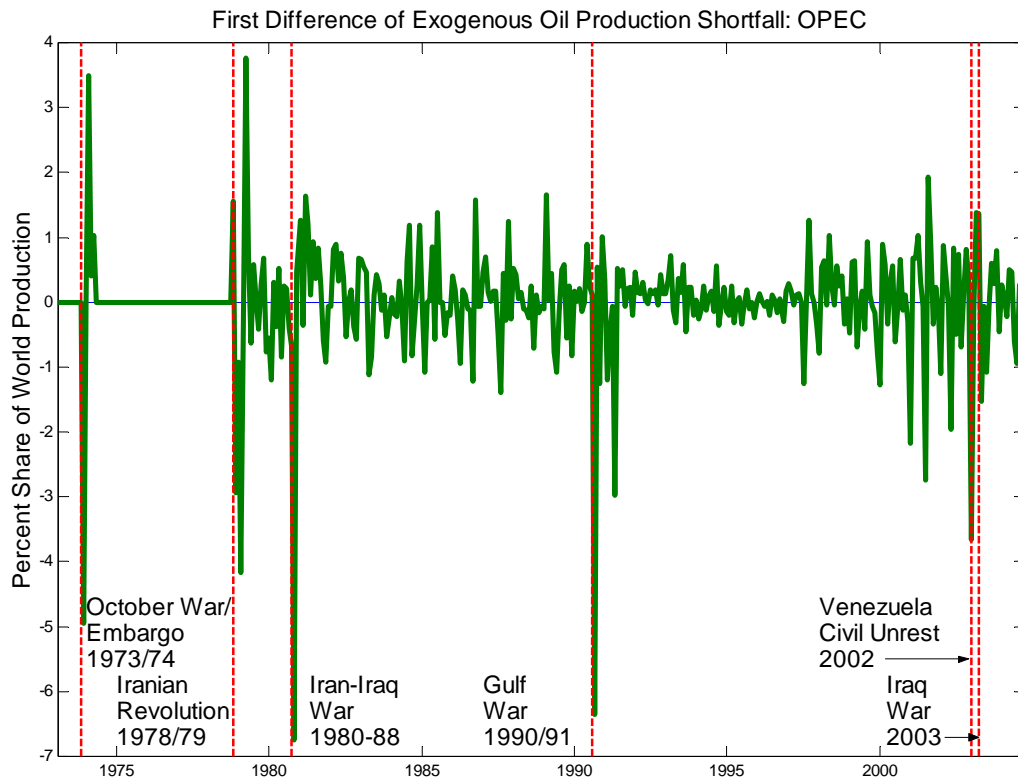
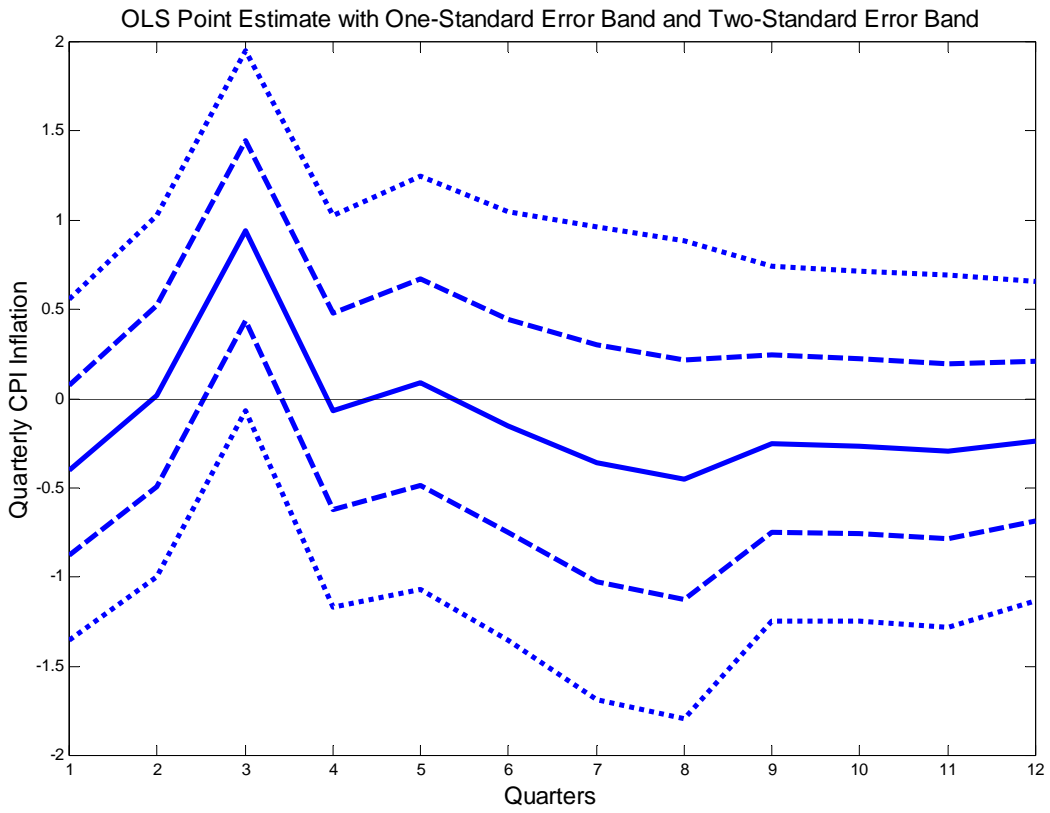
