

Retail Energy Prices and Consumer Expenditures*

Paul Edelstein[†]

Lutz Kilian[‡]

University of Michigan

University of Michigan and CEPR

June 20, 2007

Abstract

In the absence of a major disruption in spending by consumers and firms, the effects of energy price shocks on the economy will be small. In this paper, we quantify the direct effect on real consumption of (1) unanticipated changes in discretionary income, (2) shifts in precautionary savings, and (3) changes in the operating cost of energy-using durables. We also evaluate the evidence for asymmetries in the response of real consumption that would be expected, for example, if shifting expenditure patterns cause sectoral reallocations. Formal tests do not reveal compelling evidence for asymmetries in the response of consumer spending, aggregate unemployment, or consumer expectations. The absence of a reallocation effect in particular, despite a comparatively large effect of energy price shocks on the consumption of new domestically produced automobiles, is consistent with the small share of the U.S. auto industry in domestic real GDP and employment. It is also consistent with the symmetric behavior of real consumption in 1979 (when energy prices rose sharply) and in 1986 (when they fell equally sharply). This finding has important implications for theoretical models of the transmission of energy price shocks. Our analysis also sheds light on the declining importance of energy price shocks for the U.S. economy. We not only document the extent to which consumption aggregates have become less responsive to energy price shocks since the mid-1980s, but we trace the declining importance of energy price shocks relative to the 1970s to changes in the composition of U.S. automobile production and the declining overall importance of the U.S. automobile sector.

KEYWORDS: Energy prices; Consumption; Consumer Sentiment; Purchasing Power; Asymmetry; Price elasticity of energy consumption.

JEL: E21, Q43.

*We thank Dan Cooper, Lucas Davis, Ana-Maria Herrera, and Jim Hamilton for helpful comments on an earlier draft of this paper.

[†]Department of Economics, University of Michigan, 611 Tappan Street, Ann Arbor, MI 48109-1220. E-mail: pedelste@umich.edu.

[‡]Department of Economics, University of Michigan, 611 Tappan Street, Ann Arbor, MI 48109-1220. E-mail: lkilian@umich.edu.

1 Introduction

Large fluctuations in energy prices have been a distinguishing characteristic of the U.S. economy since the 1970s. Between January 2002 and July 2006, for example, the PCE price index for energy goods increased by 68% in real terms.¹ The average price of a gallon of regular grade gasoline in particular increased from \$1.11 in January 2002 to \$2.98 in July 2006.² The price of energy is but one of many prices faced by households, yet they attract a disproportionate amount of attention in the media and from policymakers and economists.³ A common perception is that energy price increases are fundamentally different from increases in prices of many other goods (such as higher education, cable television, refrigerators, or movies) 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 (see, e.g. Douglass 2007). Similarly, households have little choice but to endure higher natural gas prices, as they cannot afford to leave their homes unheated. Thus, it is widely accepted that energy prices tend to lower the purchasing power of households, as consumers devote more of their income to energy consumption, leaving less discretionary income available for the purchases of other goods and services.⁴

Despite the attention that such purchasing power losses have received in the media and in policy discussions, little is known about their magnitude, about their effect on real consumption and about the extent to which consumption patterns change in response to energy price fluctuations.⁵ Answers to these questions are important for understanding the transmission of global energy price shocks to the U.S. economy. As noted in a recent survey of the literature on oil and the macroeconomy by Hamilton (2005), the 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 shared by policymakers. For example, in a recent speech, Bernanke (2006a) stressed that an

¹See <http://bea.gov/bea>. The PCE price index for energy goods includes gasoline and motor fuels, heating oil, natural gas, and electricity

²See <http://tonto.eia.doe.gov>.

³For example, on July 27, 2006 the New York Times in reference to the Fed's Beige Book reported that "consumers were spending less in stores and auto showrooms in part because of high gasoline prices" (see Peters 2006). Similarly, a July 15, 2006 New York Times article, quoting Anthony Chan, the chief economist for J.P. Morgan Private Client Services, announced that "... the effects of rising energy prices are starting to gain some traction on the consumer spending front" (see New York Times 2006). When energy prices retreated sharply later in 2006, driven primarily by falling gasoline prices, the Wall Street Journal predicted that "lower gasoline prices should boost the annual growth rate of consumer spending a full percentage point and could lift fourth-quarter economic growth from a forecast of 3% at an annual rate, to as high as 3.7%" (see Ip 2006).

⁴For example, Nariman Behraves, chief economist of Global Insight Inc, in a January 2007 NPR interview on the effect of falling oil prices expressed the view that consumers whose "discretionary spending" had been squeezed by high energy prices, would now increase their discretionary spending (see Behraves 2007).

⁵In related work, Mehra and Peterson (2005) concluded that increases in retail oil and gasoline prices significantly decrease quarterly consumption growth over the 1961-2004 period, while price declines have insignificant effects. Cullen, Friedberg, and Wolfram (2004) investigated whether households face a "heat or eat" decision when confronted with rising home heating costs. The authors use household-level panel data to determine how consumption of nondurable goods and food responds to anticipated and unanticipated changes in home energy costs.

increase in energy prices slows economic growth primarily through its effects on consumer spending.

This paper studies the effect of energy price changes on consumer spending. There are four mechanisms by which consumer 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. This *precautionary savings effect* does not appear to have been discussed in the literature, but will play an important role in our analysis. Finally, consumption of durables that are complementary in use with energy (in that their operation requires energy) will 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 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 purchasing power) can have large effects on output and unemployment (see Hamilton 1988). A similar reallocation may occur within the same sector, as

⁶Implicit 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 transferred abroad and is not recycled in the form of higher U.S. exports. In the case of a purely domestic energy price shock (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.

⁷This last effect need not involve a reduction in the number of automobiles sold. It can also take the form of consumers switching from large energy-inefficient automobiles to small energy-efficient 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).

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 purchasing power losses 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 policymakers with the effects of energy price shocks on the automobile sector (see, e.g., Bernanke 2006b).

In this paper, we provide a comprehensive look at the evidence for these channels of transmission based on a detailed analysis of BEA data on personal consumption expenditures (PCE), Michigan Survey of Consumers data on consumer expectations, and the aggregate unemployment rate reported by the Bureau of Labor Statistics. Our analysis is complicated by the fact that some of the mechanisms described above affect consumption symmetrically, whereas other effects are asymmetric in purchasing power losses and gains. In fact, a testable implication of the leading theoretical models of how energy prices are transmitted to the economy is that the response to an energy price decrease should differ from the response to an energy price increase (see, e.g., Bernanke 1983; Hamilton 1988, 2005).⁸ Thus, evidence of symmetry in the responses to energy price changes would be inconsistent with the presence of either an uncertainty effect or a reallocation effect.

Our empirical analysis therefore is conducted in two stages. We first assess the evidence for asymmetries in the data. Having been unable to reject the assumption of symmetry, we then proceed with a model that imposes symmetry on the responses to unexpected changes in purchasing power. Having ruled out the uncertainty and reallocation effect, our empirical strategy is to identify the remaining three effects described above by quantifying the differential response of major components of real consumption to unpredictable changes in purchasing power driven by energy price fluctuations. Finally, we demonstrate that the model is capable of explaining key episodes in the data without introducing asymmetries.

The remainder of the paper is organized as follows. In section 2, we document how energy price fluctuations have affected consumers' purchasing power since 1970. Our analysis incorporates not only gasoline expenditures, but all forms of energy consumption including natural gas, electricity and heating oil. We construct time series of the expenditure shares for each form of energy and of real energy prices. Using these data, we compute a monthly time series of the losses and gains

⁸Asymmetries also have played a central role in empirical work on the effect of energy price shocks since the late 1980s (see, e.g., Mork 1989; Hooker 1996a,b; Hamilton 1996; Huntington 1998; Davis and Haltiwanger 2001; Lee and Ni 2002; Hooker 2002; Balke, Brown, and Yücel 2002).

in purchasing power associated with changes in energy prices. The measure indicates the required change in discretionary spending if households wish to purchase last month's quantity of energy goods at the current month's prices.

We show that this effect in general is ambiguous. In section 3, we investigate the evidence for asymmetries in the responses of real consumption and its major components to unanticipated changes in purchasing power driven by fluctuating energy prices. We show that the null hypothesis cannot be rejected that the estimated responses are symmetric. Our results are robust to differentiating between small and large shocks. They also are robust to using net increases and net decreases in purchasing power in the spirit of work by Hamilton (2003). While the evidence against asymmetries in real consumption responses is subject to considerable sampling uncertainty in some cases, we find strong evidence against asymmetries in the responses of expectations data from the Michigan Survey of Consumers to the same purchasing power shocks. This evidence suggests that the reallocation effect, the uncertainty effect, and any asymmetry associated with the precautionary saving effect are not a dominant feature of the real consumption data, contrary to the prevailing view. It 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, our 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 we show that they are not.

Based on this evidence, in section 4 we use a VAR model that imposes symmetry to investigate how unanticipated changes in purchasing power driven by fluctuating energy prices affect aggregate real consumption and its major components. We quantify the magnitude of the effect of changes in discretionary income and changes in precautionary savings on all forms of consumption based on estimates of the short-run price elasticity of energy demand. We also quantify the effect of changes in operating costs on the consumption of energy-using durables. We estimate that a one-time 1% increase in energy prices in a given month is associated with a decline in real total consumption of -0.15% after one year. Our results show that the decline in total consumption as well as the decline in the consumption of nondurables, services and durables excluding vehicles is larger than would be expected given the energy share in consumption and given plausible estimates of the short-run elasticity of energy demand, and large enough to reject the hypothesis that there is no effect on total consumption beyond the reduction in discretionary income. Given our estimates of the elasticity of energy demand, the discretionary income effect of a 1% increase in energy prices is at most -0.04%. This implies a marginally statistically insignificant precautionary savings effect of -

0.08% for nondurables and a statistically significant effect of -0.07% for services. The corresponding precautionary savings effect of -0.16% on durables is not statistically significant. The operating cost effect on vehicles consumption lies between -0.60% and -0.65% and is statistically significant. Expenditures on vehicles decline about four times as much as expenditures on other consumer durables. They decline more than seven times as much as expenditures on nondurables and services.

In section 5 we summarize our estimates of the direct effects, gauge the overall importance of energy price shocks for real consumption, and discuss implications of our results for the way economists think about the transmission of energy price shocks. Our results differ from commonly held beliefs about how energy price shocks are transmitted to the economy. In particular, the presence of a strong reallocation effect has been considered crucial for explaining how energy price shocks may have effects on the economy that are disproportionate to the expenditure share for energy goods. While we find compelling quantitative evidence in section 5.1 that expenditures on new domestic automobiles decline disproportionately compared to new foreign automobiles and that consumers substitute cars for light trucks (such as SUVs, minivans, or pickup trucks), consistent with the mechanism described in Hamilton (1988) and Bresnahan and Ramey (1993), our results do not support the view that there are asymmetries in the aggregate real consumption responses of the type that would be consistent with the presence of a reallocation effect. A likely explanation is that the share of the U.S. automobile industry in U.S. real value added and employment is only about 1%.

There are other reasons, however, for the popularity of models that embody allocative effects. Theoretical models that incorporate a reallocation effect are capable in principle of explaining why rising energy prices are unambiguously bad news for the economy, whereas falling energy prices may have little or no effect on the economy. This model feature is thought to be required to explain the apparent absence of a strong economic expansion following the collapse of OPEC in late 1985. Thus, the fact that we do not find evidence of a reallocation effect in the data may seem to cast doubt on our analysis. We show that this doubt is unwarranted. Using historical decompositions we analyze how consumption would have evolved in the absence of the 1986 decline in energy prices, and demonstrate in section 5.2. that there actually was an expansion in consumption without which actual economic growth would have been much weaker in 1986, that this expansion was similar in magnitude to the decline in real consumption in 1979, and that this expansion is predicted by standard symmetric models. While there indeed is an asymmetry in the real GDP growth data for 1979 and 1986 (as opposed to the real consumption growth data), as shown in section 5.3, that asymmetry appears to be driven by a decline in business investment in 1986 that was not related to the fall in energy prices, but to the 1986 Tax Reform Act. This effect was exacerbated by the

response of investment by the petroleum and natural gas industry to the collapse of OPEC in late 1985, which (for reasons detailed in the paper) is asymmetric, but in the opposite direction from the asymmetries previously discussed in the literature.

It has been widely observed that energy price shocks do not appear to affect the U.S. economy as much as we used to think. In section 5.5., we document the declining importance of energy price shocks for consumption compared to the 1970s and early 1980s and show that this decline is consistent with the changing structure of the U.S. automobile industry and its declining weight in aggregate employment and value added. The conclusions are discussed in Section 6.

2 The Effects of Changing Retail Energy Prices on Consumers' Purchasing Power

The quantitative importance of changes in energy prices for discretionary income is by no means self-evident. For one thing, the degree to which rising energy prices affect household purchasing power depends not only on how much energy prices increase, but also on the fraction of consumer expenditures devoted to energy. In this section, we construct a measure of the gains and losses in purchasing power attributable to fluctuating energy prices. Motivated by the quotes in the introduction, the thought experiment underlying this measure is that a household, when faced with rising energy prices, is unable to reduce its energy consumption in the short run, resulting in higher energy expenditures. The resulting loss in “purchasing power” causes households to curtail their discretionary expenditures on non-energy goods and services.

Let E be the quantity of energy goods consumed, N the quantity of non-energy goods consumed, P^E the price of energy goods, P^N the price of non-energy goods, and P^{PCE} the price index for personal consumption expenditures. Real energy consumption in consumption units at date t is given by:

$$e_t \equiv \frac{E_t \cdot P_t^E}{P_t^{PCE}}$$

Real total consumption at date t is the sum of real energy consumption and real non-energy consumption:

$$c_t \equiv \frac{E_t \cdot P_t^E}{P_t^{PCE}} + \frac{N_t \cdot P_t^N}{P_t^{PCE}}$$

This leaves the household with $c_t - e_t$ to spend on non-energy goods at date t .

Suppose that at date $t + 1$, the consumer requires the same quantity of energy as at date t , but

now pays P_{t+1}^E . Then the consumer's real energy consumption would be:

$$e_{t+1} \equiv \frac{E_t \cdot P_{t+1}^E}{P_{t+1}^{PCE}}$$

and he would have $c_t - e_{t+1}$ left to spend after paying his energy bill. The implied percent change in purchasing power between dates t and $t + 1$ is:

$$\frac{(c_t - e_{t+1}) - (c_t - e_t)}{c_t - e_t} = \frac{e_t - e_{t+1}}{c_t - e_t} = \frac{E_t \cdot P_t^E / P_t^{PCE} - E_t \cdot P_{t+1}^E / P_{t+1}^{PCE}}{N_t \cdot P_t^N / P_t^{PCE}}$$

Multiplying this expression by P_t^E / P_t^E and rearranging terms yields:

$$- \left(\frac{E_t \cdot P_t^E}{N_t \cdot P_t^N} \right) \cdot \left(\frac{P_{t+1}^E / P_{t+1}^{PCE} - P_t^E / P_t^{PCE}}{P_t^E / P_t^{PCE}} \right)$$

The first term in this expression is the ratio of nominal energy expenditures to nominal non-energy expenditures. Since the fraction of energy expenditures in total expenditures is small, we can approximate this ratio with the nominal energy expenditure share. The second term is the monthly percent change in the real price of energy goods. Thus, the percent change in purchasing power approximately equals the product of the nominal expenditure share and the percent rate of change in real energy prices:

$$-\eta_t^E \cdot \% \Delta (P_{t+1}^E / P_{t+1}^{PCE}),$$

where

$$\eta_t^E \equiv \frac{E_t \cdot P_t^E}{C_t \cdot P_t^{PCE}}.$$

To illustrate our approach, consider the following numerical example: Suppose that at time t , 3% of the average household's expenditures are devoted to energy consumption. In other words, the average household's nominal expenditure share, η_t^E , is equal to 3% and its expenditure share for non-energy goods is 97%. Suppose that at time $t + 1$ the real price of energy increases by 10%. If the household wishes to consume the same quantity of energy as in time t , the energy expenditure share will rise to 3.3%, and the non-energy expenditure share will fall to 96.7%. In other words, the household's discretionary income falls by 0.3 percentage points.

