ON THE IMPORTANCE OF COMMODITY AND ENERGY PRICE SHOCKS FOR THE MACROECONOMY

by

Paul S. Edelstein

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Economics)
in The University of Michigan
2007

Doctoral Committee:

Associate Professor Lutz Kilian, Chair
Professor Matthew D. Shapiro
Assistant Professor Lucas William Davis
Assistant Professor Kathy Z. Yuan

For my parents, Bernard and Susan; for my sister Marisa; and for Bonnie.

ACKNOWLEDGEMENTS

I would like to thank my advisor, Lutz Kilian, for his guidance and patience, and for always holding me to the highest professional standards.

I would like to thank the other members of my dissertation committee - Matthew Shapiro, Lucas Davis, and Kathy Yuan - for their valuable comments and suggestions in writing this dissertation.

I want to thank my Aunt Ettie and my grandfather Nathan Hocherman for their love and encouragement.

I want to thank my parents, Bernard and Susan, and my sister Marisa for supporting me throughout my life and for helping me with this achievement.

Finally, I want to thank Bonnie, who stayed in Ann Arbor far longer than she intended and stood by me through five years of graduate school. You've helped me celebrate my accomplishments and weather my setbacks and for that I am eternally grateful. I couldn't have done this without you.

TABLE OF CONTENTS

LIST OF FIGURES

<u>**Figure**</u>

LIST OF TABLES

Table

CHAPTER I

Introduction

Rising energy prices in recent years have renewed interest in the effects of global commodity price increases, and energy price shocks in particular, on the U.S. economy. Understanding how the economy evolves in response to these shocks is important to policy makers and professional economists. One reason economists and financial market analysts are interested in commodity prices is that they are thought to be good predictors of inflation. Better predictions of inflation in turn are thought to be useful for understanding the Fed's policy decisions. A number of researchers have shown that commodity prices were good predictors of inflation during the 1970s and early 1980s. However, the predictive power of commodity prices seems to have disappeared since the mid-1980s. In Chapter II, I demonstrate that this is not the case. The apparent breakdown in the predictive power of commodity prices is shown to be the result of using broad commodity price indices that often assign all commodity prices a fixed and equal weight regardless of their predictive content. Recent advances in the forecasting literature show that such weights can perform poorly in out-of-sample forecasting, and that data-based weights are preferable. Using several state-of-the-art forecast combination methods, including bagging, Bayesian shrinkage estimation, static factor models, and Bayesian model averaging, that implicitly chose the appropriate weights for individual commodity prices, I find that commodity prices are good predictors of inflation during the 1990s and early 2000s. Furthermore, commodity prices contain information about future inflation not captured by the leading principal components of a larger set of macroeconomic and financial variables. These improved inflation forecasts, however, do not necessarily improve our understanding of Fed policy decisions. Specifically, inflation forecasts based on commodity prices do not produce more accurate predictions for short-term interest rates, suggesting that commodity prices contribute little to a better understanding of monetary policy.

Economists are interested in commodity prices non only as leading indicators for inflation, but also because of their real effects on the economy. Rising prices for energy commodities such as crude oil, gasoline, and natural gas are considered particularly worrisome because they require households to devote a larger share of their budgets to energy consumption and a smaller share to non-energy consumption. Though it is widely accepted that rising energy prices decrease households' purchasing power for non-energy goods and services, little is known about the size of these resulting declines in purchasing power, or about their effects on consumer expenditures, consumer expectations, and aggregate unemployment. Chapter III (joint with Lutz Kilian), estimates their effects and characterizes the channels of transmission of energy prices to the economy. Several popular models of the transmission of energy price shocks suggest that consumption will fall when energy prices increase but remain unresponsive to energy price declines. The empirical evidence for such asymmetries provided by the existing literature, however, is not based on formal statistical tests. Using such tests, we show that there is no compelling evidence that real consumption, consumer expectations, or aggregate unemployment respond asymmetrically to unanticipated increases and decreases in energy prices. This result has immediate implications for models of the transmission of energy price shocks to the economy. Asymmetric transmission mechanisms are thought to be necessary for explaining why real GDP fell sharply in 1979 following an increase in energy prices, but was flat in 1986 following a decrease in energy prices. Using historical decompositions to analyze how consumption would have evolved in the absence of the 1986 decline in energy prices, we demonstrate that there actually was a boom in consumption in 1986, and that this boom was similar in magnitude to the decline in real consumption in 1979. While there indeed is an asymmetry in the real GDP growth data for 1979 and 1986, 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 and to the response of investment in the petroleum and natural gas industry to the collapse of OPEC in late 1985.

When we estimate the response of real consumption to an unanticipated increase in energy prices, we find that real consumption falls by more than the associated decline in discretionary income. This suggests that other transmission mechanisms, such as an increase in precautionary savings or an increase in the operating cost of energy-using durables, are necessary for explaining the overall response. Our analysis allows us to make predictions about the effects of rising energy prices on household expenditure patterns. For example, a 25 cent-per-gallon increase in the price of gasoline should cause a household with $\$4,000$ to spend per month to cut back its expenditures by \$41 one year later.

It has been suggested that energy prices are less important for the economy today than they were during the 1970s. Our analysis demonstrates that the response of real consumption to an increase in energy prices is larger for the 1970-1987 period than for the 1988-2006 period, consistent with this view. We trace 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.

Whereas Chapter III studies the effects of energy price shocks on consumption, Chapter IV (joint with Lutz Kilian) focuses on the effects of energy price shocks on business fixed investment in structures and equipment. This question is important because if there is evidence of an asymmetric response of real GDP to energy price shocks it must come from investment, not from consumption. Moreover, several leading theories that predict asymmetries in the response of durables consumption, apply to business fixed investment as well. While we do find that certain types of structures and equipment investment respond asymmetrically, the pattern of asymmetries that we observe is inconsistent with the predictions of these leading theories. Rather this evidence is an artifact (1) of the aggregation of mining-related expenditures by the petroleum industry and all other expenditures, and (2) of ignoring an exogenous shift in investment mainly caused by the 1986 Tax Reform Act. Once we control for these factors, the evidence for asymmetries mostly disappears. Furthermore, except for investment in mining structures, mining and oilfield machinery, and transportation equipment, most forms of structures and equipment investment are unresponsive to unanticipated energy price shocks. Since the 1970s, the contribution of energy prices to fluctuations in non-mining investment has been negligible.

The findings in this thesis suggest that it is necessary to revise our thinking about the transmission of energy price shocks to the economy. Researchers typically believe that frictions associated with sectoral reallocations are important channels through which energy price shocks are transmitted to the economy. These frictions in turn imply that real consumption, real investment, and aggregate unemployment will respond asymmetrically to energy price increases and decreases. We find no compelling evidence for asymmetries in the response of these variables, suggesting that sectoral reallocations are not a significant transmission mechanism. This does not mean, however, that the effects of energy price shocks are necessarily small. In addition to a discretionary income effect, we find that energy price shocks result in a sizeable precautionary savings effect and a sizeable operating cost effect on energy-using durable goods. The extent to which each of these channels is important depends on the type of expenditure in question. Besides mining-related investment in structures and equipment, the types of expenditure that matter most for the transmission of energy price shocks to the economy include consumer and firm expenditures on motor vehicles and residential fixed investment in housing.

CHAPTER II

Commodity Prices, Inflation Forecasts, and Monetary Policy

2.1 Introduction

This paper examines the usefulness of commodity prices for predicting consumer price inflation. Inflation forecasts are routinely used by professional economists for predicting monetary policy and the path of short-term interest rates. A common descriptive model of monetary policy states that the Federal Reserve increases or decreases short term interest rates in response to changes in the rate of consumer price inflation and a measure of real activity such as deviations in unemployment from its natural rate. Though there are many specifications of this “Taylor rule” one that appears frequently in the empirical literature postulates that short term interest rates are adjusted linearly in response to forecasts of future inflation and real activity. Thus, forecasts of inflation are important statistics for anyone interested in monetary policy. Since higher commodity prices are often thought to presage higher rates of inflation, the evolution of commodity prices in turn is intensively scrutinized by professional economists.

The relevance of commodity prices to monetary policymakers is not uniformly accepted. The debate revolves around the usefulness of commodity prices as a leading indicator for inflation. Conventional wisdom holds that increases in the prices of certain commodities such as gold or petroleum are indicators of higher future inflation. For example, a July 15, 2005 article in the New York Times reports that:

> Wholesale prices, an early indicator of inflation in the economy, were unchanged in June, the government reported today, indicating that most businesses are not passing on rising energy costs to consumers.

This quote illustrates an important issue. While market participants seem to look towards intermediate commodity prices as a signal for future consumer price inflation, the link does not always materialize. The indicator role of commodity prices is also debated by policymakers. For example, in a speech to the American Economic Association in January 2004, Fed governor Ben Bernanke stated that:

> ...the direct effects of commodity price inflation on consumer inflation are empirically minuscule, both because raw materials costs are a small portion of total costs and because part of any increase in the cost of materials tends to be absorbed in the margins of final goods producers and distributors. Accelerations in commodity prices comparable to or larger than the most recent one occurred following the 1981-1982 and 1990-1991 recessions, as well as in 1986-1987 and 1999, with no noticeable impact on inflation at the consumer level.

According to Fed governor Donald Kohn, however:

> Commodities and imports are only a small part of what we consume, and changes in their prices as well as in the price of energy usually do not affect measured core consumer inflation very much. But the recent situation has been notable for the size and number of price movements going in the same direction, so that even with limited pass-through of individual price movements, the total effect probably has been significant. Judging from the results of statistical models incorporating the factors we have been examining, increases in commodity, energy, and import prices together might have boosted core consumer inflation on the order of roughly 1/4 to 1/2 percentage point over the past four quarters.

There are several theoretical justifications for using commodity price inflation to predict future consumer price inflation. Commodities are often purchased in spot markets, so their prices are flexible, as opposed to final goods prices, which tend to be sticky. A macroeconomic shock will affect commodity prices contemporaneously, but will affect consumer prices only with a lag. For instance, the first signs of an increase in the aggregate demand for final goods and services will be found in commodity prices, and only subsequently in final goods prices. Alternatively, a cost shock might originate in the market for certain commodities such as oil. To the extent that these commodities are important for the domestic production of final goods, the costs will be passed through to final goods prices. But since final goods prices are sticky, they will adjust more slowly.[1]

An empirical link between commodity prices and consumer prices, however, might be difficult to detect for a number of reasons. First, higher production costs that result from higher commodity prices might not be passed on to consumers in the form of higher final goods prices. This could be the case if producers experience increased competition from free trade. As Bernanke points out in his speech, "...part of any increase in the costs of materials tend to be absorbed in the margins of final goods producers and distributors." Second, the argument for a theoretical link between commodity prices and final consumer goods prices typically postulates that a specific monetary or supply shock is driving the positive relationship. If, however, consumer price inflation is subject to several offsetting shocks at any one time, the unconditional covariance between consumer prices and commodity prices could be small and insignificant. Commodity prices would then have little predictive ability, despite the existence of a theoretical link. Thus, commodity prices' relevance to forecasters remains an open question.

A natural and conventional metric for judging the performance of a particular forecasting model is the recursive prediction mean-squared error (PMSE) in simulated out-of-sample forecasts. If adding commodity price inflation to the information set of a forecaster reduces the PMSE of the inflation forecast, then I can conclude that commodity price inflation is a leading indicator of final goods price inflation. However, a more accurate inflation forecast will not necessarily indicate whether commodity prices are useful for describing monetary policy. For example, policymakers might put little weight on the rate of commodity price inflation when forming there own projections about future consumer price

[1] See Boughton and Branson (1991) for a model of commodity and industrial prices.

inflation. This suggests another metric for gauging the usefulness of commodity prices, namely the fit of a Taylor rule. If adding commodity price inflation to the information set of the policy rule, via the consumer price inflation forecast, reduces the root mean-squared error of the model, then I can conclude that commodity price inflation is a useful predictor of monetary policy.

An early appeal to the forecasting ability of commodity prices is Sims' (1992) demonstration that the inclusion of a commodity price index in a VAR can partially resolve the often noticed "price puzzle", i.e., a short-term rise in the price level following an exogenous monetary tightening. Sims made the case that the consecutive rise in the price level and short-term interest rate was the result of an endogenous policy response to higher expected inflation. In order to identify exogenous movements in the interest rate correctly, the econometrician must control for expectations of future inflation. Sims proposed including a commodity price index for this purpose. Since then, a commodity price index has been included in numerous VAR studies of monetary policy (see, e.g., Bernanke and Mihov 1998). Subsequent research has examined the link between commodity prices and consumer price inflation. Hanson (2004), revisiting the price puzzle, demonstrates that adding certain commodity price indices to a VAR forecasting model can significantly reduce the PMSE of a price level forecast, relative to a VAR without commodity prices. Hanson finds that adding an agricultural raw materials price index to the VAR reduces the PMSE of a CPI forecast by 35% at a one year horizon. The accuracy gains, however, are much smaller at shorter horizons.

Other recent studies dispute the usefulness of commodity prices as predictors of future consumer price inflation. Blomberg and Harris (1995) demonstrate that certain commodity and producer price indices are good indicators of future inflation during the 1970s and early 1980s, but that they lose their ability to predict changes in core CPI inflation starting in the mid 1980s. They show that the importance of commodities for U.S. production wanes around this time and point to the absence of significant food and oil price shocks that could affect consumer prices. Cecchetti, Chu, and Steindel (2000) report mixed results for the forecasting ability of commodity prices. The authors add a series of individual indicators of future inflation, such as commodity price indices, M1 growth, exchange rates, interest rates, and unemployment rates, to an autoregression for CPI inflation. They demonstrate that the model with certain commodity price indices, such as the Journal of Commerce price index for industrial materials, the NAPM price diffusion index, a gold price index, or an oil price index, outperforms the autoregression more often than most other indicators over the 1975-1998 period. An autoregression, however, often outperforms these forecasting models. Cecchetti et al. also find a *negative* covariance between certain commodity price indices and final goods price inflation, leading them to dismiss most of their other results. Stock and Watson (2003) show that forecasting models with a commodity price index can increase the accuracy of a forecast for the *change* in monthly CPI inflation relative to an autoregressive benchmark for the 1971-1984 period by 21%. Adding the commodity price index to the model for the 1985-1999 period, however, decreases the forecast accuracy by 26%. Finally, Furlong and Ingenito (1996) find that the ability of a non-oil commodity price index to Granger cause CPI inflation weakens considerably around 1984.

While this literature constitutes an important first step towards uncovering the leading indicator role of commodity prices for consumer price inflation, these papers suffer from several drawbacks. First, all of these papers forecast inflation with broad commodity price indices. These indices may include some commodity prices that are irrelevant for predicting future inflation, and thus may increase the variance of the forecast. They might also aggregate the information from a set of commodity prices in a way that is less than optimal; for instance, if all commodities are given equal weight in the index. Second, many of the aforementioned papers do not use state-of-the-art forecasting techniques. For example, Hanson (2004), and Blomberg and Harris (1995) use unrestricted VAR forecasting models that are known to perform poorly in out-of-sample forecasting (see, e.g., Stock 2001). Third, a number of these papers do not evaluate commodity prices on a genuine out-of-sample basis. For example, Boughton and Branson (1991) and Blomberg and Harris (1995) estimate model parameters for the first half of their sample and use them to forecast inflation for the quarters in the second half of their sample. While their forecasts are out-of-sample, they are based on information sets that do not include the most up-to-date set of data. Cecchetti, Chu, and Steindel (2000) use a bivariate iterated multi-step ahead forecasting model, yet they only forecast future values of inflation and use ex-post realizations of commodity prices instead of their forecasted values. Their forecasts are based on data that would not have been included in a forecaster's information set. Finally, the existing literature evaluates the usefulness of commodity prices based on their ability to reduce the PMSE of an inflation forecast and largely ignores the question of whether these forecasts help us understand monetary policy. While quadratic loss is a commonly used criterion for evaluation, most users of inflation forecasts will be interested in inflation as it relates to the perceived policy reaction function of the central bank. There is no a priori reason to believe that a forecasting model that forecasts well by the PMSE criterion, also is a good predictor for the path of short-term interest rates.