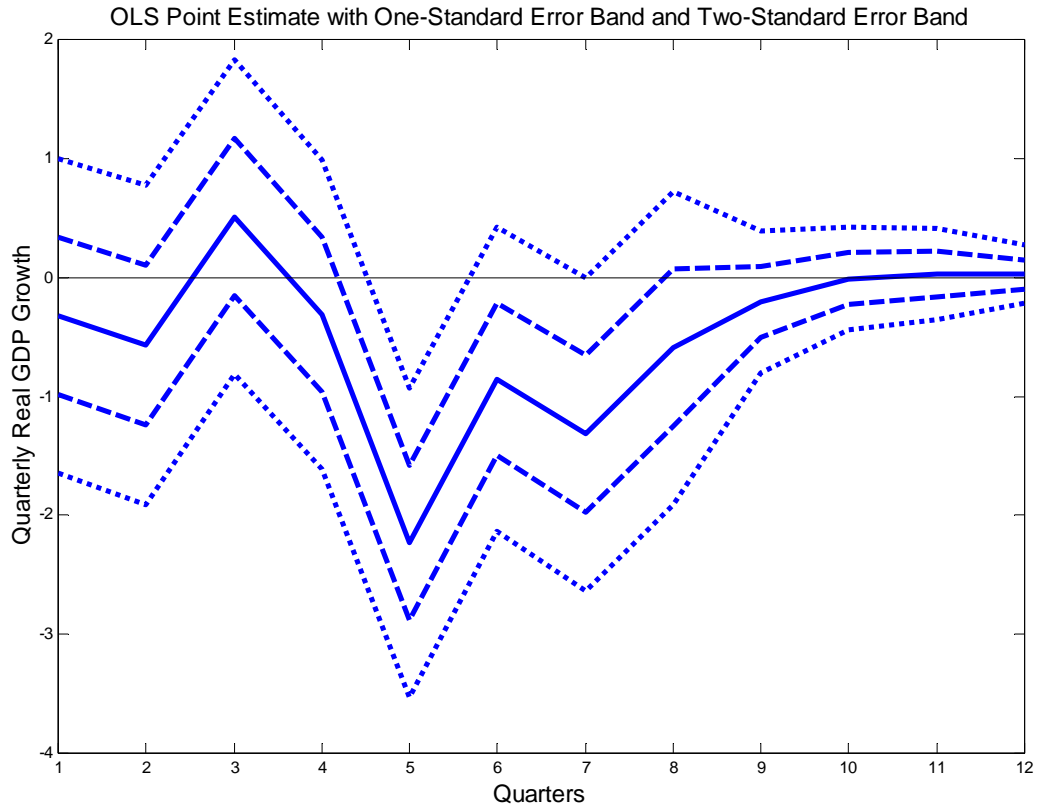