The BEA's PCE price index for energy goods is comprised of four main components: gasoline (and other motor fuels), natural gas, electricity and all other energy goods (including heating oil, coal and oil lubricants). Figure 1 shows the monthly real price indices for gasoline, natural gas, and electricity as well as an aggregate energy price index including all items. The sample period is

January 1970 - July 2006. All four indices rose from the early 1970s until the early 1980s, declined from the early 1980s until the late 1990s, and (notwithstanding some temporary reversals) have been on the rise again since 1999. Although the four energy price series share the same trend, there are important differences. For example, natural gas prices rose more gradually in the 1970s than gasoline and electricity prices because the residential market for natural gas was heavily regulated until 1979. Deregulation of this market was complete only in 1989 (see Davis and Kilian 2007). There are also price spikes in some energy markets that are not shared by other markets. For example, gasoline prices rose sharply during the Persian Gulf War of 1990/91, whereas electricity and natural gas prices did not. Similarly, there was a large spike in natural gas prices in 2000 that far exceeded changes in other energy prices. Even when the direction of the price change is the same, the magnitudes may differ. For example, between 1970 and the early 1980s, gasoline prices and natural gas prices almost doubled. In contrast, electricity prices increased only by about one third.⁹

Figure 2 shows the monthly nominal expenditure share for gasoline, electricity, and natural gas as well as the aggregate expenditure share for energy. The discrepancy between the aggregate share and these three components is the share of other energy goods such as heating oil and lubricants. Figure 2 illustrates that, despite the importance attached to energy prices in the press, the energy share in consumption has never been large compared to the expenditure share of food or housing, for example. The overall energy share was stable at about 6.5% in the early 1970s. It rose to a peak of 9.6% in 1980, but fell steadily throughout the 1980s and 1990, reaching a low of 4.1% in 1999, only to rise back to its initial level of about 6.5% by 2006. The nominal expenditure share for gasoline, which constitutes the largest individual component of real energy spending, follows a similar historical pattern as the overall energy share, whereas the shares of natural gas and of electricity are much more stable.

The three upper panels of Figure 3 show the corresponding changes in purchasing power associated with gasoline, electricity, and natural gas prices, respectively. The bottom panel shows the corresponding results for the aggregate energy price. A negative value in Figure 3 denotes an increase in prices and thus a decline in purchasing power. It means that consuming the previous month's quantity of energy at the current month's price would leave households with less money to spend on all other goods. A positive value indicates that purchasing power has increased.

While gasoline prices tend to occupy the headlines and are foremost on the minds of consumers, changing prices for other energy goods can affect consumers' purchasing power as well. For example,

⁹A striking feature of Figure 1 is that real gasoline prices in 2005 and 2006 reached their highest levels ever. This result is in sharp contrast to the data for real crude oil prices reported in Kilian (2006). The key difference is the shortage of crude oil refining capacity following Hurricane Katrina, which sharply drove up U.S. gasoline prices, while slightly lowering world crude oil prices.

an increase in gasoline prices might be accompanied by a decline in natural gas or electricity prices, at least partially offsetting the decline in purchasing power. Figure 3 shows that changes in gasoline prices, nevertheless, tend to have a disproportionate impact on household purchasing power. The contemporaneous correlation of the purchasing power loss series for energy in the last panel with the corresponding series for gasoline in the first panel is 97%. This result is consistent with the relatively high amplitude of changes in gasoline prices in Figure 1 and the relatively high expenditure share for gasoline in Figure 2. Because natural gas and electricity each comprise a smaller fraction of total expenditures, and because their prices are more stable, their impact on purchasing power is much smaller.

In the remainder of the paper, we will focus on changes in purchasing power associated with energy prices, as shown in the bottom panel. Increases in energy prices left households with less money to spend on other goods in about 46% of the months in the sample. In the mid-1990s, most changes in purchasing power were in the range between plus and minus 0.1 percentage points. Since the late 1990s, when rising global demand for oil and constraints in refining capacity began pushing up gasoline prices, however, the volatility and amplitude of purchasing power losses has greatly increased. The largest monthly loss in purchasing power (-0.66 percentage points) is associated with Hurricanes Rita and Katrina in late 2005, and the largest monthly gains in purchasing power were 0.46 percentage points in March of 1986 (associated with the collapse of OPEC in late 1985) and 0.54 percentage points in November 2005 in the wake of Hurricanes Rita and Katrina.

It is useful to put the changes in purchasing power documented in Figure 3 in perspective. A simple example helps assess the economic significance of these fluctuations. In February 2003, the average household's nominal expenditure share for energy was 5.20%. In March, energy prices rose by 4.89%, reducing households' purchasing power by about 0.25 percentage points, and the nominal expenditure share fell to 5.17%. As energy prices fell by 5.64% in April, households' purchasing power increased by 0.29 percentage points. Since the average household's annual expenditure in 2003 was \$40,817, or \$3,401 per month (see Bureau of Labor Statistics 2005), in effect, the rise in energy prices left households with \$8.50 less to spend on non-energy goods in March. The subsequent decline in energy prices in April slightly more than reversed this loss. This exercise demonstrates that an increase or decrease in monthly energy prices has only a small direct impact on a household's resources available for consuming other goods and services. The main reason is that the relatively low expenditure share for energy (6.32% on average) blunts the impact of even sharp monthly swings in energy prices.

Our measure of purchasing power losses abstracts from all margins of adjustment that might reduce the impact effect of energy price changes. In reality, the quantity of energy demanded

is unlikely to be completely inelastic. For example, a household could take one big trip to the grocery store instead of several small trips, switch to energy-efficient light bulbs, or simply adjust the thermostat in the house. Furthermore, households have the option of drawing down their savings (or of borrowing) in order to smooth their consumption when energy prices rise temporarily. Our empirical methodology in the sections below is designed to incorporate such endogenous responses. Moreover, we will provide direct estimates of the price elasticity of energy consumption in section 4.

The fact that the purchasing power losses and gains induced by energy price fluctuations are small, even under the extreme assumption of perfectly inelastic demand for energy, does not necessarily mean that they cannot have large effects on real consumption. For example, changes in precautionary savings can greatly amplify the response to an unanticipated purchasing power loss. In addition, changes in uncertainty may trigger a disproportionate response of durables consumption. Likewise, changes in the operating cost of energy-using durables such as vehicles may further amplify the overall response of real consumption. Finally, changing expenditure patterns may trigger a reallocation effect that further amplifies the response to a purchasing power loss. The next two sections will investigate the empirical support for these effects.

3 Quantifying the Effects of Purchasing Power Shocks on Real Consumption

In this section, we will present estimates of impulse responses based on bivariate vector autoregressive (VAR) models. Our analysis is intended to shed light on the empirical relevance of various channels of transmission that play an important role in theoretical work on energy price shocks. As discussed in the introduction, there are five distinct channels that might rationalize a response of real consumption to purchasing power shocks. Our objective in this paper is to assess the evidence for these effects and to quantify them when possible. Our empirical strategy is to identify the individual effects by estimating the differential responses of major components of real consumption to unpredictable changes in purchasing power driven by energy price fluctuations. We are interested in whether there is compelling statistical evidence for the existence of these effects and in how large these effects are in the data.

Our main focus in this section is on aggregate real consumption and its major components (durables, nondurables, and services). We fit separate VAR models for each consumption aggregate. Durables are further disaggregated into vehicles (defined to include automobiles, motorcycles, recre-

ational vehicles, aircraft and boats) and other durables.¹⁰ This distinction is important for assessing the operating cost effect. While the use of many durables requires some energy input, our rationale is that vehicles are much more energy intensive than, say, appliances, and that effectively operating costs will not matter much for durables other than vehicles. More disaggregated results on vehicles will be presented in section 5.

We focus on measures of purchasing power changes, as opposed to unweighted changes in energy prices, because the purchasing power changes incorporate the effects of changes in the expenditure share for energy. This allows us to avoid a potential source of structural instability in the relationship between energy prices and the economy. Each VAR model includes the purchasing power loss series described in the previous section (or a suitable transformation of that series) and the percent growth rate of the measure of real consumption of interest. The sample period is 1970.2-2006.7, unless noted otherwise. The monthly real consumption data are from the BEA's National Income and Product Accounts. Throughout the paper, we impose a VAR lag order of 6. This lag order tends to be larger than the estimates suggested by the Akaike Information Criterion conditional on an upper bound of 12 lags, which in some cases produces implausibly low lag order estimates. While our qualitative results are not sensitive to the lag order choice, given the well-known dangers of underfitting a VAR model, we adopt a conservative approach (see Kilian 2001, Ivanov and Kilian 2005).

The VAR models are identified recursively with the purchasing power loss series ordered first, implying that its innovations are not affected contemporaneously by innovations to real consumption growth.¹¹ The advantage of using a VAR model is that it isolates the linearly unpredictable component of losses in purchasing power and allows for reverse causality.¹² In the figures below, we show responses of the level of real consumption to a one-time, one standard-deviation purchasing-power shock. The maximum horizon of the impulse response functions is 18 months. All figures show point estimates as well as 90% bootstrap confidence intervals based on the bias-corrected method of Kilian (1998). Since the dominant autoregressive roots of the VAR models are near 0.8 throughout, the coverage accuracy of these intervals is likely to be very high.¹³

¹⁰In computing real consumption of other durables we make suitable adjustments to account for the fact that the BEA data are based on chain-weighted price indices.

¹¹Our approach does not allow us to differentiate between energy price changes driven by demand and by supply shocks in energy markets. The distinction between demand and supply shocks can be important, as these shocks tend to have very different effects on macroeconomic aggregates (see Kilian 2006). The impulse responses shown below hence are best viewed as an average reflecting the (unknown) composition of demand and supply shocks over our sample period.

¹²Our approach differs from that in Mehra and Peterson (2005) who treat the change in energy prices as exogenous. This assumption is at odds with historical and econometric evidence (see, e.g., Barsky and Kilian 2004, Hamilton 2005, Kilian 2006)

¹³Under the assumption that purchasing power changes driven by energy price changes are predetermined, there is no loss in generality in restricting ourselves to bivariate models. Adding more variables to the VAR model would only undermine the efficiency of the estimates, while creating identification problems, since there is no reason to believe that the remaining model structure is recursive.

3.1 On the Source of Asymmetries in the Response of Real Consumption

Our empirical analysis is complicated by the fact that some of the effects listed in the introduction are symmetric in purchasing power losses and gains, whereas other effects are asymmetric. It is useful to review the main arguments, before discussing the empirical specification. The effect of changes in discretionary income on real consumption is clearly symmetric in purchasing power increases and decreases. All else equal, consumers will tend to lower their discretionary expenditures when energy prices rise, and increase their expenditures when energy prices fall.

Similarly, the effect of changes in operating costs on the consumption of energy-using durables is likely to be symmetric. This point deserves some elaboration. Hamilton (2005) makes a case that consumers may postpone car purchases, when the price of energy rises, but will not buy a second car, when energy prices go down. This conclusion is not obvious in that there are likely to be consumers who used to think that purchasing their first (or second) car was beyond their means, and who may elect to buy a car after all, following a decline in energy prices. More importantly, even if Hamilton's conclusion were correct, it misses the point that consumers will tend to replace their existing car with a new and less energy-efficient car. For example, they may trade in their Honda Civic for an SUV. This decision is the direct mirror image of the decision to trade in one's SUV for a Honda Civic in response to higher energy prices. Thus, on a priori grounds, there is no reason to expect an asymmetry in the response of real consumption of motor vehicles.

In contrast, the uncertainty effect on the consumption of durables is asymmetric. In its strongest form, the uncertainty effect implies that unexpected increases as well as unexpected decreases of energy prices will cause a decline in real durables consumption, as households postpone major purchases, while they wait to see whether the energy price change will persist (see Bernanke 1983). Since this effect reinforces the discretionary income effect, when energy prices increase, but counteracts the discretionary income effect when energy prices decrease, it causes the responses of consumer durables to be asymmetric. A somewhat weaker form of the uncertainty hypothesis recognizes that consumers may be more risk averse in the direction of higher energy prices. In that view, when energy prices unexpectedly increase, households will postpone major purchases of durables that were already planned. When energy prices decline unexpectedly, however, there is no reason to postpone those planned purchases and real durables consumption does not respond to the energy price decline. For our purposes this distinction is immaterial. Either way, the response to an energy price increase should be larger in magnitude than the response to a decrease, causing the responses of real durables consumption to be asymmetric.

The precautionary savings effect could be symmetric or asymmetric. One view is that consumers

increase precautionary savings (and hence decrease real consumption) in response to all changes in energy prices, simply because changing energy prices create uncertainty. For example, energy expert Daniel Yergin in a recent interview reported results of a study that shows that, even though gasoline prices have dropped substantially from their highs in 2006, there is “a greater sense of insecurity and people don’t want to be caught emptying their wallet at the gasoline pump” (see Douglass 2007). This contrasts with the view that consumers in practice view falling energy prices as unambiguously good news, while energy price increases are considered a reason for concern about the future. In the latter case, one would expect a reduction in precautionary savings and an increase in real consumption in response to an unexpected decline in energy prices that mirrors the response to an unexpected energy price increase. Our empirical analysis will allow for both of these possibilities.

Finally, the reallocation of resources in response to changing expenditure patterns will cause an asymmetric response to energy price increases and decreases. This hypothesis is based on the observation that any change in the price of energy, whether it is an increase or a decrease, will cause reallocations across industries, which in the presence of frictions in capital and labor markets will cause falling output and rising unemployment. However, a fall in energy prices also raises the discretionary income of households and stimulates output and employment. Thus, the positive income effects of such a price decline will offset at least in part the negative reallocation effects. In contrast, when the price of energy increases, both effects work in the same direction, amplifying the response of output and unemployment (see, e.g. Davis and Haltiwanger 2001).

It is essential for our empirical analysis to allow for these potential asymmetries in the response of real consumption to energy price shocks. Clearly, the standard linear regression framework that treats energy price increases and energy price decreases symmetrically will only be appropriate, if we can rule out the presence of an allocative channel of transmission and of the asymmetric effects arising from uncertainty. In the remainder of section 3, we will show that there is no compelling evidence against the symmetry hypothesis, consistent with the absence of both the uncertainty effect and the reallocation effect in the real consumption data. In section 4, we will estimate the effect of changes in discretionary income, changes in precautionary savings and changes in operating costs, conditional on having ruled out all asymmetric effects.

3.2 Is There Evidence of a Reallocation Effect or an Uncertainty Effect?

The presence of a reallocation effect of energy price changes as well as the presence of an uncertainty effect in durables consumption may be detected based on the pattern of the responses of real consumption expenditures and their major components to unanticipated purchasing power gains and

losses.

Hypothesis 1: The fact that the reallocation effect will lower all forms of consumption whether energy prices increase or decrease implies that the response of real consumption aggregates to an increase in energy prices should be at least as large in absolute terms as the response to a decrease in energy prices of the same magnitude. The absence of such an asymmetric effect of energy price changes on *real consumption of nondurables and services* is evidence against the presence of a reallocation effect.

For the real consumption of durables, the analysis is complicated by the fact that there potentially is a second asymmetric effect at work.

Hypothesis 2: Uncertainty will increase whether energy prices increase or decrease. The effect of increased uncertainty on the consumption of durables is negative. Thus, the absence of an asymmetric effect of energy price changes on *real consumption of durables* is evidence against the existence of both the reallocation effect and the uncertainty effect.