In this paper I use a set of 46 individual commodity prices that allows me to determine which commodity prices are good predictors of future inflation. I apply state-of-the-art forecasting methods to combine the information content of this data set. These methods, including bagging, Bayesian model averaging, shrinkage estimation, and factor models, implicitly determine the appropriate weights for combining the 46 regressors. The forecasting models are evaluated on a recursive out-of-sample basis in order to simulate the experience of a forecaster who uses all information available to him at a particular point in time. Recursive out-of-sample forecasting is used in a number of recent papers including Stock and Watson (1999, 2002, and 2003), Massimiliano, Stock and Watson (2004), Wright (2003), Bernanke and Boivin (2003), Hanson (2004), and Inoue and Kilian (2005). I also focus on single equation, direct h-step ahead forecasting methods to avoid the problems associated with multivariate models.

Several papers, such as Blomberg and Harris (1995), Furlong and Ingenito (1996) and Stock and Watson (2003) find a structural break in the relationship between commodity prices and consumer price inflation during the early 1980s, with the forecasting power of these predictors weakening around this time. Since there is little argument in the literature that commodity prices were good leading indicators of consumer price inflation during the 1970s and early 1980s, this paper evaluates this relationship over the 1993-2004 period for which there is controversy. In the last part of this paper, I assess the usefulness of commodity prices in improving the predictive power of a monetary policy reaction function. Assuming that Fed behavior is described by a forward-looking Taylor rule, I compare actual interest rates to those implied by the model when inflation is forecasted with commodity prices, and when it is not.

The accuracy of each forecasting method employed in this paper is evaluated relative to an inflation-only model, which uses lags of inflation to predict future inflation. This basis for comparison is common to a number of research papers about inflation forecasting.[2] Yet forecasters and policymakers today have access to tremendous amounts of data that are useful for predicting future inflation, such as wages, industrial production, stock prices, and exchange rates. With this in mind, I also examine inflation forecasts for which commodity prices are added to a factor model based on a larger information set containing more than 100 macroeconomic time series.

The remainder of this paper is organized as follows. Section 2 examines forecasting methods based on individual commodity prices, based on individual price indices, and based on forecast combinations of individual commodity prices. I present rescursive PMSE results for monthly CPI inflation forecasts. I also present results for quarterly CPI inflation and GDP deflator inflation forecasts. The latter analysis is motivated by the analysis in Barsky and Kilian (2002) that shows that under certain conditions, the prices of imported commodities may have a bigger impact on CPI inflation than on GDP deflator inflation.

[2]See, for example, Massimiliano, Stock, and Watson (2004); Stock and Watson (2003); Wright (2003); and Bachmeier, Li, and Liu (2005).

Section 3 considers the alternative factor model benchmark. I examine factor model forecasts of CPI inflation, core CPI inflation, and core Personal Consumption Expenditure (PCE) price inflation. The latter two measures exclude the more volatile food and energy prices. Section 4 addresses the ability of commodity prices to describe the behavior of monetary policymakers. I compare estimated short term interest rates from a Taylor rule to actual interest rates, when commodity prices are included in the inflation forecast and when they are not. Section 5 examines the usefulness of alternative forecasting model specifications based on error correction terms involving the deviation of the logged commodity price from the logged CPI. Section 6 concludes.

2.2 Forecasting Inflation with Commodity Prices

This section addresses the predictive ability of commodity prices for inflation. Some of the questions of interest are: Does adding commodity prices to the forecasting model improve the accuracy of inflation forecasts? Which commodity prices are the best predictors? Are industrial commodity prices more informative than food and beverage prices? Do oil prices in particular help forecast inflation? Should one rely on individual commodity prices or commodity price indices? Are there benefits from using state-of-the-art forecast combination methods rather than equal-weighted indices? If so, which methods work best? Do the results depend on the definition of inflation? Do the results apply equally to long and to short horizons? Is there information in commodity prices not captured by the leading principal components of a broader set of macroeconomic and financial variables?

A commodity or commodity group's predictive ability is measured as the recursive PMSE of a forecasting method that includes these commodity prices, relative to a model without them. I perform a simulated out-of-sample forecasting exercise for different subsets of commodity prices, different h-month forecasting horizons, and various forecasting methods. The dependent variable in each case is the annualized h-period growth rate of the monthly or quarterly price series, measured as the first difference of the logarithm of the series. The predictors include one or more lags of the annualized 1-month (or in some cases 1-quarter) growth rate of consumer prices, and the corresponding growth rate of commodity prices. The commodity price data are from the IMF's International Financial Statistics database. Table 1 lists the different commodity prices and commodity price indices used in this study. The commodities in the data set are categorized into three groups: food and beverage commodities (28 commodities), materials commodities (18 commodities), and all commodities (46 commodities). The commodity prices are monthly averages of world spot market prices measured in U.S. dollars. Each forecasting method is evaluated relative to an inflation-only direct-forecasting model. I chose this benchmark based on a comparison of the forecasting accuracies of a direct-forecasting model, an iterated AR model, and an iterated AR in first differences. The lag order for the inflation-only model is selected by minimizing the Akaike Information Criterion (AIC).[3]

Let h denote the monthly forecast horizon, p the lag order for monthly inflation chosen by minimizing the AIC, and π_t^h the annualized h-month rate of consumer price inflation between time t and t+h. Then the inflation-only model used as the benchmark for the forcast accuracy comparison is given by:

$$\pi_t^h = \alpha + \rho_1 \pi_{t-1} + ... + \rho_p \pi_{t-p} + \epsilon_t$$

In the remainder of this paper, the predictive ability of commodity prices is investigated based on a number of different forecasting methods. My analysis provides a more comprehensive assessment of the predictive relationship between commodity price inflation and consumer price inflation than the results based on vector autoregressions or univariate inflation models augmented with a single commodity price or commodity price index, as in Hanson (2004), Blomberg and Harris (1995) and Cecchetti, Chu, and Steindel (2000). The accuracy of each forecasting method is measured by the ratio of its recursive PMSE rela-

[3]The recursive PMSE of the inflation-only model is smaller when the lag order is selected with the AIC than the Schwarz Information Criterion (SIC), reflecting the presence of a large moving-average component in inflation (see Inoue and Kilian (2006) for further discussion).

tive to that of the inflation-only model. A ratio less than 1 indicates that the model with commodity price inflation is more accurate, on average, than the inflation-only model. The initial parameter estimates for each forecasting model are based on data from the 1980-1992 period. Recursive forecast accuracy results are reported for the 1993-2004 period. This evaluation window is chosen because it is large enough to obtain reasonably accurate estimates of the model parameter values, yet it still provides forecasts for about 144 months.

2.2.1 Forecasting Monthly CPI Inflation with Individual Commodity Prices and Price Indices

The first set of results assesses the predictive power of each individual commodity price and of commonly used commodity price indices. For this purpose, I combine lags of consumer price inflation, with lags of the inflation rate for a single commodity price or commodity price index in the form of a distributed lag model.[4] The lag orders for consumer price inflation and commodity price inflation are selected jointly by minimizing an AIC, with an upper bound of 12 lags for each predictor. Letting $\Delta x_{i,t}$ denote the commodity price inflation rate for commodity price i at time t, the distributed lag model is given by:

$$\pi_t^h = \alpha + \rho_1 \pi_{t-1} + ... + \rho_p \pi_{t-p} + \gamma_1 \Delta x_{i,t-1} + ... + \gamma_q \Delta x_{i,t-q} + \epsilon_t$$

Results for the distributed lag model forecasts of monthly CPI inflation are presented in Table 2. It is widely believed that industrial commodity prices have predictive power for monthly CPI inflation. Indeed, certain individual metals prices, such as copper, lead, nickel, and zinc prices improve the forecast at various horizons, but the gains are not uniform across horizons or commodities.

Perhaps the series most often singled out for its alleged predictive ability for inflation is the price of crude oil. Table 2 confirms that crude oil price inflation is particularly useful for predicting CPI inflation at the 1 and 3 month forecast horizons, although at longer

[4] A distributed lag model is used in Stock and Watson (2003) to assess the predictive content of different macroeconomic and financial variables.

horizons, the accuracy gains are either modest or non-existent. Among other materials prices, hard sawn wood price inflation is useful at medium and long forecast horizons, in some cases lowering the PMSE by 12% relative to the inflation-only model, while fine wool peforms moderately well at short horizons. Food and beverage prices almost always increase the PMSE of the forecast model. Two notable exceptions are cocoa bean prices and coconut oil prices.

One of the problems with using individual commodity prices is that individual series tend to be noisy. This suggests aggregating series into commodity price indices. On the other hand, a number of existing papers have shown that commodity price indices have not been useful for forecasting CPI inflation since the mid-1980s.

Table 2 confirms that a broad non-fuel price index will make the forecasting model less accurate than an inflation-only model at most forecast horizons. These results are similar to those in Blomberg and Harris (1995) and Stock and Watson (2003). Turning to the more narrowly defined price indices, I find that the various food and beverage indices as well increase the PMSE of the forecast at every forecast horizon, sometimes by more than 40%. On the other hand, the industrial inputs price index works well at most forecast horizons, typically decreasing the PMSE of the forecast by 4%-9% compared to the inflation-only model. The agricultural raw materials and metals price indices also generate some accuracy gains, but they occur at fewer horizons and tend to be more modest than those from the industrial inputs price index. Table 2 suggests that combining information from many commodity prices can be effective in reducing noise, but it also highlights important differences in the predictive ability across different types of commodities. Even for materials prices, however, the gains in accuracy from using the industrial inputs price index are at most 10%. One possible reason is that these indices are based on equal weights. A reasonable conjecture is that methods that choose these weights based on the data may generate superior forecasts of inflation. I will consider several such methods below. As forecast combination methods depend on auxiliary parameters such as prior variances and critical values, for these methods I report only the best forecasting results after searching over alternative choices for these parameters.

The IMF data set of commodities does not include precious metals such as gold or silver. The latter commodities matter not so much because of their importance as industrial inputs, but because they are often considered a hedge against inflation. This view dates back to the 1970s. Given the absence of high and sustained inflation rates since the early 1980s, the relevance of gold and silver prices for inflation forecasting is no longer obvious.

In unreported results, I examined the predictive value of gold and silver prices. I found that neither price is a useful predictor individually. In fact, adding the price of gold to the inflation only model worsens the forecast accuracy in virtually every case. The performance of silver prices is only marginally better.

Moreover, when adding the price of gold and silver to the baseline data set of commodity prices, the forecast accuracy of the forecast combination methods considered below is worse in most cases and marginally better only in rare cases. This is not surprising since the main rationale for the predictive power of gold and silver prices does not seem to apply to the period after the mid-1980s. Thus, I choose not to report these additional results.

2.2.2 Combination Forecasts of Monthly CPI Inflation Using Multiple Commodity Prices

Unrestricted Model Forecasts

A useful benchmark for the forecast combination methods is the performance of forecasts from unrestricted models that include the most recent lag of all commodity prices in addition to lags of CPI inflation. This unrestricted model is given by:

$$\pi_t^h = \alpha + \rho(L)\pi_t^1 + \sum_{i=1}^{n} \delta_i \Delta X_{i,t-1} + \epsilon_t$$

The lag length for consumer price inflation is selected via the AIC with an upper bound of 12 lags. The model is estimated by OLS. I consider the set of materials price inflation regressors, that of food and beverage price inflation regressors, and the set of all commodity price inflation regressors. Table 3 shows that the unrestricted model outperforms the inflation-only model only when it includes materials price inflation and only at the 1-month forecast horizon. In all other cases, the inflation-only model is more accurate. Table 3 illustrates the difficulty in improving on simple equal-weighted commodity price indices. A partial explanation for these poor results is that estimating the unrestricted model with 28 or 46 regressors in addition to several lags of inflation leaves the model with few degrees of freedom. OLS estimation of the model's parameters with few degrees of freedom will leads to poor out-of-sample prediction when the number of parameters is too large.

Bayesian Shrinkage Forecasts

If the poor out-of-sample predictions obtained from the unrestricted combination forecast model are due to OLS estimation of the model, the results might be improved by imposing additional structure on the parameters. For the next combination forecast method, I estimate the unrestricted model parameters with a Bayesian shrinkage estimator. The prior mean reflects the notion that each commodity price regressor has no predictive power for the dependent variable (i.e., a coefficient equal to zero). Thus, the parameter estimates are shrunk towards zero. For the prior standard deviation of these parameters I explore alternative values $\lambda \in \{0.01, 0.02, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 1, 2, 5, 100\}$. The smaller the variance of the prior distribution, the higher the degree of shrinkage, and the more weight is placed on the prior distribution. So, as $\lambda \to \infty$, the shrinkage estimator reduces to the OLS estimator of the unrestricted model. Finally, following Wright (2003) and Inoue and Kilian (2005), the prior means for the intercept and lags of inflation are based on the fitted values of a regression of consumer price inflation on its lags and an intercept over the pre-sample period (1957-1979), with the corresponding prior variance set to infinity.

Table 3 shows that shrinkage estimation of the unrestricted model generates forecasts that are at least as accurate as the inflation-only model for all forecast horizons and all sets of regressors. Furthermore, the shrinkage forecast is more accurate than the unrestricted model estimated by OLS. For food and beverage prices and for the full set of commodity prices, using a shrinkage estimator reduces the relative PMSE by 30 to 50 percentage points compared to the unrestricted model estimated by OLS. The shrinkage forecast with materials price inflation is preferable for the 1,3, and 12 month horizons, reducing the PMSE by 10%, 5%, and 4% respectively. In all cases, a high degree of shrinkage ($\lambda \leq .2$) works best for all sets of predictors, with slightly larger degrees of shrinkage preferable for food and beverage prices and for the full set of commodity prices. I conclude that shrinkage estimation strictly dominates OLS estimation of the unrestricted model. In addition, shrinkage estimation of the unrestricted model performs at least as well, and typically better than, forecasts based on standard commodity price indices.

Bagging Forecasts

An alternative method for dealing with the over-fitting problem is to use statistical tests to determine which commodity prices are relevant to the forecasting exercise. For example, one may decide whether to include a given predictor based on a 2-sided t-test.[5] Small changes in the data, however, can overturn the results of such tests, leading to an unstable pre-test decision rule and an inflated out-of-sample forecast variance. Bagging, or bootstrap aggregation, is a method for reducing this model instability while exploiting the information content of a large number of predictors (see, e.g., Buhlmann and Yu 2002, Inoue and Kilian 2005). Bagging involves fitting the unrestricted model, creating a large set of bootstrap samples, applying the pre-test to each bootstrap sample, and generating a forecast for each bootstrap sample based on the model selected for that sample. The bagging forecast is defined as the average of these bootstrap forecasts.[6] Following Inoue and Kilian (2005) I set the number of bootstrap re-samples to 100. For the pre-test I consider the following critical values: $c \in \{1.96, 2.1701, 2.5758, 2.807, 2.8782, 3.0233, 3.2905, 4.4172, 4.5648, 4.8916, 5.0263,$

[5] Boughton and Branson (1991) use a t-test to determine which regressors to include in an aggregate commodity price index. Stock and Watson (2005a) demonstrate that model selection with pretests has a shrinkage interpretation.

[6] Inoue and Kilian (2005) demonstrate that a bagging model with a large number of macroeconomic time-series predictors can reduce the PMSE of an inflation forecast relative to an inflation-only model.

