DISCUSSION PAPER SERIES

No. 6559

THE ECONOMIC EFFECTS OF ENERGY PRICE SHOCKS

Lutz Kilian

INTERNATIONAL MACROECONOMICS

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP6559.asp

ISSN 0265-8003

THE ECONOMIC EFFECTS OF ENERGY PRICE SHOCKS

Lutz Kilian, University of Michigan and CEPR

Discussion Paper No. 6559
November 2007

Centre for Economic Policy Research
90–98 Goswell Rd, London EC1V 7RR, UK
Tel: (44 20) 7878 2900, Fax: (44 20) 7878 2999
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 6559

November 2007

ABSTRACT

The Economic Effects of Energy Price Shocks

Large fluctuations in energy prices have been a distinguishing characteristic of the U.S. economy since the 1970s. Turmoil in the Middle East, rising energy prices in the U.S. and evidence of global warming recently have reignited interest in the link between energy prices and economic performance. This paper addresses a number of the key issues in this debate: What are energy price shocks and where do they come from? How responsive is energy demand to changes in energy prices? How do consumers' expenditure patterns evolve in response to energy price shocks? How do energy price shocks affect real output, inflation, stock markets and the balance-of-payments? Why do energy price increases seem to cause recessions, but energy price decreases do not seem to cause expansions? Why has there been a surge in gasoline prices in recent years? Why has this new energy price shock not caused a recession so far? Have the effects of energy price shocks waned since the 1980s and, if so, why? As the paper demonstrates, it is critical to account for the endogeneity of energy prices and to differentiate between the effects of demand and supply shocks in energy markets, when answering these questions.

JEL Classification: E21 and Q43
Keywords: asymmetry, causality, channels of transmission, crude oil, elasticity, gasoline, price shocks and propagation

Lutz Kilian
Department of Economics
University of Michigan
238 Lorch Hall
Ann Arbor, MI 48109-1220
USA
Email: lkilian@umich.edu

Submitted 31 October 2007

1 Introduction

The price of energy is only one of many prices faced by households and firms, yet it attracts a disproportionate amount of attention in the media and from policymakers and economists. A common perception is that energy price increases are fundamentally different from increases in the prices of other goods. One reason is that energy prices at times experience sharp and sustained increases not typical of other goods and services. A second reason is that these price increases matter more than in the case of other goods because the demand for energy is comparatively inelastic. For example, most workers have to drive to work every day and thus have little choice but to acquiesce to higher gasoline prices. Similarly, households have little choice but to endure higher natural gas prices, as they cannot afford to leave their homes unheated. A third reason that energy prices are perceived to be different is that energy price fluctuations seem to be determined by forces that are exogenous to the U.S. economy such as political strife in the Middle East. A fourth reason is that major energy price increases in the past have often been followed by severe economic dislocations, suggesting a causal link from higher energy prices to recessions, higher unemployment and possibly inflation.

In this paper, I selectively review the recent literature on the effect of energy price shocks on the U.S. economy. For a complementary survey of the relationship between oil prices and the macroeconomy the reader is referred to Hamilton (2005). My objective is not to provide a comprehensive survey of the literature. Rather I wish to highlight some recent methodological developments and to outline how these developments have altered our perceptions of where energy price shocks originate, how they are transmitted, and how much they affect the economy. The paper addresses a number of the key questions in the ongoing debate about the economic effects of energy price shocks: Where do energy price shocks originate? How responsive is energy demand to changes in energy prices? How do consumers' expenditure patterns evolve in response to energy price shocks? How do energy price shocks affect real output, inflation, stock markets and the current account? Why do energy price increases seem to cause recessions, but energy price decreases do not seem to cause economic expansions? Why has there been a surge in gasoline prices in recent years? Why has this new energy price shock not caused a recession so far? Have the effects of energy price shocks waned since the 1980s and, if so, why?

The remainder of the paper is organized as follows. In section 2, I discuss the identification of exogenous shifts in energy prices with special emphasis on alternative specifications of energy price shocks and alternative frameworks for estimating the effects of energy price shocks. Section 3 provides an overview of the effects of unanticipated changes in energy prices on U.S. consumer expenditures and firms' investment expenditures. I discuss the most prominent channels of transmission and the empirical evidence in their support. I also address the question of whether there is an asymmetry in the responses to energy price increases and energy price decreases. Section 3 also contains detailed estimates of the responsiveness of energy consumption to higher energy prices and, more generally, assesses how consumers' expenditure patterns evolve in response to energy price shocks. The section concludes with a discussion of the link between crude oil prices and monetary policy. In section 4, I address the question of why the effects of energy price shocks have weakened since the second half of the 1980s. Section 5 illustrates how disentangling demand and supply shocks in energy markets can help us understand the evolution of energy prices. I focus on explaining the recent surge in U.S. gasoline prices. This section also demonstrates the differential impact of demand and supply shocks in energy markets on U.S. macroeconomic aggregates. The emphasis is on the responses of real GDP, consumer prices, real stock returns and the current account. The concluding remarks are in section 6.

2 Methodological Issues Raised by the Endogeneity of Energy Prices

It is widely accepted that energy prices in general and crude oil prices in particular have been endogenous with respect to U.S. macroeconomic conditions dating back to the early 1970s. Endogeneity here refers to the fact that not only do energy prices affect the U.S. economy, but that there is reverse causality from U.S. and more generally global macroeconomic aggregates to the price of energy. Clearly, both the supply of energy and the demand for energy depend on global macroeconomic aggregates such as global real economic activity and interest rates (see Barsky and Kilian 2002, 2004). Thus, a correlation between energy prices and macroeconomic outcomes does not necessarily imply causation. One response to this problem has been to apply statistical transformations to the price of energy to extract the exogenous component of oil prices. The leading example of this approach is the net oil price increase measure proposed by Hamilton (1996, 2003).[1]

2.1 Measures of Net Oil Price Increases

An innovative approach to dealing directly with the lack of exogeneity of the price of oil was proposed in Hamilton (2003). Building on work by Mork (1989), Lee, Ni and Ratti (1995) and Hamilton (1996), Hamilton suggested that, although the price of crude oil itself has been endogenous since 1973, a suitable nonlinear transformation of the price of oil (based on the amount by which nominal

[1] A number of alternative transformations have been suggested by, among others, Lee, Ni and Ratti (1995). Hamilton (2003) shows that these alternative measures in general produce similar results to the net increase measure.

oil prices exceed their maximum value over the previous three years) may be treated as exogenous. Such regressors have been used widely in studying the impulse responses of U.S. sectoral and macroeconomic aggregates to oil price shocks. The purpose of the statistical transformation of oil prices is to isolate the component of the price of crude oil that can be attributed to political events in the Middle East, which in turn are exogenous to global macroeconomic conditions.

Rather than simply asserting the exogeneity of the transformed oil price measure, as other authors have been content to do in similar contexts, Hamilton proceeded to establish formally the exogeneity of net oil price increases. He made the case that using the oil price changes predicted by exogenous oil supply variations as instruments in a regression of real GDP growth on lagged percentage changes in nominal oil prices results in estimates of the structural regression coefficients that look remarkably similar to the reduced form estimates obtained from regressing GDP growth on the net oil price increases instead. Hamilton concluded that the latter reduced-form relationship effectively represented a causal relationship, lending credence to the practice of treating net oil price increases as exogenous.[2]

Although Hamilton's analysis represents an important step forward in this literature, there is reason to be skeptical of the exogeneity of the net oil price increase measure because of the nature of the instruments used by Hamilton. 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). In the presence of many variables to be instrumented, as in the instrumental variable (IV) regressions presented in Hamilton (2003), weak instrument problems may be detected based on the $g_{\min}$ statistic of Cragg and Donald (1993). Critical values for a formal test of the null hypothesis of weak instruments have been compiled by Stock and Yogo (2005). Table 1 presents IV regression estimates of the type presented in Hamilton (2003) in support of the exogeneity of the net increase measure of crude oil prices. In addition to the specification favored by Hamilton (2003), which is shown in column (1), I include a number of alternative specifications involving different measures of exogenous crude oil supply shocks developed in Kilian (2007a,b), different sample periods, and different measures of nominal and real oil prices. Table 1 shows that in *all* cases the estimated $g_{\min}$-statistics are far below the least conservative critical value of about 4, suggesting that a weak instrument problem cannot be ruled out. The presence of weak instruments would render the IV estimates inconsistent and standard inference on the dynamic effects of exogenous variations in the price of oil invalid, invalidating any comparison with the reduced form evidence. Consequently, we have no credible evidence that the responses to net oil price increases correspond to those of truly

[2] For example, Hamilton (2003, p. 395) writes that nonlinear transformations "... 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).

causal models.

This evidence of weak instruments is not surprising, as it has been shown that measures of exogenous oil supply shocks driven by political events in the Middle East fail to explain much of the observed fluctuations in crude oil prices. This result is robust across alternative specifications of these exogenous political supply shocks. For example, at most one fourth of the observed increase in the price of oil in 1973/74 can be explained based on such shocks (see Kilian 2007a). In fact, recent evidence in Kilian (2007c) suggests that crude oil supply shocks, narrowly defined as unanticipated changes in world crude oil production, are far less important for understanding crude oil prices than are shocks to the demand for crude oil. This point will be discussed in more detail in section 5.

The fact that net oil price increases are not measures of exogenous oil price shocks can also be seen more directly. For example, Kilian (2007a) shows that oil price shocks detected by the nonlinear transformation proposed by Hamilton (2003) occur at times when exogenous oil supply shocks in the Middle East were conspicuously absent, and that there are major exogenous events that are not followed by oil price shocks. Kilian also shows that the same procedure applied to other industrial commodity prices generates spurious evidence of exogenous price shocks. Hence, oil price series must be treated as endogenous whether they have been transformed to net oil price increases or not.

The lack of exogeneity of the price of oil is not as much of a problem as it may seem, because exogeneity is not required for estimating the economic effects of changes in oil prices. A much weaker and more defensible assumption than strict exogeneity of the price of oil is that innovations to the oil price series (whether transformed or not) are predetermined with respect to U.S. macroeconomic aggregates. In other words, the price of oil responds to changes in macroeconomic conditions only with a delay. Under this assumption, recursively identified vector autoregressions with energy prices ordered first may be used to estimate the dynamic effect of a change in energy prices, whether the energy price has been transformed or not. Indeed, this assumption forms the basis of a number of influential papers in the literature including Bernanke, Gertler and Watson (1997) and Davis and Haltiwanger (2001) that focus on the linearly unpredictable component of the net increase in oil prices. The assumption of predeterminedness typically is inappropriate when working with annual data, but may provide a good approximation when working with quarterly and in particular with monthly data. Thus, there are clear advantages to studying the effects of energy price shocks in a high-frequency time series context compared to panel studies using data at lower frequencies.[3] The implementation of this VAR approach is discussed in more detail in section 2.5. In section 5, I will discuss in detail the economic rationale of the assumption of predetermined energy prices in light

[3] In the latter case, it becomes necessary to find suitable instruments for energy price changes. For example, Cullen, Friedberg and Wolfram (2004) use weather data as exogenous instruments for home energy costs.

of recent work stressing the important distinction between demand and supply shocks in energy markets.

2.2 The Effect of Changes in the Energy Share

It is sometimes argued that regressions of macroeconomic aggregates on unanticipated energy price changes are potentially misleading in that they fail to account for the declining share of energy in value added since the 1970s. Figure 1 shows that indeed this share has been falling from a maximum of 5% in 1981 to a low of 1% in 1998, but by 2005 the share had recovered to its original level of 3.3% in 1977.[4] Interestingly, the pattern of these fluctuations seems to reflect primarily changes in the price of crude oil rather than shifts in energy use. While not trivial, the observed fluctuations in the energy share in value added are largely immaterial for regression estimates of the effect of energy price shocks because the share does not fluctuate enough on a quarter-to-quarter basis to affect much the response estimates of U.S. macroeconomic aggregates. Weighted and unweighted quarterly energy price changes have a correlation of 99%. Thus, little is lost by studying the effect of unweighted energy price shocks, as has been demonstrated in Edelstein and Kilian (2007b). Weighting can be important, however, in constructing measures of the retail energy prices faced by households and firms, as the next subsection illustrates.

2.3 On the Choice of the Energy Price Series

Much of the work on energy price shocks has focused on the price of crude oil. This approach reflects the perception that the bulk of the fluctuations in energy prices since the 1970s has been driven by disturbances in global crude oil markets that are reflected in the price of imported crude oil and are transmitted from the crude oil market to retail energy prices. While it is true that disturbances in global crude oil markets are typically reflected in gasoline prices, this is not the only major source of shocks to gasoline prices. A good example is provided by the effects of Hurricanes Rita and Katrina in late 2005. Whereas the reduction of U.S. crude oil supply associated with these exogenous events was negligible on a global scale, as shown in Kilian (2007c), the reduction in refining capacity was not. It constituted both a decline in the demand for crude oil (associated with a fall of crude oil prices) and a decline in the supply of gasoline and other refined products, reflected in a sharp increase in gasoline prices. This example illustrates that from a consumer's point of view a direct measure of retail gasoline prices may be more relevant than a measure of crude oil prices.

[4] Following Rotemberg and Woodford (1996), the energy share in value added is approximated by the sum of nominal value added in oil and gas extraction and imports of petroleum and petroleum products, divided by nominal GDP. No data are available prior to 1977.

As of 2002, according to the BEA, gasoline accounts for 48.7% of all energy used by consumers, compared with 12.3% for natural gas and 33.8% for electricity. Since gasoline is by far the most important form of energy consumed in the United States and the one with the most volatile price, little would be lost by focusing on gasoline prices alone in studying the response of consumer expenditures. In contrast, in studying firm behavior neither gasoline nor crude oil prices will be representative of energy prices. For producers, based on 2002 BLS data, electricity makes up 40.3% of energy use, natural gas 14.5%, and unleaded gasoline only 14%, followed by diesel fuel (11.4%) and jet fuel (9.7%). The difference in weights has important implications for the magnitude of energy price shocks. For example, during the Persian Gulf War in 1990, crude oil prices rose by 83%, whereas the intermediate energy prices faced by firms only rose by 12%. This example illustrates that the choice of energy price series may matter a great deal. Different questions may require a different measure of energy prices.

2.4 Alternative Specifications of Energy Price Shocks

In studying the effects of energy price shocks, a natural baseline is the hypothesis that firms and consumers respond proportionately to a percent change in energy prices, regardless of the magnitude of the change. We will refer to this model as the percent change or C specification. There are other behavioral models, however, that involve nonlinear transformations of the price of energy. One alternative is the possibility that consumers and firms only respond to large shocks. For example, the presence of costs to monitoring energy expenditures and of costs of adjusting consumption patterns might make households reluctant to respond to small energy price changes (see Goldberg 1998). One may allow for that possibility by defining energy price shocks as increases or decreases that exceed, say, one standard deviation of the percent change in energy prices. We will refer to this model as the large percent change or LC specification. A third possibility is that consumers and firms respond only to changes in energy prices that are unprecedented in recent history. This view calls for a measure of the net energy price increase of the type proposed by Hamilton (1996, 2003), which is defined as the difference between the current price and the maximum price over the previous year (or alternatively the previous three years) if the current price is larger than the previous maximum, and zero otherwise. While Hamilton focuses on net energy price increases, Edelstein and Kilian (2007a,b) extend this measure to include net energy price declines that are constructed in a similar fashion. They document statistically significant responses of U.S. macroeconomic aggregates to both net increases and net decreases. The resulting model will be referred to as the net percent change or NC specification. Figure 2 illustrates the differences between these three specifications for U.S. retail energy prices. The example in Figure 2 (as well as all subsequent empirical analysis in this paper) is based on the real price of energy, which is the price relevant for resource allocation.[5]

An important question is how to choose between these alternative behavioral models. Edelstein and Kilian (2007a,b) show that the C specification cannot be rejected in favor of the LC specification in general. Whether the C or NC specification is more appropriate remains an active area of research. Since the NC specification is not nested in the C specification, it is not straightforward to test these specifications, or even to compare the magnitudes of the responses implied by the two specifications. In this paper, I will focus on the C specification, but note that the qualitative results would be similar under the NC specification, and indeed in most cases for all three specifications.

2.5 Alternative VAR Frameworks for Modeling Energy Price Shocks

Models that treat innovations to energy prices as predetermined with respect to macroeconomic aggregates at high frequency have been used in the literature for many years (see, e.g., Rotemberg and Woodford 1996; Blanchard and Galí 2007). A full discussion of the rationale of the assumption of predetermined energy prices can be found in section 5.3. The assumption of predetermined energy prices permits the consistent estimation of the expected response of real U.S. macroeconomic aggregates to an innovation in energy prices. In conjunction with the plausible assumption that there are no other exogenous events that are correlated with the exogenous energy price innovation, these impulse responses can be interpreted as the causal effect of the energy price innovation (see Cooley and LeRoy 1985).

A convenient vehicle for assessing the average economic effects of energy price shocks on a given macroeconomic aggregate is a recursively identified bivariate vector autoregression (VAR), in which the percent change in real energy prices is ordered first and the macroeconomic aggregate of interest is ordered second. For example, we may assess the response of real consumption to a real energy price innovation by specifying a model in the percent change of the real price of energy and the percent growth in real consumption. Generalizations to other specifications of energy price shocks are straightforward. In section 3.1 and 3.2, I use this bivariate workhorse model to quantify in detail the responses of consumption and investment expenditures to real energy price shocks. These results will be used in assessing the empirical support for the main channels of transmission discussed in the literature.

There is no loss of generality from restricting ourselves to a bivariate model under the maintained assumption of predetermined energy price innovations, if we are only interested in consistently esti-

[5] For further discussion of the distinction between nominal and real energy prices see, e.g., Hamilton (2005) and Kilian (2007a).

mating the causal effects of energy price innovations on macroeconomic aggregates. If, in addition, we want to assess the precise nature of the transmission, a bivariate model will not suffice. For example, it is common in the literature to estimate larger-dimensional VAR models of the effects of unanticipated oil price shocks that include a monetary policy-reaction function (see, e.g., Bernanke, Gertler and Watson 1997, 2004; Balke, Brown and Yücel 2002; Lee and Ni 2002; Hamilton and Herrera 2004; Jiménez-Rodriguez and Sánchez 2005; Herrera and Pesavento 2007). Such VAR models can be useful in assessing the extent to which the overall response of macroeconomic aggregates to unanticipated energy price changes is driven by the response of the central bank to the change in energy prices. Since these larger structural VAR models require further identifying assumptions (in addition to the assumption of predetermined oil prices) that may not be realistic, and since larger-dimensional models tend to be less precisely estimated than bivariate models, there is no reason to depart from the baseline bivariate VAR model for the purpose of the analysis in sections 3.1 and 3.2. Models of the reaction of monetary policy to oil price shocks will be considered in section 3.5.