By implication, an asymmetry of the hypothesized form only in the response of *real consumption of durables*, but not in the response of *real consumption of nondurables and services*, would be evidence of the presence of an uncertainty effect, but the absence of a reallocation effect.

In investigating these two hypotheses, a natural starting point is to divide purchasing power shocks into negative and positive changes, as shown in the upper panel of Figure 4. Implicit in this specification is the assumption that energy price changes have a proportionate effect on consumption. We will refer to this specification as the baseline model. The assumption of proportionate effects will be relaxed further below.

3.2.1 Baseline Model: All Purchasing Power Changes Matter

For each consumption aggregate of interest, we estimate impulse response functions based on two bivariate VAR models, one for purchasing power losses and one for purchasing power gains. The underlying shocks are scaled to be identical in magnitude, so the impulse response function can be compared across purchasing power losses and gains. Thus, if symmetry holds, we would expect the response to a purchasing power loss to be the exact mirror image of the response to a purchasing power gain, except for sampling error.¹⁴

¹⁴An alternative approach would have been to fit a regression model that involves lags of both increases and decreases and to test the equality of the coefficients at a given lag by means of a Wald test. In related work, many authors have included oil-price increases and oil-price decreases as separate variables in a single-equation model for real output growth and performed a Wald test of the equality of the coefficients (see, e.g., Mork 1989, Dotsey and

Real Consumption The first two columns of Figure 5 contrast the responses of major real consumption aggregates to unexpected losses and unexpected gains in purchasing power. First, for total consumption as well as for each of its three major components (durables, nondurables, services) the point estimates are negative in response to a purchasing power loss (or energy price increase) and positive in response to a purchasing power gain (or energy price decrease). Second, the estimated responses to a purchasing power loss are systematically larger in absolute terms than the responses to a purchasing power gain. Third, while the responses to a purchasing power loss are highly statistically significant, the responses to purchasing power gains are invariably much less precisely estimated. The 90% confidence intervals in most cases include values that would be expected if the response were symmetric.

The next two rows of Figure 5 show the corresponding responses for durables disaggregated into vehicles (defined as automobiles, pleasure boats, pleasure aircraft, motorcycles, and recreational vehicles) and other durables. The response of real vehicles consumption to losses is larger in magnitude and more statistically significant than that of real durables consumption, but the qualitative results are similar. In the case of other durables, we find that the estimated response to a purchasing power gain is larger than the response to a loss, contrary to the implications of theoretical models. Moreover, the response to a purchasing power loss is small and not statistically significant, whereas the response to a purchasing power gain is statistically significant at most horizons.

What inspection of these plots does not reveal is whether the differences in the impulse response point estimates are statistically significant. The first column of Table 1 provides p -values for Wald tests of the symmetry of the response functions for each of these aggregates. These tests are computed based on a modification of the residual-based bootstrap method of Kilian (1998) that takes proper account of the fact that the underlying VAR models are only seemingly unrelated and that preserve the contemporaneous correlations in the data across models. Table 1 shows that, in no case, do we reject the null of symmetry. All but one p -value is above 0.85 and in the one case where the p -value is only 0.61, as discussed earlier, inspection of Figure 5 reveals departures from the null of symmetry in the wrong direction in that the response to a purchasing power gain is larger than the response to a loss, contrary to the implications of standard theoretical models. While a non-rejection of the symmetry null does not establish that the null is true, Table 1 shows that there is no compelling reason to depart from the standard model imposing symmetry.

Reid 1992, Hooker 1996a, Hooker 2002). A drawback of this approach is that such tests only tell us whether the estimates of the regression slope parameters are different, whereas we really are interested in the extent to which the estimated impulse responses differ. We avoid this ambiguity by focusing directly on the differences in the impulse response estimates with respect to unexpected increases and decreases in purchasing power.

Unemployment An additional plausibility check of the results for real consumption is provided by the response of the U.S. aggregate unemployment rate to the same purchasing power shocks. In the presence of a reallocation effect in particular, one would expect a strong degree of asymmetry in that response. We investigate this conjecture based on the same VAR framework. As before, purchasing power shocks are treated as predetermined. The upper panel of Figure 6 shows a statistically significant positive response to purchasing power losses, whereas the response to a purchasing power gain is flat and surrounded by very wide confidence bands. While the point estimates are seemingly asymmetric, the first column of Table 1 shows a p -value of essentially 1 for the null of symmetric responses. Thus, the unemployment data are even less informative than the real consumption data.

Consumer Expectations The evidence presented so far is open to interpretation. While there is weak evidence of asymmetries in the impulse response point estimates of real consumption aggregates, formal statistical tests indicate that the data are essentially uninformative. We cannot tell whether the nonrejection of the symmetry null hypothesis occurs because the null is true, or whether the test may have low power to detect departures from symmetry. Nevertheless, our inability to reject the symmetry null is surprising, given the consensus in the literature that asymmetries are important. One way of assessing whether the non-rejections for real consumption and aggregate unemployment are due to low power, is to compare these results to similar response estimates for consumer expectations data from the Michigan Survey of Consumers. If the responses of real consumption and of unemployment were truly asymmetric, one would expect to see a similar asymmetry in consumer expectations data.

Our data set includes the overall index of consumer sentiment as well as individual series. One set of expectations data relates to households' perceptions about the future evolution of the economy and includes expected changes in business conditions and expected changes in unemployment. Another set of measures relates to households' precautionary savings motives and includes expected changes in households' personal financial situation and expected changes in real family income. A third set of expectations measures relates directly to households' decisions about major durables purchases and includes current buying conditions for large household goods and current buying condition for vehicles.

All series are constructed such that a fall in the index indicates a worsening of conditions from the consumer's point of view. The sample period is 1978.1-2006.5. No monthly expectations data are available prior to 1978. While the scale of the responses is not comparable, given the way survey responses are represented, the qualitative patterns and degree of symmetry of the responses are. Our identifying assumption throughout this paper is that purchasing power innovations are

predetermined with respect to innovations in consumer sentiment within the month.

The use of consumer expectations data provides important additional insights. The first two columns of Figure 7 show that the responses of overall consumer sentiment are highly symmetric in purchasing power gains and losses. The overall shape of the response function as well as the scale of the responses are very similar, and both response functions are highly significant. The same pattern of results is found for most of the relevant disaggregates in the Michigan Survey. In most cases, the responses look highly symmetric and are statistically significant. In some cases, the responses to purchasing power gains actually are larger in absolute terms than the responses to purchasing power losses rather than smaller, as predicted by theoretical models (see, e.g., current buying conditions for vehicles). The only result that might potentially be consistent with an asymmetric reaction is the response of expected changes in interest rates. While households expect interest rates to increase temporarily in the first few months following a purchasing power loss, the response to a purchasing power gain after the expected initial decrease of interest rates suggests a persistent increase. While puzzling, this persistent response is not statistically different from zero.

The first column of Table 1 provides p -values for Wald tests of the symmetry of the response functions for each of these expectations measures. In no case, do we reject the null of symmetry. All but one p -value is above 0.97. Even for the response of expected changes in interest rates the p -value for the symmetry null is essentially 1.

We conclude that there is no reason to depart from standard linear models that impose symmetry in the responses to purchasing power losses and gains. There is no compelling evidence for the reallocation effect (Hypothesis 1) or the uncertainty effect (Hypothesis 2). That conclusion, however, is conditional on the premise that consumers respond to all changes in the price of energy equally. There are two alternative models of consumer behavior that suggest that this premise may be unrealistic. Below we will investigate both of these models and show that our conclusions about the symmetry of the responses are robust.

3.2.2 Alternative Model 1: Only Large Purchasing Power Changes Matter

The baseline model postulated that consumers' responses are proportionate to the magnitude of the shock. An alternative hypothesis is that consumers only respond to large purchasing power shocks. For example, Macintyre (2006) suggests that a gas price increase of 25 cents (an increase of about 10% in gas prices as of 2006) would make consumers angry. It seems plausible that a shock of only 4 cents (an increase of about 1.5%), which roughly corresponds to the magnitude of the shocks studied in the earlier VAR analysis in this paper, might not evoke a reaction from consumers. In fact, it might go unnoticed by many households. Such behavior might be rationalized by adjustment costs.

The presence of costs to monitoring energy costs and of adjusting consumption patterns might make households reluctant to respond to small changes in purchasing power (see Goldberg 1998).

In this subsection, we investigate the possibility that real consumption, the unemployment rate and consumer expectations may react differently to purchasing-power shocks of different signs and magnitudes. We divide the observations for the purchasing power loss series into four mutually-exclusive categories: large purchasing power losses, small purchasing power losses, large purchasing power gains, small purchasing power gains. A monthly change in purchasing power is considered *large* if it is greater than 0.11%, which corresponds to one standard deviation of the purchasing power loss series. We find that 14.2% of the purchasing power losses involve large losses, 31.7% involve small losses, 47.3% involve small gains and 6.9% large gains. The time series of large purchasing power increases and large purchasing power decreases are plotted in the middle panel of Figure 4.

Columns 3 and 4 of Figure 5 show the estimated impulse responses of each of the six real consumption categories to large losses and large gains in purchasing power. The reduced form correlations are similar to those in the baseline model. All shocks have been normalized in absolute terms to have the same scale, so differences across the estimated responses are entirely due to differences in slope parameter estimates. With the exception of services, responses to large gains and large losses are of opposite signs and similar magnitudes. As in the baseline model, while the point estimates of the responses to large gains tend to be somewhat smaller than the estimated responses to large losses, they are less precisely estimated. The corresponding results for the U.S. unemployment rate are shown in the middle panel of Figure 6. The large and statistically significant increase in the unemployment rate in response to a large loss contrasts with a small and statistically insignificant decline in the unemployment rate in response to a large gain in purchasing power.

Overall, the evidence for this alternative model leaves room for doubt about the presence of asymmetries in the consumption data, not unlike in the baseline model. A more direct measure of the degree of asymmetry is again provided by the responses of consumer expectations data to the same shocks, as shown in columns 3 and 4 of Figure 7. As in the baseline model, the estimated responses for the Michigan index of consumer sentiment show no evidence of the asymmetry one might have expected based on theory. For most expectations indicators shown in Figure 7 there is a large degree of symmetry between responses to large losses and to large gains, and to the extent that there are departures from this pattern, there are as likely to be in the direction of larger responses to large gains as to large losses.

Formal statistical tests based on these responses allow us to address two distinct questions. The first question is whether there is evidence against the assumption that the responses to large purchasing power gains and large purchasing power losses are symmetric. Column 2 of Table 1

provide evidence that the symmetry hypothesis cannot be rejected in any case. The p -values range between 0.68 and 1.00 for real consumption aggregates (with the exception of other durables, similar to the baseline model) and between 0.79 and 1.00 for consumer expectations. Similarly, the p -value for the symmetry of the unemployment responses is 0.98. Hence, the earlier conclusions based on the baseline model are supported even if we relax the assumption of scale invariance.

The second question we can address is whether there is empirical support for the notion that only large price changes matter. Since the latter model can be nested in the baseline model, we can address this question by testing the equality of the impulse response functions for large and small losses on the one hand and for large and small gains on the other hand, controlling for the size of the shock. Columns 4 and 5 of Table 1 shows that we cannot reject the equality null for the responses to large and small losses. The p -values range from 0.43 to 1.00. For gains, they range from 0.95 to 1.00. Similar results hold for the unemployment rate and all measures of consumer expectations we considered. It is not the case that consumers ignore small changes in purchasing power. Thus, there is no compelling reason not to impose scale invariance as in the baseline model. The evidence against the equality null becomes even weaker if we impose symmetry before testing the equality of large and small changes (see last column of Table 1).

3.2.3 Alternative Model 2: Only Net Purchasing Power Changes Matter

We have already considered the possibilities that consumers respond proportionately to purchasing power shocks or that they respond only to large changes. Yet another view is that consumers only respond to net changes in purchasing power. The idea is that consumers will not respond at all to losses in their purchasing power from one month to the next that simply offset earlier gains in purchasing power; whereas a decline in purchasing power to levels unprecedented in recent history will change consumption behavior. This model of consumer behavior can be motivated based on a proposal by Hamilton (1996).

A measure of the net decrease in purchasing power at a given point in time can be constructed by comparing the current level of purchasing power to its minimum over the previous year. If the current level is above the benchmark, the net decrease measure is zero.¹⁵ Previous empirical studies based on net increments have focused on net energy price increases, while ignoring net energy price decreases such as the unprecedented fall in energy prices in 1986 after the collapse of OPEC or the sharp decline in energy prices following the Asian crisis of 1997/98. This restriction seems implausible on a priori grounds (see Hooker 1996b). In this paper we compute measures of both the

¹⁵Alternatively, one could have used a longer window of three years in computing the net purchasing power loss. Our main findings are robust to the length of the window.

net purchasing power loss and the net purchasing power gain. The resulting net change series are plotted in the bottom panel of Figure 4.

We estimate the VAR models already described in the previous sections. Columns 5 and 6 of Figure 5 show the results. For all real consumption aggregates, responses to net losses and net gains tend to be of opposite signs. Sometimes, as in the case of real durables consumption, they even are of similar magnitude. Typically, the response to net decreases is smoother and more precisely estimated than the response to net increases. As in the baseline model, the response of durables other than vehicles is inconsistent with theory in that net increases have larger and more significant effects. The null of symmetric responses is not rejected for any consumption aggregate. Abstracting from other durables, the p -values range from 0.42 to 1.00 (see column 3 of Table 1).

The response of the aggregate U.S. unemployment rate in the bottom panel of Figure 6 is as expected. While the response to a net loss in purchasing power is larger in magnitude than the response to a net gain, the latter estimate is of the expected negative sign and the intervals are wide enough to accommodate symmetric responses. Compared to the baseline model, the point estimates are more symmetric. Table 1 confirms that symmetry cannot be rejected. The p -value is 0.98.

As in the baseline model, the responses of consumer expectations in the last two columns of Figure 7 are generally consistent with the symmetry hypothesis. The lowest p -value is 0.75. There is strong evidence of symmetry in the response of current buying conditions for vehicles in particular, and as much evidence of larger responses to net decreases in purchasing power as evidence for smaller responses.

3.3 Discussion

The previous subsections demonstrated that, regardless of the choice of model specification, there is no compelling evidence for asymmetries in the response of consumers to positive and negative purchasing power shocks. While the estimated responses of real consumption expenditures and of the unemployment rate to net increases in purchasing power are too imprecise to permit firm conclusions, they are consistent with the symmetry hypothesis. In addition, there is striking evidence that the responses of consumer expectations exhibit a high degree of symmetry.

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 for some investment expenditures reported in Edelstein and Kilian (2007). Moreover, our results are consistent with an alternative approach to assessing the evidence for asymmetries that involves two separate tests. The first null hypothesis is that purchasing power

losses have no effect at any horizon. The second null hypothesis is that purchasing power gains have no effect at any horizon. If the first null hypothesis were systematically rejected, but not the second hypothesis, this would be direct evidence of asymmetry. We conducted such tests and did not find systematic evidence in favor of asymmetries. While it is not possible to reject either null for the major consumption aggregates or for unemployment on the full sample, for the *expected change in one's personal financial situation*, for example, or for *current buying conditions for vehicles* both null hypotheses are strongly rejected. This occurs despite the shorter sample for the consumer expectations data. Thus, low power alone is unlikely to explain our findings in favor of symmetry.