5.3267}. The AIC with an upper bound of 12 lags is used to select the lag order for CPI inflation. As in the unrestricted model, I include one lag of each $\Delta x_{i,t}$. The bagging exercise is performed on the original set of commodity price inflation regressors as well as on two sets of orthogonalized regressors. The first is based on an orthogonalization of the regressors via a Choleski decomposition. The second transformation replaces the original regressors by their principal components. For a comparison of these alternative bagging methods see Inoue and Kilian (2005). In my case, the bagging forecast results are broadly similar across the three specifications of the predictors. I therefore present results for the standard bagging forecasts on the original set of commodity price inflation regressors.

Like the Bayesian shrinkage forecast, the bagging method reduces the PMSE of the forecast relative to the inflation-only model at the 1- and 3-month horizons for all three sets of predictors. The largest accuracy gains are obtained when using materials prices, in which case the PMSE falls by 11% and 7% at the 1- and 3-month horizons, respectively. Bagging loses its predictive accuracy at the 12-month horizon with PMSE ratios greater than 1 for all sets of predictors. In all cases, bagging is more accurate than the unrestricted model, suggesting that model selection is beneficial to the forecast. The best critical value for the 2-sided t-test depends on the regressors included in the model and the forecast horizon. Bagging with materials price inflation requires a lower critical value (2.807 or 3.2905 work best) than bagging with the set of food and beverage prices or the full set of commodity price prices (4.8916, 4.5648, or 5.0263 work best).

Static Factor Model Forecasts

The static factor model provides yet another method for forecasting a single time series with a large number of predictors. The static factor model reduces the dimension of the model by combining the information content of these predictors into a small number of principal components.[7] Factors are estimated for each set of $\Delta x_{i,t}$ regressors; materials prices, food and beverage prices, and the full set of commodity prices. For each forecast

[7] For details on the estimation of the factors, see Stock and Watson (1999).

horizon and initial estimation window, lags of consumer price inflation are augmented with the first r estimated commodity price inflation factors where $r \in \{1, 2, 3, 4, 5\}$. I also estimated dynamic factor models where the lag order for the first r factors is selected jointly by the AIC or the SIC. These models typically performed worse than the static factor models in out-of-sample forecasting.

Table 3 shows that the static factor models produce more accurate forecasts than the inflation-only model for all forecast horizons when materials price regressors are used. On the other hand, when food and beverage prices or the full set of commodity prices are used, the inflation-only model is generally at least as accurate as the factor model. For the 1-month and 3-month forecast horizons, including the first three materials price inflation factors works best, reducing the PMSE by 8% relative to the inflation-only model. For the 12-month forecast horizon, including only the first materials price inflation factor is preferrable.

Equal-weighted Forecasts and Bayesian Model Averaging

Stock and Watson (2003) show that forecast accuracy can be improved by taking a simple average or the median of forecasts generated by adding a single lag of $\Delta x_{i,t}$ to the inflation-only model. Wright (2003) demonstrates that equal-weighted forecasts can be improved further by taking a weighted average of different forecasts, where the weights are based on the Bayesian posterior probability of each forecasting model conditional on assuming that each model is equally likely. As before, the lag order for inflation is selected by the AIC. The prior mean for each forecasting model is $\frac{1}{M}$, where M is the number of commodity prices in a given commodity price group. The prior for the model parameters is analagous to that used for Bayesian shrinkage estimation. The prior standard deviation for commodity price inflation regressors is $\phi \in \{0, 0.01, 0.05, 0.1, 0.2, 0.3, 0.4, 0.5, 1, 2, 5, 100\}$.[8]

Table 3 shows that Bayesian model averaging (BMA) and the equal-weighted forecast

[8]The higher the value of ϕ, the more weight is put on the data when constructing the posterior probability. When $\phi = 0$, for instance, the model reduces to an equal-weighted forecast.

are at least as accurate as the inflation-only model in all cases. Similar to Wright (2003), I find a number of cases in which the best BMA forecast is superior to the equal-weighted forecast, particularly when materials price inflation or the full set of commodity price inflation regressors are included in the model. For food and beverage price inflation, however, there is no difference between the two methods. Typically, a small degree of shrinkage ($\phi \leq .4$) works best.

2.2.2.1 Which Models Work Best?

At this point it is useful to summarize the findings for monthly CPI inflation. At short horizons, the price of oil is the best single predictor with gains relative to the inflation-only model of 8% and 7%, respectively, at the 1-month and 3-month horizon. Oil prices also have better predictive ability than even the best commodity price indices. However, modern forecast combination methods such as bagging and Bayesian model averaging tend to perform as well or better than oil prices at short horizons, provided that the set of potential predictors includes either materials prices only or all commodity prices. For example, the PMSE reduction from bagging materials prices is 11% and 7%, respectively, at the 1-month and the 3-month horizon. BMA applied to all commodity prices reduces the PMSE by 11% at the 1-month horizon. Forecast accuracy gains based on food and beverage prices alone were modest at best. These findings overturn the consensus that commodity prices have lost their predictive ability since the 1990s.

Although the forecast accuracy improvements may seem small by the standards of PMSE reductions reported in the literature for the 1970s and 1980s, it is important to keep in mind that the latter results were based on a much larger set of predictors rather than just commodity prices. Moreover, as I will show in Section 3, the large-dimensional forecasting models of the type that appeared to predict well in the 1970s and 1980s actually are considerably less accurate than the model based on commodity prices only for 1993-2004 with gains of at most 5% relative to the inflation-only model.[9]

[9] The extent of the forecast accuracy gains also is sentitive to the specification of the inflation-only model. For example, if

At longer horizons, the performance of forecast combination methods deteriorates. For example, bagging is actually less accurate than the benchmark model at the 12-month horizon, and Bayesian model averaging is only marginally superior to the inflation-only model. Moreover, there is no improvement on an equal-weighted industrial inputs price index at the 12-month horizon.

A natural question is whether the reductions in the recursive PMSE at short horizons are statistically significant. There are currently no statistical tests that would allow us to answer this question. The root of the problem is that I compare forecasting methods rather than forecasting models. Thus, a given forecasting method will tend to select different models at each point in time, invalidating the distribution theory for standard tests of equal predictive ability.

2.2.3 Forecasting Quarterly CPI Inflation and GDP Deflator Inflation

Barsky and Kilian (2002) demonstrate that an increase in the price of imported commodities such as oil will unambiguously increase the price of gross output, but could increase *or* decrease the price of value added. This suggests that higher oil prices should presage an increase in measures of the price of gross output, such as the CPI, but not necessarily in measures of the price of real value added, such as the GDP deflator. Thus, the relative merits of other commodity prices for forecasting GDP deflator inflation and CPI inflation are unclear.

To control for the effects of forecasting a quarterly series (GDP deflator) instead of a monthly series (CPI), I compare the forecasting results for quarterly deflator inflation to the results for quarterly CPI inflation. The dependent variable in these forecasts is the annualized, h-quarter growth rate in the GDP deflator or the CPI. $\Delta x_{i,t}$ is the annualized, 1-quarter growth rate of commodity price i between time t and *t+1.* I forecast at horizons of 1-6 quarters. Because the GDP deflator is available only on a quarterly basis, the time

I replace the specification that appeared most accurate for my dataset by an iterated model of the type used as a benchmark in other studies or if I replaced the AIC used in lag-order selection by the SIC, the relative gains from using commodity prices for prediction may appear much larger.

series is substantially shorter than the monthly CPI series. Estimating many of the models with low initial estimation windows results in poor parameter estimates and inaccurate out-of-sample forecasts. For this reason, and to keep the forecast evaluation period consistent with the evaluation period for the monthly CPI inflation forecasts, I report the quarterly forecasting results for the 1995-2004 period.

2.2.3.1 Forecasting with Individual Commodity Prices and Price Indices

The results for the distributed lag model are shown in Tables 4a and 4b. Table 4b shows no evidence that any of the commodity price indices are useful for predicting quarterly CPI inflation at any horizon. There also is little evidence that food and beverage prices individually have systematic predictive ability. On the other hand, there is limited evidence that materials may help in predicting quarterly CPI inflation at the 1-quarter horizon. The best performers are copper, tin, crude oil, lead, and fine wool. Only crude oil and hard sawn wood, however, seem to have substantial predictive ability at longer horizons.

Table 4a shows the corresponding results for quarterly GDP deflator inflation. Here there is evidence that at least the non-fuel price index as well as the industrial inputs price index and the agricultural raw materials index have limited predictive power at horizons of 1- and 2-quarters. That predictive ability is also reflected in individual materials and food and beverage series. Again crude oil and hard sawn wood as well as hard logs stand out among materials for having predictive power at longer horizons, while the predictive ability of food and beverage prices, to the extent that there is any, is largely limited to the first two quarters.

Of particular interest is a comparison of the predictive ability of crude oil prices for the two inflation series. The accuracy gains from crude oil prices tend to be larger for GDP deflator inflation at the 2-, 3- and 4-quarter forecast horizons, and larger for quarterly CPI inflation at the 1-, 5-, and 6-quarter forecast horizons. Thus, I find that the price of crude oil may be a good leading indicator for both inflation measures, at least for selected horizons. These results do not contradict the theoretical arguments presented by Barsky and Kilian (2002) regarding the link between oil prices and inflation measures following an exogenous oil supply shock. Barsky and Kilian concluded that, since most U.S. oil is imported, higher oil prices arising from exogenous oil supply shocks should lead to higher rates of gross output price inflation, but not to higher rates of value added price inflation, which seems consistent with their empirical evidence from the 1970s. In contrast, the evidence in this paper is for a period (1995-2004) that is recognized for high levels of demand in the oil market. Oil price inflation, GDP deflator inflation, and CPI inflation are all expected to rise following a demand shock, so there is no apparent ordering of the accuracy gains across inflation measures.

2.2.3.2 Combination Forecasts with Many Commodity Prices

For the other forecasting methods, I present the results in the same way as for the monthly CPI inflation forecasts, except that I restrict myself in Table 5 to the 1-, 2-, and 4- quarter horizons. The results in Table 5 show that forecast combination methods based on materials work best. The most favorable results are obtained for the shrinkage method, for bagging and for the factor models. The BMA method is more erratic. However, none of the forecast combination methods dominates a distributed lag forecast based on crude oil prices or the price of hard sawn wood at the 4-quarter horizon. By implication, this means that these methods also dominate forecasts based on commodity price indices. In summary, suitable forecast combination methods, as in the case of monthly inflation, appear useful for forecasting quarterly inflation. This result holds for both measures of inflation, although the extent of the gains varies with the inflation measure, the forecast combination method and the horizon.

2.3 Full Information Set Forecasts

The exercise in the preceding sections demonstrated that combination forecasts of CPI inflation that include lags of certain commodity price inflation regressors are often superior to inflation-only forecast models. While inflation-only benchmark models are often used in the forecasting literature to evaluate the predictive ability of various modeling techniques, these types of models are rarely used by professional forecasters. Their most obvious drawback is a failure to exploit the large number of macroeconomic and financial predictors that are available for forecasting inflation. Recent advances in computing power allow economists to analyze scores of data sets and estimate complex models at very low cost. Yet the loss of degrees of freedom caused by the inclusion of too many regressors in a model can severely limit our ability to fully exploit these gains. In recent years, factor models have emerged as a popular method for harnessing the information content of a large set of predictors, while limiting the dimension of the forecasting models. Stock and Watson (2002) and Bernanke and Boivin (2003) show that factor model forecasts, with factors estimated from a set of 215 macroeconomic predictors, outperform AR forecasts and VAR forecasts of inflation, unemployment, and industrial output. Given the appeal of these models, professional forecasters might want to know whether adding commodity prices to the information set of such a factor model forecast will improve its out-of-sample performance.

In this section, I first evaluate a static factor model forecast, relative to the inflation-only model, that uses factors estimated from a set of 112 macroeconomic time series used in Stock and Watson (2005b).[10] This data set includes series on real output; income; employment; hourly earnings; retail, manufacturing, and trade sales; consumption; housing starts and sales; inventories; stock prices; exchange rates; interest rates; interest rate spreads; and money stock and credit market conditions, but not commodity prices. I then add various subsets of the commodity price inflation regressors to the Stock and Watson data set and re-estimate the factors. In addition to materials prices, food and beverage prices, and the full set of commodity prices used in the previous exercises, I also add the price of crude oil and materials prices excluding the price of crude oil to the Stock and Watson regressors. The relative PMSEs are reported in Table 6. Since all ratios have

[10] The dataset from Stock and Watson (2005b) included 132 time series. In order to avoid duplicating series, I remove 20 variables that include producer prices or commodity prices.

the same denominator, the PMSE of the inflation-only model, a direct comparison of the predictive power of each set of predictors is possible. To conserve space, I report results for the static factor model when the first lag of the first factor is included. In addition to forecasts of CPI inflation, I also report results for forecasts of core CPI inflation and core PCE price inflation.

The first row of Table 6 shows that when only the 112 Stock and Watson predictors are used, the static factor model improves the accuracy of the CPI inflation and core CPI inflation forecasts, relative to the inflation-only model, but not the core PCE price inflation forecasts. Adding materials price to the Stock and Watson data improves the factor model forecasts of CPI and core CPI inflation. The size of the accuracy gains increase with the forecast horizon. For instance the factor model with just the 112 Stock and Watson regressors reduces the PMSE of the CPI inflation forecast by 5% at the one month horizon, and 3% at the 3-month and 12-month horizons. After adding materials prices to the information set, the accuracy gains from the factor models increase by 2 percentage points at the 1-month horizon, by 6 percentage points at the 3-month horizons, and by 8 percentage points at the 12-month horizon. This shows that commodity prices convey information not captured by the principal components of standard macroeconomic and financial datasets.

In contrast, relative to the larger set of predictors, oil prices alone are not informative. Adding the price of oil to the information set does not reduce the PMSE of the factor model.[11] Adding food and beverage prices or adding the full set of commodity prices to the information set does not improve the accuracy of the factor models either. This result again underscores the unique role of materials prices in predicting inflation.

[11] For factor models with two and three factors, adding just the price of oil may raise or lower the forecast accuracy of the factor model, depending on the horizon.

2.4 Commodity Prices and Federal Reserve Behavior

In this section, I investigate the usefulness of commodity prices for describing the behavior of monetary policymakers. Commodity prices reduce the PMSE of inflation forecasts, but it is unclear whether this forecast accuracy gain translates into better predictions of monetary policy. Ultimately, professional economists are interested in whether commodity prices employed in conjunction with the various forecasting techniques from the previous sections produce inflation forecasts that can better predict the time path of short term interest rates. It is well known that the Fed operates in what Bernanke and Boivin (2003) refer to as a "data-rich environment " monitoring perhaps thousands of economic time series when conducting monetary policy. Any description of Fed policy must incorporate this vast amount of information. The question examined in this section is whether including or excluding commodity prices changes an ex-post description of monetary policy. Specifically, I examine the implications of adding commodity prices to the information set of a policy reaction function by estimating a forward-looking Taylor rule with inflation forecasts and forecasts of real activity that either include or exclude commodity prices as predictors. I then compare the implied interest rate from the model to the actual interest rate. If including the commodity price regressors in the information set of the forward-looking Taylor rule moves the implied interest rate closer to the actual rate, then commodity prices are valuable for describing Fed behavior.

I take as a benchmark reaction function a forward-looking Taylor rule similar to the one used in Bernanke and Boivin (2003). The model is given by:

$$r_t = \alpha + \beta E_t(\pi_{t+k}) + \delta[E_t(U_{t+j}) - U^*_{t+j}] + \rho(L)r_t + \epsilon_t$$

In equation (1), r_t is the Fed Funds rate at time t. $E_t(\pi_{t+k})$ is the forecast of the inflation rate at time t+k, conditional on information available at time t. $E_t(U_{t+j})$ is the forecasted unemployment rate for time t+j conditional on information available at time

t. U^*_{t+j} is a time varying NAIRU. In order to capture possible interest rate smoothing by the Fed, the model also includes lags of the Fed Funds rate. ϵ_t represents any episode or variable to which the Fed might respond that is orthogonal to expected future inflation and unemployment.