VAR models based on the assumption of predetermined energy prices are not without their limitations. Notably, they do not distinguish between energy price innovations driven by supply shocks or demand shocks. The latter distinction can be crucial because different demand or supply shocks in energy markets tend to cause different responses of macroeconomic aggregates. In section 5, I will discuss an alternative VAR approach that makes this distinction apparent, and I will make the case that bivariate VAR models based on the assumption of predetermined retail energy prices, as used in sections 3.1 and 3.2, are well identified even in light of the insights provided by alternative VAR models highlighting the distinction between various demand and supply shocks in energy markets. The resulting impulse responses can be interpreted as an estimate of the economic effects of an average energy price shock during the sample period in question.

3 The Effect of Energy Price Shocks on the Economy

The standard approach to modeling energy price shocks has been to focus on the effects of an exogenous increase in the price of imported crude oil. Such terms-of-trade shocks traditionally have been thought to matter for the domestic economy through their effects on production decisions (see, e.g., Kim and Loungani 1992; Backus and Crucini 2000). In this view, oil is treated as an intermediate input in domestic production. How imported oil enters the production function for domestic value added is one of the most studied and least resolved issues in empirical macroeconomics (Backus and Crucini 2000).

There are well known problems in explaining a decline in real GDP based on this intermediate input cost or supply channel, however. The first problem is that the interpretation of crude oil as an intermediate input in the value added production function is questionable if we think of oil as an imported commodity. Under standard assumptions, imported oil enters the production function of domestic gross output, but it does not enter the production function of domestic value added (see, e.g., Rotemberg and Woodford 1996). Since gross output is separable in value added and imported energy, holding capital and labor fixed, oil price shocks do not move value added. Hence, oil price shocks by definition cannot be interpreted as productivity shocks for real GDP (see Barsky and Kilian 2004). The second problem is that, to the extent that oil prices affect domestic output, under standard assumptions their impact should be bounded by the cost share of oil in gross domestic production, which is known to be very small.

There are three distinct proposals in the literature for dealing with these problems. All three involve major modifications of the baseline supply shock model to generate quantitatively important effects of oil price shocks on real GDP. The first proposal is Rotemberg and Woodford's (1996) model which relies on large and time-varying markups to generate large effects of oil price shocks on real GDP. The second proposal is Atkeson and Kehoe's (1999) putty-clay model which appeals to capital-energy complementarities in production. The third proposal is due to Finn (2000). Finn establishes that in a perfectly competitive model, in which energy is essential to obtaining a service flow from capital, there may be a large effect of oil price shocks on real GDP. In all three models, the supply channel of the transmission of energy price shocks may be quantitatively important, yet there is no consensus which, if any, of these models has empirical support. For example, it remains to be shown that mark-ups in the U.S. economy are as large and as time-varying as required for the Rotemberg and Woodford model to generate large effects on value added. Likewise, it remains to be shown that changes in capacity utilization in response to oil price shocks are indeed as important and pervasive in the real world as they are in Finn's model. Similarly, the microeconomic evidence on the existence and quantitative importance of capital-energy complementarities is mixed at best. A second unresolved issue is whether these models can account for a large share of business cycle fluctuations in real GDP. A third issue is that all three models postulate that oil prices follow an exogenous stochastic process, an assumption that is at odds with both the data and standard economic models of the oil market, as discussed in section 2.

These caveats are important because in the absence of an empirically supported model of the supply channel, there is no reason to expect global oil price shocks to exert major effects on the domestic economy. In part in response to these challenges, another branch of the literature has developed that focuses on the reduction in the demand for goods and services triggered by energy price shocks rather than treating energy price shocks as aggregate supply shocks for the U.S. economy (or as productivity shocks for U.S. domestic production). In this alternative view, the primary channel of transmission is on the demand side of the economy. For example, in a recent survey on the effects of energy price shocks, Hamilton (2005) stresses that a key mechanism whereby energy price shocks affect the economy is through a disruption in consumers' and firms' spending on goods and services other than energy. This view is consistent with anecdotal evidence of how oil price shocks affect U.S. industries (see Lee and Ni 2002). It is also shared by many policymakers. There is a widespread perception that an increase in energy prices slows economic growth primarily through its effects on consumer spending (see, e.g., Bernanke 2006a). In the remainder of section 3, I outline the economic rationale for the demand channel of transmission and assess its empirical support. I focus on the response of real consumption first, before considering real investment expenditures in section 3.2.

3.1 How Do Consumer Expenditures Respond to Higher Energy Prices?

3.1.1 The Channels of Transmission

The literature has focused on four complementary mechanisms by which consumption expenditures may be directly affected by energy price changes. First, higher energy prices are expected to reduce discretionary income, as consumers have less money to spend after paying their energy bills.[6] All else equal, this *discretionary income effect* will be the larger, the less elastic the demand for energy, but even with perfectly inelastic energy demand the magnitude of the effect of a unit change in energy prices is bounded by the energy share in consumption. Second, changing energy prices may create uncertainty about the future path of the price of energy, causing consumers to postpone irreversible purchases of consumer durables (see Bernanke 1983, Pindyck 1991). Unlike the first effect, this *uncertainty effect* is limited to consumer durables. Third, even when purchase decisions are reversible, changes in uncertainty may have an effect on all forms of consumption, as consumers increase their *precautionary savings* in response to higher energy prices. Such a response may reflect, for example, the increased likelihood of becoming unemployed. Finally, consumption of durables that are complementary in use with energy (in that their operation requires energy) will tend to decline even more, as households delay or forego purchases of energy-using durables. This *operating cost effect* is more limited in scope than the uncertainty effect in that it only affects specific consumer

[6] Implicit in this view is the assertion that higher energy prices are primarily driven by higher prices for imported energy goods, and that at least some of the discretionary income lost from higher prices of imported energy goods is transfered abroad and is not recycled in the form of higher U.S. exports. In the case of a purely domestic energy price shocks (such as a shock to U.S. refining capacity), it is less obvious that there is an effect on aggregate discretionary income. First, the transfer of income to the refiner may be partially returned to the same consumers in the form of higher wages or higher stock returns on domestic energy companies. Second, even if the transfer is not returned, higher energy prices simply constitute an income transfer from one consumer to another that cancels in the aggregate.

durables. It should be most pronounced for motor vehicles (see Hamilton 1988).[7]

In addition, there may be *indirect* effects related to the changing patterns of consumption expenditures. A large literature has stressed that shifts in expenditure patterns driven by the uncertainty effect and operating cost effect amount to allocative disturbances that are likely to cause sectoral shifts throughout the economy (see, e.g., Davis (1987) and Hamilton (2005) for a review). For example, it has been argued that reduced expenditures on energy-intensive durables such as automobiles may cause the reallocation of capital and labor away from the automobile sector. As the dollar value of such purchases may be large relative to the value of the energy they use, even relatively small changes in energy prices (and hence in the purchasing power of consumers) can have large effects on output and employment (see Hamilton 1988). A similar reallocation may occur within the same sector, as consumers switch toward more energy-efficient durables (see Hamilton 1988; Bresnahan and Ramey 1993). In the presence of frictions in capital and labor markets, these intersectoral and intrasectoral reallocations will cause resources to be unemployed, thus causing further cutbacks in consumption and amplifying the effect of higher energy prices on the real economy. This indirect effect could be much larger than the direct effects listed earlier, and is considered by many economists to be the primary channel through which energy price shocks affect the economy (see, e.g., Davis and Haltiwanger (2001) and Lee and Ni (2002) and the references therein). Concerns over reallocation effects also help explain the preoccupation of policy makers with the effects of energy price shocks on the automobile sector (see, e.g., Bernanke 2006b).

Unlike the other three effects discussed above, the uncertainty effect and the reallocation effect necessarily generate asymmetric responses of macroeconomic aggregates to energy price increases and decreases. The asymmetry arises because these effects amplify the response of macroeconomic aggregates to energy price increases, but reduce the corresponding response to falling energy prices. In sections 3.1 and 3.2, I will deliberately abstract from the possible presence of such asymmetric effects. As will be discussed in section 3.3, there is no compelling statistical evidence against the symmetry hypothesis, and the symmetric model appears to fit the real consumption data rather well. The historical reasons for the prominence of asymmetric models in empirical and in theoretical work on oil price shocks, and the reasons why the apparent evidence of asymmetries is likely to be spurious, are discussed in section 3.3.

[7] This last effect need not involve a reduction in the number of automobiles sold. It can also take the form of consumers substituting small energy-inefficient automobiles for large energy-inefficient automobiles. If the latter automobiles tend to be lower priced, aggregate real consumption of automobiles may fall, even when the number of automobiles sold does not (see Bresnahan and Ramey 1993).

3.1.2 Estimating the Energy-Price Elasticities of Energy Demand

A central question in the literature is how price elastic the demand for energy is. Such estimates play an important role in assessing the potential impact of a carbon tax, for example, in assessing alternative regulatory policies, and in understanding the transmission of energy price shocks. The response of real consumption may be estimated by applying the bivariate regression model described in section 2 to various forms of energy consumption. All energy prices have been weighted by the nominal expenditure share on energy to obtain a measure of the gains and losses of households' purchasing power associated with energy price fluctuations (see Edelstein and Kilian 2007a). The sample period is 1970.2-2006.7. The results are expressed as elasticities with respect to the price of energy, evaluated at the average energy share. The upper panel of Table 2 shows that consumption of all forms of energy declines in response to energy price increases. The elasticity estimate for total energy consumption is -0.45 with error bounds of -0.27 and -0.66, but there are important differences across different forms of energy.

The strongest responses are observed for gasoline and for heating oil and coal. Contrary to the conventional wisdom, *gasoline* consumption responds immediately to unanticipated energy price increases reaching an elasticity of -0.48 after one year. The strikingly large response of -1.87 for heating oil and coal is likely due to households' ability to store heating oil in tanks. This storage feature allows households to delay purchases of new heating oil when the price of heating oil is high and to fill the tank completely when prices are low.[8] In contrast, electricity and natural gas are inherently unstorable, and gasoline may not be stored for safety reasons beyond the tank capacity of a car. Indeed, the declines in electricity consumption and in natural gas use are smaller and not statistically significant.

It is useful to put these estimates into perspective by comparing them with estimates based on other methodologies. Using a structural model, Reiss and White (2005) arrive at an estimate of the short-run price elasticity of electricity demand of -0.39. The point estimate in Table 2 is only -0.15 after twelve months, but the 95% confidence interval for the elasticity estimate includes -0.39. Dahl and Sterner (1991), in a comprehensive survey, report estimates of the short-run price elasticity of gasoline demand between -0.08 and -0.41. It is not clear, however, how well identified their estimates are. Our point estimate of -0.48 is larger with 95% error bands of -0.32 and -0.69, respectively.[9]

[8] For a discussion of this storable goods feature see Dudine, Hendel, and Lizzeri (2006).

[9] Recently, Hughes, Knittel, and Sperling (2007) have reported a sharp decline in the elasticity of gasoline demand when comparing estimates for 1975-1980 and for 2001-2006. If we split our sample in half, we observe a similar decline in the price elasticity of overall energy demand, but there are important differences across different forms of energy. Whereas the elasticity of gasoline demand and the elasticity of demand for heating oil and coal have declined consistent with the finding in Hughes et al., that of electricity demand has actually increased, and the elasticity of natural gas consumption has remained unchanged.

These estimates suggest that the demand for energy is not as unresponsive to energy prices as is sometimes believed.

3.1.3 Estimating the Energy-Price Elasticities of Non-Energy Consumption

The lower panel of Table 2 summarizes the corresponding elasticities of major non-energy real consumption aggregates. All estimates have the expected negative sign and most are statistically significant. Table 2 demonstrates that the overall elasticity of -0.15 is driven mainly by a reduction in vehicles purchases. The elasticity of demand for vehicles is -0.84. It can be shown that this estimate in turn primarily reflects reduced consumer demand for cars. There is little evidence of reduced demand for other durables.

It is useful to put these results in perspective. The sharp rise in gasoline prices in recent years has renewed interest in the question of how much higher energy prices affect consumer expenditures. Our analysis allows us to assess the overall effect of such a price increase on household consumption. Suppose, for example, that gasoline prices unexpectedly increase by 25 cents per gallon (which translates into a 6.85% increase in the overall price of energy, assuming all other energy prices remain unchanged). Then, one year later, a typical household with $4000 to spend per month will cut back its consumer expenditures by $41 based on the full-sample estimates (or by $20 based on the post-1987 estimates). This example illustrates that it takes repeated surprise increases in gasoline prices to generate large effects on household consumption.

It is instructive to disaggregate further the responses of consumption aggregates to higher energy prices. First, consider the likely consequences of an energy price increase based on the discretionary income effect alone. Given that households may choose to borrow or to dissave as a short-run response to higher energy prices, it is quite possible for the impact effect of such a shock on consumption to be smaller than consumers' loss in purchasing power, even when energy demand is inelastic. Such consumption smoothing is likely to be short-lived, however, and in the long run the response should be bounded by the magnitude of the purchasing power loss. Hence, a reasonable upper bound on the discretionary income effect is the initial reduction in purchasing power,

In practice, the long-run response could be much smaller than this bound to the extent that demand for energy declines over time, as households increasingly utilize extensive and intensive margins of adjustment in response to purchasing power losses driven by higher energy prices. It stands to reason that such efforts at energy conservation will increase over time. Beyond simple remedies such as driving less or changing the thermostat, households will gradually upgrade their home heating and insulation systems or trade in their gas-guzzling car for a more energy-efficient vehicle. Thus, a tighter bound may be obtained by taking account of the elasticity of energy demand obtained in the upper panel of table 2. Taking account of the price elasticity of energy demand in Table 2, the reduction in discretionary income associated with a 1% increase in energy prices is bounded by -0.04%. It can be shown that the response of total consumption is too large to be accounted for by a discretionary income effect alone.

Where does the excess decline come from? A decline in nondurables and services consumption by more than this bound can be interpreted as an indication that consumers reduce expenditures because they increase their precautionary savings. The data imply a marginally statistically insignificant precautionary savings effect of -0.08 percentage points for nondurables and a statistically significant effect of -0.07 percentage points for services. The corresponding precautionary savings effect of -0.16 percentage points implied for durables other than vehicles is not statistically significant. The excessively high response of vehicles compared to other durables, in contrast, is an indication of the presence of an operating cost effect. The operating cost effect on vehicles consumption lies between -0.60 and -0.65 percentage points and is statistically significant. Taken together, these three effects imply a larger reduction of real consumer expenditures in response to an unanticipated energy price increase than would be expected based on the small share of energy in consumption, but the overall effect on aggregate consumption is nevertheless small.

3.1.4 Are Oil Price Shocks Different from Other Terms-of-Trade Shocks?

Consumer expenditures depend on retail energy prices (such as the prices of gasoline, heating oil, electricity or natural gas) rather than the price of imported crude oil. Whereas shocks to the price of imported crude oil can be interpreted as terms-of-trade shocks, this is not necessarily the case for retail energy price shocks. One difference is that an unanticipated change in retail energy prices may be triggered by a terms of trade shock, but not every retail energy price shock implies a terms-of-trade shock. Moreover, the extent to which an oil price shock triggers a retail energy price shock depends on the degree of pass-through. A second difference is that the effect of retail energy price shocks on consumption depends on the loss in purchasing power associated with the energy price shock, which in turn depends on the expenditure share of energy. Many imported commodities such as copper are not being consumed directly, making it impossible to determine the loss in consumer purchasing power associated with an unanticipated increase in the price of copper.

Despite these differences, the analysis in section 3.1.3 has implications for the literature on terms-of-trade shocks (see, e.g., Backus and Crucini 2000). The results in Tables 2 and 3 on how the demand channel operates suggest important differences between oil price shocks and other terms-of-trade shocks that can be illustrated based on two examples. First, consider bananas as an example of an imported commodity that can be purchased by consumers. The effect on consumption of an unanticipated increase in the price of imported bananas (to the extent that it will be reflected in the retail price of bananas) will be bounded by the expenditure share for bananas. In contrast, a similar shock for imported crude oil (to the extent that it will be reflected in retail energy prices) will elicit an additional response beyond the limit set by the expenditure share. Not only does the operating cost effect amplify the effect of higher prices for imported oil (controlling for the expenditure share) compared with higher prices for imported bananas, but households are more likely to increase their precautionary savings, given the track record of previous oil price shocks, than they would be if the price of imported bananas increased.

Whereas bananas are a consumption good, most commodities of interest in the terms-of-trade literature are industrial and cannot be consumed directly. An example is imported copper. The supply shock channel would suggest that the effect of a terms-of-trade shock in copper and in crude oil should be similar, controlling for the expenditure shares, as both copper and oil are complementary inputs in the manufacturing of many products. The demand channel, in contrast, suggests an additional effect of oil price shocks on the economy because motor vehicles are complementary in use with oil (through the link from oil prices to retail energy prices), but not complementary in use with copper. This line of reasoning suggests that shocks to the price of imported crude oil should have larger effects, all else equal, than shocks the prices of other industrial commodity imports.

The conclusion that oil price shocks are different from other terms-of-trade shocks, of course, depends on the interpretation of oil price shocks as demand shocks for the domestic economy. Notwithstanding the clear conceptual differences between oil and other imported commodities, it is important not to overstate the quantitative importance of this distinction. Specifically, the elasticity estimates in Tables 2 and 3 suggest that, for shocks of the same magnitude, the differences in the effect of oil price shocks and other terms-of-trade shocks on aggregate expenditures, while large in relative terms, will be small in absolute terms.

3.1.5 How Do Expenditure Patterns Change in Response to Energy Price Shocks?

Section 3.1.3 investigated the response of major real consumption aggregates to the gains and losses in purchasing power associated with energy price shocks. It is equally important to understand how individual expenditure items respond to such shocks. Shifts in expenditure patterns play an important role in theories of the demand channel of the transmission of energy price shocks. They are central to the existence of a reallocation effect, for example. Despite the preponderance of anecdotal evidence on how energy price shocks affect specific types of expenditures (such as gasoline consumption, purchases of SUVs, or dining out), to date there exists virtually no quantitative evidence of these effects.[10]

A detailed investigation of the responses of a large number of disaggregated expenditure items is beyond the scope of this paper. Regression analysis along the lines of section 3.1.3 (not reported to conserve space) confirms that there is evidence of shifts in expenditure patterns, although the responses may differ greatly by expenditure item. For example, the data show that restaurant and lodging expenditures are adversely affected by unanticipated energy price increases, as are sales of airline tickets. Automobile purchases are by far the most responsive expenditure item. Purchases of other durables such as furniture or appliances, by comparison, are far less sensitive to energy price fluctuations. The most surprising result is the low degree of responsiveness of nonessential expenditures on entertainment, sports and other leisure activities. Expenditures on public transportation and on food at home are among the few expenditure items that benefit from unanticipated higher energy prices.

Given the central importance of the automobile industry in many accounts of the transmission of energy price shocks, this subsection focuses on the responses of the components of motor vehicle consumption. The extent to which consumers buy smaller and more energy-efficient cars in response to purchasing power losses, in particular, plays a central role in policy discussions of the effect of higher oil prices (see Bresnahan and Ramey 1993). Figure 3 presents selected impulse response functions that shed light on this margin of adjustment. The impulse response functions were estimated in the same manner as in previous subsections with the results normalized to represent a 1% increase in energy prices.