The evidence against asymmetries runs counter to common beliefs and suggests that several mechanisms that feature prominently in theoretical models of the transmission of energy price shocks are not quantitatively important. First, we conclude that the reallocation effect modeled in Hamilton (1988) and emphasized by Davis and Haltiwanger (2001) and Lee and Ni (2002) cannot be detected in aggregate real consumption or consumer expectations data. Second, there is no apparent effect of rising uncertainty on durables consumption. This is true whether we focus on vehicles consumption, other durables consumption or total durables consumption. Such effects played a central role in the closely related analysis of Bernanke (1983). Third, the absence of significant asymmetries across all consumption aggregates also suggests that the precautionary savings effect, if it exists, only exists in symmetric form.

Based on these findings, for the remainder of the paper we will impose symmetry in studying the effects of changes in purchasing power on real consumption and its components. Of particular interest is the question of how large the effects of changes in operating costs, precautionary savings and discretionary income are. Since there is no direct way of testing the net change model against the baseline model of consumer behavior outlined above, we will present the results for the baseline model, augmented by selected additional results for the net change model.

4 How Large are the Effects of Changes in Discretionary Income, Precautionary Savings and Operating Costs?

Having found no compelling evidence of an uncertainty or reallocation effect, in the remainder of the paper we work with bivariate VAR models that impose symmetry on the effect of changes in purchasing power. Our objective is to quantify the effects of changes in discretionary income, precautionary savings and operating costs. Since all these effects are symmetric in energy price increases and decreases, we quantify them by comparing the responses of different real consumption

aggregates to the same purchasing power shock in a standard linear VAR framework. Again we focus on testable implications of the effects outlined in section 3:

Hypothesis 3: The effect of unanticipated changes in discretionary income is bounded by the loss in purchasing power derived under the assumption of inelastic energy demand. This bound may be tightened by taking account of the short-run price elasticity of energy demand.

Hypothesis 4: Evidence of a disproportionately large response of *real consumption of nondurables, services, and durables other than vehicles* to purchasing power shocks is an indication of a precautionary savings effect.

Hypothesis 5: Suppose that only vehicles are subject to the operating cost effect. Further suppose that the precautionary savings effect on durables consumption is the same for vehicles and other durables. Then, in the absence of an operating cost effect, the response of *real vehicles consumption* and the response of *real consumption of other durables* to energy price changes should be equal. Thus, the difference between these two responses allows us to quantify the operating cost effect.

Figure 8 shows responses of the level of real consumption aggregates to a one-time, minus one standard-deviation shock to purchasing power. Since an unanticipated decline in purchasing power is followed by additional declines, the own-response of purchasing power reaches its maximum decline only after eight months. A one standard-deviation shock in purchasing corresponds approximately to a 1.5% increase in energy prices, which translates to a 0.096 percentage point reduction in purchasing power. A 1.5% increase in energy prices may seem small, but it is not, given that our data are monthly. Historically, the average change in prices has been 1.2% in absolute terms. The largest monthly energy price increase in our sample is 11% and occurred following Hurricane Katrina in September of 2005. This corresponds to a purchasing power loss of -0.66% evaluated at the 2005 energy share.

4.1 Discretionary Income Effect

It is instructive to consider the expected consequences of a one standard deviation shock to purchasing power 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 0.096 percentage points, 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. 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. First of all, households may attempt to reduce energy consumption. 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.

These responses may be estimated using regressions analogous to those underlying Figure 8 for various forms of energy consumption. Figure 9 shows that consumption of all forms of energy declines, but there are some unexpected patterns. Contrary to conventional wisdom, *gasoline* consumption responds immediately to unanticipated purchasing power losses. The impact response is -0.57%. Most of the adjustment takes place on impact. The response is highly significant at all horizons with a maximum impact of -0.73%. In contrast, the consumption of *heating oil and coal* takes somewhat longer to adjust fully, but is more elastic in the long run in response to purchasing power losses. After half a year, the response reaches its maximum impact of about -2.28%. The negative response of *electricity* and of *natural gas* is much smaller and statistically insignificant at all horizons with a long run response of -0.25% and -0.53%, respectively. The strikingly large response of heating oil and coal is likely to be 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 up the tank completely when prices are low.¹⁶ In contrast, electricity and natural gas are inherently unstorable, and gasoline may not be stored for safety reasons.¹⁷

The overall response of *energy* consumption is -0.43% on impact and -0.70% after 18 months, and is statistically significant at all horizons. The fact that a 1.5% increase in energy prices (corresponding to a one standard deviation shock to purchasing power) reduces real energy consumption by -0.43% on impact, suggests that the discretionary income effect on consumption should be bounded by -0.055% rather than -0.096% (as discretionary income falls by only 57% of the initial loss in purchasing power that would have occurred under completely inelastic energy demand). Thus, if the response of real consumption were driven entirely by the loss of discretionary income, the response

¹⁶For a discussion of this storable goods feature see Dudine, Hendel, and Lizzeri (2006).

¹⁷It is useful to put these estimates into perspective. Using a structural model, Reiss and White (2005) arrive at an estimate of the short-run price elasticity of electricity demand of -0.39. Our point estimate is -0.24 after seven months. The 90% 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. Our point estimates range from -0.38 to -0.48, depending on the horizon. The bounds of the 90% confidence interval for this elasticity estimate are -0.27 and -0.66, respectively. The short-run elasticity estimate for all energy consumption combined is -0.28 on impact, which does not seem unreasonably high. The interval estimate is [-0.14, -0.42].

should fall between 0% and -0.055%.

How does this prediction compare to the response of total consumption in Figure 8? Total real consumption falls immediately. The response is highly statistically significant and stabilizes after about ten months. After one year, a 1.5% increase in energy prices (which leads to -0.055% decline in discretionary income after accounting for the drop in energy consumption) causes a -0.23% reduction in the level of real consumption relative to the original level. The fact that the point estimate is four times as large as the upper bound on the discretionary income effect suggests that part of the response must be associated with an increase in precautionary savings or an increase in operating costs of energy-using durables. We clearly reject the null that the discretionary income effect alone can explain the response of total real consumption. Moreover, these estimated responses are economically significant. One way of assessing the economic significance of these estimates is to consider a shock equal in size to the largest shock observed in our sample. A Katrina-sized shock in energy prices would imply a fall of -1.58% in real consumption.

It is useful to decompose total real consumption by type of consumption (durables, nondurables, services). Figure 8 shows that each series experiences a statistically significant and persistent decline following an unexpected decline in purchasing power due to rising energy prices. Real durables consumption experiences by far the biggest decline, falling by -0.73% in the long run. This result suggests that durables play an important role in the transmission of such shocks.¹⁸ Real services consumption and real nondurables consumption decline only by about -0.15% and -0.17%, respectively. The confidence intervals allow us to reject the notion that the bound of -0.055 percentage points implied by the discretionary income effect is consistent with the data with the exception of nondurables, for which the response is marginally insignificant.

4.2 Operating Cost Effect

The disproportionately large responses of real consumption relative to the discretionary income effect of Hypothesis 3 is puzzling at first sight. How can a fairly small reduction in purchasing power following an energy price increase generate such large reductions in real consumption? There are two possible explanations for a response of real consumption that exceeds the bound set by the discretionary income effect. One possibility is the presence of an operating cost effect for energy-using durables. This point may be investigated further by decomposing durables into vehicles and

¹⁸We would also expect households to reduce their expenditures on home improvements and other forms of household residential fixed investment, given the durables goods nature of this form of investment. The latter series is only available at quarterly frequency. Additional results (not shown) confirm that the response of real residential fixed investment after six quarters is large and highly significant, not unlike that of consumer durables. As for durables, the estimated response is significantly larger than the any reasonable bound on the discretionary income effect.

durables other than vehicles. The size of the operating cost effect corresponds to the difference in the responses of vehicles relative to other durables. It can be shown that this excess response of real vehicles consumption reaches its maximum after two months with a statistically significant decline of -0.99 percentage points at the 10% level, consistent with Hypothesis 5. After 8 months, it stabilizes at a statistically significant level of about -0.91 percentage points. Thus, the operating cost effect amounts to a decline of 0.99 percentage points of real vehicles consumption in the short run and a decline of 0.91 percentage points in the long run.

4.3 Precautionary Savings Effect

A second possible reason for a disproportionate response of real consumption is the presence of a precautionary savings effect. In this view, household consumption responds to anticipated increases in unemployment (or declines in real income) that are caused by the initial loss of purchasing power. In other words, households respond not only to the immediate loss of discretionary income, but they also respond in anticipation of the delayed effects on unemployment and real household income triggered by such a shock. As the probability of becoming unemployed increases, households increase their precautionary savings at the expense of consumption. This effect is not limited to consumer durables. Households may choose to reduce nonessential consumption of services and nondurables as well, and the reductions need not be spread evenly across all forms of consumption, but depend on how essential a given expenditure item is.

An estimate of this precautionary savings effect can be obtained by comparing the response of consumption excluding vehicles to the bound set by the discretionary income effect. Consider, for example, durables other than vehicles in Figure 8. We cannot reject that this response is bounded by -0.055 percentage points, resulting in a statistically insignificant bound on the precautionary savings effect of 0.24 percentage points (the difference between the point estimate and the upper bound on the discretionary income effect). This compares to a marginally statistically insignificant upper bound of 0.12 percentage points for nondurables and a statistically significant upper bound of 0.10 percentage points for services.

One of the central motivations for a precautionary savings effect is the fear of becoming unemployed in the foreseeable future. If precautionary savings are triggered by fears of unemployment, one would also expect an increase in unemployment in response to unanticipated losses in purchasing power. Figure 10 shows a positive response of the U.S. unemployment rate to an unanticipated loss in purchasing power. After 18 months the unemployment rate rises by 1.36%. The response is marginally statistically significant at some horizons. This evidence is weakly consistent with the

view that there is a precautionary savings motive in response to purchasing power shocks.

While the evidence in Figure 10 is not clear-cut, it seems plausible that consumers - rightly or wrongly - have come to associate unanticipated losses of purchasing power with increases in unemployment. Thus, it makes sense to investigate the plausibility of this explanation further based on the response of consumer expectations about changes in unemployment rather than the response of actual unemployment. In the next subsection, we will address this point based on data from the Michigan Survey of Consumers. We will provide evidence that a loss in purchasing power due to rising energy prices has a significant and negative impact on a wide array of consumer expectations and attitudes that can explain the excess response of real consumption to unanticipated purchasing power losses. We will focus not only on expectations of changes in unemployment, but also a number of alternative indicators of future economic conditions that are relevant to household consumption decisions. We will show that rising energy prices tend to make consumers pessimistic about the state of the economy and about their own personal financial situation. They cause consumers to expect worsening future economic conditions, and they heighten concerns about current buying conditions.

Deteriorating consumer confidence is likely to be an important additional link in the relationship between energy prices and household consumption. The importance of this channel has also been recognized by policymakers. For example, Bernanke (2006b) in a recent speech on the U.S. economic outlook stressed that “recent declines in energy prices ... have boosted household purchasing power and *consumer confidence* [emphasis added]”.

4.4 The Effects of Purchasing Power Shocks on Consumer Sentiment

Rising energy prices are often associated with pessimism and uncertainty about one’s financial situation and the broader economy. In a report on the February 2006 Survey of Consumers at the University of Michigan, Richard Curtin noted that “the February loss in confidence was due to higher energy costs, higher interest rates, and a heightened concern about potential future increases in the unemployment rate”. The same report stated that one-in-five families cited higher prices, mainly for energy, as the cause of their decreased living standards. Such attitudes could lead to a decline in non-energy consumption, even if the discretionary income effect of the purchasing power loss is miniscule. If households are fearful of the economic outlook, they may curtail their consumption on a variety of goods and services driven by a precautionary savings motive. In order to investigate this conjecture we estimate the impulse response functions for a set of measures of consumer confidence and consumer expectations to a negative one-time, one standard-deviation purchasing-power shock. The impulse response functions are estimated in the same fashion as in the previous sections.

The indices for consumer sentiment, expected change in one's personal financial situation, and expected changes in business conditions measure the difference between the number of respondents who expect a better situation and the number who expect a worse situation. A decline in the index suggests that more respondents expect a worsening situation, fewer expect a better situation, or both. Similarly, a decline in the index for buying conditions for large household goods and vehicles suggests that an increasing proportion of respondents think it is currently a bad time to make these purchases. A decline in the indices for unemployment, interest rate, and real family income expectations suggests that a greater proportion of survey respondents expect more unemployment, higher interest rates, and a decline in real family income, respectively.

Figure 11 shows the response of each sentiment series to the purchasing power shock. A one standard deviation fall in purchasing power decreases the overall index of consumer sentiment by 1.6 points. The fall in the index is immediate. While the index begins to rise again a few months after the shock, it remains below its initial level even 18 months after the shock. The observed decline in consumer sentiment compares to a standard deviation of 12.3 for this series, suggesting that an unusually large shock such as Hurricane Katrina, all else equal, could move consumer sentiment nearly one standard deviation away from its mean.

The indices for expected changes in one's personal financial situation and for general business conditions fall by 1.4 and 2.3 points, respectively, suggesting that an increasing number of people expect general business conditions and their personal financial situation to deteriorate over the coming year in response to an unanticipated loss in purchasing power. Whereas the response of the expected change in one's personal financial situation is quite persistent and statistically significant even after 18 months, the response of the expected change in general business conditions, while larger in magnitude, reverts back to zero more quickly and is statistically insignificant after only four months. Given this evidence, one would expect households to cut back on nonessential consumption and to increase precautionary savings.

Of particular interest in judging the empirical plausibility of an operating cost effect is the response of expectations about current buying conditions for durables. Figure 11 shows that the index for buying conditions for large household goods falls by 1.9 points. An even larger decrease is observed for vehicles. The latter index falls by 2.8 points. This implies that a shock of the size associated with Hurricane Katrina would move the vehicle index by nearly one standard deviation. The relatively strong reaction of the index for buying conditions for vehicles in particular is qualitatively consistent with theories stressing energy complementarities in use.

The following results in Figure 11 provide additional insight into why rising energy prices cause households to curb their consumption. Increased pessimism about buying conditions in response

to purchasing power losses is associated with expectations of higher unemployment, higher interest rates, and lower real family income. First, the index for expected changes in unemployment falls by 2.1 points, indicating that an increasing number of people expect higher unemployment. This response is consistent with households perceiving an increased risk of unemployment, as required for the existence of a precautionary savings effect. Second, the index for expected interest rates falls by 1.1 points, indicating that an increasing number of people expect higher interest rates in the future. This suggests another channel of transmission. To the extent that consumers (rightly or wrongly) expect interest rates to rise in response to higher energy prices, their expected liabilities would increase as credit card rates and mortgage rates increase, making it necessary to cut back on consumption. This second channel, however, is short-lived and the responses are largely insignificant, suggesting that it is of minor importance. Third, the index for changes in expected real family income falls by 1.2 points, indicating that a greater number of survey respondents expects real family income to fall in the future. These results are fully consistent with the view that the effect of purchasing power shocks on real consumption operates in part through changes in precautionary savings and through changes in the operating cost of vehicles.

5 Summary and Implications

Based on the evidence presented so far, there are three distinct effects of an unanticipated change in purchasing power driven by higher energy prices. Consider a loss of purchasing power corresponding to an unanticipated 1% increase in energy prices in a given month. First, taking account of the response of fuel consumption to higher energy prices, the effect of the resulting changes in discretionary income after one year can be bounded by -0.04%. In addition, expectations of deteriorating economic conditions are associated with an increase in precautionary savings that causes real consumption of nondurables to fall by an additional -0.08%, services by an additional -0.07% and durables by an additional -0.16% (although only the effect on services is statistically significant). This evidence is important in that it helps explain how the effect of energy price changes may be larger than would be expected based on the small share of energy in consumer expenditures. Finally, rising operating costs will cause real consumption of vehicles to fall by an additional -0.60% in the long run. The short-run response of vehicles consumption may be as high as -0.65%. Combining these effects, our baseline model predicts a reduction of total consumption of -0.15% after one year.¹⁹

¹⁹These estimates presume that households respond proportionately to changes in energy prices. As discussed in section 3, an alternative assumption would be that households respond symmetrically to net changes in purchasing power. In that case, it is not clear how to bound the discretionary income effect, which in turn makes it impossible to gauge the precautionary savings effect. Moreover, the estimated responses from this alternative model cannot be compared directly to those from the baseline model, since the nature of the shocks is different. Nevertheless, we find the

5.1 Automobile Consumption: How Much Does Fuel Economy Matter?

We have already shown that declines in real vehicles consumption are one of the main causes for the disproportionate fall in real total consumption in response to purchasing power losses. This result is consistent with the automobile sector playing a crucial role in the transmission of energy price shocks, as conjectured by Hamilton (1988) and Bresnahan and Ramey (1993), among others. In this subsection, we focus on the overall response of new automobile consumption and on changes in its composition in response to energy price shocks.