The forecasts of inflation used in the Taylor rule are based on the static factor model forecasts discussed in Section 3. In the baseline case, the factors are estimated as the first r principal components from the set of 112 macroeconomic time series used in Stock and Watson (2005b), excluding commodity prices and producer prices. To determine the relevance of commodity prices, I add the 18 materials price inflation series to the Stock and Watson data set. Thus this exercise will determine whether commodity prices are useful for describing monetary policy above and beyond all other relevant data that are available to policymakers. In addition to these two cases, I also consider adding oil prices and all materials prices except the price of crude oil to the Stock and Watson regressors. Since the leading indicator properties of commodity prices for future unemployment are beyond the scope of this paper, the forecasting model for the unemployment rate includes only the 112 Stock and Watson predictors excluding commodity prices and producer prices. Thus the Taylor rules differ *only* in the information set used to forecast inflation.[12]

The Fed Funds rate data are from the Federal Reserve Economic Data (FRED) website at the Federal Reserve Bank of St. Louis. The series is a monthly average of the effective nominal Fed Funds rate. The unemployment rate is the monthly civilian unemployment rate released by the Bureau of Labor Statistics. The natural rate of unemployment is measured as a 5-year moving-average for each month. As in the previous sections, the AIC is used to determine the lag order for inflation and unemployment. To this I add one lag of each of the first three principal components of the predictors to the model. As in Bernanke and Boivin (2003), the forecast horizon is set to 12 months for inflation and to 6 months for the unemployment rate. I include two lags of the Fed Funds rate to account

[12] I also computed all results with commodity prices included in the unemployment rate forecast. The estimated coefficients were somewhat smaller and the MSEs somewhat larger, but the basic patterns across specifications of the model were the same as those reported.

for interest rate smoothing.

Following Bernanke and Boivin (2003), I estimate the Taylor rule by OLS. In Table 7, I report the estimated *long-run* responses to the inflation forecast and unemployment rate forecast, $\frac{\beta}{1-\rho(L)}$ and $\frac{\delta}{1-\rho(L)}$, with Newey-West standard errors in parentheses. There is some uncertainty as to the appropriate measure of inflation to include in the Taylor rule. Based on anecdotal evidence in Meyer (2004), I estimate alternative Taylor rules using forecasts of CPI inflation, core CPI inflation, as well as core PCE price inflation. In Table 8, I report the root mean-squared error (RMSE) for each estimate of the Taylor rule. If adding commodity prices to the forecasting model for inflation and unemployment reduces the RMSE of the resulting forward-looking Taylor rule, then commodity prices are useful for describing monetary policy. Since the dimension of the model is constant across specifications, ranking the models by their RMSEs is equivalent to ranking them according to a suitable information criterion. Inoue and Kilian (2006) show that this method will consistently select the best forecasting model among a set of candidate models under weak conditions. In other words, this exercise ranks the candidate Taylor rules according to their out-of-sample predictive power. The estimation period covers each month from February 1990 until December 2002, covering most of Alan Greenspan's chairmanship of the Federal Reserve Board.

Looking at the results in Table 7, I find that the estimates of the long-run response of the nominal Fed Funds rate to forecasted CPI inflation is positive in all 4 cases. The estimated response is largest (1.16) when commodity prices are included in the forecast. Furthermore, only when materials prices are used does the estimated reaction to forecasted inflation satisfy the Taylor Principle that the Fed must change interest rates more than one-for-one with inflation in order to affect the real interest rate and aggregate demand. Therefore, to the extent that a good description of monetary policy should exhibit the Taylor Principle, adding commodity prices is beneficial. I find a similar pattern for the core PCE price Taylor rule. The estimated response to long-run inflation is greater than one in all cases for the core CPI Taylor rule, but largest when materials prices are included in the forecast. The estimated long-run response to forecasted unemployment is negative in all cases and statistically significant.

The results in Table 8 show that including any subset of materials prices in the information set for the forward-looking Taylor rule decreases the RMSE of the model. The model fits the data best when only the price of crude oil is included, and the overall best fit of the Taylor rule is obtained when the Fed is assumed to respond to forecasts of CPI inflation (RMSE=0.1519) as opposed to core CPI inflation or core PCE price inflation.

In practice, Economists who monitor monetary policy and the Fed Funds market are interested in predicting the direction of change in the Fed Funds rate from month to month. A model that more accurately predicts the level of the Fed Funds rate, may or may not more accurately predict the sign of the change in the Fed Funds rate. In Table 9, I report the probability that the Taylor rule correctly predicts the change in direction of the monthly Fed Funds rate. All models correctly predict the sign of the change slightly more than half the time. Adding commodity prices to the inflation forecast only marginally improves the performance of the Taylor rule in this instance.

To summarize, commodity prices are shown to help describe monetary policy is two ways. First, when materials prices are included in the forecasting model for inflation, the estimated response of the interest rate to future inflation satisfies the Taylor Principle. Second, including materials prices in the forecasting model reduces the RMSE of the Taylor rule compared to the case when commodity prices are excluded from the forecast. A Taylor rule with forecasts of CPI inflation and the price of crude oil in the information set provides the best predictions of the actual monthly Fed Funds rate out of all the models examined in this paper.

The accuracy gains, however, are quite small and have little economic significance, obliviating formal tests for the equality of fit. The largest reduction in the RMSE is about 1.5%. Furthermore, adding commodity prices only marginally improves our ability to predict the direction of change in the Fed Funds rate from one month to the next. Thus, the usefulness of commodity prices for improving our understanding of monetary policy seems negligible.

2.5 Sensitivity Analysis

In this paper, I have followed the convention of focusing on percent changes in commodity prices in assessing the predictive content of these variables. This approach ignores the fact that commodity prices ($X_{i,t}, i = 1, ..., 46$) and consumer prices (P_t) are likely to be cointegrated, in which case one would expect the error correction term ($log(X_{i,t}) - log(P_t), i = 1, ..., 46$) to have predictive power for inflation. I explored this possibility in sensitivity analysis not reported in the paper. I considered (1) forecasting models including only the error correction terms in addition to lags of inflation, and (2) forecasting models including both the error correction terms and percentage changes of commodity prices in addition to lags of inflation.

I find that using the error correction term instead of percentage changes in commodity prices may substantially improve or substantially worsen forecasting performance. There is evidence that forecasts based on the error correction term for most industrial commodity prices perform well, whereas agricultural raw materials do not. The performance of food and beverages is erratic. Error correction terms for price indices in general do not work well. A similar pattern emerges when I use the error correction term in addition to percent changes in commodity prices.

In contrast, when applying forecast combination or shrinkage methods, including both percentage changes and error correction terms never results in more accurate forecasts than just including percentage changes. Compared to methods based on percentage changes only, those based on error correction terms only are inferior in most cases. Moreover, methods based on percentage changes only, tend to be more robust across forecast horizons.

In particular, for the factor models based on large information sets that form the basis of the analysis of Taylor rules in section 5, the model based on percentage changes in materials prices continues to be more accurate than all other versions.

2.6 Conclusion

Though higher commodity prices are often thought to presage higher rates of inflation (and hence are intensively scrutinized by professional economists) the existing literature has found it difficult to establish such a link empirically, at least for the 1990s and 2000s. In this paper, I provide a comprehensive analysis of the predictive content of 46 commodity prices and seven commodity price indices for 1993-2004.

The analysis reveals several patterns. Among the individual commodity prices with the highest predictive power is the price of crude oil, which improves the accuracy of CPI inflation forecasts mainly for the first three months, and the price of hard sawn wood, which has predictive ability at longer horizons. In contrast, individual food and beverage prices typically are not useful for forecasting monthly CPI inflation.

One of the problems with using individual commodity prices is that individual series tend to be noisy. I show that combining information from many commodity prices into a price index can be effective in reducing this noise, but there are important differences in the predictive ability across different types of indices. An index of industrial input prices generates the largest reduction in the forecast PMSE of all the price indices that I evaluate. At longer horizons, adding an industrial input price index to lags of inflation will reduce the PMSE by up to 10%. At shorter forecast horizons, however, the price of crude oil is a better predictor. Other commodity price indices are typically not informative.

A reasonable conjecture is that statistical methods that choose the weights underlying these indices based on the data rather than imposing fixed weights may generate superior forecasts of inflation. Indeed, my analysis shows that suitable forecast combination methods reduce the PMSE of the CPI inflation forecast by up to 11% relative to an inflation-only model. These accuracy gains are concentrated mainly at short horizons and are more systematic across horizons than those obtained from the industrial materials price index. The best forecasting technique depends on the forecast horizon, but bagging, shrinkage, and factor models provide some of the best combination forecast results. While the price of crude oil is the best single predictor at short forecast horizons, a bagging forecast with materials prices or a Bayesian model average forecast with the full set of commodity prices results in even larger accuracy gains.

In addition to examining monthly and quarterly CPI inflation forecasts, I also study forecasts of GDP deflator inflation. While several individual commodity prices and price indices have some predictive power at short horizons, only a few commodity prices, notably the price of crude oil and hard sawn wood, improve forecast accuracy at longer horizons. As in the case of monthly CPI inflation, suitable forecast combination methods improve the accuracy of quarterly GDP deflator inflation compared to models using individual commodity prices or price indices, but systematic gains are limited to the first two quarters.

Forecasting models with commodity prices are superior not only to inflation-only models. I find that there is information in commodity prices not captured by the leading principal components of a broader set of macroeconomic and financial variables. In particular, adding materials prices to a set of 112 macroeconomic and financial variables improves the accuracy of factor model forecasts of CPI inflation and core CPI inflation. Forecasts of core PCE price inflation, however, are not improved by the addition of commodity prices. The PMSE reductions are particularly large at longer forecast horizons. In contrast, the price of crude oil alone is not informative relative to the larger set of predictors.

Professional economists are interested in accurate forecasts of inflation mainly because these forecasts are thought to be informative about future monetary policy and the path of short-term interest rates. I show that including the set of materials prices in the forecasting model for inflation indeed improves the fit of a forward-looking Taylor rule compared to the case when materials prices are excluded. However, the RMSE reduction is only about 1% and thus economically insignificant. Moreover, adding commodity prices to the inflation forecast only slightly improves the ability of the Taylor rule to correctly predict the sign of the monthly change in the Fed Funds rate. I therefore conclude that while commodity prices are useful predictors of U.S. consumer price inflation, their usefulness for improving our understanding of the Federal Reserve's policy decisions is negligible.

Table 2.1: Series Description

Price Indices	Materials Prices	Food and Beverage Prices
Non-Fuel Price Index	Coal	Bananas
Food and Beverage Price Index	Crude Oil (petroleum)	Barley
Food Price Index	Uranium	Beef
Beverage Price Index	Aluminum	Cocoa Beans
Industrial Inputs Price Index	Copper	Coffee, other mild arabicas
Agricultural Raw Materials Price Index	Lead	Coffee, robusta
Metals Price Index	Nickel	Coconut Oil
	Tin	Fishmeal
	Zinc	Groundnuts (peanuts)
	Cotton	Lamb
	Hides	Maize (corn)
	Soft Logs	Olive Oil
	Hard Logs	Oranges
	Rubber	Palm Oil
	Hard Sawn Wood	Swine (pork)
	Soft Sawn Wood	Poultry (chicken)
	Wool, coarse	Rice
	Wool, fine	Fish (salmon)
		Shrimp
		Soybean meal
		Soybean Oil
		Soybeans
		Sugar, European import price
		Sugar, free market
		Sugar, U.S. import price
		Sunflower oil
		Tea
		Wheat

Table 2.2: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Distributed Lag Model - Monthly CPI Inflation (1993-2004)

horizon	1	3	6	9	12	15	18
Price Indices							
Non-Fuel	1.03	1.02	0.97	0.96	1.00	1.17	1.19
Food and Beverage	1.02	1.02	1.01	1.01	1.02	1.19	1.29
Food	1.01	1.05	1.17	1.16	1.16	1.44	1.50
Beverage	1.06	1.11	1.06	1.12	1.20	1.08	1.08
Industrial Inputs	1.02	0.96	0.91	0.95	0.96	0.96	0.96
Agricultural Raw Materials	1.02	1.00	0.99	0.98	1.08	1.14	1.13
Metals	1.04	1.00	0.95	0.98	0.99	0.96	0.99
Materials							
Coal	1.00	0.99	1.00	1.01	0.99	1.01	1.01
Crude Oil (Petroleum)	0.92	0.93	0.99	0.99	1.10	1.02	1.02
Uranium	1.01	1.01	1.01	1.01	1.00	1.01	1.01
Aluminum	1.03	1.01	0.99	0.99	1.00	1.00	0.99
Copper	1.06	0.98	0.97	0.96	0.98	0.98	0.99
Lead	1.00	0.98	0.91	0.98	1.03	1.00	1.02
Nickel	1.01	0.94	0.91	0.95	0.99	0.97	1.00
Tin	1.01	1.00	0.99	0.98	1.00	1.00	1.01
Zinc	0.99	0.96	0.92	0.97	0.99	1.03	1.05
Cotton	0.99	1.05	1.01	0.99	1.05	1.19	1.19
Hides	1.00	1.00	0.99	0.98	1.12	1.18	1.02
Soft Logs	1.00	1.00	1.01	1.05	1.01	1.02	1.02
Hard Logs	1.01	1.04	1.02	1.03	1.03	1.03	1.02
Rubber	1.02	1.02	1.00	1.00	1.07	1.05	1.05
Hard Sawn Wood	1.02	1.03	1.03	0.95	0.88	0.88	0.95
Soft Sawn Wood	1.02	1.11	1.04	1.01	1.03	1.02	1.02
Wool, coarse	1.03	1.00	1.01	1.03	1.02	1.02	1.05
Wool, fine	1.00	0.99	0.95	0.97	1.11	1.01	1.03
Food and Beverage							
Bananas	1.00	1.13	1.23	1.02	1.03	1.00	1.00
Barley	1.00	1.03	1.16	1.12	1.06	1.00	1.18
Beef	0.94	1.04	1.14	1.12	1.26	1.33	1.32
Cocoa Beans	1.02	1.02	0.94	0.89	0.94	0.96	1.04
Coffee, other mild arabicas	1.07	1.11	1.12	1.22	1.26	1.17	1.14
Coffee, robusta	1.17	1.29	1.30	1.33	1.44	1.24	1.16
Coconut Oil	1.00	0.98	0.99	0.97	0.99	0.94	0.96
Fishmeal	1.00	0.99	1.16	1.18	1.05	1.14	1.17
Groundnuts (peanuts)	1.00	1.00	1.00	1.01	1.02	1.02	1.04
Lamb	1.01	1.07	1.05	1.05	1.00	1.01	1.02
Maize (corn)	1.01	1.01	1.00	1.00	1.04	1.22	1.24
Olive Oil	0.99	1.01	1.01	0.96	1.00	1.00	1.01
Oranges	1.01	1.00	1.01	1.00	1.00	1.01	1.01
Palm Oil	1.00	1.05	1.24	1.05	1.04	1.02	1.01
Swine (pork)	0.99	0.99	1.01	1.01	1.02	1.02	1.00
Poultry (chicken)	1.01	1.00	1.02	1.03	1.02	1.11	1.25
Rice	1.04	1.01	1.00	1.23	1.25	1.30	1.16
Fish (salmon)	1.01	1.01	1.00	1.02	1.14	1.15	1.07
Shrimp	1.13	1.06	1.02	0.97	0.99	0.99	1.01
Soybean meal	1.01	1.00	0.99	0.99	1.28	1.42	1.49
Soybean Oil	1.02	1.01	1.03	1.00	1.03	1.02	0.99
Soybeans	1.01	1.02	1.00	0.99	1.03	1.03	1.17
Sugar, European import price	1.02	1.01	1.04	1.09	1.08	1.08	1.02
Sugar, free market	1.06	1.15	1.18	1.23	1.36	1.41	1.41
Sugar, U.S. import price	1.04	1.05	1.11	1.10	1.38	1.27	1.18
Sunflower oil	1.04	1.04	1.05	1.06	1.07	1.05	1.05
Tea	1.03	1.00	1.18	1.30	1.31	1.26	1.21
Wheat	1.01	1.05	1.14	1.16	1.28	1.25	1.29