Motor Vehicle Consumption An unanticipated 1% increase in energy prices causes a highly significant drop of -0.76% in the overall consumption of *motor vehicles and parts*. Figure 3 allows us to examine more closely different types of vehicles. Consumption of *pleasure boats* declines by -1.25% after 18 months. The response is persistent and highly statistically significant at most horizons. Consumption of *pleasure aircraft* declines by -1.05%. The response is persistent, and statistically significant at the 16% level at virtually all horizons. Consumption of *recreational vehicles* drops sharply and highly significantly in the short run, reaching a low of -1.58%, and remains statistically significant at the 16% level at long horizons. In contrast, consumption of *motorcycles* does not

[10] A detailed investigation of the responses of a large number of expenditure items (not reported to conserve space) confirms that there is evidence of shifts in expenditure patterns, although the responses may differ greatly by expenditure item. For example, the data show that restaurant and lodging expenditures are adversely affected by energy price shocks, as are sales of airline tickets. Automobile purchases are by far the most responsive expenditure item. Purchases of other durables such as furniture or appliances, by comparison, are far less sensitive to energy price fluctuations. The most surprising result is the low degree of responsiveness of nonessential expenditures on entertainment, sports and other leisure activities. Expenditures on public transportation and on food at home are among the few expenditure items that benefit from unanticipated higher energy prices.

change nor does consumption of *motor vehicle rentals* (not shown). While these results are generally consistent with the overall response of vehicles consumption, the combined consumption share of all these vehicles of 0.47% is small. Clearly, the bulk of the vehicles response is driven by automobile consumption.

If we are interested in whether there is an effect from reduced demand for automobiles on the automobile industry, the relevant metric is the effect of unanticipated increases in energy prices on the demand for new automobiles. Figure 3 shows that elasticity of demand for *new automobiles* is about -0.71, but the estimate is only partially statistically significant at the 2.5% level. This elasticity estimate is close to the fuel cost elasticity of -0.5 reported in Goldberg (1998) based on a structural model and micro data.

A possible explanation of the surprisingly low level of statistical significance of the response of new car purchases is that the effects of energy price shocks are not so much driven by an overall reduction in the demand for cars, but by an increase in the demand for energy-efficient small cars at the expense of energy-inefficient large cars. This view seems to fit not just the 1970s, but also the 2000s, as SUVs and pick-up trucks became increasingly unattractive to consumers. While we do not have data on the consumption of automobiles broken down by energy efficiency, we can contrast the real consumption of *new domestic automobiles* with that of *new foreign automobiles*.[11] To the extent that U.S. automobile manufacturers tend to produce less energy-efficient cars, a disproportionate decline in the consumption of domestically produced new cars would be evidence in favor of a shift in demand. Figure 3 shows a strong decline in new domestic automobile consumption that is significant at the 16% level at most horizons and at the 2.5% level at some horizons. In contrast, consumption of new foreign automobiles initially increases, consistent with the perception that there is increased demand for foreign, more energy efficient cars. The increase is statistically significantly at the 16% level. After four months, consumption of new foreign cars slumps as well, although the effect is not as persistent, statistically significant only at the 16% level, and smaller than for domestic autos. It can be shown that the excess response of the consumption of domestically produced automobiles over its foreign-produced counterpart is statistically significant for months 2, 3 and 4. The excess decline reaches its maximum of -0.88% after two months. The long run response is -0.63% and not statistically significant. An important question is how economically significant the decline in automobile consumption is. What the data tell us is that a permanent shock of the magnitude associated with Hurricane Katrina could wipe out 10.3% of the domestic demand for U.S. automobiles.

The consumption data on new automobiles do not include light trucks or trucks. A different

[11] Domestic cars are defined by the BEA to include cars assembled in the United States, Canada, or Mexico.

approach to determining the importance of shifts among different types of automobiles is to focus on unit sales reported by the BEA. While these data ignore the price of a given car (and hence differences in quality), they do allow us to assess whether purchases of light trucks (including minivans, SUVs, and pickup trucks) respond differently to unanticipated losses in purchasing power than purchases of regular automobiles. There has been much discussion of the softening market for SUVs in recent years. Figure 3 shows a mostly significant decline in *unit auto sales* at the 16% level (consistent with the evidence on new auto consumption). The decline in *unit light truck sales* is much larger and significant even at the 2.5% level at most horizons, whereas the smaller decline in *unit heavy truck sales* is significant only at the 16% level. The latter two responses approach -1.05% and -0.86%, respectively. This evidence strengthens the case for the operating cost channel. Assuming that all producers of light trucks are equally affected by such a shock, a permanent shock associated with an event such as Hurricane Katrina would reduce the number of light trucks sold by about 11.2%, making this channel economically significant for U.S. companies such as Ford, GM and Chrysler, which devote between 35% and 80% of their production to trucks.[12]

3.2 How Do Investment Expenditures Respond to Higher Energy Prices?

As noted in Hamilton (2005), energy price shocks may be transmitted not only through cutbacks or shifts in consumer expenditures, but through similar adjustments in firms' investment expenditures. There are two main channels by which energy price shocks may affect nonresidential investment. One channel is that an increase in the price of energy raises the marginal cost of production. This cost channel depends on the cost share of energy. A second channel is through reduced demand for the firm's output, as consumer expenditures fall in response to rising energy prices. For example, Herrera (2007) studies a linear-quadratic inventory model that links shifts in consumer demand in response to energy price shocks to real economic activity. There is also a direct link from reduced demand to cutbacks in nonresidential investment in equipment and structures (see Edelstein and Kilian 2007b).

The response of nonresidential fixed investment need not be symmetric in energy price changes. For example, changes in energy prices are thought to create uncertainty about future energy prices, causing firms to postpone irreversible investment decisions (see Bernanke 1983 and Pindyck 1991). This uncertainty affect has implications for both supply-side and demand-side accounts of the transmission of energy-price shocks. Specifically, firms may respond to uncertainty about future production costs or to uncertainty about future sales and revenue. In either case, when energy prices rise,

[12] This information was obtained from unit sales and production data on the company websites.

the uncertainty effect will reinforce the decline in firms' investment expenditures due to reduced consumer demand and higher energy costs. When energy prices fall, in contrast, the uncertainty effect counteracts the increase in investment expenditures driven by lower costs and increased consumer demand, dampening the increase in investment spending. As in the case of consumption, however, there is no compelling evidence of asymmetries in the responses of investment expenditures to energy price shocks with the exception of some subcomponents of equipment investment. We hence will defer the discussion of asymmetries until section 3.3.

Table 3 summarizes estimates of the elasticity of investment expenditures after one year. All estimates are based on quarterly investment aggregates reported by the BEA for 1970.II-2006.IV. The analysis again is based on bivariate VAR models as discussed in section 2. The estimated VAR models include quarterly dummies for 1986 for reasons discussed in section 3.3. In reporting the results it is useful to distinguish investment related to mining activities (which in the U.S. are largely accounted for by mining for crude oil, coal and natural gas) from other investment. Whereas mining-related investment will tend to be stimulated by higher energy prices, other investment expenditures will tend to decline, if they respond at all. The net effect is indeterminate and depends on the weight of each component.

Table 3 shows that the energy price elasticity of total nonresidential investment expenditures is only -0.16 and statistically insignificant. Among the subcomponents, only the response of *mining structures* and *mining and oil field machinery* is large and statistically significant. The positive elasticity of 0.09 for *structures* is driven largely by the aggregation of *mining structures* with other structures. Excluding mining, the response is virtually zero. The response of *equipment* has the expected negative sign and is somewhat larger with an elasticity of -0.30, but is statistically insignificant. It can be shown that this estimate is mainly driven by transportation equipment, consistent with the results for durables consumption. Overall, there is no evidence that energy prices exert a large effect on total nonresidential investment expenditures. For comparison, the last row of Table 3 includes the elasticity of residential investment expenditures on structures. The estimate of -1.02 is considerably larger than for comparable nonresidential structures and statistically significant. We conclude that energy price shocks make themselves felt primarily through reduced demand for cars and new houses.

3.3 Are the Responses Asymmetric in Energy Price Increases and Decreases?

A large literature has stressed the potential importance of asymmetric responses of U.S. macroeconomic aggregates to energy price shocks (see, e.g., Balke, Brown and Yücel 2002; Davis and Haltiwanger 2001; Ferderer 1996; Hooker 1996a,b, 2002; Hamilton 1996, 2003; Huntington 1998; Lee and Ni 2002; Lee, Ni and Ratti 1995; Mork 1989). It is important to distinguish the idea of using a nonlinear transformation of the energy price (such as the LC or NC specification), which by itself does not preclude symmetry in the responses to energy price increases and decreases, from the question of whether the responses to a change in energy prices are asymmetric in the sign of the energy price change. It has been common in empirical work to impose the restriction that the economy does not respond at all to energy price decreases. This view is implicit in the use of the net oil price increase measure, for example, which assigns zero weight to net declines in the price of energy. This specification involves an extreme form of asymmetry. More generally, asymmetries could involve a weaker response to energy price decreases than to energy price increases.

Interest in asymmetries dates back to the late 1980s, when it became apparent that the sharp decline in crude oil prices in 1986 was not followed by a major economic expansion. Given the presumption that the equally sharp increase in crude oil prices in 1979 caused a major economic decline, the absence of a major economic expansion in 1986 seemed to provide iron-clad evidence of the need to allow for asymmetries in the response to crude oil price increases and decreases. This observation has led a number of researchers to incorporate asymmetries into models of the transmission of energy price shocks to the economy such as the uncertainty effect on irreversible investment decisions described in Bernanke (1983) and Pindyck (1991) or the reallocation effect of Hamilton (1988).

Until recently, however, the hypothesis of symmetric responses has not been formally tested.[13] Edelstein and Kilian (2007a, 2007b) perform formal statistical tests for the presence of asymmetries in the response of nonresidential fixed investment to energy price shocks of different sign, but the same magnitude. They allow for a variety of different measures of energy price shocks including percent changes in energy prices, large percent changes and net percent changes. Symmetry in this

[13] One common approach has been to include oil price increases and decreases as separate variables in a single-equation model for output growth, and to perform a Wald test for the equality of the coefficients on the lags of these variables (see, e.g., Mork 1989, Dotsey and Reid 1992, Hooker 1996, and Hooker 2002). A drawback of this approach is that this test only alerts us to differences in the slope coefficients, whereas we are really interested in whether the impulse responses to positive and negative energy price shocks are different. Another common approach has been simply to inspect the point estimates of the impulse response functions without formal testing (see, e.g., Davis and Haltiwanger 2001). That comparison, however, tells us nothing about the statistical significance of the difference. Nor do tests for pointwise statistical significance of the differences in the impulse response function constitute a formal test of the symmetry assumption.

context means that the sum of the impulse response function to energy price increases and of the impulse response function to energy price decreases is jointly equal to zero at all horizons. Based on a comprehensive set of monthly real consumption aggregates and disaggregates, they are unable to reject the null hypothesis of symmetry for any of the major expenditure aggregates.

While the evidence against asymmetries in real consumption responses is subject to considerable sampling uncertainty in some cases, the tests are not without statistical power, as indicated by the rejections of symmetry reported in Edelstein and Kilian (2007b). Moreover, the estimated responses of expectations data from the Michigan Survey of Consumers to the same purchasing power shocks tend to be fairly symmetric. Together, this evidence suggests that the reallocation effect, the uncertainty effect, and any asymmetry associated with the precautionary saving effect discussed in section 3.1.1 are not a dominant feature of the real consumption data. It is of particular interest that there is no evidence of a reallocation effect, despite the evidence of shifts in expenditure patterns documented in section 3.1.5. A possible explanation is the relatively small share of the U.S. automobile industry in employment.

This interpretation is also consistent with the lack of statistical evidence against the symmetry hypothesis in the responses of the U.S. unemployment rate. Notwithstanding some important methodological differences, these results are qualitatively consistent with the plant-level net employment change responses estimated in Davis and Haltiwanger (2001). Both studies show asymmetric point estimates. The chief difference is that Davis and Haltiwanger did not investigate whether these asymmetries are statistically significant, whereas Edelstein and Kilian (2007a) show that they are not.

An immediate implication of this result is there should have been a boom in consumption in 1986. This implication is largely consistent with the data, further supporting the symmetry hypothesis. In 1979, purchasing power declined by 1.69% due to energy price increases, whereas in 1986 purchasing power increased by 1.43% due to energy price decreases. Thus one would expect the effect on real consumption to be roughly symmetric. The symmetric VAR model implies that rising energy prices (all else equal) lowered real consumption by -1.92% in 1979, and raised it by +2.02% in 1986, making these effects nearly symmetric. By comparison, actual real consumption growth in 1979 was -2.20% relative to its mean, whereas in 1986 it was +1.44%. Thus, energy prices alone are capable of explaining a substantial part of observed real consumption growth in 1979 and 1986.

Nevertheless, the perception of an asymmetry in economic performance between 1979 and 1986 is correct. The observed behavior of real consumption growth in 1979 and 1986 contrasts sharply with that of real GDP growth. Real GDP growth was -1.81% relative to its mean in 1979 *and* -0.31% relative to its mean in 1986. Thus, the asymmetry alluded to earlier does exist in real GDP growth, but is not reflected in real consumption growth. A comparison of the 1979 and 1986 growth rates of real GDP and its components reveals that the asymmetry originates in nonresidential investment in equipment and structures. In 1979, these investment expenditures grew by -2.80% and +7.54% relative to the mean, respectively, whereas in 1986 they grew by -4.65% and -16.35%. The behavior of firms' investment expenditures in 1986 contrasts sharply with that of private residential fixed investment and of durables consumption.

Edelstein and Kilian (2007a) suggest that an exogenous drop in nonresidential fixed investment expenditures in 1986 was mainly responsible for the low rate of real GDP growth in 1986. A natural candidate for such an exogenous shift in investment expenditures is the 1986 Tax Reform Act, which sharply raised the effective tax rate for many corporations by severely curtailing deductions for capital expenditures and by eliminating the investment tax credit. For most types of equipment, the repeal of the investment tax credit, which became effective in the first quarter of 1986, amounted to the elimination of a 10% subsidy on investment. This fact helps explain the sharp drop in nonresidential fixed investment expenditures on equipment in 1986.[14] The even larger drop in nonresidential fixed investment in structures is unlikely to be explained by the repeal of the investment tax credit alone because it was offset by other changes in the tax code and because business investment dropped even in sectors that were not subject to the investment tax credit prior to 1986 (see, e.g., Auerbach 1987).

Further disaggregation of the BEA data reveals that the decline in nonresidential investment in structures is concentrated in two components. The first component is *commercial structures (including office space)* and *manufacturing structures*, which account for 21% and 6% of total real nonresidential investment in structures, respectively. A likely explanation is that the elimination of real estate tax shelters as part of the 1986 Tax Reform Act contributed to the observed 17% drop (relative to the average growth rate) in these two components in 1986 (see *Survey of Current Business* 1987, p. 4). The second component is nonresidential investment in *mining exploration, shafts and wells*. That component accounts for about 11% of all nonresidential investment in structures and mainly comprises investments in the petroleum, natural gas and coal mining industry. In fact, one third of the total decline in real business investment in structures can be accounted for by the dramatic 65% drop in this component in 1986 below the average growth rate. While one would expect some decline in investment in these industries in response to falling energy prices, this particular drop was swifter and larger than the corresponding increase in investment in the domestic petroleum and natural gas industry observed after 1979. This asymmetric reaction is consistent with the view that the market treated the breakdown of OPEC in late 1985 as an exogenous shock and responded more

[14] For details of the timing of the 1986 Tax Reform Act see Wakefield (1987).

strongly than it would have based on the fall of energy prices alone. The evidence is also consistent with the view that there were limited investment opportunities in the domestic petroleum, natural gas and coal mining industry after 1979, making the response of this component of real GDP growth inherently asymmetric (but in the opposite direction of the asymmetries previously discussed in the literature on oil and the macroeconomy).

Hence, there are good reasons for the existence of an asymmetry between 1979 and 1986 in the real GDP growth data. The Tax Reform Act of 1986 and the unprecedented fall in investment in the oil and gas industry also help explain why real consumption did not grow quite as much in 1986 as predicted by a linear econometric model on the basis of falling energy prices alone and why unemployment remained higher than it would have been otherwise. Ignoring this exogenous shift in nonresidential fixed investment, given the short sample, may bias VAR estimates of the responses of nonresidential investment and cause them to look asymmetric in small samples, even when the true responses are not. Indeed, the responses of nonresidential investment in equipment and structures to energy price shocks estimated on the 1970-2006 period appear asymmetric, and the symmetry null can be rejected in several cases. The nature of the asymmetries, however, in many cases departs sharply from the predictions of commonly used economic models of the transmission of energy price shocks.

These spurious and economically implausible asymmetry results can be traced to two problems. One problem is the exogenous shift in investment in 1986 and can be addressed by introducing quarterly intercept dummies for 1986. The other problem is that the aggregation of mining-related investment by domestic producers of coal, natural gas and crude oil with other forms of investment expenditures may generate perverse asymmetries (in the sense that the observed asymmetries are seemingly at odds with the predictions of economic models). Excluding investment in mining-related activities and including the 1986 dummies removes all of the evidence of asymmetries in structures. Similarly, there is no compelling evidence of asymmetries in the responses of aggregate nonresidential investment in equipment. There is a marginal rejection of symmetry at the 10% level for only one of the components of equipment investment. We conclude that there is no compelling evidence of asymmetries for either consumer expenditures or investment expenditures, lending credence to the symmetric response estimates presented in sections 3.1 and 3.2.

3.4 Oil Price Shocks and Monetary Policy

Starting with Bernanke, Gertler and Watson (1997), there has been interest in the extent to which the response of the U.S. economy to crude oil price shocks is driven by the endogenous response of monetary policy, as opposed to the direct effect of oil price shocks on the economy. Motivating this line of work was the perception that the magnitude of the recessions following major oil price increases is too large to be caused by rising oil prices alone. The alternative explanation was proposed that the Federal Reserve chooses to raise interest rates in anticipation of the higher inflation expected as a result of oil price increases, thereby aggravating the relatively benign economic downturn normally expected in response to higher oil prices. This perception has spurred interest in VAR models of the effects of oil price shocks that incorporate monetary policy reaction functions. The inclusion of a monetary policy reaction function results in much larger VAR systems than the models discussed so far, exemplified by the 7-variable system specified in Bernanke et al. (1997).