Figure 9: Exogenous Oil Supply Shocks: OPEC



**Figure 10: Dynamic Effects of a 10% World Oil Supply Disruption
Baseline Counterfactual**



**Figure 11: Counterfactual U.S. Economic History 1973-2004
Without Exogenous Oil Supply Shocks**





Figure 12: Oil Price Shocks Measured by Net Oil Price Increase (Relative to Previous Three Years)

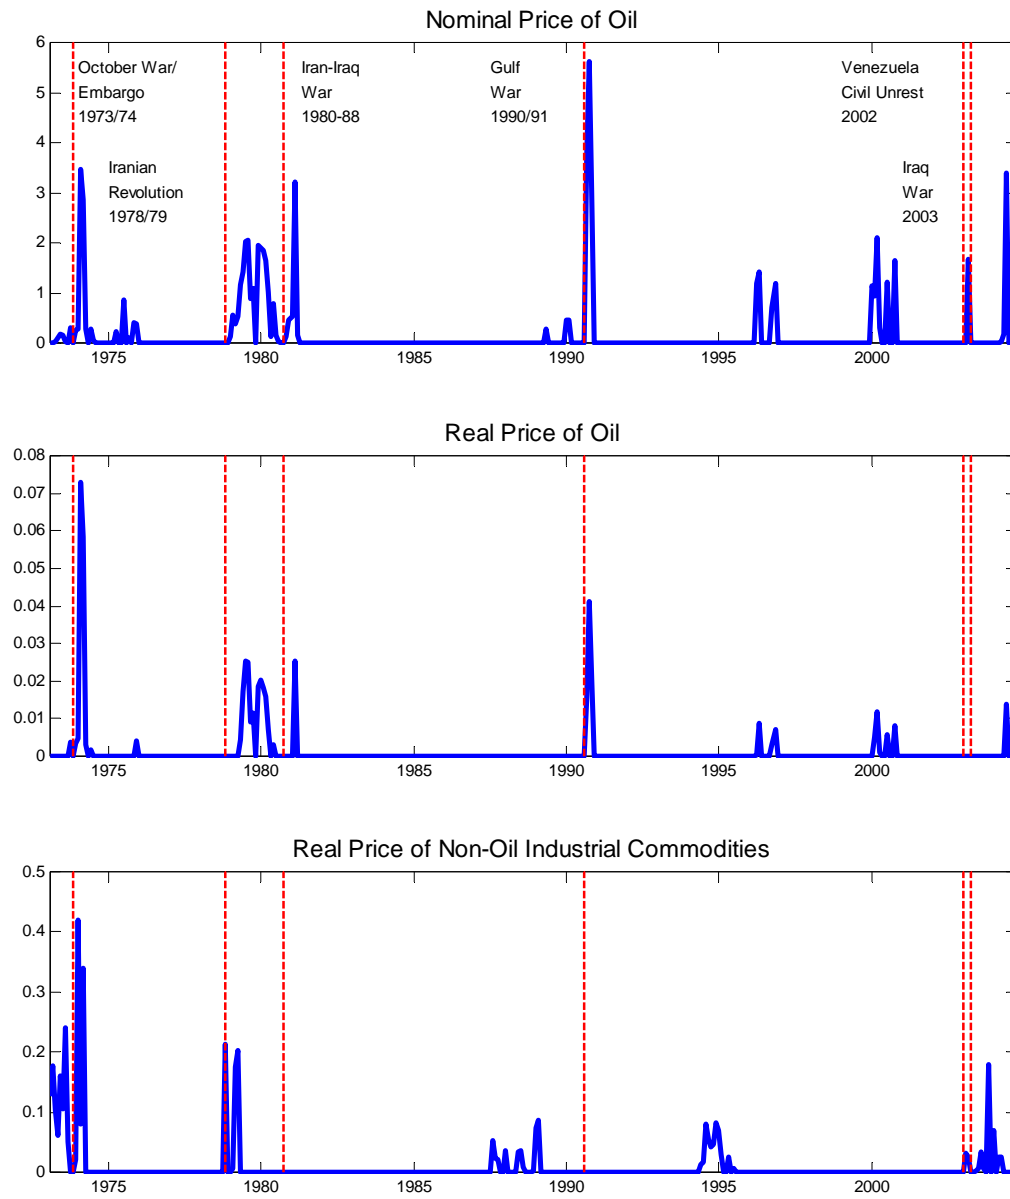


Figure 13: Comparison with Hamilton's Quantitative Dummy Series


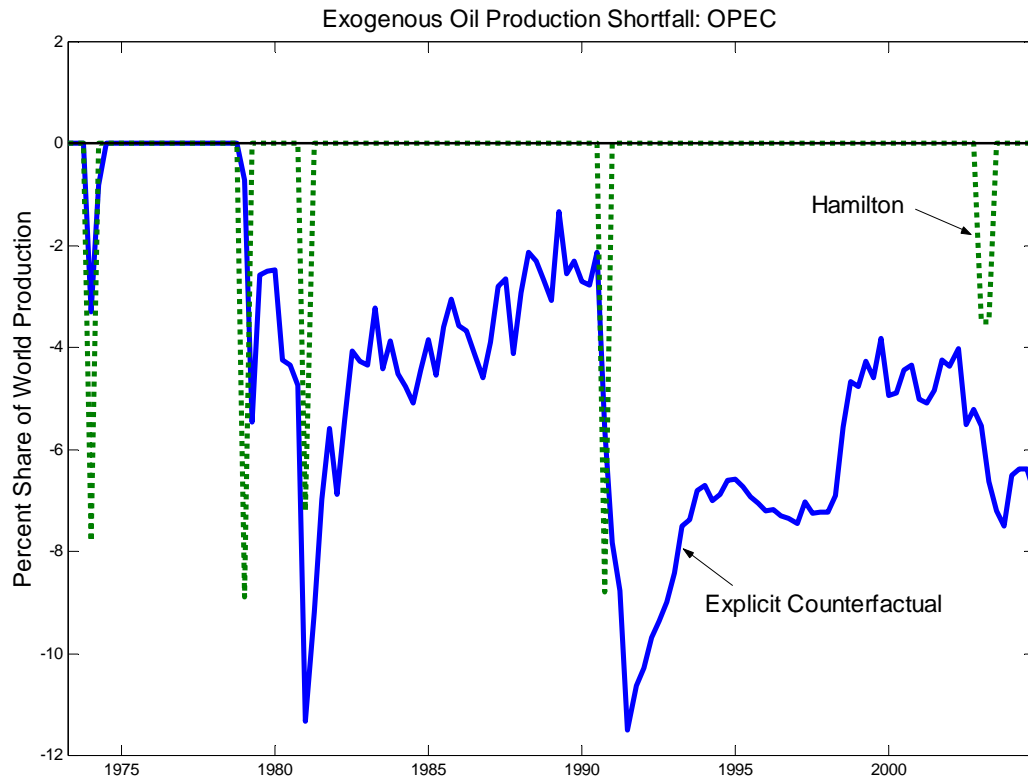


Figure 14: Actual Oil Price Changes and Oil Price Changes Explained by Exogenous Oil Supply Shocks