Table 2.3: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Combination Forecasts - Monthly CPI Inflation (1993-2004)

	Materials Prices			Food and Beverage Prices			All Commodity Prices		
	1m	3m	12m	1m	3m	12m	1m	3m	12m
Unrestricted	0.92	1.06	1.12	1.29	1.48	1.59	1.24	1.54	1.53
Shrinkage	0.90	0.95	0.96	0.97	0.98	1.00	0.95	0.96	0.98
Standard Bagging	0.89	0.93	1.04	0.96	0.98	1.08	0.90	0.94	1.06
Factor									
r=1	0.99	0.93	0.96	1.00	1.00	1.02	1.00	1.00	1.02
r=2	0.95	0.92	0.99	1.02	1.03	1.03	1.02	1.02	1.03
r=3	0.92	0.92	0.99	1.03	1.01	1.04	1.01	0.98	1.03
Equal-Weighted Forecast	0.95	0.99	0.98	0.97	0.98	1.00	0.96	0.96	0.98
BMA	0.90	0.94	0.97	0.97	0.98	1.00	0.89	0.95	0.98

Table 2.4a: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Distributed Lag Model - Quarterly GDP Deflator Inflation (1995-2004)

horizon	1	2	3	4	5	6
Price Indices						
Non-Fuel	1.08	0.99	0.99	1.12	1.34	1.35
Food and Beverage	1.01	1.12	1.05	1.16	1.15	1.04
Food	1.19	1.17	1.29	1.42	1.43	1.30
Beverage	1.02	1.44	1.38	1.49	1.29	1.21
Industrial Inputs	0.98	0.90	1.05	1.16	1.43	1.27
Agricultural Raw Materials	0.95	1.35	1.51	1.74	1.01	1.00
Metals	1.01	1.00	1.00	1.33	1.87	1.39
Materials						
Coal	1.21	1.39	1.27	1.41	1.34	1.09
Crude Oil (Petroleum)	1.01	0.89	0.77	0.71	0.78	0.79
Uranium	1.02	1.19	1.26	1.12	1.04	1.04
Aluminum	1.04	1.02	0.96	1.08	1.31	1.27
Copper	1.00	1.09	1.39	2.02	2.65	2.35
Lead	1.04	0.96	1.28	2.20	3.73	4.56
Nickel	1.04	0.93	1.14	1.20	1.43	1.43
Tin	1.07	1.06	1.02	1.02	1.06	1.10
Zinc	1.04	0.99	0.96	1.08	1.36	1.29
Cotton	1.06	1.05	0.98	1.19	1.38	1.19
Hides	0.84	0.86	0.85	1.01	1.04	0.96
Soft Logs	1.24	1.10	1.00	1.01	1.12	1.00
Hard Logs	1.02	0.95	0.98	0.95	0.92	0.98
Rubber	1.41	1.54	1.15	0.98	1.23	1.18
Hard Sawn Wood	1.02	0.84	0.75	0.78	0.67	0.71
Soft Sawn Wood	1.05	1.04	1.01	1.00	0.96	0.94
Wool, coarse	1.20	1.44	1.46	1.61	1.97	2.15
Wool, fine	1.10	1.15	1.41	1.89	2.43	2.46
Food and Beverage						
Bananas	1.37	1.50	1.07	1.00	0.98	1.06
Barley	1.08	1.15	1.19	1.13	1.14	1.23
Beef	1.18	1.26	1.27	1.35	1.33	1.28
Cocoa Beans	1.04	1.24	1.17	1.35	1.45	1.46
Coffee, other mild arabicas	1.06	1.03	1.05	1.12	1.13	1.11
Coffee, robusta	1.40	1.43	1.11	1.17	1.33	1.36
Coconut Oil	1.15	1.02	1.01	1.01	0.99	1.02
Fishmeal	1.08	1.18	1.18	1.24	1.26	1.22
Groundnuts (peanuts)	0.95	0.99	0.99	1.04	1.02	1.00
Lamb	1.15	1.55	1.77	1.11	1.03	1.06
Maize (corn)	1.17	1.14	1.14	1.33	1.28	1.30
Olive Oil	1.04	1.02	1.02	1.03	1.04	1.03
Oranges	0.91	0.90	0.97	1.03	1.08	1.02
Palm Oil	1.01	1.03	1.01	1.00	0.99	1.03
Swine (pork)	1.07	1.05	1.01	1.03	1.07	1.04
Poultry (chicken)	1.13	0.95	0.97	1.00	1.02	1.08
Rice	1.06	1.16	1.27	1.15	1.19	1.19
Fish (salmon)	1.01	1.05	0.98	1.12	1.16	1.10
Shrimp	1.32	1.31	1.30	1.56	1.54	1.13
Soybean meal	0.94	1.01	1.42	1.71	2.02	1.93
Soybean Oil	1.12	1.11	1.21	1.27	1.50	1.58
Soybeans	1.11	1.13	1.10	1.08	1.04	1.09
Sugar, European import price	1.03	0.97	0.99	1.00	0.94	0.97
Sugar, free market	1.19	1.22	1.30	1.28	1.34	1.33
Sugar, U.S. import price	1.04	1.32	1.19	1.18	1.38	1.31
Sunflower oil	1.06	1.14	1.15	1.24	1.25	1.11
Tea	0.95	1.30	1.42	1.74	1.95	2.16
Wheat	0.95	1.15	0.98	1.02	0.99	1.02

Table 2.4b: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Distributed Lag Model - Quarterly CPI Inflation (1995-2004)

horizon	1	2	3	4	5	6
Price Indices						
Non-Fuel	1.17	1.51	1.28	1.69	1.96	2.08
Food and Beverage	1.13	1.31	1.02	1.04	1.03	1.11
Food	1.23	1.37	1.11	1.26	1.35	1.60
Beverage	1.10	1.18	1.07	1.07	1.11	1.08
Industrial Inputs	1.20	1.25	1.17	1.51	1.53	1.55
Agricultural Raw Materials	1.74	1.33	1.45	1.92	2.07	2.33
Metals	1.08	1.44	1.40	1.76	1.72	1.68
Materials						
Coal	1.41	1.23	1.43	1.38	1.41	1.50
Crude Oil (Petroleum)	0.80	1.21	0.82	0.74	0.73	0.76
Uranium	1.21	1.43	1.15	1.18	1.16	1.21
Aluminum	0.98	1.33	1.15	1.31	1.27	1.32
Copper	0.91	2.03	1.86	1.97	2.00	1.97
Lead	0.96	1.25	1.47	1.98	2.28	2.72
Nickel	1.12	1.26	1.18	1.24	1.41	1.35
Tin	0.83	1.02	0.98	0.99	1.00	0.96
Zinc	1.30	1.51	1.30	1.27	1.31	1.39
Cotton	1.00	1.15	0.99	1.17	1.10	1.09
Hides	1.28	1.50	1.17	1.22	1.05	0.95
Soft Logs	1.11	1.04	1.05	1.02	1.03	1.00
Hard Logs	1.46	1.41	1.09	1.13	1.06	1.04
Rubber	1.25	1.25	1.03	1.02	1.07	1.10
Hard Sawn Wood	1.13	1.12	0.81	0.75	0.78	0.89
Soft Sawn Wood	1.08	1.11	1.00	1.00	1.01	0.99
Wool, coarse	1.02	1.13	1.12	2.27	2.79	2.95
Wool, fine	0.94	1.70	1.74	1.93	2.07	1.83
Food and Beverage						
Bananas	1.52	2.55	2.21	2.19	1.94	1.87
Barley	1.05	1.62	1.43	1.29	1.23	1.24
Beef	1.11	1.22	1.65	1.83	1.93	2.02
Cocoa Beans	1.09	1.40	1.67	1.87	1.84	1.88
Coffee, other mild arabicas	1.09	1.17	1.09	1.14	1.13	1.08
Coffee, robusta	1.31	1.37	1.17	1.25	1.21	1.18
Coconut Oil	0.99	1.12	0.96	0.95	0.82	0.88
Fishmeal	1.06	1.64	1.62	1.57	1.68	2.01
Groundnuts (peanuts)	1.00	1.06	1.09	1.03	1.01	1.01
Lamb	1.21	1.07	1.08	1.03	1.04	1.06
Maize (corn)	1.23	1.21	1.27	1.36	1.40	1.40
Olive Oil	1.26	1.35	1.15	1.04	1.08	1.05
Oranges	0.87	0.99	1.03	0.99	1.02	1.01
Palm Oil	1.38	1.44	1.55	1.29	1.02	1.01
Swine (pork)	1.07	1.15	1.00	1.00	1.00	1.03
Poultry (chicken)	1.00	1.07	0.99	1.03	1.03	0.97
Rice	0.98	1.04	1.07	1.04	1.15	1.06
Fish (salmon)	1.01	1.00	1.35	1.57	1.39	1.34
Shrimp	1.01	1.01	1.02	1.17	1.31	1.54
Soybean meal	1.57	2.08	1.84	1.95	1.90	1.94
Soybean Oil	1.06	1.49	1.07	1.02	1.45	1.57
Soybeans	1.48	1.12	1.08	1.07	1.05	1.07
Sugar, European import price	0.96	1.20	1.11	1.05	1.00	0.90
Sugar, free market	1.13	1.79	2.28	2.38	2.25	2.27
Sugar, U.S. import price	1.34	2.01	1.87	1.30	0.99	0.97
Sunflower oil	1.05	1.34	1.09	1.05	1.03	1.05
Tea	1.78	1.84	2.10	1.42	1.89	1.72
Wheat	1.14	1.33	1.36	1.30	1.32	1.26

Table 2.5: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Quarterly GDP Deflator Inflation and Quarterly CPI Inflation Forecasts

	Materials Prices			Food and Beverage Prices			All Commodity Prices		
	1q	2q	4q	1q	2q	4q	1q	2q	4q
Quarterly GDP Deflator Inflation									
Unrestricted	1.04	1.00	1.23	-	-	-	-	-	-
Shrinkage	0.83	0.85	0.93	-	-	-	-	-	-
Standard Bagging	0.92	0.84	0.96	-	-	-	-	-	-
Factor									
r=1	0.93	0.93	0.97	1.01	0.98	1.03	1.00	0.98	1.05
r=2	0.93	0.93	0.96	0.97	0.99	1.05	0.99	0.97	1.10
r=3	0.95	0.91	0.89	0.93	0.98	1.13	0.88	0.91	1.04
Equal-Weighted Forecast	0.92	1.00	1.03	0.91	0.95	0.99	0.96	1.00	1.10
BMA	0.91	0.92	1.00	0.91	0.95	0.99	0.93	0.97	1.07
Quarterly CPI Inflation									
Unrestricted	1.38	1.58	1.27	-	-	-	-	-	-
Shrinkage	0.82	0.92	0.86	-	-	-	-	-	-
Standard Bagging	0.83	0.92	0.87	-	-	-	-	-	-
Factor									
r=1	0.67	0.91	0.87	1.01	1.03	1.07	0.99	1.04	1.06
r=2	0.72	0.91	0.89	1.04	1.06	1.06	1.00	1.05	1.08
r=3	0.78	0.91	0.84	1.06	1.07	1.12	1.04	1.05	1.08
Equal-Weighted Forecast	0.89	1.04	0.92	0.88	1.10	0.93	0.96	1.19	1.08
BMA	0.87	1.00	0.91	0.88	1.10	0.93	0.93	1.16	1.06

Table 2.6: Recursive Out-of-Sample PMSE Relative to Inflation-Only Model: Full Information Set Forecasts

	CPI Inflation			Core CPI Inflation			Core PCE Inflation		
	1m	3m	12m	1m	3m	12m	1m	3m	12m
No Commodity Prices	0.95	0.97	0.97	1.00	0.98	0.98	1.00	1.02	1.02
With Materials Prices	0.93	0.91	0.89	1.01	0.96	0.93	1.02	1.01	1.01
With Food and Beverage Prices	1.00	1.00	1.01	1.02	1.00	1.00	1.00	1.00	1.00
With All Commodity Prices	1.00	1.00	1.01	1.03	1.00	1.00	1.00	1.00	1.00
With Materials Prices Excluding Oil	0.94	0.93	0.90	1.01	0.95	0.93	1.01	1.01	1.01
With Oil Prices	0.95	0.96	0.96	1.00	0.98	0.98	1.00	1.02	1.02

Table 2.7: Taylor Rule Parameter Estimates

	CPI Inflation		Core CPI Inflation		Core PCE Inflation	
	$\frac{\beta}{1-\rho(L)}$	$\frac{\delta}{1-\rho(L)}$	$\frac{\beta}{1-\rho(L)}$	$\frac{\delta}{1-\rho(L)}$	$\frac{\beta}{1-\rho(L)}$	$\frac{\delta}{1-\rho(L)}$
No Commodity Prices	0.46	-0.31	1.29	-0.32	0.97	-0.31
	(0.67)	(0.11)	(1.05)	(0.09)	(0.99)	(0.09)
With Materials Prices	1.16	-0.25	1.46	-0.26	1.04	-0.31
	(0.61)	(0.08)	(0.90)	(0.07)	(0.97)	(0.09)
With Oil Prices	0.97	-0.28	1.36	-0.31	0.97	-0.31
	(0.74)	(0.10)	(1.10)	(0.08)	(0.97)	(0.09)
With Materials Prices Excluding Oil	1.11	-0.25	1.41	-0.27	1.02	-0.32
	(0.62)	(0.08)	(0.91)	(0.07)	(0.98)	(0.09)

Note: Taylor rule evaluated over the 1990-2004 period.

Table 2.8: RMSE of the Taylor Rule

	CPI Inflation	Core CPI Inflation	Core PCE Inflation
No Commodity Prices	0.1541	0.1533	0.1542
With Materials Prices	0.1526	0.1521	0.1540
With Oil Prices	0.1519	0.1532	0.1542
With Materials Prices excluding Oil	0.1529	0.1523	0.1541

Note: Taylor rule evaluated over the 1990-2004 period.

Table 2.9: Probability that the Taylor Rule Correctly Predicts the Direction of Change in the Fed Funds Rate

	CPI Inflation	Core CPI Inflation	Core PCE Inflation
No Commodity Prices	0.51	0.52	0.52
With Materials Prices	0.52	0.54	0.52
With Oil Prices	0.52	0.53	0.52
With Materials Prices excluding Oil	0.52	0.52	0.52

Note: Taylor rule evaluated over the 1990-2004 period.

BIBLIOGRAPHY

[1] Bachmeier, L.J., Li, Q., and D. Liu (2005), "Can Oil Prices Improve Economic Forecasting? An Empirical Investigation," mimeo.

[2] Bajaj, V. (2005), "Inflation Remains Flat Despite Energy Costs," *The New York Times*, July 15.

[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] Bernanke, B.S., and I. Mihov (1998), "Measuring Monetary Policy," *Quarterly Journal of Economics*, 113, 869-902.

[5] Bernanke, B.S., and J. Boivin (2003), "Monetary Policy in a Data-Rich Environment," *Journal of Monetary Economics*, 50, 525-546.

[6] Bernanke, B.S. (2004), "Monetary Policy and the Economic Outlook" speech to the American Economic Association, January 4.

[7] Blomberg, S.B., and E.S. Harris (1995), "The Commodity-Consumer Price Connection: Fact or Fable?" *FRBNY Economic Policy Review*, October, 21-38.

[8] Boughton, J.M., and W.H. Branson (1991), "Commodity Prices as a Leading Indicator of Inflation," in K. Lahiri and G. Moore (eds.), *Leading Economic Indicators: New Approaches and Forecasting Records*, 305-336.