Based on their VAR model estimates, Bernanke et al. document the cumulative effect of oil price increases on real output relative to trend, including both the direct effect of higher oil prices and the effect associated with the monetary policy response to higher oil prices. The historical decompositions in Bernanke et al. suggest that the oil price shock of 1973/74 overall did not contribute much to the sharp decline in real output relative to trend in 1974/75, consistent with evidence that the Fed was tightening monetary policy well before the oil price shock and did so in response to rising industrial commodity prices which were viewed as an indication of rising future inflation (also see Barsky and Kilian 2002). During the 1979-1983 episode, the model predicts more of a decline in real output below trend than actually occurred from mid-1980 through late 1981, but it does not explain well the sharp decline in real output relative to trend starting in mid-1981, which seems due to an autonomous tightening of monetary policy under Paul Volcker. Similarly, oil prices are only a contributing factor for the decline in real output relative to trend in 1991/92. Subsequently, Hamilton and Herrera (2004) have observed that the magnitude of the effect of oil price shocks on real output is sensitive to the lag order. Allowing for additional lags results in somewhat larger declines in real output after the two oil price shocks of the 1970s, but only after the 1990/91 shock does the price of oil explain the bulk of the decline in real output below trend.[15]

Having documented the effects of these three oil price shocks, Bernanke et al. propose a thought experiment in which the Fed instead pursues a policy of holding the Fed Funds rate constant in the relevant sample periods. They show that the resulting path of real output would have been considerably less recessionary and attribute the difference in outcomes to the endogenous policy response to higher oil prices. A direct implication of the Bernanke et al. analysis is that the recessionary consequences of an oil price shock in principle could have been avoided at the cost of higher inflation by simply holding constant the Fed Funds rate. This implication has been challenged

[15] These results are difficult to compare to other results in the literature because real output consists of monthly real GDP interpolated based on industrial production data and subsequently detrended by a cubic spline.

by Hamilton and Herrera (2004) who suggest that the monetary expansion required to stabilize the Fed Funds rate and to prevent a decline in real output below trend would be implausibly large. Such an expansion certainly would involve a change in interest rates outside of historical experience and hence would make the analysis subject to the Lucas critique, casting doubt on the validity of the analysis (see Bernanke, Gertler and Watson 2004).

More generally, coming up with a suitable counterfactual for assessing the quantitative importance of the monetary policy response to oil price increases remains a challenge. Bernanke et al. consider the counterfactual of a constant Fed Funds rate following an oil price shock. This counterfactual by construction rules out not only a monetary policy response to oil price shocks, but it prevents the Fed from responding to *any* shocks in the economy, whether related to oil prices or not. While this specific counterfactual potentially sheds light on the overall importance of systematic monetary policy, it does not allow us to assess the extent to which the Fed had amplified the recessionary impact of an oil price shock. For example, it precludes a response of the Federal Reserve to commodity price increases. This assumption is problematic because, as Bernanke et al. point out, the Fed's endogenous response to commodity price increases in 1973 (starting well before the oil price increase of late 1973) may explain a substantial part of the difference between the counterfactual and the policy outcome. Likewise, the oil price increases in 1979 coincided with rising commodity prices, making it difficult to attribute the policy response to any one price shock. Thus, the important question of how much the Fed's endogenous response to higher oil prices contributed to the subsequent economic declines still remains unresolved.[16]

Apart from sampling error, the real output responses to oil price increases obtained in this framework ought to be identical to the results of a bivariate VAR model in net oil price increases and (detrended) real output. While there is no advantage to using the framework of Bernanke et al. to learn about the response of real output compared with the bivariate framework of section 2, there may be disadvantages to the extent that higher-dimensional VAR models involve additional identifying assumptions that may prove incorrect. Indeed, there is no agreement on how to identify the remaining shocks of the VAR system. For example, Herrera and Pesavento (2007) under an alternative set of assumptions obtain much smaller estimates of the effects of endogenous monetary policy responses. Moreover, larger VAR systems require more data for estimation, and in practice must be estimated on data from the pre-1973 period, which is characterized by very different institutional arrangements in the global crude oil market, as discussed in Hamilton (1983) and Barsky and Kilian

[16] More recent studies have aimed to address this question in the context of theoretical macroeconomic models (see. e.g., Leduc and Sill 2004, Carlstrom and Fuerst 2006, Kormilitsina 2007, Dhawan and Jeske 2007a). While the answer depends to a large extent on the definition of the counterfactual and on the modeling assumptions, the emerging consensus is that monetary policy may be partially responsible for the observed declines in real GDP.

(2002).

4 Has the U.S. Economy Become Less Responsive to Energy Price Shocks?

It has been widely observed that energy price shocks do not appear to affect the U.S. economy as much as they used to (see, e.g., Herrera and Pesavento 2007).[17] This observation can be substantiated by comparing responses of consumption aggregates estimated on the first half (1970.2-1987.12) and the second half (1988.1-2006.7) of the sample used in sections 3.1.3 and 3.1.5. After normalizing the scale of the impulses to be the same across the two subsamples to make the magnitudes of the impulse responses comparable, striking differences emerge. Evaluated at the average energy share for the full sample, the response of total real consumption to an unanticipated 1% energy price increase drops from -0.30% after 18 months in the first half of the sample to -0.08% in the second half. The corresponding decline for durables is from -0.84% to -0.24%. Vehicles consumption declines from -1.31% in the first half to -0.49% in the second half of the sample. The decline in durables consumption excluding vehicles shrinks from -0.44% in the first half of the sample to -0.01% in the second half. The response of nondurables shrinks from -0.29% to -0.02% and that of services from -0.18% to -0.07%. A similar reduction occurs in the response of real residential fixed investment (not shown). The response drops from -4.7% to -1.3%. Finally, the rise in unemployment associated with an unanticipated purchasing power loss drops from 1.53% to 0.36%.

There are several possible explanations for the declining importance of energy price shocks. One conjecture is that this result is related to the declining share of energy in consumption in the late 1980s and 1990s. Since our results are based on innovations to purchasing power changes rather than innovations to unweighted the energy price changes, they already control for changes in the expenditure share of energy, eliminating this explanation. A second conjecture is that the variability of purchasing power shocks may have declined in the second half of the sample. Our analysis shows that in fact the variability of both total changes and linearly unpredictable changes in purchasing power has *increased* in the second half of the sample. The innovation standard deviation increased from 0.08 to 0.11. The average size of positive innovations increased from 0.056 to 0.076, and the average size of negative innovations increased from -0.049 to -0.073. Moreover, both the maximum and the minimum of the innovations increased.

[17] A weakening of the statistical relationship between oil prices and the U.S. economy in the mid-1980s has been noted as early as Hooker (1996b, p. 222) and Davis and Haltiwanger (2001, p. 482). There also is a widely held view among policymakers that the surges in oil prices in the 1970s and 1980s had much more prononced economic effects than the more recent increases (see, e.g., Bernanke 2004).

A third and more plausible explanation is that the structure of the U.S. automobile industry has changed. In the 1970s, U.S. auto manufacturers were simply not producing any small, energy-efficient cars, leaving consumers no choice but to buy small cars from abroad. Thus, the U.S. auto industry was hit particularly hard by rising energy prices and falling demand for large cars (see, e.g., Bresnahan and Ramey 1993, Davis and Haltiwanger 2001). In contrast, by the late 1980s and 1990s the differences between domestic and foreign auto producers had been greatly reduced, as domestic auto manufacturers offered small and energy efficient cars of their own, while foreign manufacturers were beginning to branch out into the market for jeeps, SUVs, vans and pickup trucks. Thus, the U.S. auto industry became relatively less vulnerable to energy price increases than in the 1970s.

This point is illustrated by comparing the responses of new domestic and foreign automobiles in the two subsamples. Whereas in the first subsample expenditures on new domestic automobiles in response to a 1% increase in energy prices drop by -2.8% after two months and by -1.7% after 18 months, in the second half the short-run response drops to -0.7% and the long-run response to -0.3%. The strongly significant short-run decline in the first sample is only marginally significant in the second sample. In contrast, in the first half of the sample, after one month expenditures on new foreign automobiles rise significantly by 1.3%, followed by an insignificant decline of -0.99% after five months and a long-run response of -0.3%. In the second half of the sample, the initial increase in the response has become small and insignificant, the decline after 5 months has shrunk to -0.4% and the long-run response to -0.1%. While it is still true that the consumption of new domestic autos is more responsive to energy price shocks than the consumption of new foreign autos, the differences are much smaller than they used to be.

There is also a fourth and complementary explanation. As the U.S. automobile industry restructured itself after the energy price increases of the 1970s, the share of domestically produced automobiles in total U.S. real expenditures on new cars declined (from 88% in 1970 to 60% in 1988 and 57% in 2006), as did the employment share of the industry (from a peak of 1.3% in 1973 to 0.9% in 1988 and 2005).[18] Thus, the relative importance of the auto industry for the U.S. economy and the potential for spillovers from the automobile industry to other sectors has declined relative to the 1970s, further reducing the precautionary savings effect.

Yet another possibility taken up in the next section is that the nature of the energy price shocks has evolved and that recent energy price shocks have been qualitatively different from earlier shocks. It will be shown that an energy price increase driven by strong global demand for industrial commodities (including crude oil), for example, may have far less adverse consequences for U.S. real output

[18] See http://bea.gov/bea. There are no data on the share of the automobile industry in real value added prior to 1987. The current share of 1.1% is only slightly lower than in 1987.

than the same energy price increase driven by adverse global oil supply shocks or by expectations-driven shocks to the precautionary demand for oil. This distinction matters because positive shocks to global demand at the same time stimulate U.S. real output growth, as they raise the price of oil, yielding a positive stimulus on balance, at least initially. Thus, the origin of energy price increases matters.

5 Disentangling Demand and Supply Shocks in Energy Markets

5.1 A Joint Model of the Global Crude Oil and the U.S. Gasoline Markets

There is an important distinction between the price of gasoline and other motor fuels in the U.S. and the price of crude oil in global markets that is often ignored in discussions of the impact of higher energy prices. This section makes explicit the relationship between demand and supply shocks in these two markets. Building on a structural VAR model of the global crude oil market proposed in Kilian (2007c), I explore the implications of a joint VAR model of the global market for crude oil and the U.S. market for motor gasoline. The data frequency is monthly and the model includes 12 lags. The model includes (in the order listed) the percent change of world production of crude oil, a suitably detrended index of global real economic activity as it relates to industrial commodity markets (which may be thought of as a measure of the business cycle in global industrial commodity markets), the real price of imported crude oil, the real price of gasoline and other motor fuels in the U.S. and the percent growth rate of the quantity of gasoline and motor fuels consumed in the U.S. The sample period is 1973.2-2006.12.[19] I postulate that these variables are driven by five structural shocks: (1) crude oil supply shocks (*oil supply shocks*); (2) shocks to the demand for *all* industrial commodities in global markets (*aggregate demand shocks*), which may be thought of as innovations to the business cycle in global industrial commodity markets; (3) demand shocks that are specific to the global crude oil market (*oil-market specific demand shocks*); (4) shocks to the supply of gasoline in the U.S. (exemplified by *refinery shocks*); (5) shocks to the U.S. demand for gasoline (*gas demand shocks*). Whereas the aggregate demand shock is designed to capture shifts in the demand for all industrial commodities (including crude oil) driven by the global business cycle as well as structural shifts in the demand for industrial commodities such as the emergence of industrialized economies in Asia, the oil-market specific demand shock is designed to capture shifts in the price of oil driven

[19] The date are obtained from the Department of Energy and the BEA. Data availability constraints preclude the extension of the sample beyond 2006.12. For further discussion of the data the reader is referred to Kilian (2007c).

by higher precautionary demand associated with fears about future oil supply shortfalls.[20]

Identifying Assumptions in the 5-Variable VAR Model The identifying assumptions for the 5-variable VAR model are: (1) world crude oil production does not respond within the month to demand shocks in the crude oil market, nor does world crude oil production respond to demand or supply shocks in the U.S. gasoline market within the same month; (2) oil-market specific demand shocks do not affect, within the month, the business cycle in global industrial commodity markets; (3) while shocks to the supply of or demand for crude oil may have an immediate effect on gasoline prices, the demand for crude oil remains unaffected within the same month by demand and supply specific to the U.S. gasoline market; (4) shocks to the demand for gasoline that are orthogonal to shocks to the demand for crude oil do not affect the price of gasoline within the same month; (5) shocks to the supply of gasoline (such as refineries shutting down due to accidents or changes in the regulatory environment, as discussed in Muehlegger 2006), however, may affect the price of gasoline within the same month. These identifying assumptions imply a recursive structure for the innovations in the structural VAR model. Figure 4 shows the impulse response estimates for a horizon of up to 24 months. The impulse response confidence intervals have been constructed using a recursive-design wild bootstrap (see Gonçalves and Kilian 2004). The qualitative pattern of the response estimates conforms with basic economic theory.

Assumptions (1) and (2) represent a partial equilibrium model of the global crude oil market. The model postulates that the stochastic supply curve for crude oil is vertical in the short run and does not respond to demand shifts within the month. This assumption is reasonable because supply decisions are made based on expectations of medium-term demand. Since changing supply is costly, and innovations to demand will have a negligible effect on expected trend growth in demand, supply will only respond to demand shocks with a delay. The supply curve may be shifted by production disruptions in the Middle East and other exogenous events. The short-run demand curve is downward sloping. It is being shifted by innovations to global aggregate demand and by innovations to oil-specific demand. Thus, all three shocks are allowed to affect the real price of oil within the month. The real oil price innovation is simply a weighted average of the crude oil demand and crude oil supply innovations.

Assumptions (4) and (5) imply that all innovations in gasoline prices not associated with changes in crude oil prices are driven by gasoline supply shocks. The gasoline market is modeled as follows:

[20] In the context of a theoretical model of the spot and futures market for crude oil, Alquist and Kilian (2007) show that measures of the percent spread of the oil futures price over the current spot price of oil may also be used to measure the precautionary demand component of the real spot price. Their empirical work shows that such measures are highly correlated with the fluctuations in the spot price of oil driven by the precautionary demand shocks as identified by VAR models of the type discussed here.

First, refineries process crude oil to produce gasoline (and other refined products). Thus, crude oil price increases represent a supply shock for refineries and the model postulates that such shocks are being passed on to gasoline prices within a given month. Second, the model stipulates that in the aggregate the gasoline price is determined by the U.S. refineries and that owners of U.S. gas stations in essence charge that price (possibly with a fixed markup). The owner of the gas station will sell gasoline at that price regardless of the ups and downs of demand for gasoline. The spare capacity of the gas station's underground tank allows the owner to absorb such fluctuations, as long as fluctuations in demand are moderate. In short, the model embodies the assumption that the supply of gasoline is perfectly elastic in the short run, conditional on the price set by the refinery. Third, the demand curve for gasoline is downward sloping, so shifts in the demand curve will not affect the price within the month. This is assumption (5). Since changes in the demand for gasoline depend on the composition of the fleet of vehicles, on the average distance of commuters from their job, and other factors that evolve only slowly, we would expect demand for gasoline to evolve smoothly over time, except for a noise component. As refineries set the price of gasoline, they keep in mind expected demand for gasoline. In essence, they project out the trend growth of gasoline demand. To the extent that owners of gas stations report back to refineries that they expended the contents of their underground tank more quickly than anticipated, refineries may infer that demand has shifted and will adjust the gasoline price. Such signals are noisy, however, and refineries will be slow to take action until a trend emerges. As is well known in econometrics, our ability to detect such a "trend change" will depend on the time span of data, rather than the frequency with which gas stations report in, which makes the delay assumption (5) economically reasonable. On the other hand, refinery supply shocks (such as a refinery fire or refineries shutting down because of a hurricane) will be taken into account immediately by refineries and will be reflected in a higher gasoline price within the month. This fact is embodied in assumption (4). An event such as a refinery fire would amount to an upward shift of the horizontal supply curve.

Finally, assumption (3) postulates that the price of crude oil is predetermined with respect to gasoline prices. In other words, the global crude oil market can be treated as block recursive with respect to the U.S. gasoline market within the month. This assumption amounts to stipulating that U.S. gasoline demand shocks and U.S. gasoline supply shocks will not affect the world crude oil price within the same month, but only with a delay of at least a month. In the case of unanticipated increases in U.S. gasoline demand I already made the case that such shocks will not be reflected in higher gasoline prices within the month, and hence by construction do not cause higher refinery output, higher U.S. demand for crude oil and higher global crude oil prices within the month. In the case of U.S. gasoline supply shocks such as a refinery shutdown (which would reduce global demand for crude oil), this assumption may be defended by observing that world crude oil supply only responds with a lag to crude oil demand shocks, as discussed earlier. Likewise, it is not plausible that a U.S. gasoline supply disruption through its effect on U.S. real consumption would affect global demand for industrial commodities within the first month because the global aggregate demand shock is predetermined with respect to U.S. real consumption, as discussed earlier. A U.S. gasoline supply disruption on impact will have a direct effect on the global demand for crude oil specifically, but that direct effect is likely to be negligible, given the magnitudes of the refinery outages we have observed historically. More important is the indirect effect working through a reduction in U.S. real consumption, but that reduction will have negligible effects on the estimate of the trend component of oil demand and hence negligible effects on precautionary demand.

Results The first row of Figure 4 shows that an unanticipated permanent reduction in world crude oil supplies is partially offset by subsequent increases in crude oil production. The global real price of crude oil rises temporarily, as does the real price of gasoline in the U.S., but the latter response is much weaker. U.S. gasoline consumption drops, as the price of gasoline rises.[21] An unanticipated increase in global demand for industrial commodities, as shown in the second row, causes a temporary increase in world crude oil production and a persistent increase in the real price of crude oil. The response of the real price of crude oil reaches its maximum only with a delay. There is also a persistent, but much smaller increase in the real price of U.S. gasoline. U.S. real consumption of gasoline declines.

The third row focuses on the responses to oil-market specific demand shocks. Such shocks typically arise from an increase in the precautionary demand for crude oil driven by concerns about future crude oil supply shortfalls (see Kilian 2007c; Alquist and Kilian 2007). They reflect the market's expectations of both future demand and future supply developments, and may reflect concerns that never materialize. Large shifts in oil-market specific demand occurred, for example, in 1979, when the Iranian Revolution, the Iranian hostage crisis and the Soviet invasion of Afghanistan coincided with strong global demand for oil. Large shifts also occurred following the collapse of OPEC in late 1985 and after the invasion of Kuwait in 1990. Figure 4 shows that an unanticipated increase in the precautionary demand for crude oil would be associated with an immediate and sharp increase in both crude oil and gasoline prices. Again the response of gasoline prices is smaller. There is an immediate and persistent fall of real gasoline consumption, but no increase in world crude oil

[21] The model does not distinguish between crude oil supply shocks driven by exogenous political events in the Middle East, as discussed in Kilian (2007a,b) and other exogenous shocks to crude oil production. This distinction could be easily incorporated into the VAR framework above, but is largely immaterial in the present context, as shown in the working paper version of Kilian (2007c).

production.

An unanticipated reduction in the U.S. capacity to refine crude oil, as shown in the fourth row, is associated with a sharp increase in U.S. gasoline prices, but a temporary fall in world crude oil prices. The decline in real gasoline use is immediate and fairly persistent. There is also evidence that world crude oil production falls with some delay. Finally, an unanticipated increase in U.S. demand for gasoline, as shown in the last row, will cause higher gasoline prices, will have no significant effects on the real price of crude oil and will have little effect on global crude oil production.