We first show that the response of vehicles is primarily driven by the response of automobile consumption. Figure 12 shows that a reduction of purchasing power of -0.096% (corresponding to a one-standard deviation shock to purchasing power or a 1.5% change to energy prices evaluated at the average energy share) causes the consumption of *pleasure boats* to decline by -1.9% in the long run. The response is persistent and significant. Consumption of *pleasure aircraft* declines by -1.6%. The response is persistent, but only marginally significant. Consumption of *recreational vehicles* drops sharply and significantly in the short run, reaching a low of -2.4%, but becomes insignificant in the long run. In contrast, consumption of *motorcycles* is not changing nor are motor vehicle rentals. 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 purchasing power losses on the demand for new automobiles. Figure 12 shows that consumption of *new automobiles* declines sharply, reaching -1.08% in the long run, but the response is barely statistically significant. This translates into a short-run elasticity of demand of about -0.71, which is close to the fuel cost elasticity of -0.5 reported in Goldberg (1998) based on a structural model and micro data.

One possible explanation is that the sectoral reallocation is not so much driven by an overall reduction in the demand for cars, but by an increase in 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 consumption of *new domestic automobiles* with that of *new foreign automobiles*.²⁰ To the extent that

same qualitative patterns as in Figure 8. The responses of total consumption and of all major consumption components are negative and statistically significant. The estimated response of durables far exceeds that of nondurables and services. When durables are broken down further into vehicles and other durables, the latter response is relatively small and statistically insignificant, whereas the former response is large and highly statistically significant.

²⁰Domestic cars are defined by the BEA to include cars assembled in the United States, Canada, or Mexico.

U.S. automobile manufacturers tend to produce less energy-efficient cars, as was certainly the case in the 1970s, considering the larger share of pickup trucks and SUVs in U.S. automobile production, a disproportionate decline in the consumption of domestically produced new cars would be evidence in favor of a shift in demand. Figure 12 shows a strong and highly significant decline in new domestic automobile consumption. In contrast, consumption of new foreign automobiles initially increases, albeit insignificantly. After four months, consumption of new foreign cars slumps as well, although the effect is not as persistent, insignificant, and smaller in the long run 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 -1.34% after two months. The long run response is -0.95% 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 energy price increase 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 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 consumption of light trucks (including minivans, SUVs or pickup trucks) responds differently to unanticipated losses in purchasing power than regular automobiles. There has been much discussion of the softening market for SUVs in recent years.²¹ Figure 12 shows no significant decline in *unit auto sales* (consistent with the evidence on new auto consumption), but a significant decline in both *unit light truck sales* and *unit heavy truck sales*, with long-run responses of -1.6% and -1.3%, 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 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.²²

5.2 Understanding the 1986 Episode: Where is the Boom?

Our results suggest that Hamilton (1988) was correct that expenditures on consumer durables that are complementary in use to energy (such as cars) are sensitive to even small energy price fluctua-

²¹For example, Douglass (2007) discusses the cooling appetite for large sports utility vehicles, as gasoline prices are rising.

²²This information was obtained from unit sales and production data on the company websites.

tions. We showed that indeed there is a strong decline in the real consumption of motor vehicles in response to unanticipated purchasing power losses. This decline accounts for much of the anomalous response of consumer durables and generates increased aggregate unemployment. However, there is no evidence in the real consumption or consumer expectations data that changes in the demand for vehicles cause a sectoral reallocation effect that amplifies the effect of energy price increases and counteracts the effect of energy price decreases, as postulated in Hamilton's model. This conclusion may seem surprising, but it is consistent with the fact that only vehicles consumption experiences a dramatic decline in response to losses in purchasing power and that the U.S. automobile industry only accounts for about 1% of aggregate U.S. employment and 1% of real U.S. GDP. Thus, even if there is a drastic decline in the demand for U.S. automobiles, the effect on other parts of the economy is likely to be small in scale, which may account for our inability to detect a reallocation effect in the data.

There are other reasons, however, for the popularity of models that embody a reallocation effect. One reason is that there was no noticeable economic expansion after the sharp fall in crude oil prices in 1986. This evidence seems hard to reconcile with the perception that rising crude oil prices in 1979 contributed to a sharp economic downturn, unless one appeals to a model with asymmetric responses to energy price changes (see, e.g., Balke, Brown, Yücel 2002; Gramlich 2004). The theoretical model of Hamilton (1988) and subsequent empirical work by Davis and Haltiwanger (2001) and others, seemed to provide an explanation for this puzzling asymmetry. A common view is that this asymmetry in the data simply necessitates the existence of a large reallocation effect. Although this view is appealing upon casual inspection of the data, it misses important pieces of the puzzle.

It is useful to be explicit about the counterfactual. The implicit premise in this literature is that it suffices to compare economic performance before and after the collapse of OPEC. Clearly, however, we need to compare what actually happened in 1986 to what would have happened without the sharp fall in energy prices (rather than to the status-quo-ante). This question may be answered based on historical decompositions of the real consumption data. Historical decompositions measure the cumulative effect of the historically observed sequence of purchasing power shocks on the level of real consumption at each point in time. Quantifying this effect is important because it conveys information that cannot be gleaned from impulse response estimates. Sometimes energy price increases come in clusters, and at other times energy price increases may alternate with decreases. The cumulative effect on real consumption is a weighted average of the entire history of shocks up to a given point in time. We compute this effect based on the bivariate VAR model estimates of section 4. Figure 13 shows the actual (demeaned) real consumption growth rates and the consumption

growth rates predicted based on the cumulative effect of the purchasing power shocks alone. The difference between the two series measures the extent to which consumption growth is not explained by purchasing power shocks. To improve the readability of the plot, we have converted all growth rates to quarterly averages (upper panel) and annual averages (lower panel). The quarterly series has been annualized.

It is instructive to compare the annual results for 1979 and 1986. 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. As shown in Figure 13, the VAR model implies that rising energy prices (all else equal) lowered real consumption growth by -1.92% in 1979, and raised it by +2.02% in 1986, making these effects nearly symmetric. Moreover, 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.

Figure 13 also provides two additional insights. First, energy price shocks were responsible for substantial declines (defined as an effect on the real consumption growth rate in excess of -0.65%) in consumption growth in 1974, 1979/80, 1990 and 2004/05, but they also caused large increases in real consumption growth (defined as an effect on the growth rate in excess of +0.65%) in 1986, 1991, 1998 and 2001. Second, a substantial part of real consumption growth is not associated with energy prices. Notably, the pattern of excess consumption growth in the 1970s is consistent with go-and-stop monetary policies of the type described in Barsky and Kilian (2002) and the unusually low growth in 1980-82 and 1990-91 is at best partially explained by energy prices and suggests an important role for monetary policy under Paul Volcker. Likewise, the unusually high growth of 1984-85, 1996-1999, and 2004 cannot be attributed mainly to energy prices.

5.3 Toward an Alternative Explanation of the 1986 Episode

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 (see Table 2). Thus, the asymmetry alluded to earlier does exist in real GDP growth, but is not reflected in real consumption growth. The comparison of the 1979 and 1986 growth rates of real GDP and its components in Table 2 reveals that the source of the asymmetry in real GDP growth lies in private investment. More specifically, real nonresidential investment in equipment and structures are the two key components that caused real GDP growth in 1986 to be so low. In 1979, they grew by -2.80% and +7.54% relative to their means, 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 consumption.

There are two potential explanations for this pattern. One explanation is that energy price shocks have asymmetric effects on firms' fixed investment expenditures.²³ Such an explanation seems implausible for several reasons. First, while there is some apparent evidence of asymmetries in the point estimates of the nonresidential fixed investment responses (not shown), the type of asymmetry found in these responses (and of business investment in structures in particular) does not conform to what we would expect to see if these responses were driven by the uncertainty effect of Bernanke (1983). Specifically, the response to purchasing power losses is near zero, whereas the response to gains is strongly negative. Second, there is no statistically significant evidence against symmetry in the nonresidential fixed investment responses. For all specifications already considered in section 3, we fail to reject the null of symmetric responses to positive and negative shocks with p -values in the range from 0.66 to 0.91.²⁴ Third, the mechanisms commonly discussed in support of an asymmetric response of nonresidential fixed investment (such as the Bernanke (1983) uncertainty effect on durables) should apply equally to consumer durables and to firms' purchases of durables. If firms' fixed investment responds much more asymmetrically to energy price changes than durables consumption and real residential fixed investment, then it must do so for unrelated reasons. It is unclear what economic mechanism would explain such an asymmetry in the responses of firms' fixed investment expenditures. Finally, one would expect an asymmetric effect on nonresidential fixed investment to be reflected in similarly asymmetric responses of aggregate unemployment and therefore consumer expectations, consumer expenditures and residential fixed investment. We have already shown that there is no compelling evidence to support this conjecture.

An alternative and more plausible explanation is that a drop in firms' investment expenditures not related to the preceding fall in energy prices caused the 1986 expansion to fizzle. That explanation is consistent with the fact that the growth rate of firms' real investment expenditures fell much more in 1986 than in 1979 (see Table 2). Such a pattern is inconsistent with conventional explanations of asymmetric investment responses. As we pointed out in section 3, the response to purchasing power gains cannot be larger in absolute terms than the response to purchasing power losses and typically will be smaller. Even in the complete absence of direct effects on investment expenditures, the effect of an increase in energy prices and a decrease in energy prices of the same size should be

²³Some industry-level evidence of the effect of energy price shocks on inventory investment has been presented in recent work by Herrera (2006). The overall importance of such channels at the aggregate level remains an active area of research.

²⁴For the quarterly VAR models, our default choice is a lag order of 2. Our analysis is based on the premise that innovations to purchasing power changes (or their transformations) can be treated as predetermined with respect to the quarterly macroeconomic aggregates. For a more detailed analysis of the response of nonresidential fixed investment expenditures to energy price shocks see Edelstein and Kilian (2007).

of the same magnitude in absolute terms. In contrast, the data show that firms' real investment in structures actually increased by 7.54% relative to trend in 1979 (which is completely at odds with the uncertainty effect), but fell by -16.35% relative to trend in 1986. Investment expenditures on equipment declined in both years, but they declined much more in 1986 than in 1979, which again casts doubt on the presence of an asymmetric effect.

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 1986.²⁵

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 space (including office space)* and *manufacturing structures*, which account for 21 percent and 6 percent 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 percent 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 percent 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 percent drop of this component in 1986 (relative to 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 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

²⁵For details of the timing of the 1986 Tax Reform Act see Wakefield (1987). We thank Chris House for providing us with the detailed investment tax credit data used in House and Shapiro (2006) and constructed by Jorgenson using methods detailed in Jorgenson and Yun (1991).

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).

Thus, 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 the econometric model on the basis of falling energy prices alone and why unemployment remained higher than it would have been otherwise.

5.4 Discussion

The evidence presented in this section implies that the standard view in the literature of how energy price shocks affect the U.S. economy has to be reconsidered. The conventional wisdom is that a fall in energy prices will have only weak effects on output and employment, as the increase in aggregate demand will be offset by the reallocation effect of changing expenditure patterns. In contrast, when the price of energy increases, both effects work in the same direction, amplifying the response of output and unemployment. This explanation rationalized both a sharp contraction following energy price increases, and the absence of an economic expansion following energy price decreases. Without an allocative channel, however, one of the chief mechanisms whereby energy price increases in theory can create large economic downturns is inoperative, making it more difficult to rationalize the economic downturns of 1974 and 1979/80 based on adverse energy price shocks. Despite the absence of an allocative effect, the cumulative effect of energy price shocks on real consumption can be important. As we have shown, both the precautionary savings effect and the operating cost effect help elevate the level of the responses of real consumption beyond the limits sets by the energy share in consumption and the elasticity of energy demand. Nevertheless, energy price shocks are by no means the dominant explanation of real consumption growth in 1974, 1979-81, 1990, or for that matter since 2003.

5.5 Has Real Consumption Become Less Responsive to Purchasing Power Shocks?

It has been widely observed that energy price shocks do not appear to affect the U.S. economy as much as we used to think.²⁶ Figure 14 quantifies this phenomenon by comparing responses

²⁶A weakening of the statistical relationship between oil prices and the U.S. economy in the mid-1980s has been noted, for example, by Hooker (1996b, p. 222) and Davis and Haltiwanger (2001, p. 482). There is also a widely held view among policymakers that the surges in oil prices in the 1970s and 1980s had much more pronounced economic

of consumption aggregates estimated on the first half (1970.2-1987.12) and the second half of our sample (1988.1-2006.7). Otherwise, the models are identical to the models used in section 4. The scale of the impulses has been normalized to be the same across the two samples as for the full sample. Figure 14 shows that, compared to the first half of the sample, in the second half the long-run response of total real consumption drops from -0.46% to -0.12%. The corresponding decline for durables is from -1.27% to -0.37%. The response of vehicles consumption declines from -1.99% in the first half to -0.74% in the second half of the sample. The decline in durables consumption excluding vehicles shrinks from -0.67% in the first half of the sample to -0.02% in the second half. The response of nondurables shrinks from -0.44% to -0.03% and that of services from -0.28% to -0.10%. A similar reduction occurs in the response of real residential fixed investment (not shown). The long-run response drops from -7.1% to -2.0%. Finally, the rise in unemployment associated with an unanticipated purchasing power loss drops from 2.32% to 0.55%.

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 that we have documented in section 2. This conjecture is not obvious since the share of energy has been rising recently, as shown in Figure 2. More importantly, since our results are based on innovations in purchasing power changes rather than innovations in energy price changes, they already control for changes in the expenditure share of energy.

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 actually 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.

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 with 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 effects than the more recent increases (see, e.g., Bernanke 2004).

1970s.

This point is illustrated by comparing the responses of new domestic and foreign automobiles in each subsample (see Figure 15). Whereas in the first subsample expenditures on new domestic automobiles drop by -4.2% after two months and by -2.6% in the long run, in the second half the short-run response drops to -1.1% and the long-run response to -0.5%. 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 2.0%, followed by an insignificant decline of -1.5% after five months and a long-run response of -0.5%. 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 -0.6% and the long-run response to -0.2%. 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 re-structured 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).²⁷ 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.

6 Conclusion

There is an ongoing interest in understanding the effects of energy price shocks on the economy. Some of the most compelling evidence on this question has come from documenting the adjustments of output and employment at the industry and plant level to energy price shocks (see, e.g., Davis and Haltiwanger 2001, Lee and Ni 2002). This type of work was motivated by the view that unexpected energy price changes represent allocative disturbances because of their effects on expenditure patterns of households and firms. Such shocks in principle can have dynamic effects that far exceed the importance of energy for the economy as measured by the energy share in expenditures.

It is widely accepted that in the absence of a major disruption in spending by consumers and firms, the effects of energy price shocks on the economy will be small. In this paper we studied in

²⁷See <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.

detail the response of personal consumption expenditures to unanticipated purchasing power shocks triggered by fluctuating energy prices. On the basis of the evidence presented in this paper, we concluded that the standard view of how energy price shocks affect the U.S. economy has to be reconsidered.