[9] Buhlmann, P., and B. Yu (2002), "Analyzing Bagging," *Annals of Statistics*, 30, 927-961.

[10] Cecchetti, S.G., Chu, R.S., and C. Steindel (2000), "The Unreliability of Inflation Indicators," *FRBNY Current Issues in Economics and Finance*, 6, 1-6.

[11] Clarida, R., Gali, J., and M. Gertler (2000), "Monetary Policy rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, 115, 147-180.

[12] Furlong, F., and R. Ingenito (1996), "Commodity Prices and Inflation," *Federal Reserve Bank of San Fransisco Economic Review*, 2, 27-46.

[13] Hanson, M.S. (2004), "The 'Price Puzzle' Reconsidered," *Journal of Monetary Economics*, 51, 1385-1413.

[14] Inoue, A., and L. Kilian (2005), "How Useful is Bagging in Forecasting Economic Time Series? A Case Study of U.S. CPI Inflation," forthcoming *Journal of the American Statistical Association.*

[15] Inoue, A., and L. Kilian (2006), "On the Selection of Forecasting Models," *Journal of Econometrics*, 130, 273-306.

[16] Kohn, D.L. (2004), "The Outlook for Inflation." Speech to the National Economists Club, June 4.

[17] Massimiliano, M., Stock, J.H., and M.W. Watson (2004), "A Comparison of Direct and Iterated Multistep AR Methods for Forecasting Macroeconomic Time Series," *Journal of Econometrics*, 135, 499-526.

[18] Meyer, L.H. (2004), *A Term at the Fed*, HarperBusiness, p.42.

[19] Sims, C.A. (1992), "Interpreting the Macroeconomic Time Series Facts," *European Economic Review*, 36, 975-1000.

[20] Stock, J.H., "Forecasting Economic Time Series," ch. 27 in B. Baltagi (ed.), *Companion in Theoretical Econometrics*, Basil Blackwell (2001), 562-584.

[21] Stock, J.H. and M.W. Watson (1999), "Forecasting Inflation," *Journal of Monetary Economics*, 44, 293-335.

[22] Stock, J.H. and M.W. Watson (2002), "Macroeconomic Forecasting Using Diffusion Indexes," *Journal of Business and Economic Statistics*, 20, 147-162.

[23] Stock, J.H. and M.W. Watson (2003), "Forecasting Output and Inflation: The Role of Asset Prices," *Journal of Economic Literature*, 41, 788-829.

[24] Stock, J.H. and M.W. Watson (2005a), "An Empirical Comparison of Methods for Forecasting Using Many Predictors," mimeo, August, Harvard University.

[25] Stock, J.H. and M.W. Watson (2005b), "Implications of Dynamic Factor Models for VAR Analysis," mimeo, June, Department of Economics, Harvard University.

[26] Wright, J.H. (2003), "Forecasting U.S. Inflation by Bayesian Model Averaging," International Finance Discussion Papers, Number 780, Board of Governors of the Federal Reserve System.

CHAPTER III

Retail Energy Prices and Consumer Expenditures

3.1 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

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

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

[3] For 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).

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

This paper studies the effects 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

[4]For example, Nariman Behravesh, 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 Behravesh 2007).

[5]In 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.

[6]Implicit in this view is the assertion that not all of the revenue generated by higher energy prices is recycled in the form of higher spending, an assumption we will discuss in more detail in section 3.

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 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

[7]This 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).

(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

[8] Asymmetries 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).

for each form of energy and of real energy prices. Using these data, we compute a monthly time series of the losses and gains 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.

Section 3 contains an illustrative model of the effect of higher energy prices on consumption. In section 4, 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 5 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 6, we provide a comprehensive analysis, based on more than 130 disaggregate consumption series, of how consumers adjust their expenditures in response to unanticipated losses in purchasing power. This analysis provides several additional insights. First, we quantify the extent to which consumers substitute small, energy-efficient automobiles for large, energy-inefficient automobiles. There is compelling quantitative evidence 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). The likely reason that the large effect of purchasing power losses on purchases of new domestically-produced automobiles does not appear to cause a strong reallocation effect, as conjectured by Hamilton (1988), is that the share of the automobile industry in value added and employment is only about 1%. Second, we document the margins of adjustment used by households in dealing with purchasing power losses. This allows us to identify the extent to which industries such as tourism or the airline industry will be affected by higher energy prices. Third, we identify patterns of substitution that are of interest from a public policy point of view. There is clear evidence of consumers substituting low-cost alternatives for high-priced goods. For example, rather than dining out, households eat more at home, which in turn is reflected in changing patterns of food consumption. Similarly, consumers substitute public for private higher education. Some responses are temporary (such as the increasing use of mass transit in urban areas), whereas others appear very persistent (such as the decline in personal airline travel).

In section 7 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. Theoretical models that incorporate this 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. In contrast, we do not find evidence of a reallocation effect in the data. It may seem that our findings must be wrong because they cannot account for the absence of a strong economic boom in 1986. This is not the case. Using historical decompositions to analyze how consumption would have evolved in the absence of the 1986 decline in energy prices, we demonstrate in section 7.1. that there actually was a boom in consumption without which actual economic growth would have been much weaker in 1986, that this boom was similar in magnitude to the decline in real consumption in 1979, and that this boom 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 7.2., 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 (see, e.g., Herrera and Pesavento 2007). In section 7.4., 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 8.

3.2 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:

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

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

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

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, $\%\Delta pp_{t+1}$, approximately equals the product of the nominal expenditure share and the percent rate of change in real energy prices:

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

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

[9] A 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.

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, 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 5.

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. Sections 4 and 5 will investigate the empirical support for these effects.

3.3 An Illustrative Model of Energy and Non-Energy Consumption

This section presents a simple framework that illustrates the intuition for why higher real energy prices affect real consumption. The model abstracts from the fact that the consumer demand for energy is derived from a household production function. Consider a representative consumer who consumes energy (e) and a composite non-energy good (n). Let P_N be the real price of the composite non-energy good, P_E the real price of energy, and I real household income. Furthermore, suppose that the consumer's income is supplemented with a fraction θ of economy-wide energy revenues $P_E \cdot E$. The representative consumer faces the following constrained optimization problem:

$$\begin{aligned} \max \qquad & u = [\alpha n^\rho + \beta e^\rho]^{\frac{1}{\rho}} \\ s.t. \qquad & I + \theta P_E E = P_N \cdot n + P_E \cdot e \end{aligned}$$

where I, P_N,and P_E have been expressed in terms of the PCE numeraire good.

The LaGrangian for this constrained maximization problem is:

$$L = [\alpha n^\rho + \beta e^\rho]^{\frac{1}{\rho}} + \lambda[I + \theta P_E \cdot E - P_E \cdot e - P_N \cdot n]$$

For (n^*, e^*) to be the optimal consumption bundle, the following first-order conditions must be met:

$$\frac{1}{\rho}[\alpha n^\rho + \beta e^\rho]^{\frac{1}{\rho}-1}\rho\alpha n^{\rho-1} - \lambda P_N = 0$$

$$\frac{1}{\rho}[\alpha n^\rho + \beta e^\rho]^{\frac{1}{\rho}-1}\rho\beta e^{\rho-1} - \lambda P_E = 0$$

$$I + \theta P_E \cdot E - P_E \cdot e - P_N \cdot n = 0$$

$$e^* = E^*$$

The first two first-order conditions do not depend θ because income from energy is treated as a lump sum transfer, so consumers do not take into account the increase in income associated with an increase in the price of energy.

Define $r = \rho/(\rho - 1)$. Then the Marshallian demand function for non-energy consumption can be expressed as:

$$n^*(I, P_N, P_E) = \frac{I \cdot P_N^{r-1}}{(1-\theta)\frac{\alpha}{\beta}^{r-1} P_E^r + P_N^r}.$$

The cross-partial derivative of n^* with respect to P_E is:

$$\frac{\partial n^*}{\partial P_E} = (-r) \cdot \frac{P_N^{r-1} \cdot (1-\theta) \cdot (\frac{\alpha}{\beta})^{r-1} \cdot P_E^{r-1} \cdot I}{[(1-\theta) \cdot (\frac{\alpha}{\beta})^{r-1} \cdot P_E^r + P_C^r]^2}.$$

The second term in this expression is greater than or equal to zero for all parameter values. Thus, the sign of $\partial n^*/\partial P_E$ is determined by $-r = \rho/(1-\rho)$.

The effect of changes in the real price of energy on non-energy consumption depends on two parameters (see Table 1). The first parameter is θ, which measures the degree to which revenues from energy consumption are recycled through the economy as income. If $\theta = 1$, meaning that every dollar of energy revenue is recycled as income, then $\partial n^*/\partial P_E = 0$. Non-energy consumption does not respond to changes in real energy prices at all because an additional dollar spent on energy becomes an additional dollar of income. For $\theta \in [0, 1)$, the term $1 - \theta$ in the expression for the cross-partial derivative is positive. In this case, the sign of $\partial n^*/\partial P_E$ depends on the sign of ρ. There are three cases:

1. $\rho = 0 \Rightarrow -r = 0$. In this case, $\partial n^*/\partial P_E = 0$. This is the Cobb-Douglas case, in which utility is separable in energy and non-energy consumption. Non-energy consumption is independent of the real price of energy.

2. $\rho \in (0, 1) \Rightarrow -r > 0$. In this case, $\partial n^*/\partial P_E > 0$. The limit of this case is $\rho = 1$, in which case utility is linear in energy and non-energy consumption.

3. $\rho < 0 \Rightarrow -r < 0$. In this case, $\partial n^*/\partial P_E < 0$, and consumers reduce their non-energy expenditures in response to energy price increases.

This analysis demonstrates that the common perception that higher energy prices lower consumption is based on the implicit view that $0 \leq \theta < 1$ and $\rho < 0$. Policy discussions sometimes seem to postulate implicitly that $\theta = 0$, which would be true only if there were no recycling of the revenues generated by higher energy prices. A more plausible assumption is that $\theta > 0$. For example, higher energy prices driven by higher prices for imported energy goods will stimulate increased exports of U.S. goods, at least partially offsetting the loss in consumer income. Similarly, in the case of a purely domestic energy price shock such as a shock to U.S. refining capacity, the transfer of income to the refiner may be partially returned to consumers in the form of higher wages or higher stock returns on domestic energy companies. At the same time, it seems unlikely that $\theta = 1$, because increased exports are unlikely to offset completely the reduction in income from higher prices for imported energy goods, and because different income recipients may have different marginal propensities to consume.

For $0 < \theta < 1$, the effect of higher energy prices on non-energy consumption depends on ρ. The negative income effect of an increase in energy prices on n^* will dominate the positive substitution effect only if $\rho < 0$. When $\rho > 0$, in contrast, n^* will increase. This example shows that the response of non-energy consumption to energy price shocks is ambiguous a priori. In the next section, we investigate the extent of the empirical support for the common notion that energy price increases lower real consumption.

3.4 Quantifying the Effects of Purchasing Power Shocks on Real Consumption

In this section, we 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, recreational 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 6.

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

[10] In 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.

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 section 2 (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]

[11] Our 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.

[12] Our 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).

[13] Under 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.

3.4.1 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 4, 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 5, 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.

3.4.2 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.

3.4.2.1 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]

[14] An 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 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.

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 2 provides p-values for Wald tests of symmetry in 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 2 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 2 shows that there is no compelling reason to depart from the standard model imposing symmetry.

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 2 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 2 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.

3.4.2.2 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 gains and large purchasing power losses 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 increases and large purchasing power decreases are symmetric. Column 2 of Table 2 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 2 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 2).

3.4.2.3 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 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

[15] Alternatively, 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.

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 2).

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 2 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.

3.4.3 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.

3.5 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 4:

> 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.

Figures 8 and 9 displays dynamic responses to a one-time, one standard-deviation unanticipated decline in purchasing power. A one standard-deviation shock 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 shock size 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.

Figure 8 shows how consumers' purchasing power evolves in response to an unanticipated decline in purchasing power. The initial decline in purchasing power is followed by additional declines in subsequent months. The maximum loss of purchasing power is reached after eight months. Thereafter, the response becomes flat. Figure 9 shows the corresponding responses of the level of real consumption aggregates.

3.5.1 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 9 for various forms of energy consumption. Figure 10 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%. Virtually all 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, 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 -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

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

[17]It 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 estimate ranges 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].

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 should fall between 0% and -0.055%.

How does this prediction compare to the response of total consumption in Figure 9? 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, non-durables, services). Figure 9 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,

[18]We would also expect households to reduce their expenditures on home improvements and other forms of household residential fixed investment, given the durable-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.

for which the response is marginally insignificant.

3.5.2 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 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.

3.5.3 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 9. 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 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 11 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 11 is not clear-cut, it seems plausible that consumers - rightly or wrongly - have come to associate unanticipated losses in 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 using 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]".

3.5.4 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 Surveys 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 12 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 12 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 12 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 in 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, following rising 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.

3.6 How Do Expenditure Patterns Change in Response to Purchasing Power Shocks?

The previous two sections investigated the response of consumption aggregates and of consumers' expectations to unanticipated losses in purchasing power. It is equally important to understand how individual expenditure items respond to such shocks. Little is known about the trade-offs consumers face when confronted with purchasing power shocks and about the choices they make, although there is anecdotal evidence that we will draw on below. In this section, we document in detail how individual expenditure items respond to unanticipated losses in purchasing power triggered by higher energy prices. We estimate impulse response functions of monthly real consumption for about 130 different goods and services to a negative one-time, one standard deviation purchasing power shock. The impulse response functions were estimated in the same manner as in previous sections, and are presented in Figure 13.

The disaggregate analysis of this section is useful for several reasons. First, despite the preponderance of anecdotal evidence on how energy price shocks affect specific types of expenditures (such as gasoline consumption, purchases of SUVs, or dining out), to date there exists virtually no quantitative evidence of these effects. Direct evidence of the presence or absence of such shifts in expenditure patterns is also important because it plays a central role in models of the reallocation effect. Second, a detailed analysis of motor vehicle consumption will shed light on the extent to which consumers buy smaller and more energy-efficient cars in response to purchasing power losses, a mechanism that plays an important role in policy discussions of the effect of higher oil prices (see Bresnahan and Ramey 1993). Third, disaggregate results are of interest to industry analysts. For example, there is independent interest in assessing the effects of purchasing power shocks on the tourism industry or the airline industry. Finally, these results are of interest from a public policy point of view. Energy price shocks may have important effects on expenditures on higher education, the use of prescription drugs, or the choice of diets. Our analysis complements related survey-based results in the literature on public health.[19]

3.6.1 Motor Vehicle Consumption

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. Here we investigate all forms of consumption related to motor vehicles. Figure 13 shows that a reduction in purchasing power of -0.096% (corresponding to a one-standard deviation shock) causes a highly significant drop of -1.16% in the consumption of *motor vehicles and parts.* It is useful to consider different types of vehicles. Figure 13 shows that consumption of *pleasure boats* declines 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* does not change, nor does the consumption of 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 13 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

[19] In related work, Cullen et al. (2004) use panel data from the Consumer Expenditure Survey to characterize one specific trade-off faced by consumers, namely the "heat versus eat " decision.

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 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 13 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, largely 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 foreign-produced automobiles 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 shock of the magnitude associated with Hurricane Katrina could wipe out 10.3% of the domestic demand for U.S. automobiles.