In summary, Figure 4 demonstrates that demand and supply shocks in the global crude oil market and in the U.S. gasoline market have distinctly different effects on global crude oil production, gasoline consumption, and the prices of crude oil and gasoline, making it important to differentiate between price shocks driven by one or another of these demand and supply shocks.

5.2 What Has Been Behind the Recent Surge in U.S. Motor Gasoline Prices?

The model of section 5.1 may also be used to shed light on the behavior of U.S. gasoline prices in recent years. The average price of a gallon of regular grade gasoline increased from $1.11 in January 2002 to $2.98 in July 2006.[22] Figure 5 identifies the cumulative effect of each of the structural demand and supply shocks identified in the previous subsection on the real price of gasoline.

The first row of Figure 5 shows that overall crude oil supply shocks have had a negligible effect on gasoline prices and, if anything, have lowered gas prices slightly since early 2005. The bulk of the increase in gas prices has been associated with steady pressure from increasing global demand for crude oil, along with other industrial commodities, as shown in the second row. Since early 2003, strong demand for industrial commodities all by itself has been pushing gas prices far above their average value.

In addition, there has been upward pressure on real gasoline prices from shifts in precautionary demand since late 2005, as shown in the third row of Figure 5. The initial rise occurred following Hurricanes Rita and Katrina (possibly reflecting a misinterpretation of the effects of these events on the crude oil market). The subsequent even larger buildup in 2006 is more likely linked to concerns about Iran and Iraq and the continued strength of the world economy. In contrast, the effect of the Iraq War in early 2003 on U.S. gasoline prices was fairly small and short-lived.

An important test of the plausibility of the identifying assumptions is the behavior of gasoline and crude oil prices following Hurricanes Rita and Katrina in late 2005. As discussed earlier, the

[22] See http://tonto.eia.doe.gov.

primary effect of this exogenous event was not the reduction in U.S. crude oil production (which was negligible on a world scale), but the reduction of crude oil refining capacity in the Gulf of Mexico. Given that other U.S. refineries were already operating close to capacity at the time, this event constituted a major unanticipated reduction of the supply of gasoline in the U.S., which would be expected to raise the price of gasoline sharply. The fourth row of Figure 5 indeed shows a sharp increase of U.S. gasoline prices driven by adverse refinery shocks in late 2005. Only a year later, the price seems to have stabilized again. In contrast, the effect of exogenous shocks to gasoline demand, as shown in the last row, has been negligible throughout this period.

We conclude that the recent build-up in gasoline prices consists of three components whose relative contribution to the real price of gas has varied over time: (1) strong and persistent demand in global industrial commodity markets (consistent with a strong global expansion of advanced economies and with the integration of emerging economies in the global economy); starting in late 2005, the upward pressure on gasoline prices was also aided by (2) precautionary demand shocks specific to the oil market and (3) adverse supply shocks in the U.S. refining industry. There is no evidence that the recent build-up of gasoline prices has been associated with production decisions by OPEC or other crude oil supply shocks.[23]

This example complements related evidence in Kilian (2007c) on the relative contribution of demand and supply shocks to the real price of crude oil since 1978. This evidence suggests that efforts to link crude oil price increases to crude oil production shortfalls alone are doomed to failure, given the overriding importance of shocks to the demand for crude oil not just in the most recent period, but also during earlier oil price shock episodes. This, of course, does not preclude that crude oil production shortfalls may play a more important role in the future. If there is a shortfall of crude oil production in some country, much depends on the duration of this shortfall and on the ability of other oil-producing countries to offset the shortfall. The fact that, in the past, global oil production has tended to recover or even to increase following oil supply shocks is no guarantee that additional supplies will be forthcoming when needed in the future.

[23] It is important to stress that this analysis is based on a linear symmetric model. A number of studies have documented evidence of asymmetry in the response of the price of gasoline to crude oil price increases and crude oil price decreases using daily, weekly or bi-weekly data, although symmetry does not appear to have been tested formally (see, e.g., Borenstein, Cameron and Gilbert 1997). While the evidence is by no means clearcut, it does appear that gas stations are quicker to raise gasoline prices in response to crude oil price increases than they are to lower gasoline prices in response to crude oil price decreases. While the VAR model used in this section is much richer in detail than those studies, it does not allow for such asymmetries. Further investigation of the evidence of asymmetries at the monthly frequency is an important area of research, all the more so, since Davis and Hamilton (2004) and Douglas and Herrera (2007) in closely related work have presented formal evidence in favor of asymmetries in the relationship between daily gasoline prices as quoted on NYMEX and daily wholesale gasoline prices.

5.3 Reexamining the Assumption of Predetermined Gasoline Prices

The analysis in section 3.1 was based on the conventional assumption that innovations to retail energy prices such as gasoline prices can be treated as predetermined with respect to U.S. consumption growth. This type of assumption is standard in the literature, but it is useful to reexamine its rationale based on the joint model of the global crude oil market and the U.S. retail gasoline market discussed in section 5.1.

5.3.1 On the Predeterminedness of Global Crude Oil Prices for U.S. Consumption Growth

Since the innovations to the real price of oil can be expressed as a weighed average of the three oil demand and oil supply shocks in the VAR model of section 5.1, if we can show that each of these shocks is predetermined with respect to real U.S. consumption growth, their weighted average will also be predetermined. For this purpose it is useful to assess the contemporaneous correlation of each of the three demand and supply shocks with innovations to U.S. real personal consumption growth. The latter innovations may be constructed from an autoregressive model for monthly consumption growth. Based on the AR(6) specification, the correlation of the latter innovations with oil supply shocks is 0.03, the corresponding correlation with aggregate demand shocks is 0.00 and that with precautionary demand shocks is -0.04. Similar results are obtained with other lag order choices.

While the identifying assumption of predeterminedness is not testable, these empirical correlations in conjunction with other information help assess the plausibility of the assumption of predetermined oil demand and oil supply shocks. First, consider the global crude oil supply shock. The case for treating this shock as predetermined with respect to real U.S. consumption growth follows immediately from the identifying assumption of section 5.1 that the global supply of crude oil does not respond within a given month to shocks to the global demand for crude oil. Thus, the empirical correlation of 0.03 can be interpreted as capturing the positive, but negligible instantaneous effect of unexpectedly high oil supplies on U.S. real consumption.

Next consider the global aggregate demand shock. If positive innovations to U.S. consumption growth stimulated global demand for industrial commodities, we would expect the correlation between innovations to U.S. consumption growth and global aggregate demand shocks to be positive. To the extent that causation also runs the other way from global aggregate demand shocks to innovations to U.S. consumption growth, that correlation should also be positive. Hence, the fact that the contemporaneous correlation of these innovations is essentially zero in the data, establishes that neither causal link can be quantitatively important at the monthly frequency. Hence, we may treat oil price increases driven by global aggregate demand shocks as predetermined with respect to U.S. consumption growth.

Finally, for the precautionary demand shock the argument is more complicated. Positive innovations to U.S. consumption growth may raise precautionary demand; yet unanticipated increases in precautionary demand will increase the price of oil and hence may lower U.S. consumption growth on impact. Since it is conceivable that these two effects may offset one another at least in part, we cannot use the low empirical correlation of these innovations to argue that there is no feedback from U.S. real consumption innovations to the precautionary demand shock. Nevertheless, a plausibility check suggests that innovations to monthly consumption growth in the U.S. have a negligible effect on global precautionary demand. Precautionary demand reflects concerns about the shortfall of future oil supply relative to expected demand for oil at medium-term or long-term horizons. An innovation to monthly U.S. consumption growth will have negligible effects on the estimate of that trend component under any conceivable model of trend growth in demand, and hence negligible effects on precautionary demand. Hence, we can be confident that the identifying assumption that oil price increases are predetermined for monthly U.S. consumption growth is economically plausible, consistent with the standard approach in the literature. This conclusion follows directly from the fact that oil price innovations are by construction a weighted average of innovations to oil supply, aggregate demand and precautionary demand for oil, and that each component is predetermined individually.[24]

5.3.2 On the Predeterminedness of U.S. Gasoline Prices for U.S. Consumption Growth

There are two channels through which an unanticipated increase in U.S. real consumption may affect the price of gasoline within the month. The first channel is based on the fact that refiners process imported crude oil, making the price of gasoline dependent on the global price of crude oil. As has already been shown, however, unanticipated increases in U.S. real consumption do not cause the global price of crude oil to increase within the month, so this first channel can be ignored. The second channel is that unanticipated increases in non-gasoline real U.S. consumption are likely to be associated with an increase in U.S. demand for gasoline as well. This channel can be ruled out based on the identifying assumption that the U.S. price of gasoline does not respond within the month to an unanticipated increase in the U.S. demand for gasoline. Hence, we know that the contemporaneous correlation between innovations to U.S. real gasoline prices and innovations to U.S. real consumption growth (or any of its components) reflects a causal link from real gasoline

[24] This chain of reasoning can be extended to other macroeconomic aggregates. For example, Kilian (2007c) discusses the asssumption that global crude oil prices are predetermined for quarterly U.S. real GDP growth along similar lines.

prices to real consumption. Analogous arguments also apply in regard to other retail energy prices such as heating oil, electricity, or natural gas.

5.3.3 Synthesis

The preceding subsections established that it is reasonable to treat innovations to energy prices as predetermined, provided the data are measured at a high enough frequency. With the usual caveats, the predeterminedness of energy prices permits the consistent estimation of the causal effect of an innovation in energy prices on real U.S. consumption (see Cooley and LeRoy 1985). The fact that the energy price innovation is driven by various demand and supply shocks does not alter this result because the weighed sum of these shocks is identical to the energy price innovation.

The additional insight provided by the model in section 5.1 is that we should not interpret these causal responses as reflecting only the direct effect of energy price innovations on real consumption. Suppose, for expository purposes, that the energy price innovation is driven entirely by an unanticipated increase in aggregate demand in global industrial commodity markets. Such as shock has both a direct stimulating effect on the U.S. economy and an indirect retarding effect operating through higher crude oil prices and hence higher retail energy prices. The empirical estimates of the responses to the retail energy price increase by construction will capture the net effect of this intervention on real U.S. consumption at each horizon. More generally, to the extent that changes in energy prices are driven by different demand and supply shocks in energy markets, the impulse response estimates will appropriately differ across subperiods, reflecting the average composition of the demand and supply shocks underlying changes in the price of energy over each sample period. The apparent decline in the responsiveness of the U.S. economy to energy price shocks, documented in section 4, indeed can be attributed in part to the changing composition of demand and supply shocks. As the sample size increases, the impulse responses will converge to the population expectation and will reflect the long-run average composition of demand and supply shocks in energy markets.

It may seem that an approach based on identifying demand and supply shocks in energy markets would in general be more informative than regressions exploiting the predeterminedness of energy prices. This is not the case. Rather the two empirical approaches are complementary. One reason is that various hypotheses arising from the theoretical literature have been stated explicitly in terms of responses to price innovations and would be untestable in the framework of section 5.1. The second reason is that consumer expenditures will depend on demand and supply shocks in all relevant energy markets (crude oil, gasoline, heating oil, electricity, natural gas, coal, etc). While the VAR model in section 5.1 integrates the global crude oil market with the U.S. retail gasoline market, it is not clear how to generalize the identifying assumptions to a model including several retail energy markets.

Thus, for now, there is no substitute for the approach taken in section 3.1.

5.4 The Differential Impact of Demand and Supply Shocks in Global Oil Markets

5.4.1 Real GDP and CPI Inflation

It is useful to relate demand and supply shocks identified in a VAR model of the global crude oil market of the type discussed above to U.S. real GDP and consumer prices. Under the assumption that these shocks are predetermined with respect to U.S. macroeconomic aggregates, consistent response estimates may be obtained from regressions of real GDP growth and inflation, respectively, on a constant and a distributed lag of the shock in question (see Kilian 2007c). Figure 6 from Kilian (2007c) shows the estimated responses of real GDP growth, the level of real GDP, CPI inflation and the level of consumer prices to selected shocks. Figure 6 illustrates that each demand and supply shock in global oil markets generates a unique pattern of responses.

The first column of Figure 6 shows that an unexpected global crude oil supply disruption leads to a persistent and statistically significant decline in the level of real GDP. The response is statistically significant at the 10% level for the first nine quarters. In contrast, the response to an unanticipated increase in global aggregate demand for industrial commodities shows an initial increase of U.S. real GDP, followed by a decline below the original level of real GDP after one year. This pattern is consistent with the view that such a shock has both direct and indirect effects on the U.S. economy that work in opposite directions. In the short run, an unanticipated expansion of the business cycle in global commodity markets directly stimulates the U.S. growth. It also raises the price of oil thereby indirectly slowing U.S. growth. Initially, the direct positive effect is large enough to offset the negative effect working through higher oil prices. Over time, the stimulus from the global economy weakens and the growth-retarding effect working through higher oil prices begins to dominate. After seven quarters the decline in U.S. real GDP is statistically significant at the 10% level. Finally, oil-market specific shocks (such as an increase in precautionary demand for crude oil) cause a persistent decline in real GDP. Unlike the decline triggered by an oil supply disruption, the decline triggered by oil-market specific increases in demand reaches its maximum only after three years. After four quarters, the decline is statistically significant at the 10% level.

The second column of Figure 6 shows the corresponding responses of consumer prices measured by the CPI. Adverse crude oil supply shocks cause a statistically insignificant increase in CPI inflation on impact, but almost no increase in the price level. Aggregate demand shocks cause a delayed increase in the price level. The response is significant at the 10% level after two quarters. An oil-market specific demand shock causes a large and even more statistically significant increase in the price level at all horizons.

This evidence also helps us understand why the consequences of the increase in crude oil and gasoline prices since 2002 have been relatively benign so far. As Figure 5 shows, much of the increase in U.S. gasoline prices was fueled by strong demand for crude oil driven by a booming world economy. Such demand pressures have two effects: (1) they stimulate U.S. economic activity directly and (2) they indirectly slow down U.S. real economic activity through their effect on the price of energy. As shown in Figure 6, especially in the short run, the expansionary effects of an aggregate demand shock for industrial commodities will offset the adverse consequences of higher gasoline prices. Only with some delay will U.S. real GDP decline, as the price of energy rises and the economic stimulus from higher global demand weakens. Thus, following several such shocks, the economy will tend to remain quite resilient, and seemingly unaffected by higher oil prices.

5.4.2 Stock Markets

The same distinction between different oil demand and oil supply shocks also matters for understanding the response of the U.S. stock market to oil price shocks. Using a similar VAR methodology, Kilian and Park (2007) show that the responses of real U.S. stock returns to oil price shocks differ substantially, depending on the underlying causes of the oil price increase. On average, 13 percent of the variation in aggregate stock returns can be attributed to the shocks that drive the crude oil market (most of which is driven by demand shocks), but the contribution of each shock varies over time, making it necessary to understand the origins of a given oil price increase before its consequences on aggregate U.S. stock returns can be assessed.

This point is illustrated in Figure 7 which shows the responses of U.S. stock prices to each of the three demand and supply shocks in global oil markets already discussed in the preceding subsections. The conventional wisdom that higher oil prices necessarily cause lower returns is seen to apply only to oil-market specific demand shocks such as increases in the precautionary demand for crude oil that reflect concerns about future oil supply shortfalls. In contrast, positive shocks to the global aggregate demand for industrial commodities cause both higher real oil prices and higher stock prices. Hence, higher oil prices need not be bad news for the stock market. Finally, shocks to the global production of crude oil, while not trivial, are far less important for understanding changes in stock prices than shocks to global demand for industrial commodities and shocks to the precautionary demand for crude oil. Given the evidence that recent increases in the price of crude oil have been driven primarily by strong global demand for all industrial commodities, this evidence helps explain the apparent resilience of the U.S. stock market to higher oil prices so far. Kilian and

Park also show that shocks to the precautionary demand for crude oil provide an explanation for the negative association between stock returns and inflation found in previous studies of the postwar period (see, e.g., Kaul and Seyhun 1990, Hess and Lee 1999), whereas other shocks in the crude oil market do not.

Finally, this study uncovers interesting differences across industries. For example, shares in the petroleum and natural gas industry as well as the gold and silver mining will appreciate in response to a positive oil-market specific demand shock, while the automobile industry and the retail sector will experience a persistent and significantly negative response to the same shock. In contrast, if the same increase in the price of crude oil is driven by innovations to global real economic activity, the cumulative returns of all four industries will increase during the first year after the shock, but especially that of petroleum and natural gas stocks. A systematic analysis of industry returns suggests considerably stronger and often more significant responses to demand shocks for crude oil than to oil supply shocks, although the degree of sensitivity varies across industries. Outside of the energy sector, the strongest responses to demand shocks are found in industries such as the automobile industry, the retail industry, and tourism-related industries such as restaurants and lodging, consistent with the view that oil price shocks are primarily shocks to the demand for goods and services rather than their supply. The energy intensity of industries is not an important factor in explaining differences in the responses of real stock returns across manufacturing industries.

5.4.3 External Accounts

A similar methodology may also be applied to study the implications of oil price shocks for the U.S. balance of payments. The analysis in Kilian, Rebucci and Spatafora (2007) sheds some light on this question. First, while any demand or supply shock in the global crude oil market that raises the price of crude oil will push the U.S. oil trade balance into deficit, the timing and the magnitude of the response of the oil trade balance depends on the source of the shock. A positive aggregate demand shock, for example, tends to generate a statistically significant oil trade deficit that reaches its maximum only after about two years. In contrast, unanticipated disruptions to oil supply generate an immediate but short-lived and statistically insignificant oil trade deficit. Positive shocks to the precautionary demand for oil tend to cause a sustained and statistically significant oil trade deficit.

Second, demand and supply shocks in the global oil market also affect the non-oil trade balance. For example, aggregate demand shocks tend to be associated with a highly statistically significant non-oil trade surplus that offsets the oil trade deficit, giving rise to an overall U.S. trade surplus and current account surplus. Negative oil supply shocks are not associated with a statistically significant improvement in the non-oil trade balance. Rather the non-oil trade balance turns negative after two years, resulting in an overall trade deficit and current account deficit that becomes statistically significant three years after the shock. Finally, positive precautionary demand shocks are not associated with statistically significant changes in the U.S. non-oil trade balance. Their overall effect is a trade deficit as well as a current account deficit, the magnitude of which is not statistically significant.

6 Concluding Remarks

In recent years, our understanding of the nature of energy price shocks and their effects on the economy has evolved dramatically. Only a few years ago, the prevailing view in the literature was that at least the major crude oil prices increases were exogenous with respect to the U.S. economy and that these increases were associated with political disturbances in the Middle East. This view has not held up to scrutiny. Today, we know that simple statistical transformations of the price of oil are not sufficient to identify oil price increases driven by exogenous crude oil supply shocks.