First, whereas asymmetric responses play a central role in theoretical models of the transmission of energy price shocks, we found no compelling statistical evidence of asymmetries. Specifically, we found no evidence for an uncertainty effect on the consumption of durables of the type discussed by Bernanke (1983) in the context of investment decisions nor did we find evidence of the reallocation effect stressed in Hamilton (1988) and Bresnahan and Ramey (1993). While we did find evidence of changing expenditure patterns, which is a necessary condition for an reallocation effect, we showed that aggregate consumer spending and its major components as well as consumer expectations data do not respond in the directions one would expect if the allocative channel of transmission were quantitatively important. The apparent absence of a reallocation effect on real consumption, despite comparatively large effects of purchasing power shocks on the consumption of new domestically produced automobiles in particular, is consistent with the small share of the U.S. automobile industry in domestic real GDP and employment.

Second, we showed that despite the absence of the uncertainty and reallocation effects, the responses of real consumption aggregates are larger than suggested by the effects of unanticipated changes in discretionary income alone. The excess responses can be attributed to shifts in precautionary savings and to changes in the operating cost of energy-using durables. We quantified each of these effects. The combined effect of a one-time 1% increase in energy prices in a given month is a statistically significant decline in real total consumption of -0.15% a year later.

Third, energy price shocks were shown to have contributed substantially both to the decline in consumption in 1979, amidst sharply rising energy prices, and to its recovery in 1986 after the collapse of OPEC. This result runs counter to the conventional wisdom that the U.S. economy's response to the decline in energy prices in 1986 was muted, whereas its response to the 1979 energy price increases was strong. We showed that this perception is at odds with the largely symmetric pattern of real consumption growth in 1979 and 1986. This pattern contrasts sharply with the asymmetric pattern of real GDP growth in 1979 and 1986. Further data analysis suggested that an exogenous drop in nonresidential fixed investment expenditures related to the 1986 Tax Reform Act was mainly responsible for the low rate of real GDP growth in 1986. 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 (for reasons detailed in the paper) appears asymmetric in the opposite direction from the asymmetries previously discussed in the literature on oil and the macroeconomy.

Fourth, our analysis shed light on the declining importance of energy price shocks for the U.S. economy. We documented the extent to which consumption aggregates have become less responsive to energy price shocks since the mid-1980s. The effect of an unanticipated 1% increase in energy prices on total real consumption one year later drops from -0.30% in the first half of the sample to only -0.08% in the second half. We traced the declining importance of energy price shocks relative to the 1970s to changes in the composition of U.S. automobile production and to the declining overall importance of the U.S. automobile sector.

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 and permanently 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). If a typical household spends \$200 a month on gasoline at the January 2007 price of \$2.29 per gallon, this would raise the household's gasoline bill by almost \$22 a month, if the household continued to consume the same amount of gasoline. In response to such a shock, a typical household with about \$4000 to spend per month will have cut back its expenditures one year later by \$35 based on the full-sample estimates (or by \$17 based on the post-1987 estimates). Most of the adjustment will take place in the first six months following the gasoline price increase. Given a share of consumption in GDP of about 72%, this implies that, all else equal, real GDP will have fallen by 0.63% one year after the shock. This example illustrates that it takes repeated surprise increases in gasoline prices to generate large effects on household consumption.

References

- [1] Auerbach, A.J. (1987), “The Tax Reform Act of 1986 and the Cost of Capital,” *Journal of Economic Perspectives*, 1, Summer, 73-86.
- [2] 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.
- [3] 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.
- [4] Barsky, R.B., and L. Kilian (2004), “Oil and the Macroeconomy Since the 1970s,” *Journal of Economic Perspectives*, 18, Fall, 115-134.
- [5] Behraves, N. (2007), Global Insight Inc., “Warm Winter Pushing Oil Prices Down,” interviewed by Renée Montagne for National Public Radio, transcript of interview, January 16.
- [6] Bernanke, B.S. (1983), “Irreversibility, Uncertainty, and Cyclical Investment,” *Quarterly Journal of Economics*, 98, 85-106.
- [7] Bernanke, B.S. (2004), “Oil and the Economy,” remarks at the Distinguished Lecture Series, Darton College, Albany, Georgia, October 21.
- [8] Bernanke, B.S. (2006a), “Energy and the Economy,” speech to the Economic Club of Chicago, June 15.
- [9] Bernanke, B.S. (2006b), “The Economic Outlook,” remarks before the National Italian American Foundation, New York, November 28.
- [10] 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.
- [11] Bureau of Economic Analysis (1987), “The Business Situation,” *Survey of Current Business*, January.
- [12] 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.
- [13] Curtin, R.T. (2006), *Survey of Consumers*, University of Michigan, March 10.

- [14] Dahl, C., and T. Sterner (1991), "Analyzing Gasoline Demand Elasticities: A Survey," *Energy Economics*, 13, 203-210.
- [15] Davis, L.W., and L. Kilian (2007), "The Allocative Cost of Price Ceilings: Lessons from the U.S. Residential Market for Natural Gas," mimeo, Department of Economics, University of Michigan.
- [16] Davis, S.J. (1987), "Allocative Disturbances and Specific Capital in Real Business Cycle Theories," *American Economic Review Papers and Proceedings*, 77, 326-332.
- [17] Davis, S.J., and J. Haltiwanger (2001), "Sectoral Job Creation and Destruction Responses to Oil Price Changes," *Journal of Monetary Economics*, 48, 465-512.
- [18] Douglass, E. (2007), "U.S. Motorists Cutting Back a Bit," *The L.A. Times*, January 25.
- [19] 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.
- [20] Dudine, P., Hendel, I., and A. Lizzeri (2006), "Storable Good Monopoly: The Role of Commitment," *American Economic Review*, 96, 1706-1719.
- [21] Edelstein, P., and L. Kilian (2007), "The Response of Business Fixed Investment to Energy Price Changes: A Test of some Hypotheses about the Transmission of Energy Price Shocks," mimeo, Department of Economics, University of Michigan.
- [22] Goldberg, P.K. (1998), "The Effects of the Corporate Average Fuel Efficiency Standards in the U.S.," *Journal of Industrial Economics*, 46, 1-33.
- [23] Gramlich, E.M. (2004), "Oil Shocks and Monetary Policy," remarks by Governor Edward M. Gramlich at the Annual Economic Luncheon, Federal Reserve Bank of Kansas City, Kansas City, Missouri, September 16.
- [24] Hamilton, J.D. (1988) "A Neoclassical Model of Unemployment and the Business Cycle," *Journal of Political Economy*, 96, 593-617.
- [25] Hamilton, J.D. (1996), "This is What Happened to the Oil Price-Macroeconomy Relationship," *Journal of Monetary Economics*, 38, 215-220.
- [26] Hamilton, J.D. (2005), "Oil and the Macroeconomy," forthcoming: *New Palgrave Dictionary of Economics*.

- [27] Herrera, A-M. (2006), "Inventories, Oil Shocks and Macroeconomic Behavior," mimeo, Department of Economics, Michigan State University.
- [28] Hooker, M.A. (1996a), "What Happened to the Oil Price-Macroeconomy Relationship?" *Journal of Monetary Economics*, 38, 195-213.
- [29] Hooker, M.A. (1996b), "This Is What Happened to the Oil Price-Macroeconomy Relationship: Reply," *Journal of Monetary Economics*, 38, 221-222.
- [30] 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.
- [31] House, C., and M.D. Shapiro (2006), "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," mimeo, Department of Economics, University of Michigan.
- [32] Huntington, H.G. (1998), "Crude Oil Prices and U.S. Economic Performance: Where Does the Asymmetry Reside?," *Energy Journal*, 19, 107-148.
- [33] Ip, G. (2006), "Falling Oil Prices May Spell Relief for Consumers," *The Wall Street Journal*, September 13.
- [34] Ivanov, V., and L. Kilian (2005), "A Practitioner's Guide to Lag Order Selection for VAR Impulse Response Analysis," *Studies in Nonlinear Dynamics and Econometrics*, 9(1), 2.
- [35] Jorgenson, D.W., and K.-Y. Yun (1991). *Tax Reform and the Cost of Capital*. Oxford University Press.
- [36] Kilian, L. (1998), "Small-Sample Confidence Intervals for Impulse Response Functions," *Review of Economics and Statistics*, 80(2), 218-230.
- [37] Kilian, L. (2001), "Impulse Response Analysis in Vector Autoregressions with Unknown Lag Order," *Journal of Forecasting*, 20(3), 161-179.
- [38] Kilian, L. (2006), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market," mimeo, Department of Economics, University of Michigan.
- [39] 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.
- [40] Macintyre, D. (2006), U.S. Department of Energy, "Gasoline Price Trends Could Move Upwards Again," interviewed by Scott Horsley for National Public Radio, transcript of interview, December 29.

- [41] Mehra, Y.P., and J.D. Peterson (2005), "Oil Prices and Consumer Spending," *Federal Reserve Bank of Richmond Economic Quarterly*, 91, 53-72.
- [42] 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.
- [43] Peters, J.W. (2006), "Growth Slowing Across the Country, Fed Regional Reports Say," *The New York Times*, July 27.
- [44] Pindyck, R.S. (1991), "Irreversibility, Uncertainty and Investment," *Journal of Economic Literature*, 29, 1110-1148.
- [45] Reiss, P.C., and M.W. White (2005), "Household Electricity Demand Revisited," *Review of Economic Studies*, 72, 853-883.
- [46] "Retail Sales, Expected to Rise, Instead Fell in June," *The New York Times*, July 15, 2006.
- [47] U.S. Bureau of Labor Statistics (2005), "Consumer Expenditures in 2003," Report 986. Available at: <http://www.bls.gov/cex/csxann03.pdf>.
- [48] Wakefield, J.C. (1987), "The Tax Reform Act of 1986," *Survey of Current Business*, Bureau of Economic Analysis.

Table 1: Specification Tests for Impulse Response Functions**Bootstrap p -Values**

Symmetry Tests

Equality Tests

	Gains vs Losses	Large Gains vs Large Losses	Net Gains vs Net Losses	Large Losses vs Small Losses	Large Gains vs Small Gains	Large Changes vs Small Changes
Total Consumption	0.998	0.999	0.694	0.750	1.000	0.917
Durables	0.988	0.995	0.491	0.695	1.000	0.881
Nondurables	1.000	1.000	1.000	0.999	0.982	0.974
Services	0.855	0.675	0.948	0.998	0.954	1.000
Vehicles	0.909	0.932	0.418	0.434	1.000	0.901
Other Durables	0.610	0.570	0.315	0.988	0.975	0.978
Unemployment	0.996	0.983	0.983	0.977	0.968	0.940
Consumer Sentiment	0.995	0.991	0.799	0.644	1.000	0.964
Expected Change in Personal Financial Situation	0.856	0.791	0.753	0.385	1.000	0.929
Expected Change in Business Conditions	0.974	0.859	0.989	0.417	0.911	0.970
Buying Conditions for Large Household Goods	0.995	0.994	0.986	0.996	0.996	1.000
Buying Conditions for Vehicles	0.995	1.000	1.000	0.999	1.000	0.991
Expected Change in Unemployment	1.000	1.000	0.948	0.991	0.999	0.999
Expected Change in Interest Rates	0.998	0.998	0.980	0.981	0.994	1.000
Expected Change in Real Family Income	0.966	0.830	0.951	0.876	0.986	1.000


NOTES: The p -values were computed based on a version of the residual-based bootstrap method of Kilian (1998) that preserves the correlations across VAR models.

Table 2: Annual Growth Rates for Demeaned Real GDP and its Components (Percent)


	1979	1986
GDP	-1.81	-0.31
Consumption	-1.58	0.90
Imports	-5.29	1.50
Exports	4.21	4.40
Government Expenditures and Investment	-0.95	2.98
Gross Private Investment	-8.47	-9.92
Private Fixed Investment	-3.31	-4.17
Private Nonresidential Fixed Investment	0.65	-8.94
Equipment	-2.80	-4.65
Structures	7.54	-16.35
Private Residential Fixed Investment	-12.41	7.83
Change in Inventories	-5.50	-72.88

SOURCE: Bureau of Economic Analysis.

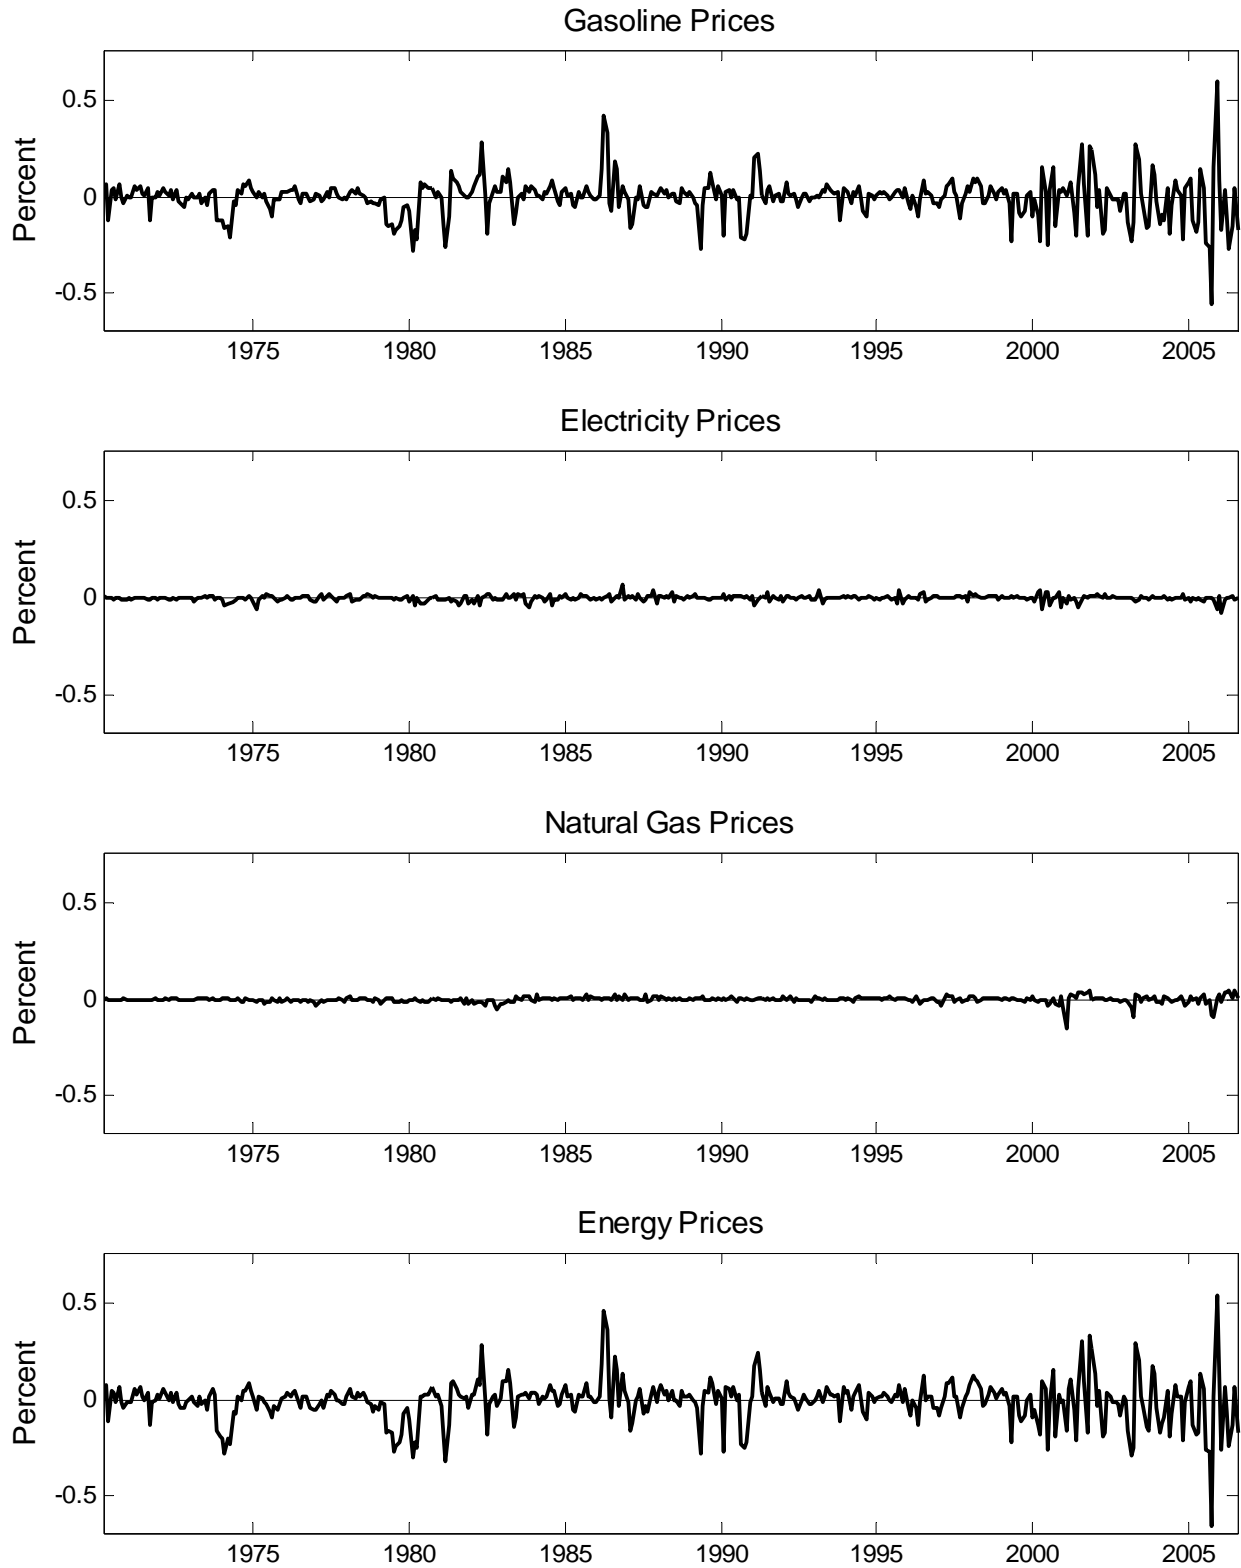
**Figure 1: Monthly Real PCE Price Indices
1970.1-2006.7**