The consumption data on new automobiles do not include light trucks or heavy trucks. A different approach to determining the importance of shifts among different types of automobiles is to focus on unit sales data 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

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

has been much discussion of the softening market for SUVs in recent years.[21] Figure 13 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]

3.6.2 Other Consumption

Faced with a reduction in purchasing power and an increased probability of becoming unemployed, there are several margins of adjustment available to households. One possibility is to delay large expenses such as vehicle repairs or major medical treatments that are not life-threatening. Another possibility is to economize on expenditures that have no immediate payoff such as vehicle maintenance or preventive health care. A third possibility is to forego nonessential purchases such as luxury or fashion items. For example, consumers may choose to spend less on jewelry, expensive clothing, vacations, dining out, lottery tickets, or entertainment. A fourth possibility is to delay major purchases such as appliances or furniture. A fifth possibility is to search for lower-priced alternatives such as public instead of private higher education or public transportation instead of private transportation. In discussing the relevance of these margins of adjustment it makes sense to group together goods and services as follows:

3.6.2.1 Other Durable Goods

The fact that payment plans for major durables constitute a long-term commitment suggests that households will cut back on durables purchases in response to unanticipated

[21] For example, Douglass (2007) discusses the cooling appetite for large sports utility vehicles, as gasoline prices are rising.
[22] This information was obtained from unit sales and production data on the company websites.

losses in purchasing power. Thus, we would expect to see declines in durables purchases such as appliances, computers, furniture, video and audio goods, or household furnishings. Many of these goods, such as video and audio goods, computers, and even new appliances and furniture, are luxuries which households might wish to forego when their purchasing power falls. Our results provide only very limited support for this common view, however. For example, we find a persistent but marginally statistically insignificant decline of -0.38% in *major appliances*, but only a small and statistically insignificant decline of -0.1% in *small electronic appliances*. There is no evidence of a decline in expenditures on *computers and software*. For *video and audio goods* we find a marginally significant decline of -0.67%.

Expenditures on both *furniture* and *outdoor equipment* show persistent but statistically insignificant declines with a delay of several months, reaching -0.41% and -0.55%, respectively, after 18 months. Declines in expenditures on *tools, hardware and supplies* as well as *china, glassware and utensils* are small and not statistically significant. Likewise, expenditures on *bicycles* and on *sporting equipment* as well as expenditures on *books and maps* show only insignificant declines, and expenditures on *household furnishings* show virtually no response.

Jewelry and watches are typically viewed as luxury items. As such, consumption of jewelry and watches is likely to be curtailed when purchasing power falls. This is exactly what we observe. The impulse response function for jewelry and watches shows a persistent decline in consumption of about -0.67% following the shock that is statistically significant at all horizons.

Overall, we conclude that the consumption of durables other than motor vehicles is not nearly as responsive to unanticipated shocks to purchasing power as one might have conjectured and typically only imprecisely estimated, consistent with the results in section 5.

3.6.2.2 Food

While overall food expenditures do not change significantly, there are some interesting patterns at the disaggregate level. One of the obvious margins of adjustment is for consumers to spend less money dining out. We find that consumption of *food in restaurants* very quickly and sharply declines by -0.25%. The decline is highly significant and very persistent. The overall pattern differs from that for *meals in limited service restaurants* (fast food). The latter response exhibits an immediate, but purely temporary and insignificant drop of about -0.19%. That drop is completely reversed after four months and followed by a flat response.

The counterpart of these responses is a short run increase in expenditures on the consumption of food at home. We find a statistically significant increase in the consumption of *food and beverages at home* of 0.15% in the short run and an insignificant increase of 0.09% in the long run. Even more persistent and statistically significant is the increase in expenditures on *alcohol consumption at home* which reaches a peak of 0.41% after one month. This increase suggests that people consume alcohol at home rather than go out for a drink. There is also a partially statistically significant and persistent decline in expenditures on *school lunches* (-0.31%).

A second margin of adjustment involves substituting less expensive foods for high priced foods. For example, expenditures on *fresh fruits* and *fresh vegetables* decline by -0.35% and -0.49%, respectively, although the decline is not statistically significant. There is also a smaller and partially statistically significant decline in the consumption of *processed fruits and vegetables* (-0.23%). A partially statistically significant decline in the consumption of *fish and seafood* (-0.55%) contrasts with a statistically significant increase in the consumption of *pork* (0.90%) and *poultry* (0.73%), indicating that consumers substitute pork and poultry for more expensive fish and seafood. In addition, there is a statistically insignificant decline in expenditures on *nonalcoholic drinks* in general (-0.45%) and of *tea and coffee* in particular (-0.76%). Food items such as *cereals*, *beef and veal*, *eggs*, and *dairy* *products* increase insignificantly, consistent with increased food consumption at home. The short-run increase in *bakery products* by 0.15% is statistically significant. The fact that overall *food* expenditures do not decline suggests that ultimately attempts to reduce food consumption are unsuccessful.

3.6.2.3 Other Nondurables

Seemingly nonessential expenditures on *toys, dolls and games* show only insignificant declines. Likewise, expenditures on *sports supplies* show virtually no response. Expenditures on *tobacco* decline after a few months and reach -0.20%, but the decline is statistically insignificant. Similarly, expenditures on *cosmetics and perfume* show a decline of -0.37%, but the decline is not statistically significant.

When it comes to clothing, the line between essential and nonessential expenditures is blurred. Total expenditures on *clothing and shoes* decline by about -0.25%, but the decline is statistically significant only for months 3, 4 and 5 after the shock. The decline in expenditures on *women's clothes* is greater in the short run, but smaller in the long run, compared to that for *men's clothes.* The decline in expenditures on *infants' clothing*, while large in magnitude with an estimated peak response of -0.37%, is not statistically significant at any horizon.

3.6.2.4 Travel, Transportation, and Tourism

It seems natural to expect that consumers will reduce their nonessential travel expenditures in response to unanticipated losses in purchasing power as well as choose more economical means of transportation. For example, it is widely believed that higher energy prices are associated with reduced consumer demand for airline travel between cities. Personal (as opposed to business) airline travel often is a luxury which consumers are likely to forego when their purchasing power falls. Indeed, we find that consumption of *airline tickets* falls by -0.79% with a delay of three months. The reduction is very persistent and highly statistically significant. This result implies that a Katrina-sized shock would reduce personal airline travel by about -5.4%.

Expenditures on *intercity railway travel* increase persistently by 0.5%, but the response is not statistically significant. Expenditures on *intercity bus travel* rise sharply upon impact, reaching a highly statistically significant peak of 1.5% after one month before the response falls to 0.7% and becomes insignificant. This pattern is clear evidence of a shift in expenditure patterns away from private transportation.

As far as local transportation is concerned, we find insignificant declines in expenditures on *taxi cabs* and in *bridge, tunnel, ferry and road tolls.* Like in the case of intercity bus travel, there is a sharp and statistically significant increase in expenditures on *mass transit* (local bus, metro, subway) on impact, but that increase is purely temporary and insignificant after three months. The temporary nature of the response suggests that consumers switch to mass transit at first, but soon decide that the added inconvenience is not worth the savings.

Combining local and intercity transportation, we find a highly significant decline in expenditures on *transportation* after half a year that reaches -0.52% after 18 months. As consumers reduce their travel, one would expect expenditures on tourism to be hit hard. For example, according to *The Detroit News*, "Michigan's dismal economy and record setting gasoline prices are the main culprits in an expected decline in the state's tourism industry this year" (see Greenwood 2006). Expenditures on tourism may fall for several reasons. First, recreation and vacations may be viewed as luxuries, so households will cut back on this type of consumption before other forms of consumption such as food, clothing, household supplies, or education. Second, the complementarity with automobile travel makes recreation and vacation consumption more expensive when energy prices, and specifically gasoline prices, rise. Consistent with this line of reasoning we find a large and persistent decline in expenditures on *hotels and motels* of -0.39%, although that decline is not statistically significant.

3.6.2.5 Education

The decline in total expenditures on *higher education* of about -0.14% is largely statistically insignificant. Further disaggregation suggests a shift in expenditure patterns, however, along the lines conjectured earlier. Whereas expenditures on *private higher education* decline insignificantly by -0.23%, expenditures on *public higher education* rise insignificantly by 0.14%, in line with the view that consumers substitute lower-priced public universities for higher-priced private universities. Whereas expenditures on *elementary and secondary schools* decline only slightly and not significantly, expenditures on *nursery schools* show a marginally significant decline of -0.75% in the long run. This pattern is consistent with the view that an increase in unemployment reduces the need for nursery schools. Expenditures on *vocational schools* show a highly significant decline of -0.74% in the long run. The relatively strong (if insignificant) response of expenditures on higher education suggests the presence of a precautionary savings motive. We conclude that households reduce their investment in forms of education that may be considered discretionary such as nursery schools, vocational training, or private higher education.

3.6.2.6 Health Care

The response of household health care expenditures to unanticipated losses in purchasing power may be mitigated by a number factors. First, many forms of health care are necessities that are unlikely to be reduced in response to purchasing power losses. Second, households with health insurance face a very low marginal cost to consuming health care goods and services, and are therefore unlikely to face a trade off between health care and energy consumption. While some households pay for health insurance out of their own pocket, and might decide to reduce or eliminate their coverage in order to meet higher energy costs, many households benefit from employer-sponsored or government-sponsored coverage and would not be required to do this. Thus, one would expect a loss in purchasing power to reduce only those health care items that are typically not covered by health insurance or require a large deductible.

The data show that expenditures on health care typically only react with some delay to unanticipated losses in purchasing power. There is a clear, but largely insignificant decline in expenditures on *eye exams* of -0.47%, consistent with a reduction of expenditures on preventive health care. Similarly, while expenditures on *physicians* remain constant, expenditures on *dentists* decline by -0.16%, albeit not significantly so. The response of expenditures on *medical care* is largely flat. The decline in expenditures on *prescription drugs* as well as on *nonprescription drugs* of up to -0.20% and -0.13%, respectively, is not statistically significant. Expenditures on *home health care* do not decline significantly either. Overall, we find only weak evidence of consumers adjusting their health care expenditures, consistent with the notion that most of these expenditures are necessities and largely covered by health insurance.

3.6.2.7 Entertainment and Leisure

Somewhat surprisingly, there is no significant decline in total expenditures on *recreation*. The response is largely flat, negligible in magnitude and statistically insignificant. More specifically, we find no significant effect on the consumption of *movies*, *bowling and billiard*, *casino gambling* and only insignificant declines for *recreational camps*, *sightseeing*, *spectator sports*, and *spectator amusements*. Thus, entertainment expenditures seem remarkably robust to shocks to purchasing power. We observe statistically significant declines only for purchases of *lottery tickets* (-1.2%) and expenditures on *theater tickets* (-0.62%).

These results contradict common perceptions. For example, we would expect fewer consumers to go to the movies, both because going to the movies involves driving and because watching movies in theaters is relatively expensive compared to renting movies or watching them on television. Similarly, the lack of responsiveness of expenditures on sightseeing and spectator amusements and sports (known to be expensive in many cases), does not fit the view that unanticipated purchasing power losses hurt the tourism and entertainment industry.

3.6.2.8 Other Services

An important component of services is telecommunications. While households may opt not to eliminate cable television service, they can chose to downgrade their service by eliminating premium channel options. Telephone service is more of a necessity, but households can save money by foregoing options such as call-waiting and caller-id or by substituting land lines for cell phones. Furthermore, cell phone subscribers can save money by eliminating optional service features and reducing their minutes. In contrast to these hypotheses, our results show no significant decline in *cable tv* expenditures or *local and cellular telephone service*, suggesting that households view these services as essential. Nor is there a significant change in expenditures on *internet service providers.*

However, we find a persistent decline in expenditures on *personal business* that is statistically significant for the first six months. The long-run decline in overall personal business expenditures is -0.22%. Particularly responsive are expenditures on *brokerage charges and investment counseling* which decline by -1.39%. The decline is partially statistically significant.

Another service on which households may choose to economize is insurance services. Households may economize by letting contracts lapse, by switching to less expensive providers, and by choosing higher deductibles or lower levels of coverage. We find no statistically significant or large responses of expenditures on *auto insurance* or *household insurance.*

A last group of services relates to maintenance and upkeep. We find a reduction of about -0.30% in expenditures on *vehicle tires and parts* that is partially statistically significant. There is also a highly statistically significant drop of -0.56% in the consumption of *vehicle repair* services, consistent with the notion that consumers delay repairs when faced with an unanticipated loss of purchasing power. We also find a marginally significant decline of -0.52% of expenditures on *personal care* such as *repair services*, *dry-cleaning*, *barbershops*, *beauty parlors and health clubs.*

3.7 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). The latter effect is important in that it explains how the effect of energy price changes may be larger than would be expected based on 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.[23]

3.7.1 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 fluctuations. 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

[23]These estimates presume that households respond proportionately to changes in energy prices. As discussed in section 4, 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 same qualitative patterns as in Figure 9. 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.

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 5. Figure 14 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 14, 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 14 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.

3.7.2 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* was -0.31% relative to its mean in 1986 (see Table 3). 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 3 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.[24] 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 4, we fail to reject the null of symmetric responses to positive and negative

[24] Some 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.

shocks with p-values in the range from 0.66 to 0.91.[25] 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 boom 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 3). Such a pattern is inconsistent with conventional explanations of asymmetric investment responses. As we pointed out in section 4, 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 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.

[25] For 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).

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.[26]

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

[26] For 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).

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 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.

3.7.3 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 channel, 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.

3.7.4 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.[27] Figure 15 quantifies this phenomenon by comparing responses 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 5. The scale of the impulses has been normalized to be the same across the two samples as for the full sample. Figure 15 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%. Vehicles consumption declines by -1.99% in the first half of the sample and by -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%. These responses are in general consistent with reductions in consumption that one would expect from the discretionary income effect. A similar reduction occurs in the response of quarterly 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

[27] A 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 prononced economic effects than the more recent increases (see, e.g., Bernanke 2004).

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 1970s.

This point is illustrated by comparing the responses of new domestic and foreign automobiles in each subsample (see Figure 16). 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 restructured itself after the energy price increases of the 1970s, the share of domestically produced automobiles in total U.S. real expenditures on new cars declined (from 88% in 1970 to 60% in 1988 and 57% in 2006), as did the employment share of the industry (from a peak of 1.3% in 1973 to 0.9% in 1988 and 2005).[28] 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.

3.8 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 result in 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

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

and firms, the effects of energy price shocks on the economy will be small. In this paper we studied in 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 problems 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 based on a detailed analysis of more than 130 expenditure items, 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% one 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 sheds 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.

Table 3.1: Predicted Response of Non-Energy Consumption to an Increase in Energy Prices

	$\theta = 1$	$\theta \in [0, 1)$
$\rho = 0$	No Response	No Response
$\rho \in (0, 1)$	No Response	Positive Response
$\rho < 0$	No Response	Negative Response

In the model of section 3, θ is the fraction of aggregate energy revenues that is transferred to U.S. consumers as supplemental income. ρ is the elasticity of substitution between energy and non-energy consumption in a CES utility function.