Moreover, it has been shown that direct measures of exogenous shocks to the production of crude oil have low explanatory power for crude oil prices. This evidence suggests that attempts to link major oil price increases to disruptions of crude oil production *alone* will not be successful. At the same time, the surge in crude oil prices since 2002 has demonstrated that large and sustained increases in oil prices may be driven primarily by demand for crude oil, especially when the ability to increase crude oil production in the near future is limited. This observation is important because it suggests that oil demand shocks may have played a central role in explaining earlier episodes of oil price shocks as well. Indeed, all major oil price shocks have coincided with capacity constraints in crude oil production and strong demand for crude oil. This point was first made in Barsky and Kilian (2002, 2004), but until recently it has not been possible to quantify the relative importance of demand and supply shocks in the global crude oil market.

Recent advances in the literature allow us to answer this question. The analysis reviewed in this paper suggests that, while no two oil price shocks are alike, most oil price shocks since the 1970s have been driven by a combination of strong global demand for industrial commodities (including crude oil) and expectations shifts that increase precautionary demand for crude oil specifically. These expectations shifts reflect the market's perception of the likelihood of a future shortfall in the supply of oil. The likelihood of a future shortfall of crude oil production is driven by expectations about future demand for crude oil as well as future supplies of crude oil. The nature of the concerns of the market may evolve over time. For example, the threat of an oil embargo or of a Soviet invasion of Iran no longer preoccupies the market today, but the possibility of political upheaval in Saudi Arabia seems more real now than in 1974 or 1979. Likewise, concerns about military action against Iran which were nil in 1974, rose sharply in 1979, then all but vanished, but recently have made a comeback. It is important to keep in mind that these expectations need not be realized in the observed sample period, similar to the phenomenon of a peso problem in foreign exchange markets. It is also important to stress that expected supply disruptions alone are not enough to cause precautionary demand to increase. It is tight supply in conjunction with strong demand for crude oil that causes expectations shifts. For example, at times in the 1980s about 30 oil tankers were attacked in the Persian Gulf in given month, yet the price of oil continued to fall, reflecting the abundant supply of crude oil elsewhere in the world and the low state of global demand for crude oil.

One of the striking findings of the recent literature is that precautionary demand shocks driven by expectations shifts, unlike other oil demand and oil supply shocks, may have immediate and large effects on the U.S. economy. In many ways, they resemble the types of shocks that the earlier literature associated with exogenous political events in the Middle East. These political events indeed matter, but not so much through their effect on crude oil production, but through their effect on expectations of future crude oil production disruptions. A case in point is the invasion of Kuwait in 1990. The reason the price of crude oil skyrocketed in mid-1990 was not so much the cessation of crude oil production in Iraq and Kuwait, but rather the concern that Iraq may invade Saudi Arabia and occupy the Saudi oil fields, causing a much larger oil supply disruption. As we know, this never happened, but it explains the sharp increase in oil prices in mid-1990 (over and above what would have been expected based on the physical reduction of crude oil supply at that point), and it explains the subsequent sharp fall in crude oil prices after the U.S. had moved enough troops to Saudi Arabia in late 1990 to forestall the occupation or destruction of the Saudi oil fields.

In short, we have a much better understanding today of how oil price shocks may arise. There also has been tremendous progress in understanding how energy price shocks affect the U.S. economy. Much of the earlier literature was preoccupied with the effect of changes in the price of crude oil. One recent insight is that there is an important distinction between retail energy prices such as motor gasoline and the price of primary energy goods such as crude oil. As we have shown in this paper, shocks to U.S. refining capacity may explain a substantial component of the price of gasoline not captured by crude oil prices. In fact, gasoline and crude oil prices may move in opposite directions. Thus, it is essential to focus on retail (or intermediate) energy prices in studying the response of consumers (or firms) to higher energy prices. It can also be important to focus on a broad measure of retail or intermediate energy prices. For example, the magnitude of real energy price shocks faced by firms is much smaller in general than the corresponding shocks to crude oil prices, owing to the large share of electric power available at stable prices. In fact, in 1974, crude oil prices rose twice as much as intermediate energy prices. Even more strikingly, in 1990, crude oil prices rose by 83%, whereas intermediate energy prices only rose by 12%.

The traditional view of oil price shocks has been that they act as *aggregate supply* shocks in a traditional textbook model or as *technology* shocks in a modern dynamic stochastic general equilibrium model. Despite some important advances, the nature of this supply channel of transmission and its quantitative importance remains an open issue. An increasingly popular alternative view in the literature, discussed in section 3, is that oil price shocks affect the economy primarily through their effect on consumer expenditures and firm expenditures instead. In this view, higher energy prices cause both a reduction in *aggregate demand* in traditional parlance and a shift in expenditures which in turn causes a ripple effect throughout the economy, as firms adjust their production plans. Models of the demand channel of transmission have the merit of being consistent with anecdotal evidence that oil price shocks are typically perceived as adverse demand shocks at the industry level. Some of these models also hold the promise of generating potentially much larger effects than would be expected based on the small share of energy in consumption. Finally, some models of the demand channel (such as the sectoral shifts model) seem capable of rationalizing apparent asymmetries in the response of the economy to oil price increases and oil price decreases.

The evidence that emerges from the recent literature is that some of the channels of transmission that collectively are referred to as the demand channel in section 3 indeed matter in practice, whereas others do not appear to be quantitatively important. In particular models that imply asymmetric responses to energy price increases and decreases were shown to lack empirical support. A large literature has been devoted to studying apparent asymmetries in the response of the economy to energy price increases and energy price decreases. This literature was motivated by the fact that sharply higher energy prices in 1979 appeared to be followed by a recession, whereas sharply lower energy prices in 1986 were not followed by a major economic expansion. This evidence seemed to call for theoretical models capable of explaining asymmetric responses to energy price increases and energy price decreases. As discussed in this paper, there is reason to believe that the profession may have misinterpreted this evidence. In fact, there is no evidence of asymmetries in real consumption growth and the absence of an increase in investment expenditures in 1986 appears to be driven by an exogenous decline in business investment in 1986, related not to the fall in energy prices but to the 1986 Tax Reform Act. This effect was exacerbated by the response of investment in the petroleum and natural gas industry to the collapse of OPEC in late 1985, which far exceeded the response one would have expected to a decline in energy prices alone. Moreover, composition effects from aggregating investment expenditures related to petroleum, coal and natural gas mining and all other investment expenditures helped generate an apparent asymmetry in the growth of aggregate investment. Hence, the apparent asymmetry in the real GDP growth data seems to be largely a statistical artifact. As the evidence reviewed in this paper suggests, despite asymmetric point estimates in some cases, there is no compelling statistical evidence against the symmetry hypothesis.

This result has important implications for demand-driven models of the transmission of energy price shocks. Models of asymmetric transmission mechanisms such as the uncertainty effect of Bernanke (1983) or the reallocation effect of Hamilton (1988) have been widely used in the empirical literature on oil prices to explain the apparent breakdown of the linear relationship between real GDP growth and oil prices in the mid-1980s. The lack of evidence against the symmetry hypothesis suggests that neither the uncertainty effect nor the reallocation effect are quantitatively important in the data. We conclude that there is no compelling reason to abandon linear models that impose symmetry on the response to energy price increases and energy price decreases. Using such models allowed us to quantify the effect of retail energy price shocks on consumer and business investment expenditures. We documented that the demand channel of the transmission of oil prices is indeed more important than the small share of energy in expenditures would suggest. The estimated elasticities for total consumption and total nonresidential investment are -0.15 and -0.16, respectively, or about four times as high as the share argument would suggest. Nevertheless, the overall responses of total consumption and of total nonresidential investment as measured by the energy price elasticities are still fairly small, and of limited importance in explaining business cycle fluctuations. Evidence of larger elasticities was found only for specific expenditure items. It was shown that the bulk of the economy's response is associated with reduced demand for vehicles and reduced residential demand for houses.

An interesting observation in the recent literature is that the effects of energy price shocks have weakened since the second half of the 1980s. For example, the one-year energy price elasticity of total real consumption drops from -0.30% prior to 1987 to only -0.08% after 1987. It can be shown that this phenomenon is not associated with the evolution of the share of energy in consumer expenditures or in value added nor is it caused by a decline in the volatility or magnitude of energy price shocks. Rather it can be explained in part by changes in the composition of U.S. automobile production and to the declining overall importance of the U.S. automobile sector. It is also related to the nature of recent energy price shocks. There has been much speculation as to why the recent surge in gasoline prices in particular has not so far caused a major recession. Part of the answer is that much of that increase was driven by strong global demand for industrial commodities. Such demand shocks have both direct effects on the U.S. economy and indirect effects working through higher oil prices. In the short run, the positive effects on the U.S. economy will dominate, whereas in the long run real growth will tend to be below average, as energy prices remain high and the economic stimulus from higher global demand weakens. This response pattern differs sharply from the effect of higher energy prices driven primarily by shocks to the precautionary demand for oil, for example.

The distinction between higher energy prices driven by one shock or another has far-reaching implications, as each shock has different effects on the U.S. economy and on the real price of energy. We illustrated this point for several U.S. macroeconomic aggregates including real GDP, consumer prices, real stock returns and the current account. One implication of this analysis is that conventional estimates of the response to unanticipated energy price changes are best thought of as the response to an average energy price shock and may be sensitive to the choice of sample period, as the composition of the underlying demand and supply shocks evolves over time.

Notwithstanding the many insights the recent literature has yielded, there is still more to be learned about how energy price shocks are transmitted throughout the economy. Future empirical work with disaggregate industry or plant level data augmented by structural models is likely to be promising. A recent example of such work is Herrera (2007). The challenge will be to combine a deeper understanding of the nature of energy price shocks with an explicit model of firm decisions and interactions. One difficulty with such extensions is the absence of disaggregate real GDP data. Many empirical studies have therefore relied on disaggregate gross output data such as measures of industrial production (see, e.g., Lee and Ni 2002, Herrera 2007). This distinction matters because gross output may respond quite differently to energy price shocks than measures of value added such as real GDP (see, e.g., Barsky and Kilian 2002). This fact makes it difficult to relate conclusions of studies based on gross output to standard macroeconomic models based on value added production functions.

There is also considerable scope for developing full-fledged dynamic stochastic general equilibrium (DSGE) models that incorporate global and domestic energy markets. Building on the early contributions of Hamilton (1988), Kim and Loungani (1992), Rotemberg and Woodford (1996), Atkeson and Kehoe (1999), Backus and Crucini (2000), and Finn (2000), among others, there has been renewed interest in DSGE models of the effects of energy price shocks recently. In addition to the extensive work on the relationship between oil prices and monetary policy discussed earlier, a number of additional channels of transmission have been explored. For example, Polgreen and Silos (2006) model the effects of oil price shocks on the skill premium in labor markets. Wei (2003) studies the relationship between oil prices and stock markets. Dhawan and Jeske (2007b) model the energy consumption of households and firms within the context of a DSGE model. Bodenstein, Erceg and

Guerrieri (2007) investigate the role of financial risk sharing in a two-country DSGE model of the external adjustments caused by oil price shocks. Wen and Aguiar-Conraria (2006) stress the role of externalities in the propagation of oil price shocks. Notwithstanding this flurry of activity, existing DSGE models with few exceptions have remained extremely simplistic in treating crude oil prices as exogenous driving processes and in avoiding aggregation issues. Clearly, the development of these models is still at an early stage. More refined models are likely to yield important additional insights and to put earlier results in perspective.

References

[1] Alquist, R., and L. Kilian (2007), "What Do We Learn from the Price of Crude Oil Futures," mimeo, Department of Economics, University of Michigan.

[2] Atkeson, A., and P.J. Kehoe (1999), "Models of Energy Use: Putty-Putty versus Putty-Clay," *American Economic Review*, 89, 1028-1043.

[3] Auerbach, A.J. (1987), "The Tax Reform Act of 1986 and the Cost of Capital," *Journal of Economic Perspectives*, 1, Summer, 73-86.

[4] Backus, D.K., and M.J. Crucini (2000), "Oil Prices and the Terms of Trade," *Journal of International Economics*, 50, 185-213.

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

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

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

[8] Bernanke, B.S. (1983), "Irreversibility, Uncertainty, and Cyclical Investment," *Quarterly Journal of Economics*, 98, 85-106.

[9] Bernanke, B.S. (2004), "Oil and the Economy," remarks at the Distinguished Lecture Series, Darton College, Albany, Georgia, October 21.

[10] Bernanke, B.S. (2006a), "Energy and the Economy," speech to the Economic Club of Chicago, June 15.

[11] Bernanke, B.S. (2006b), "The Economic Outlook," remarks before the National Italian American Foundation, New York, November 28.

[12] Bernanke, B.S., M. Gertler, and M. Watson (1997), "Systematic Monetary Policy and the Effects of Oil Price Shocks," *Brookings Papers on Economic Activity*, 1, 91-142.

[13] Bernanke, B.S., M. Gertler, and M. Watson (2004), "Oil Shocks and Aggregate Macroeconomic Behavior: The Role of Monetary Policy: A Reply," *Journal of Money, Credit, and Banking*, 36, 286-291.

[14] Blanchard, O.J., and J. Galí (2007), "The Macroeconomic Effects of Oil Shocks: Why are the 2000s so Different from the 1970s?" NBER Working Paper No. 13368.

[15] Bodenstein, M., C.J. Erceg, and L. Guerrieri (2007), "Oil Shocks and External Adjustment," mimeo, Federal Reserve Board.

[16] Borenstein, S., A.C. Cameron, and R. Gilbert (1997), "Do Gasoline Prices Respond Asymmetrically to Crude Oil Price Changes?" *Quarterly Journal of Economics*, 112, 305-339.

[17] Bresnahan, T., and V. Ramey (1993), "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry," *American Economic Review Papers and Proceedings*, 83, 213-218.

[18] Bureau of Economic Analysis (1987), "The Business Situation," *Survey of Current Business*, January.

[19] Carlstrom, C.T., and T.S. Fuerst (2006), "Oil Prices, Monetary Policy, and Counterfactual Experiments," *Journal of Money, Credit, and Banking*, 38, 1945-1958.

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

[21] Cragg, J.G., and S.G. Donald (1993), "Testing Identifiability and Specification in Instrumental Variable Models," *Econometric Theory*, 9, 222-240.

[22] Cullen, J.B., L. Friedberg and C. Wolfram (2004), "Consumption and Changes in Home Energy Costs: How Prevalent is the 'Heat or Eat' Decision?", mimeo, Department of Economics, University of California at San Diego.

[23] Dahl, C., and T. Sterner (1991), "Analyzing Gasoline Demand Elasticities: A Survey," *Energy Economics*, 13, 203-210.

[24] Davis, S.J. (1987), "Allocative Disturbances and Specific Capital in Real Business Cycle Theories," *American Economic Review Papers and Proceedings*, 77, 326-332.

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

[26] Davis, M.C., and J.D. Hamilton (2004), "Why Are Prices Sticky? The Dynamics of Wholesale Gasoline Prices," *Journal of Money, Credit, and Banking,* 36, 17-37.

[27] Dhawan, R., and K. Jeske (2007a), "Taylor Rules with Headline InflationL A Bad Idea!" mimeo, Federal Reserve Bank of Atlanta.

[28] Dhawan, R., and K. Jeske (2007b), "Energy Price Shocks and the Macroeconomy: The Role of Consumer Durables," mimeo, Federal Reserve Bank of Atlanta.

[29] Dotsey, M., and M. Reid (1992), "Oil Shocks, Monetary Policy, and Economic Activity," *Economic Review of the Federal Reserve Bank of Richmond*, 78, 14-27.

[30] Douglas, C., and A.M. Herrera (2007), "Why are Gasoline Prices Sticky? A Test of Alternative Models of Price Adjustment," mimeo, Department of Economics, Michigan State University.

[31] Dudine, P., Hendel, I., and A. Lizzeri (2006), "Storable Good Monopoly: The Role of Commitment," forthcoming: *American Economic Review*.

[32] Edelstein, P., and L. Kilian (2007a), "Retail Energy Prices and Consumer Expenditures," mimeo, Department of Economics, University of Michigan.

[33] Edelstein, P., and L. Kilian (2007b), "The Response of Business Fixed Investment to Energy Price Changes: A Test of Some Hypotheses about the Transmission of Energy Price Shocks," forthcoming: *B.E. Journal of Macroeconomics*.

[34] Ferderer, J.P. (1996), "Oil Price Volatility and the Macroeconomy: A Solution to the Asymmetry Puzzle," *Journal of Macroeconomics*, 18, 1-16.

[35] Finn, M.G. (2000), "Perfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit and Banking*, 32, 400-416.

[36] Goldberg, P.K. (1998), "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S.," *Journal of Industrial Economics*, 46, 1-33.

[37] Gonçalves, S., and L. Kilian (2004), "Bootstrapping Autoregressions in the Presence of Conditional Heteroskedasticity of Unknown Form," *Journal of Econometrics*, 123, 89-120.

[38] Hamilton, J.D. (1983), "Oil and the Macroeconomy Since World War II," *Journal of Political Economy*, 91, 228-248.

[39] Hamilton, J.D. (1988) "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96, 593-617.

[40] Hamilton, J.D. (1996), "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics*, 38, 215-220.

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

[42] Hamilton, J.D. (2005), "Oil and the Macroeconomy," forthcoming in S. Durlauf and L. Blume (eds), *The New Palgrave Dictionary of Economics*, 2nd ed., Palgrave MacMillan Ltd.

[43] Hamilton, J.D., and A.M. Herrera (2004), "Oil Shocks and Aggregate Economic Behavior: The Role of Monetary Policy: Comment," *Journal of Money, Credit, and Banking*, 36, 265-286.

[44] Herrera, A-M. (2007), "Inventories, Oil Shocks and Macroeconomic Behavior," mimeo, Department of Economics, Michigan State University.

[45] Herrera, A.-M., and E. Pesavento (2007), "Oil Price Shocks, Systematic Monetary Policy and the 'Great Moderation'," mimeo, Department of Economics, Michigan State University.

[46] Hess, P. J. and B. S. Lee (1999), "Stock Returns and Inflation with Supply and Demand Disturbances," *Review of Financial Studies*, 12, 1203-1218.

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

[48] Hooker, M.A. (1996b), "This Is What Happened to the Oil Price-Macroeconomy Relationship: Reply," *Journal of Monetary Economics*, 38, 221–222.

[49] Hooker, M.A. (2002), "Are Oil Shocks Inflationary? Asymmetric and Nonlinear Specifications Versus Changes in Regime," *Journal of Money, Credit, and Banking*, 34, 540-561.

[50] Hughes, J.E., C.R. Knittel, and D. Sperling (2007), "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand," forthcoming: *Energy Journal*.

[51] Huntington, H.G. (1998), "Crude Oil Prices and U.S. Economic Performance: Where Does the Asymmetry Reside?," *Energy Journal*, 19, 107-148.

[52] Jiménez-Rodriguez and Sánchez (2005), "Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries," *Applied Economics*, 37, 201-228.