**Figure 2: Monthly Nominal Expenditure Share
1970.1-2006.7**



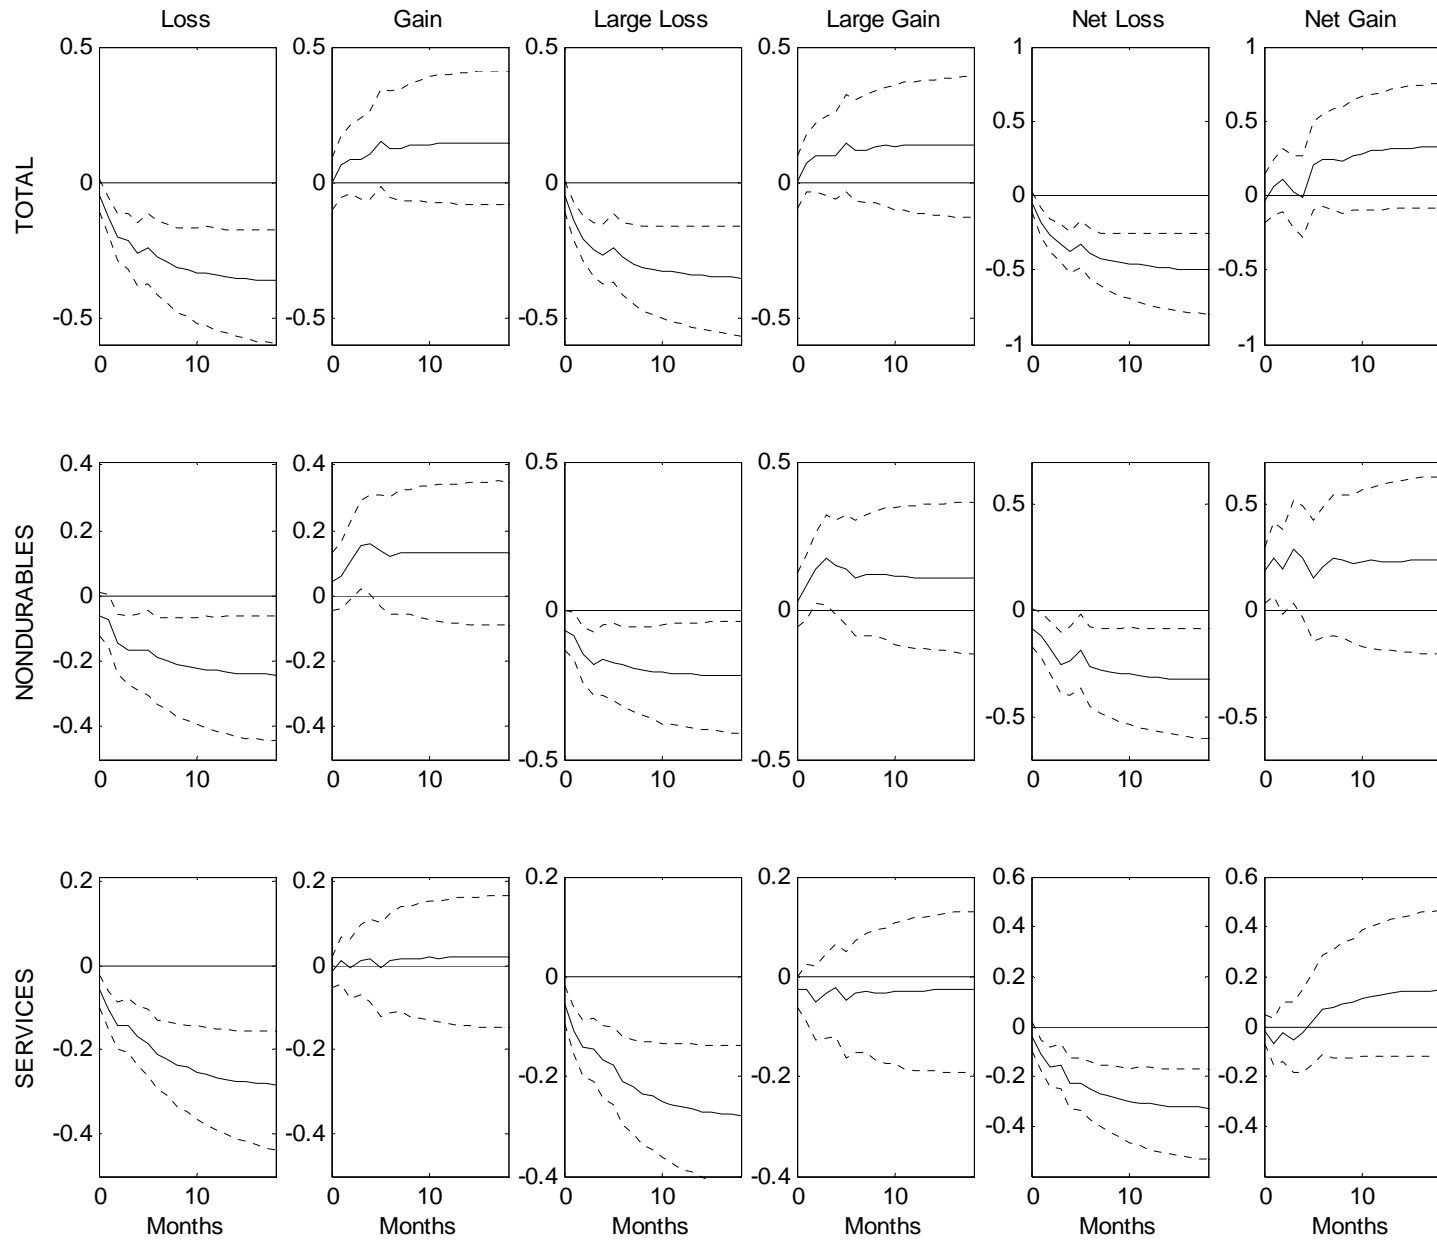
**Figure 3: Monthly Loss in Purchasing Power
1970.2-2006.7**



**Figure 4: Alternative Measures of Purchasing Power Gains and Losses
1970.2-2006.7**



**Figure 5: Response of Real Consumption to Purchasing Power Shocks
1970.2-2006.7**



**Figure 6: Response of Unemployment Rate to Purchasing Power Shocks
1970.2-2006.7**

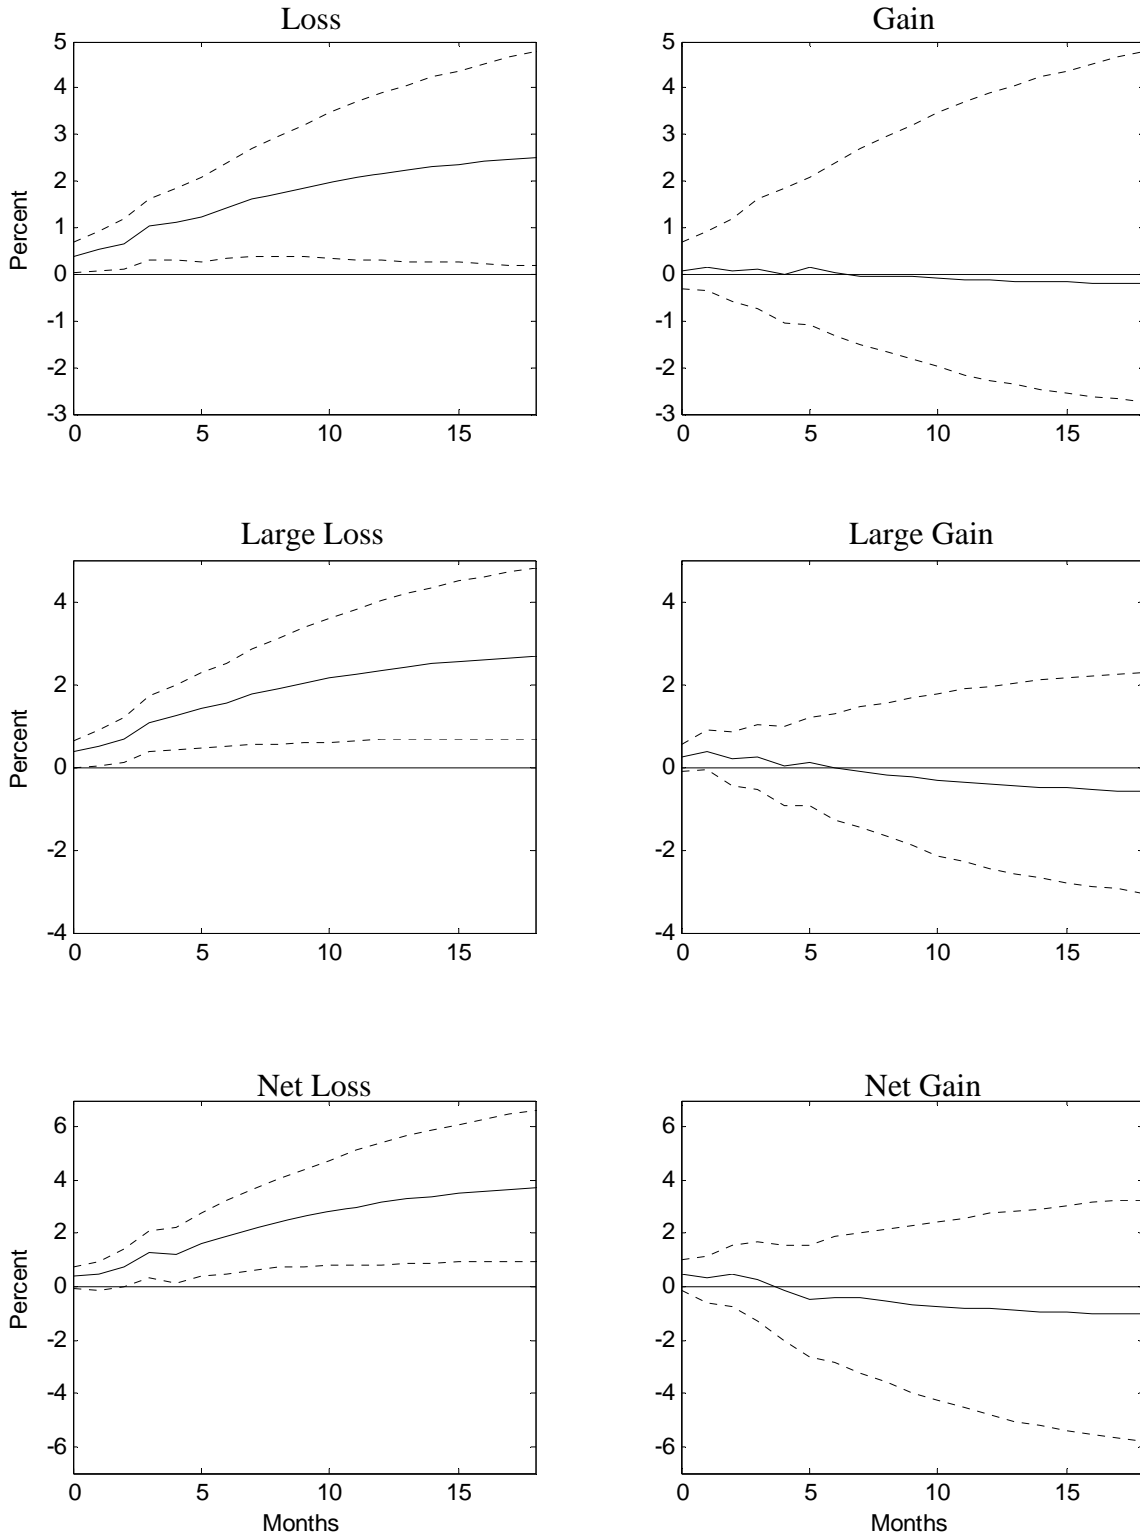




Figure 7 (continued):