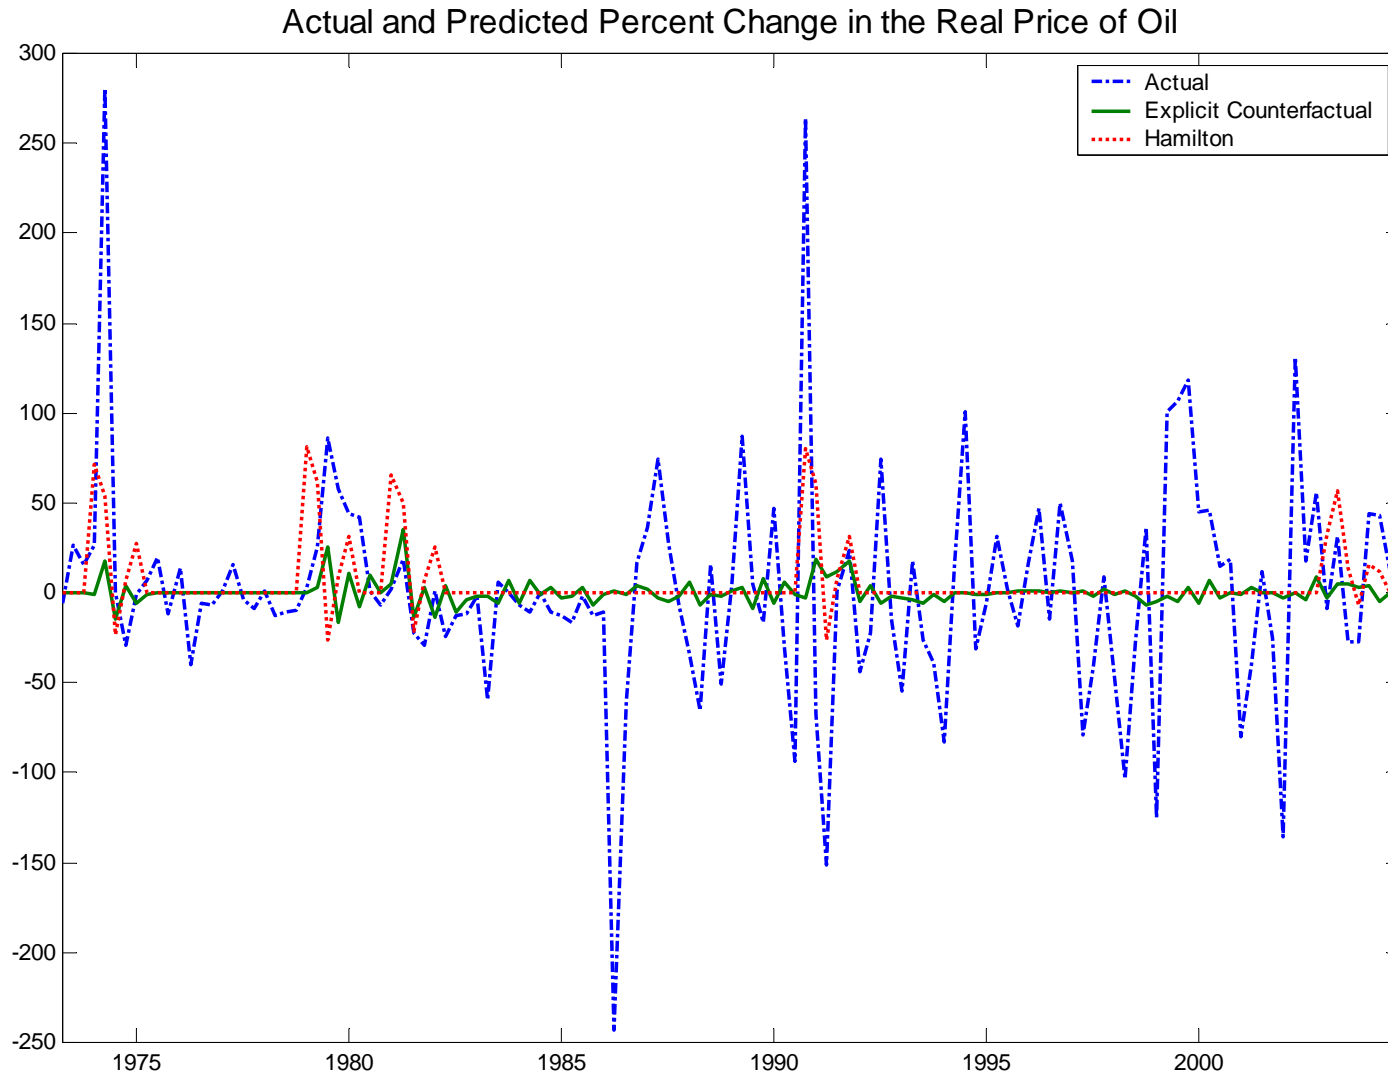


Figure 15: Actual Oil Price Changes and Oil Price Changes Explained by Exogenous Oil Supply Shocks after Including Dummy for Imminent Military Threat to Saudi Oil Fields in Last Two Quarters of 1990





Figure 16: Dynamic Effects of a 10% Oil Price Increase



**Figure 17: Dynamic Effects of a 10% World Oil Supply Disruption
Quantitative Dummies**