[53] Kaul, G., and H.N. Seyhun (1990), "Relative Price Variability, Real Shocks and the Stock Market," *Journal of Finance*, 45, 479-496.

[54] Kilian, L. (2007a), "Exogenous Oil Supply Shocks: How Big Are They and How Much Do They Matter for the U.S. Economy?" forthcoming: *Review of Economics and Statistics*.

[55] Kilian, L. (2007b), "A Comparison of the Effects of Exogenous Oil Supply Shocks on Output and Inflation in the G7 Countries," forthcoming: *Journal of the European Economic Association*.

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

[57] Kilian, L., and C. Park (2007), "The Impact of Oil Price Shocks on the U.S. Stock Market," mimeo, Department of Economics, University of Michigan.

[58] Kilian, L., A. Rebucci, and N. Spatafora (2007), "Oil Shocks and External Balances," mimeo, Department of Economics, University of Michigan.

[59] Kim, I.-M. and P. Loungani (1992), "The Role of Energy in Real Business Cycle Models," *Journal of Monetary Economics*, 29, 173-189.

[60] Kormilitsina, A. (2007), "Optimal Monetary Policy and Oil Price Shocks," mimeo, Department of Economics, Duke University.

[61] Leduc, S., and K. Sill (2004), "A Quantitative Analysis of Oil-Price Shocks, Systematic Monetary Policy, and Economic Downturns," *Journal of Monetary Economics*, 51, 781-808.

[62] 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, 823-852.

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

[64] Mork, K.A. (1989), "Oil and the Macroeconomy. When Prices Go Up and Down: An Extension of Hamilton's Results," *Journal of Political Economy*, 97, 740-744.

[65] Muehlegger, E.J. (2006), "Gasoline Price Spikes and Regional Gasoline Content Regulations: A Structural Approach," mimeo, JFK School, Harvard University.

[66] Pindyck, R.S. (1991), "Irreversibility, Uncertainty and Investment," *Journal of Economic Literature*, 29, 1110-1148.

[67] Polgreen, L. and P. Silos (2006), "Crude Substitutions: The Cyclical Dynamics of Oil Prices and the College Premium," *Federal Reserve Bank of Atlanta Working Paper 2006-14*.

[68] Reiss, P.C., and M.W. White (2005), "Household Electricity Demand Revisited," *Review of Economic Studies*, 72, 853-883.

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

[70] 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, 518-529.

[71] Stock, J.H., and M. Yogo (2005), "Testing for Weak Instruments in Linear IV Regression," in D.W.K. Andrews and J.H. Stock (eds.), *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, Cambridge: Cambridge University Press, 80–108.

[72] Wakefield, J.C. (1987), "The Tax Reform Act of 1986," *Survey of Current Business*, Bureau of Economic Analysis.

[73] Wei, C. (2003), "Energy, the Stock Market, and the Putty-Clay Investment Models," *American Economic Review*, 93, 311-323.

Table 1: Instrumental Variable Regressions for U.S. Real GDP Growth

Regressand: Δgdp_t	Exogenous oil supply shocks measured as quantitative dummies as in Hamilton (2003)								Exogenous oil supply shocks as defined in Kilian (2007a)				
	1947.II-2001.III		1947.II-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	×	×	×	×	×	×	×	×	×	×	×	×	×
Δgdp_{t-1}	×	×	×	×	×	×	×	×	×	×	×	×	×
Δgdp_{t-2}	×	×	×	×	×	×	×	×	×	×	×	×	×
Δgdp_{t-3}	×	×	×	×	×	×	×	×	×	×	×	×	×
Δgdp_{t-4}	×	×	×	×	×	×	×	×	×	×	×	×	×
Δnp_{t-1}^{oil}	×	-	×	-	×	-	×	-	×	-	-	-	-
Δnp_{t-2}^{oil}	×	-	×	-	×	-	×	-	×	-	-	-	-
Δnp_{t-3}^{oil}	×	-	×	-	×	-	×	-	×	-	-	-	-
Δnp_{t-4}^{oil}	×	-	×	-	×	-	×	-	×	-	-	-	-
Δrp_{t-1}^{oil}	-	×	-	×	-	×	-	×	-	×	×	×	×
Δrp_{t-2}^{oil}	-	×	-	×	-	×	-	×	-	×	×	×	×
Δrp_{t-3}^{oil}	-	×	-	×	-	×	-	×	-	×	×	×	×
Δrp_{t-4}^{oil}	-	×	-	×	-	×	-	×	-	×	×	×	×
$g_{\min}$	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. np_t^{oil} and rp_t^{oil} stand for the nominal and real price of crude oil, respectively. × marks regressors included in the final-stage regression. Columns (1)-(6) are based on the PPI for domestic crude oil as reported by the BLS and used in Hamilton (2003); columns (7)-(13) are based on the price of imported crude oil as used in Barsky and Kilian (2004). Columns (1) through (8) are based on the quantitative dummy measure of exogenous oil supply shocks of Hamilton (2003), as extended by Kilian (2007a). Columns (9) and (10) are based on the alternative measure of exogenous oil supply shocks introduced in Kilian (2007a,b). The last three columns are based on variations of this measure. Column (11) excludes the Saudi production response; column (12) drops the 1973 Arab oil embargo; column (13) includes Saudi Arabia in the benchmark starting in 1974. The last line shows the $g_{\min}$ -statistic of Stock and Yogo (2005). The least conservative critical value for the null of weak instruments is about 4.

Table 2: One-Year Energy Price Elasticities
U.S. Consumer Expenditures
1970.2-2006.7

Total Energy Consumption	**-0.45**
Electricity	-0.15
Gasoline	**-0.48**
Heating Oil and Coal	**-1.47**
Natural Gas	-0.33
Total Consumption	**-0.15**
Nondurables	**-0.11**
Services	**-0.10**
Durables	**-0.47**
Vehicles	**-0.84**
Other Durables	-0.19

NOTES: Full-sample estimates based on the purchasing power loss associated with a change in weighted retail energy prices. The elasticities have been computed based on the average share of energy in the sample period. All results are based on estimates in Edelstein and Kilian (2007a). Boldface indicates statistical significance at the 5% level.

Table 3: One-Year Energy Price Elasticities
U.S. Investment Expenditures
1970.II-2006.IV

Total Nonresidential Investment	-0.16
Structures	0.09
Structures Excluding Mining	0.03
Mining	**1.39**
Equipment	-0.30
Equipment Excluding Mining and Oil Field Machinery	-0.30
Mining and Oil Field Machinery	**2.13**
Residential Investment in Structures	**-1.02**

NOTES: Full-sample results with quarterly dummies for 1986. Estimates based on percent change in intermediate producer energy prices. All results are based on estimates in Edelstein and Kilian (2007b). Boldface indicates statistical significance at the 5% level. Expenditures on mining are mainly related to the exploration for and extraction of crude oil, natural gas and coal.

Figure 1: U.S. Energy Share in Value Added
1977-2005

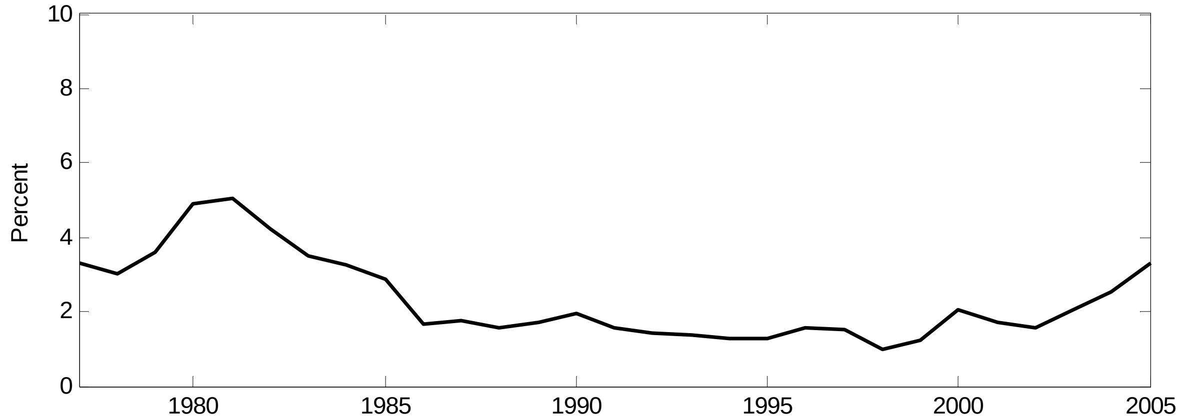

NOTES: Source: Edelstein and Kilian (2007b) based on BEA data.

Figure 2: Alternative Specifications of Energy Price Shocks
U.S. Retail Energy Prices
1970.2-2006.7

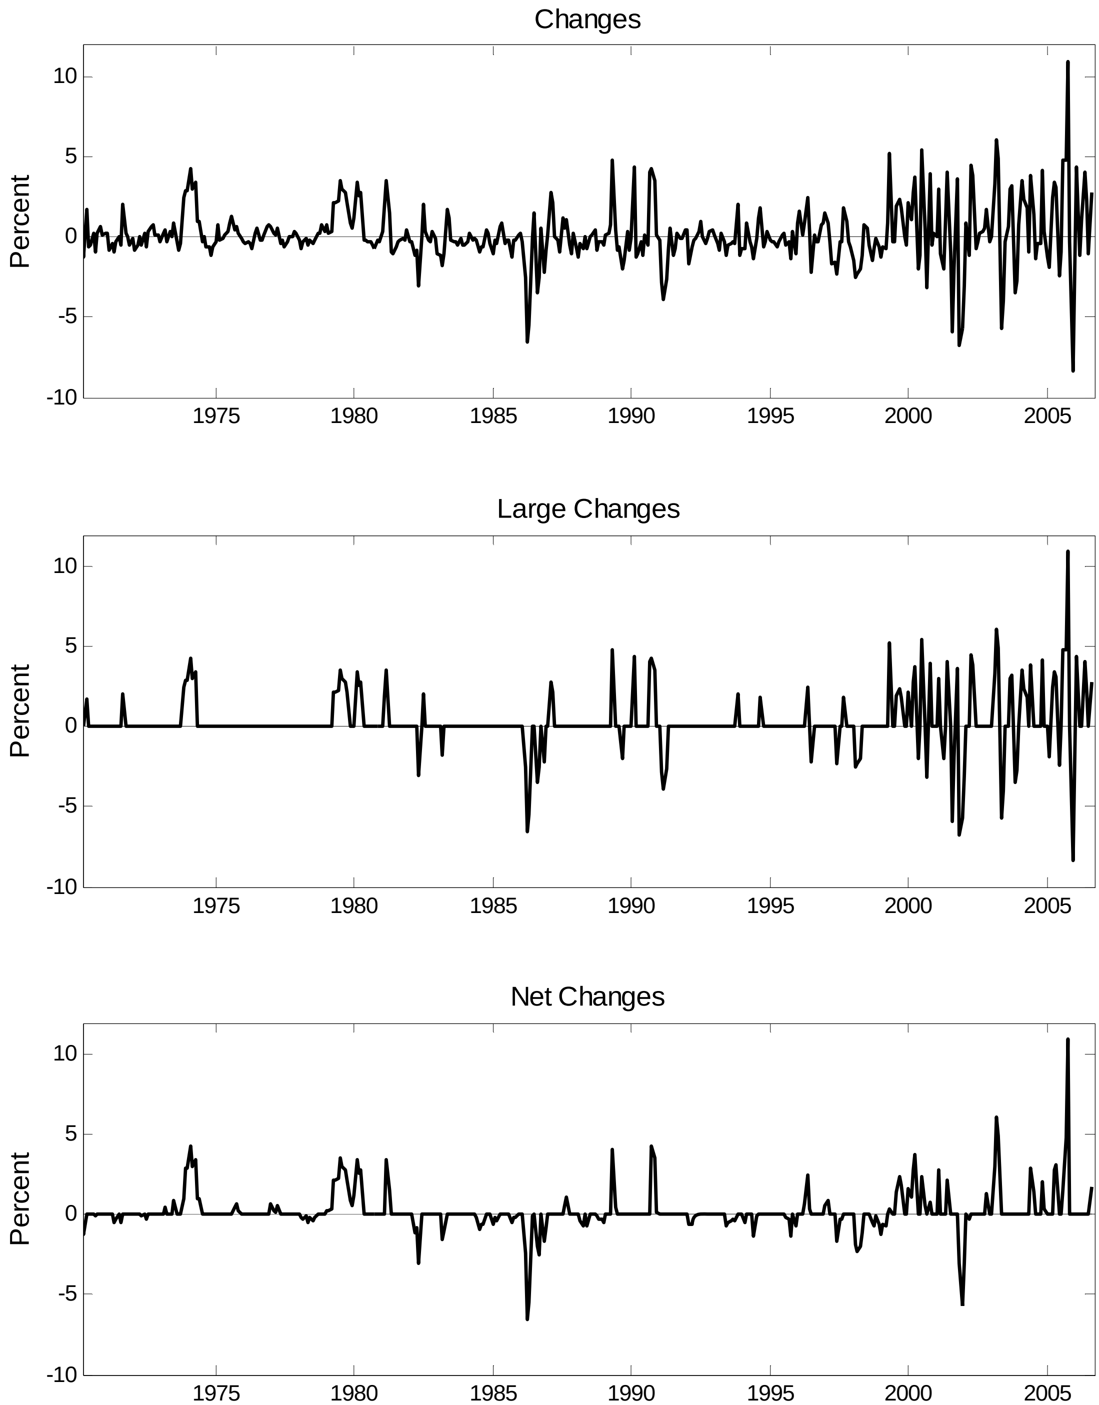

NOTES: Based on PCE price index for consumer energy prices as reported by the BEA. Source: Edelstein and Kilian (2007a).

Figure 3: Response of Real Consumption by Expenditure Item
OLS Point Estimates with One- and Two-Standard Error Bands
1970.2-2006.7

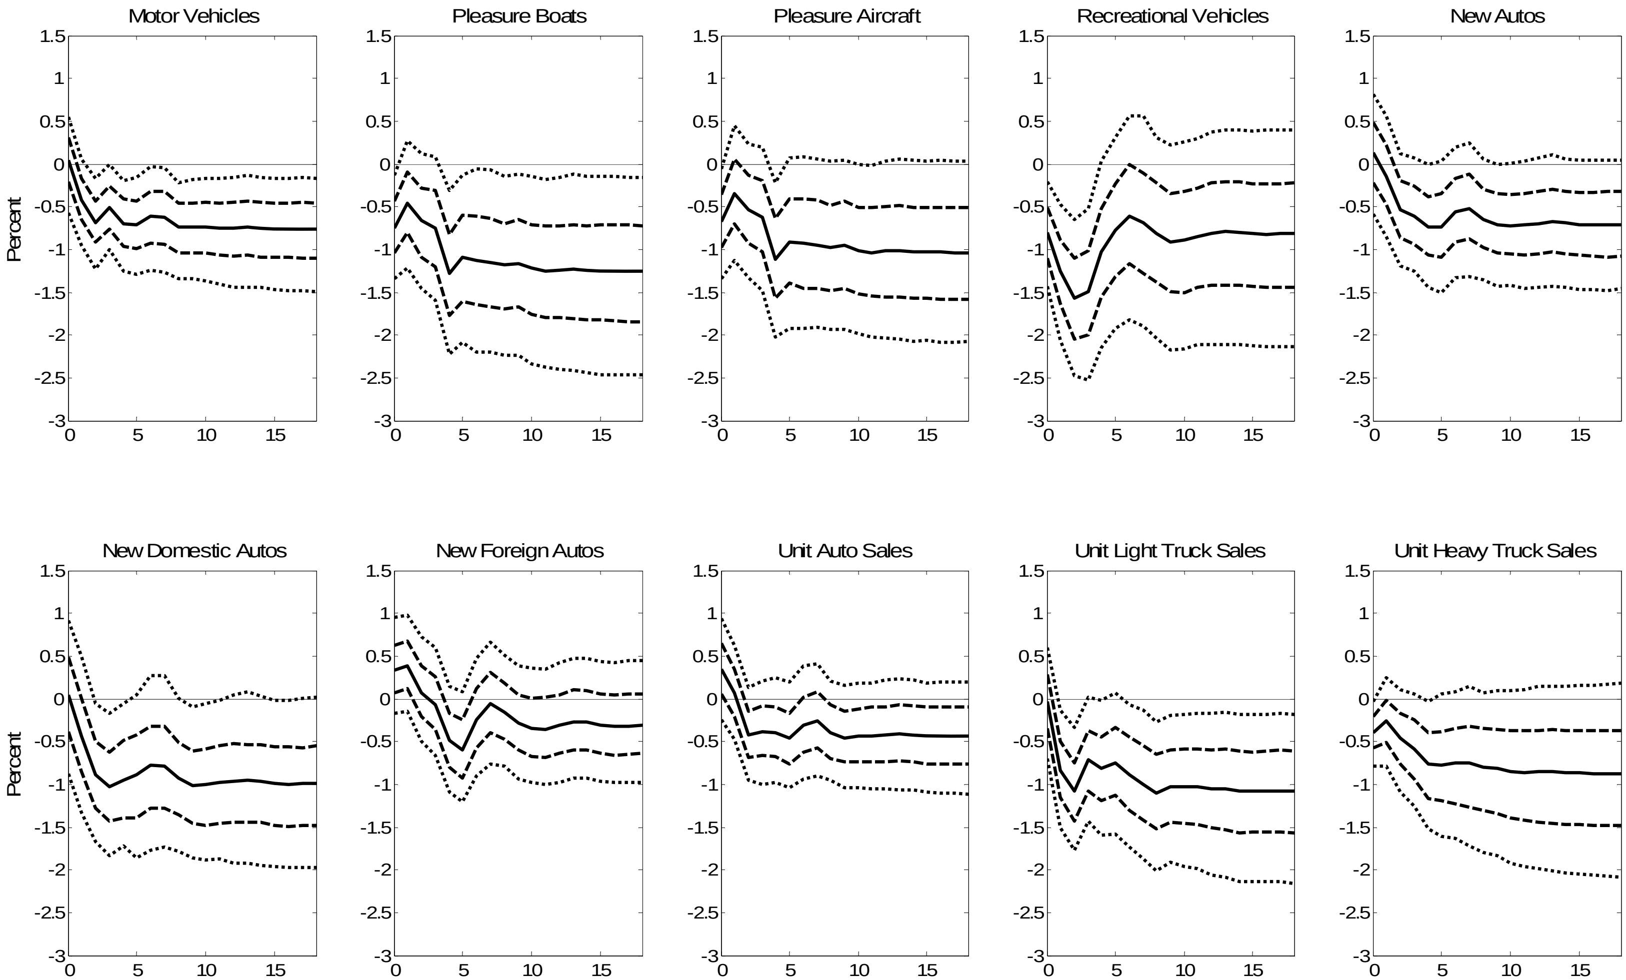

NOTES: Based on PCE price index for consumer energy prices as reported by the BEA. Source: Edelstein and Kilian (2007a).