Loss Gain Large Loss Large Gain Net Loss Net Gain


Current Buying Conditions for Vehicles




Expected Change in Unemployment



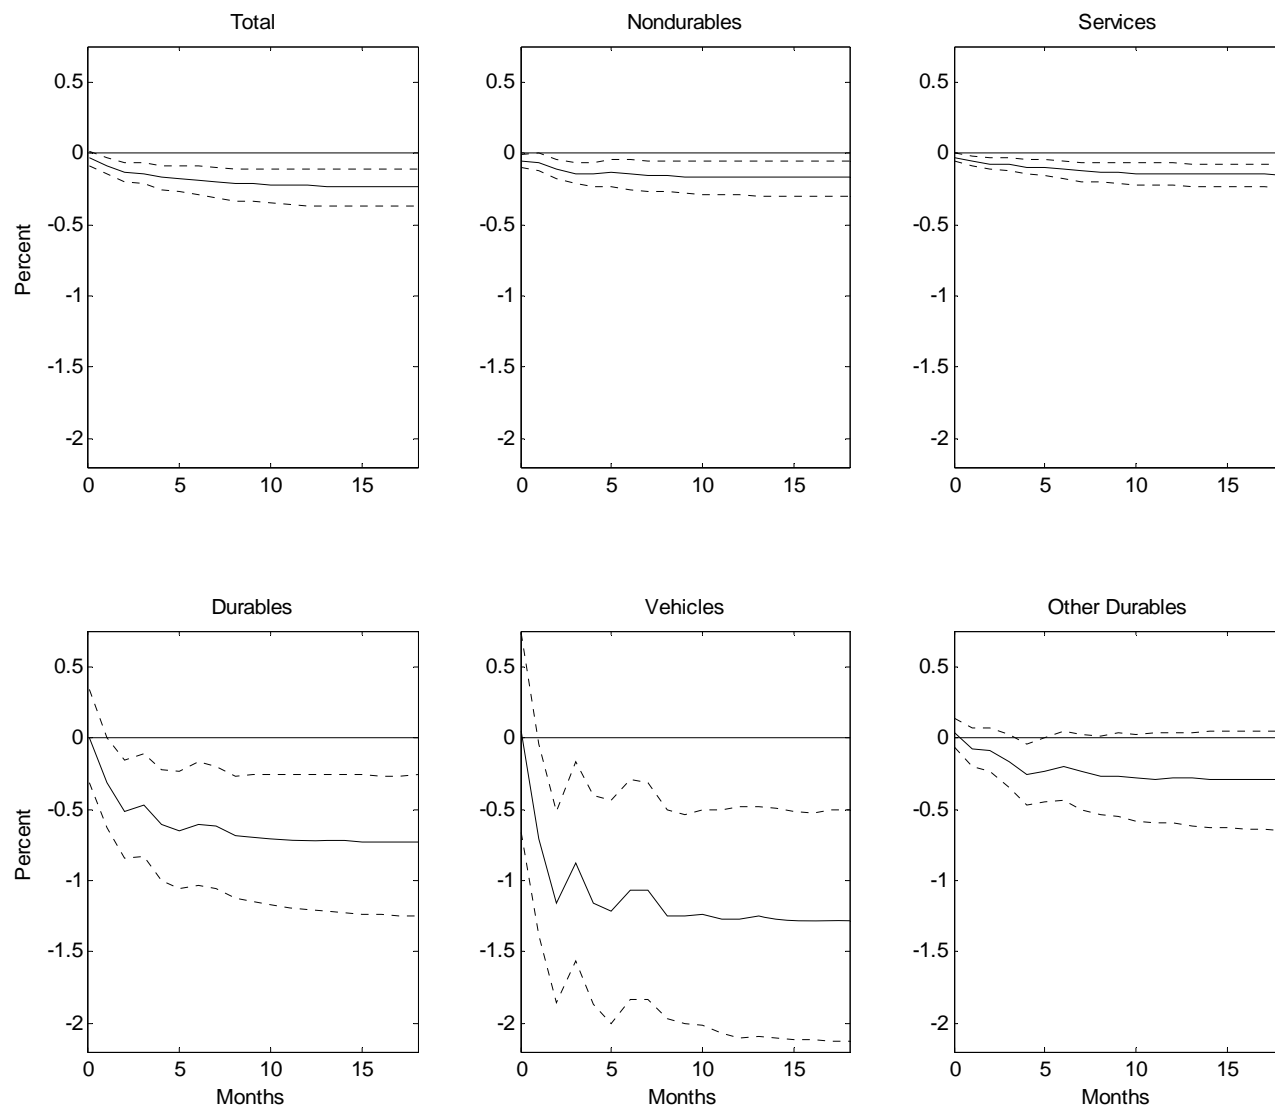
Expected Change in Interest Rates




Expected Change in Real Family Income




**Figure 8: Responses of Real Consumption to Purchasing Power Shocks
1970.2-2006.7**



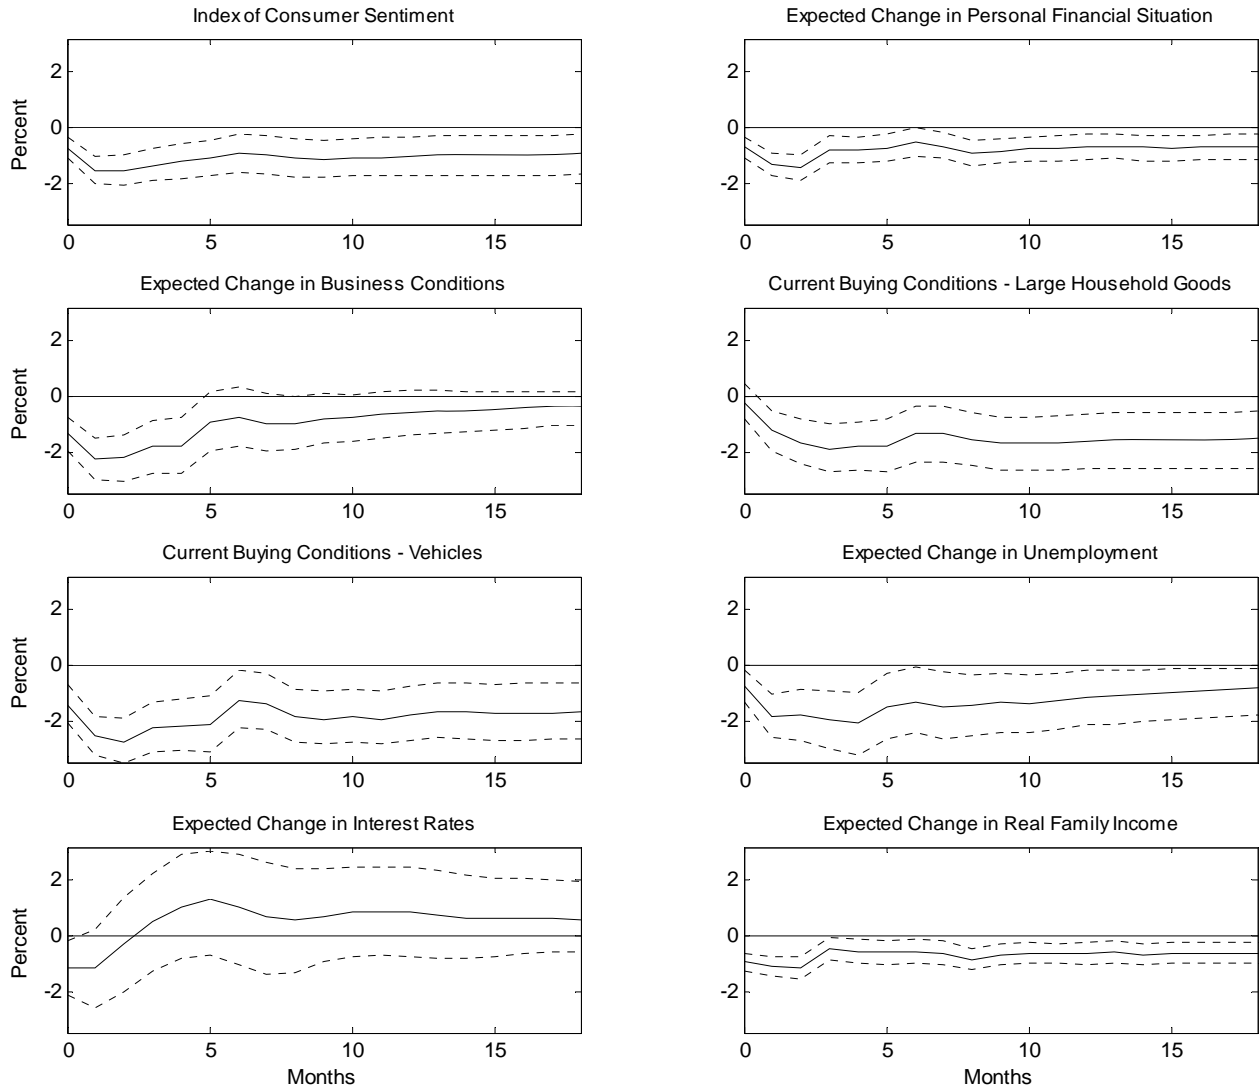
**Figure 9: Response of Real Energy Consumption to Purchasing Power Shock
1970.2-2006.7**



**Figure 10: Response of Unemployment Rate to Purchasing Power Shock
1970.2-2006.7**



**Figure 11: Response of Consumer Expectations to Purchasing Power Shocks
1978.1-2006.5**



**Figure 12: Response of Real Consumption by Expenditure Item
1970.2-2006.7**

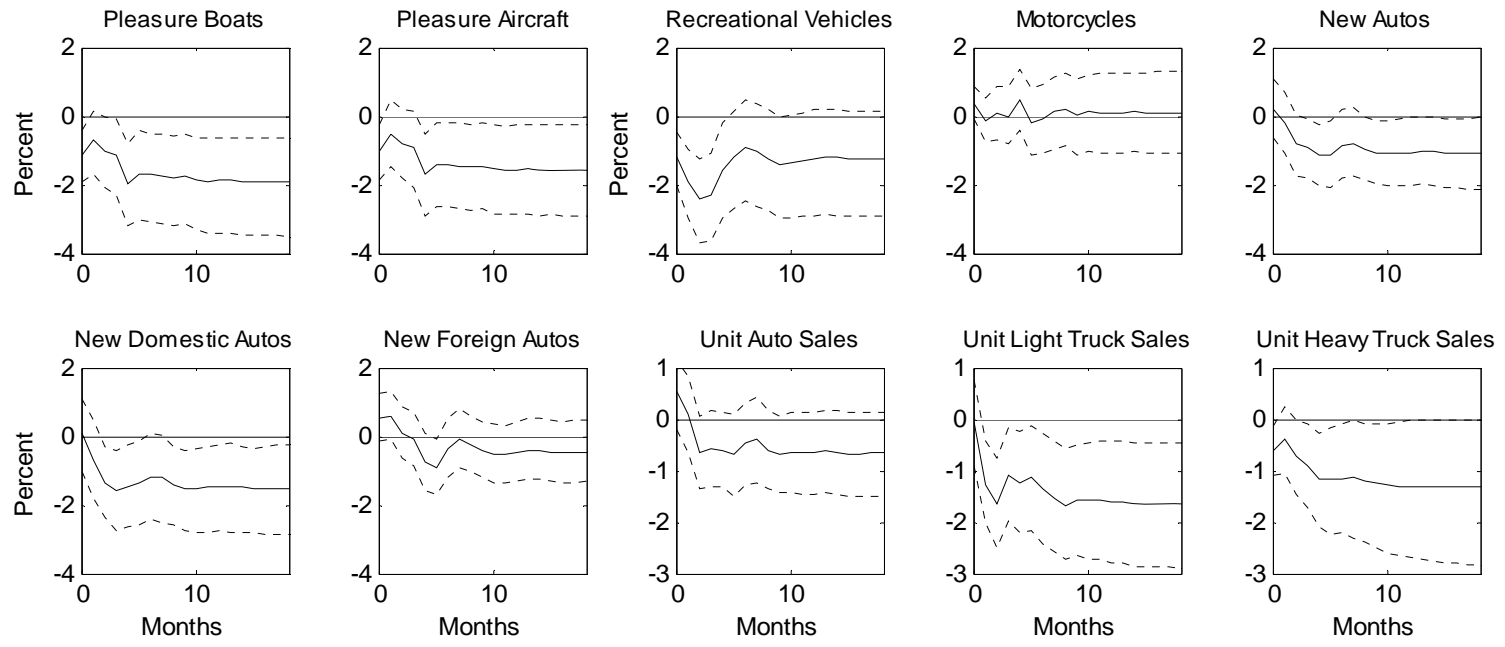


Figure 13: Contribution of Purchasing Power Shocks to Real Consumption Growth





Figure 14: Selected Responses by Sample Period

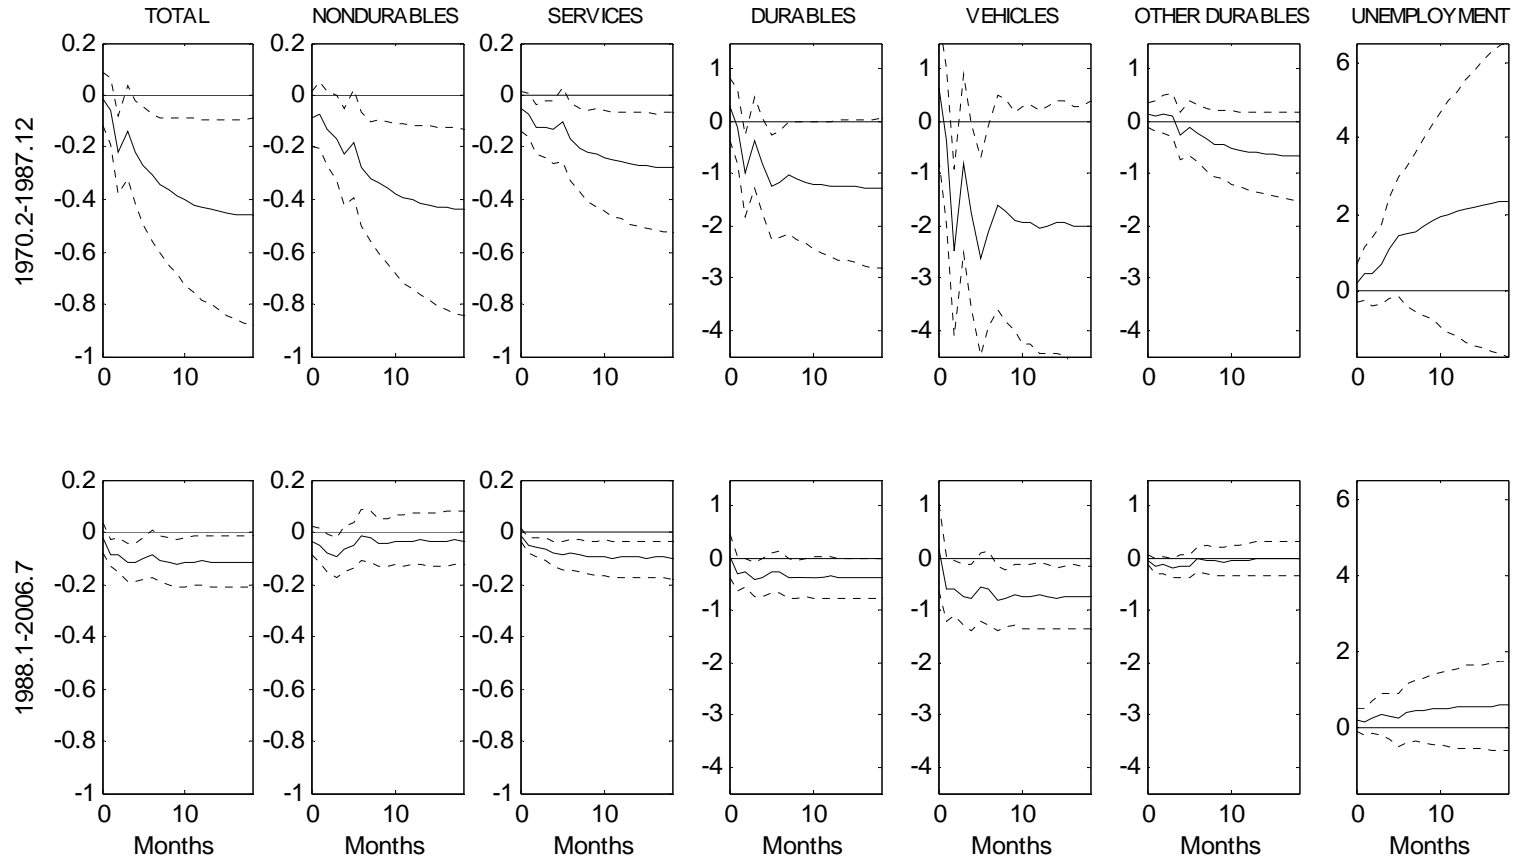


Figure 15: Response of New Auto Consumption to Purchasing Power Shocks