Figure 4: Responses to One-Standard Deviation Structural Shocks
Point Estimates with One and Two-Standard Error Bands

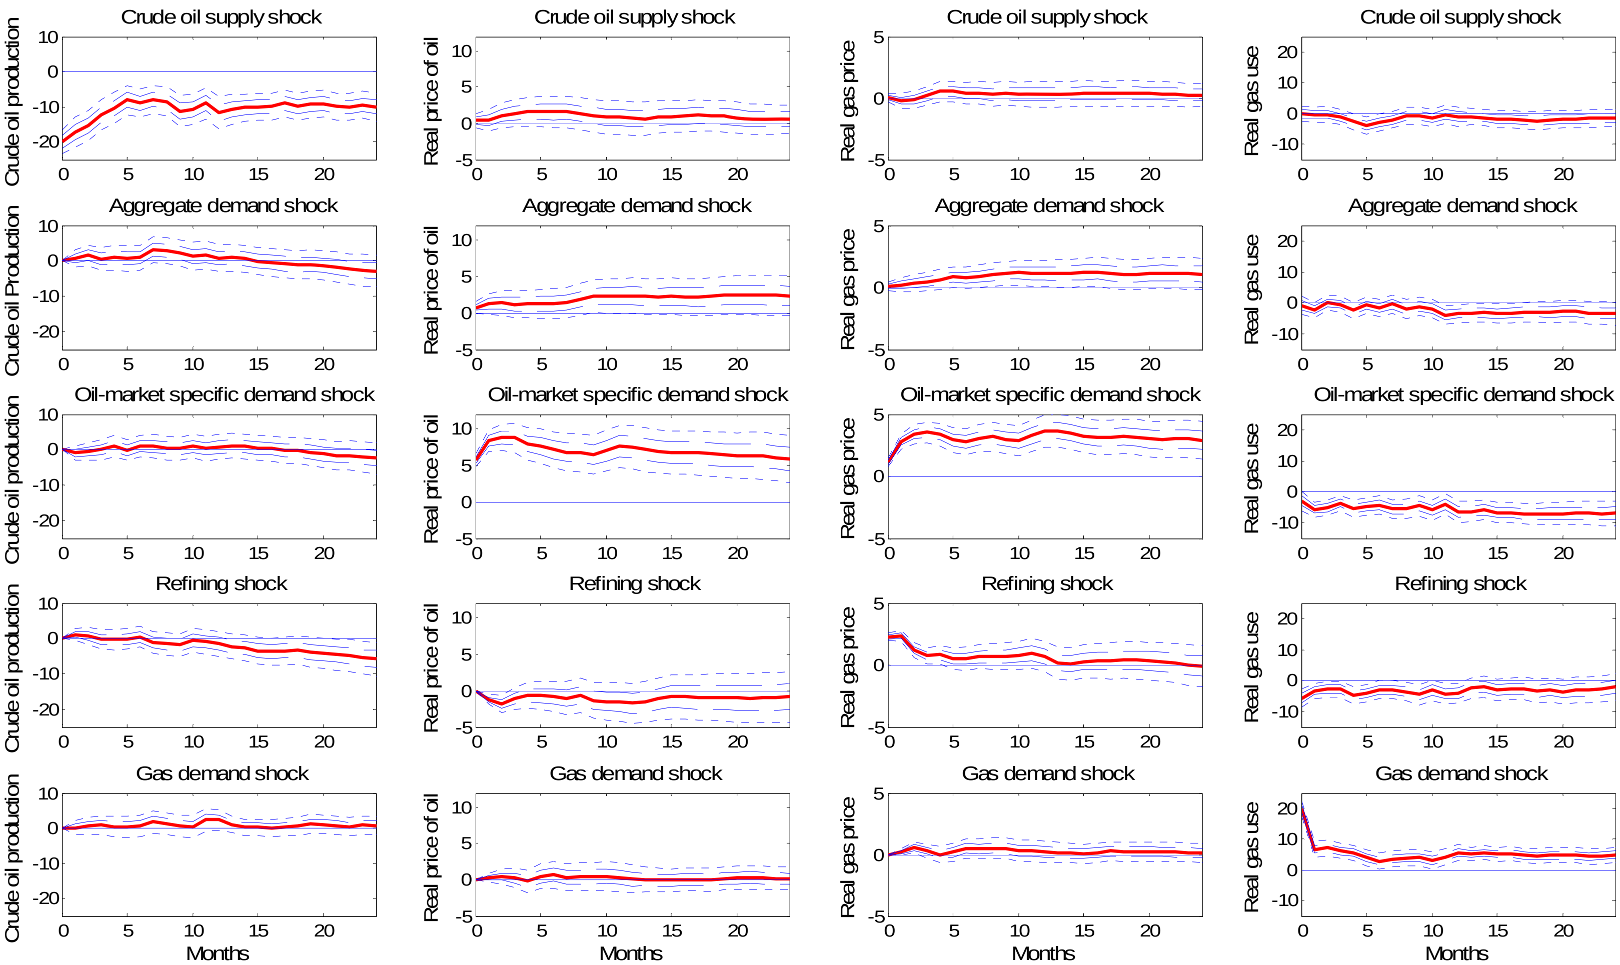

NOTES: Estimates based on the VAR(12) system described in text. The confidence intervals were constructed using a recursive-design wild bootstrap.

Figure 5: Historical Decomposition of Fluctuations in the Real U.S. Price of Gasoline 2002.1-2006.12

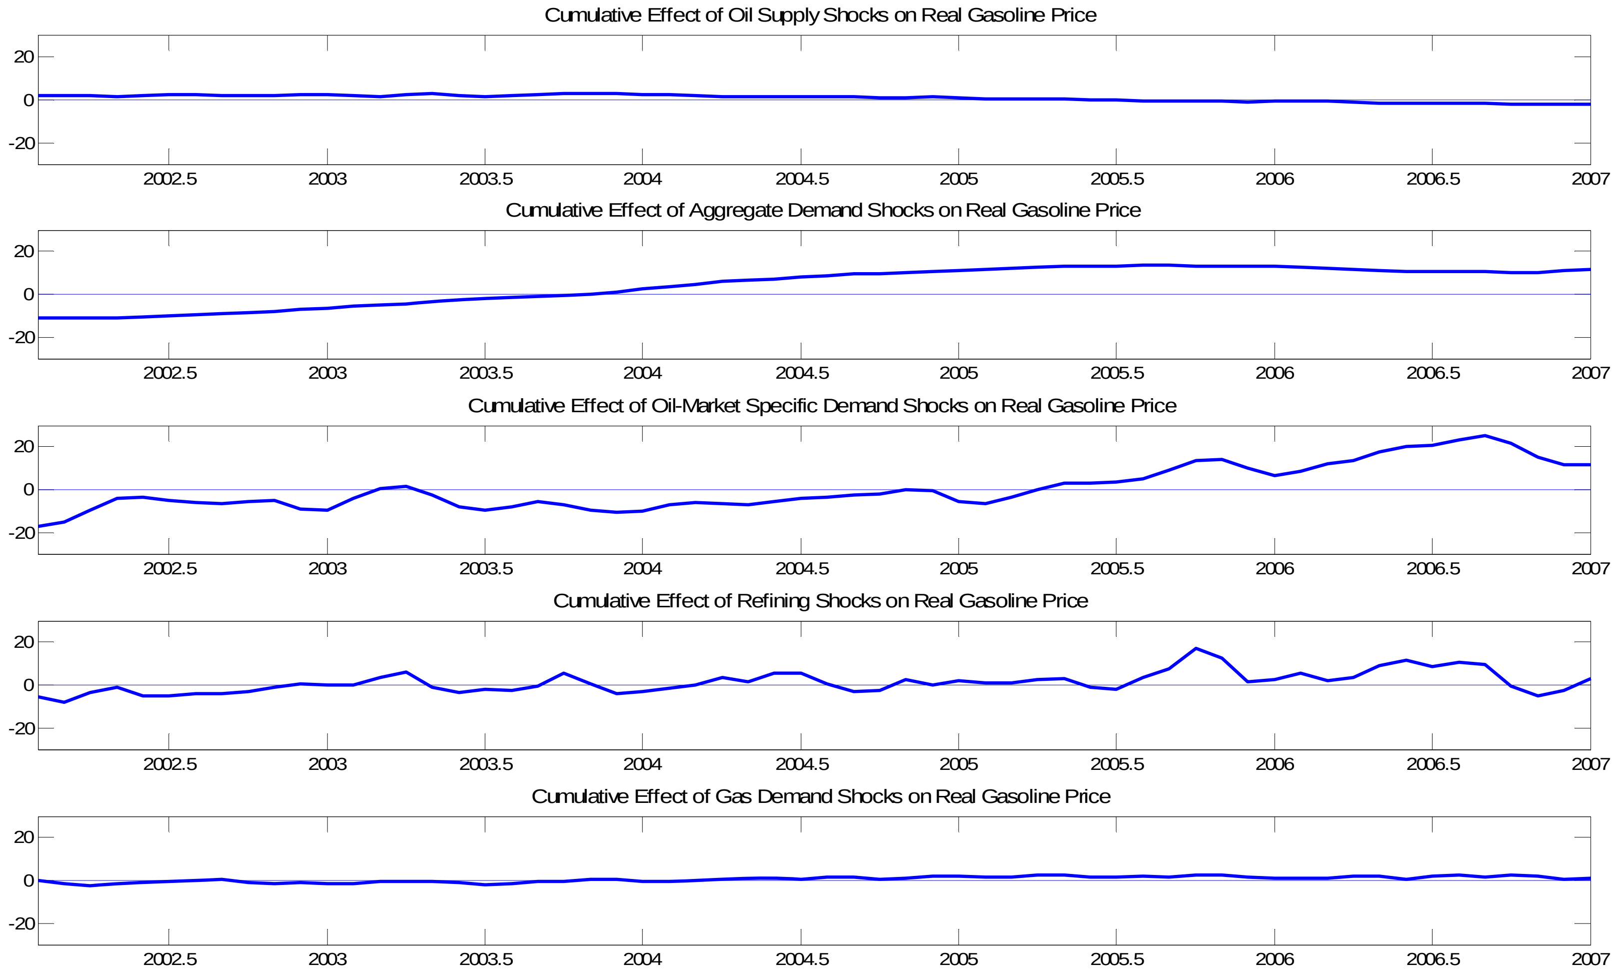

NOTES: Estimates based on the VAR(12) system described in text.

Figure 6: Responses of U.S. Real GDP and Consumer Prices to Selected Shocks
Point Estimates with One and Two-Standard Error Bands

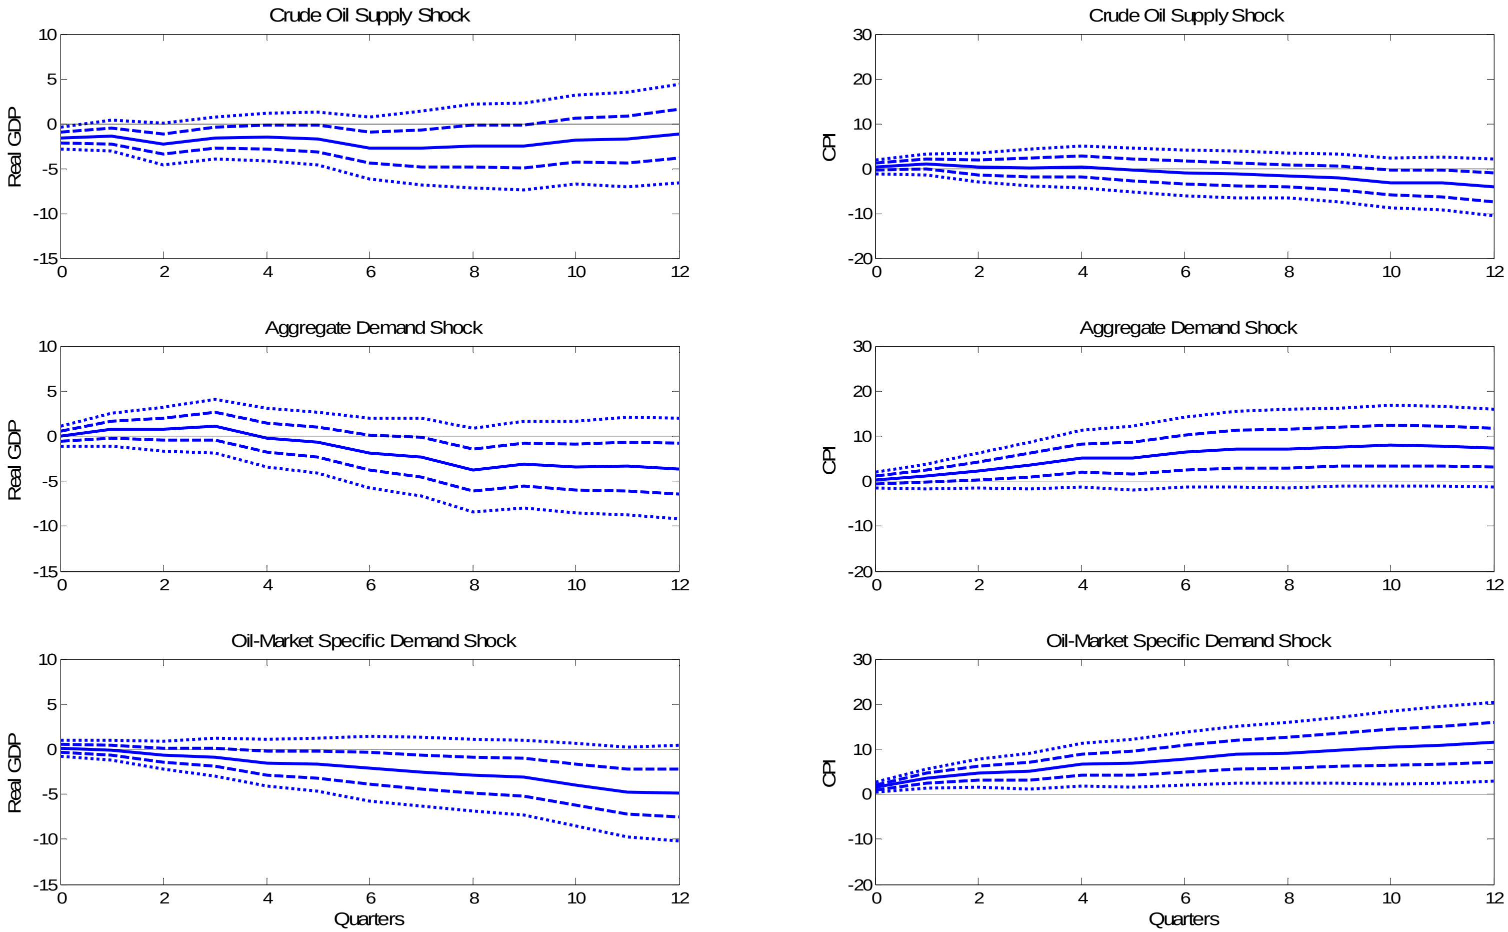

NOTES: Estimates based on the distributed lag model as described in Kilian (2007c). Confidence intervals based on block bootstrap methods.
Source: Kilian (2007c).

Figure 7: Cumulative Impulse Reponses of Real U.S. Stock Returns
Point Estimates with One and Two-Standard Error Bands

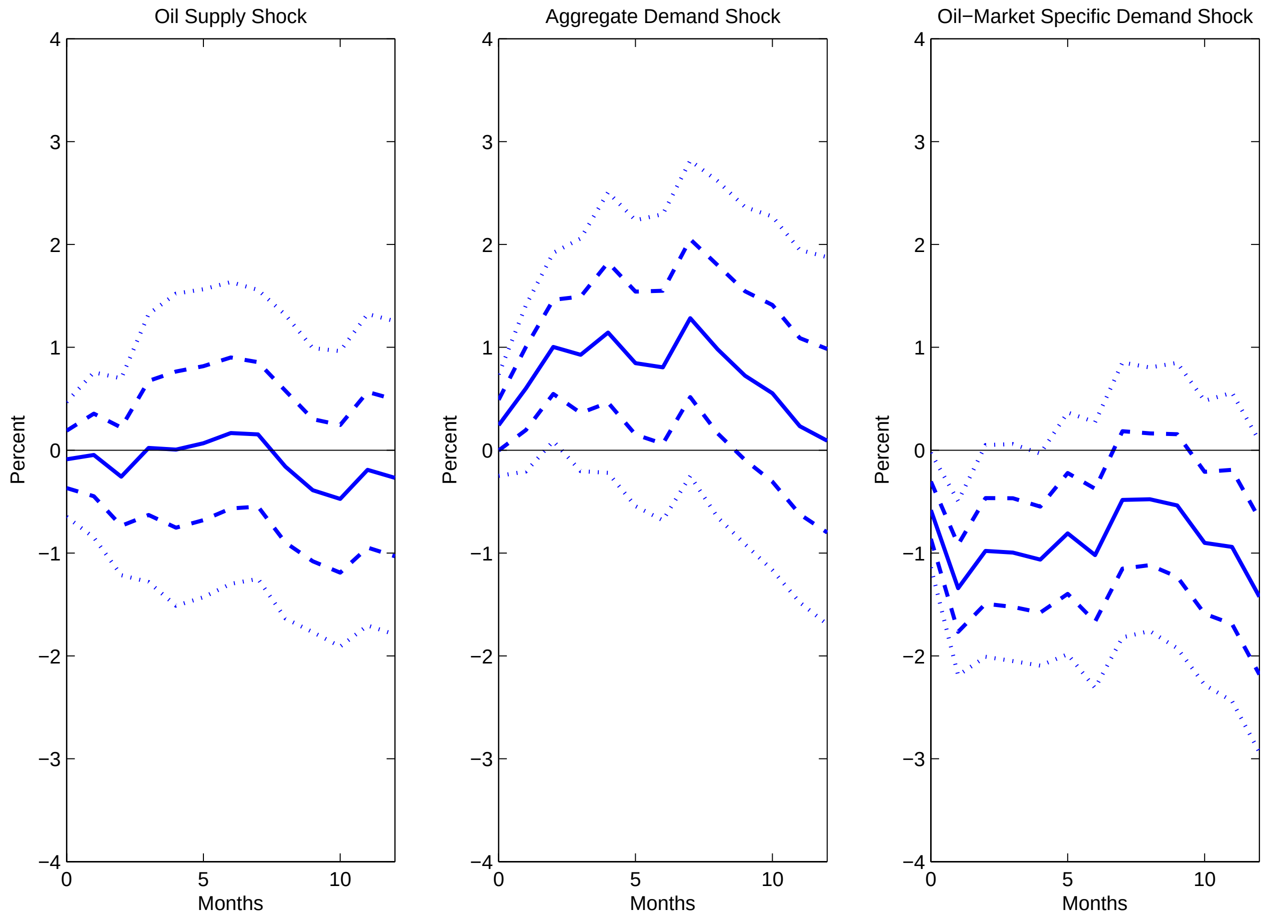

NOTES: Source: Kilian and Park (2007).