Table 3.2: 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 3.3: 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 3.1: Monthly Real PCE Price Index: 1970.1-2006.7

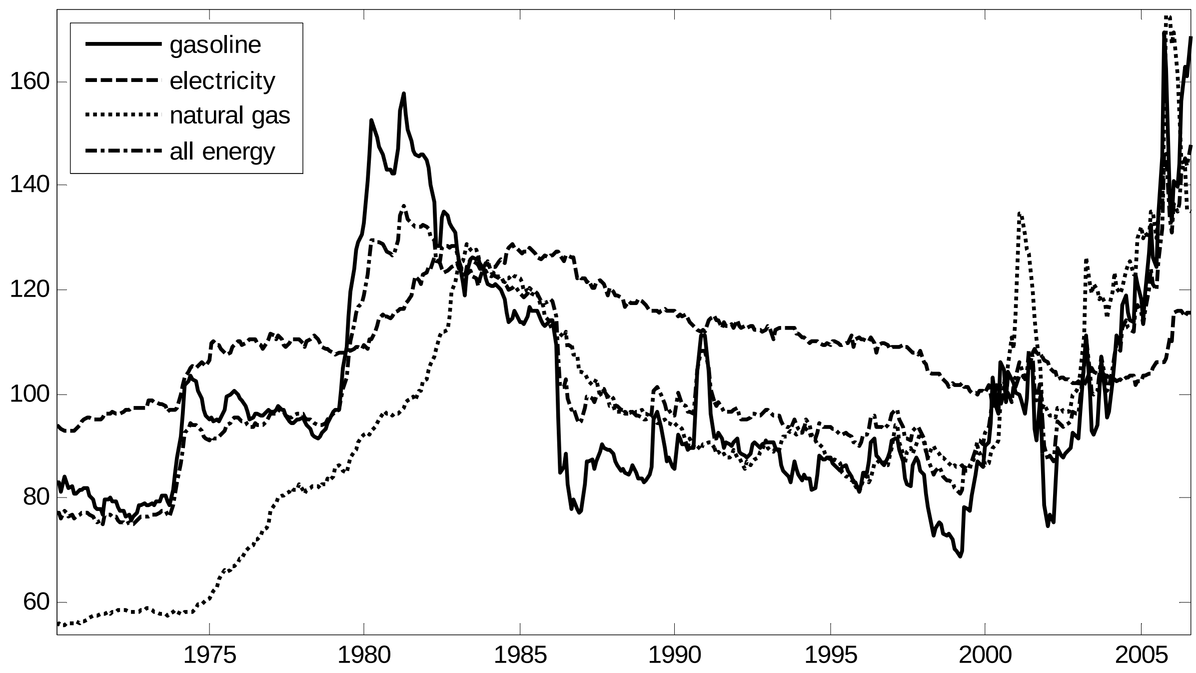

Figure 3.2: Monthly Nominal Expenditure Share: 1970.1-2006.7

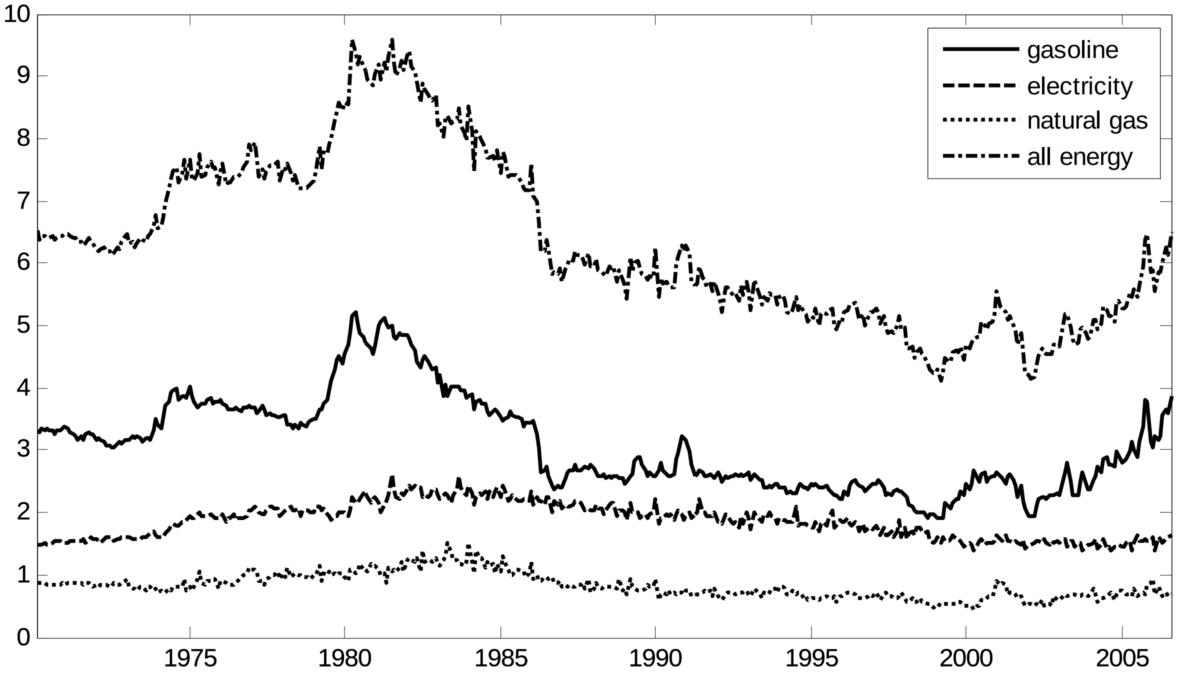

Figure 3.3: Monthly Loss in Purchasing Power: 1970.2-2006.7

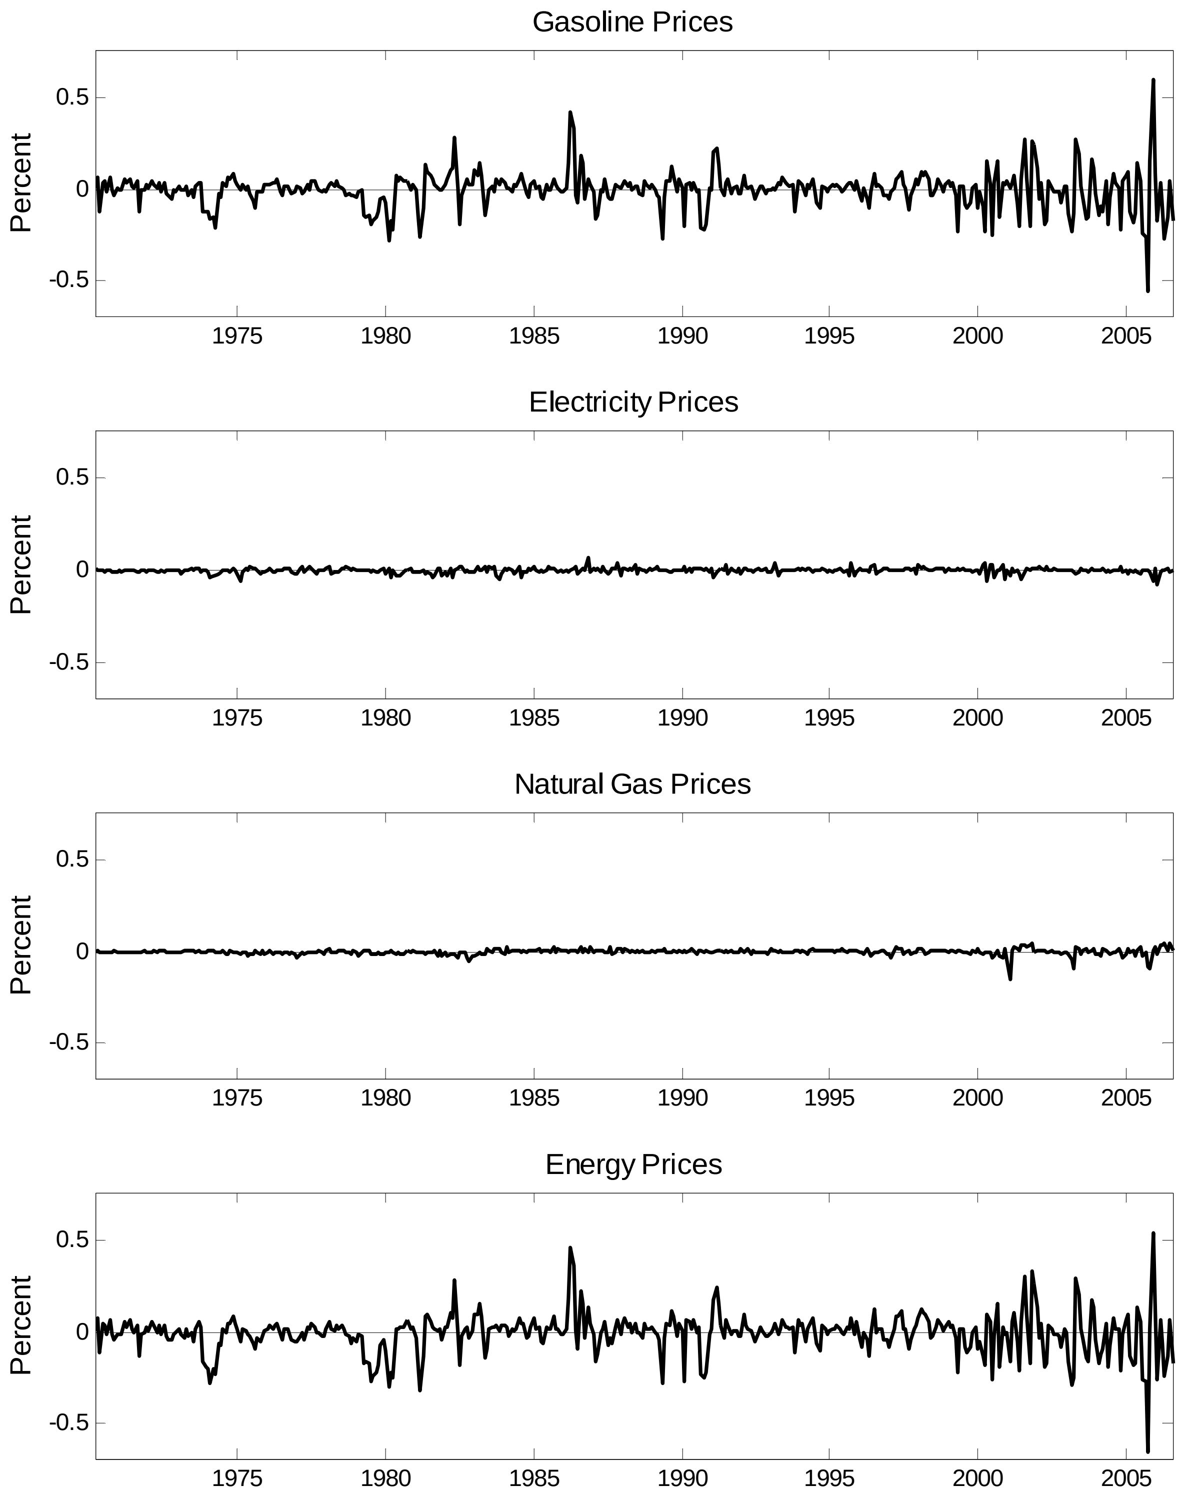

Figure 3.4: Alternative Measures of Purchasing Power Gains and Losses: 1970.2-2006.7

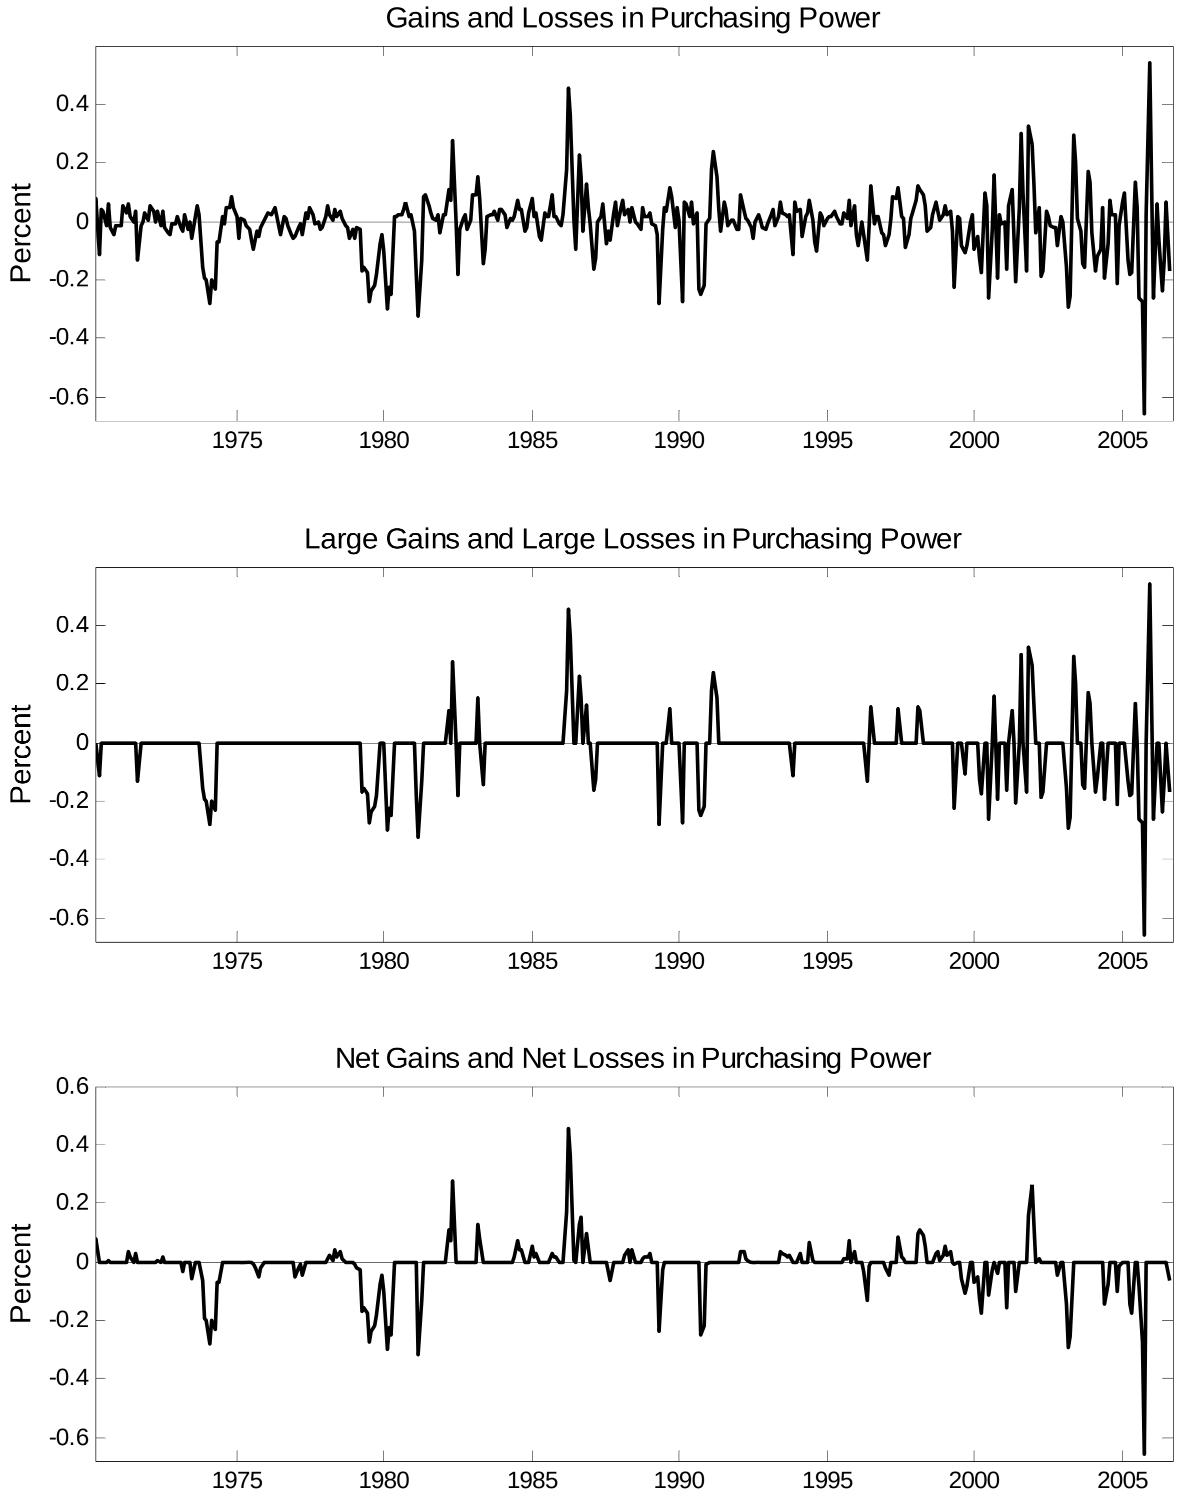

Figure 3.5: Response of Real Consumption to Purchasing Power Shocks: 1970.2-2006.7

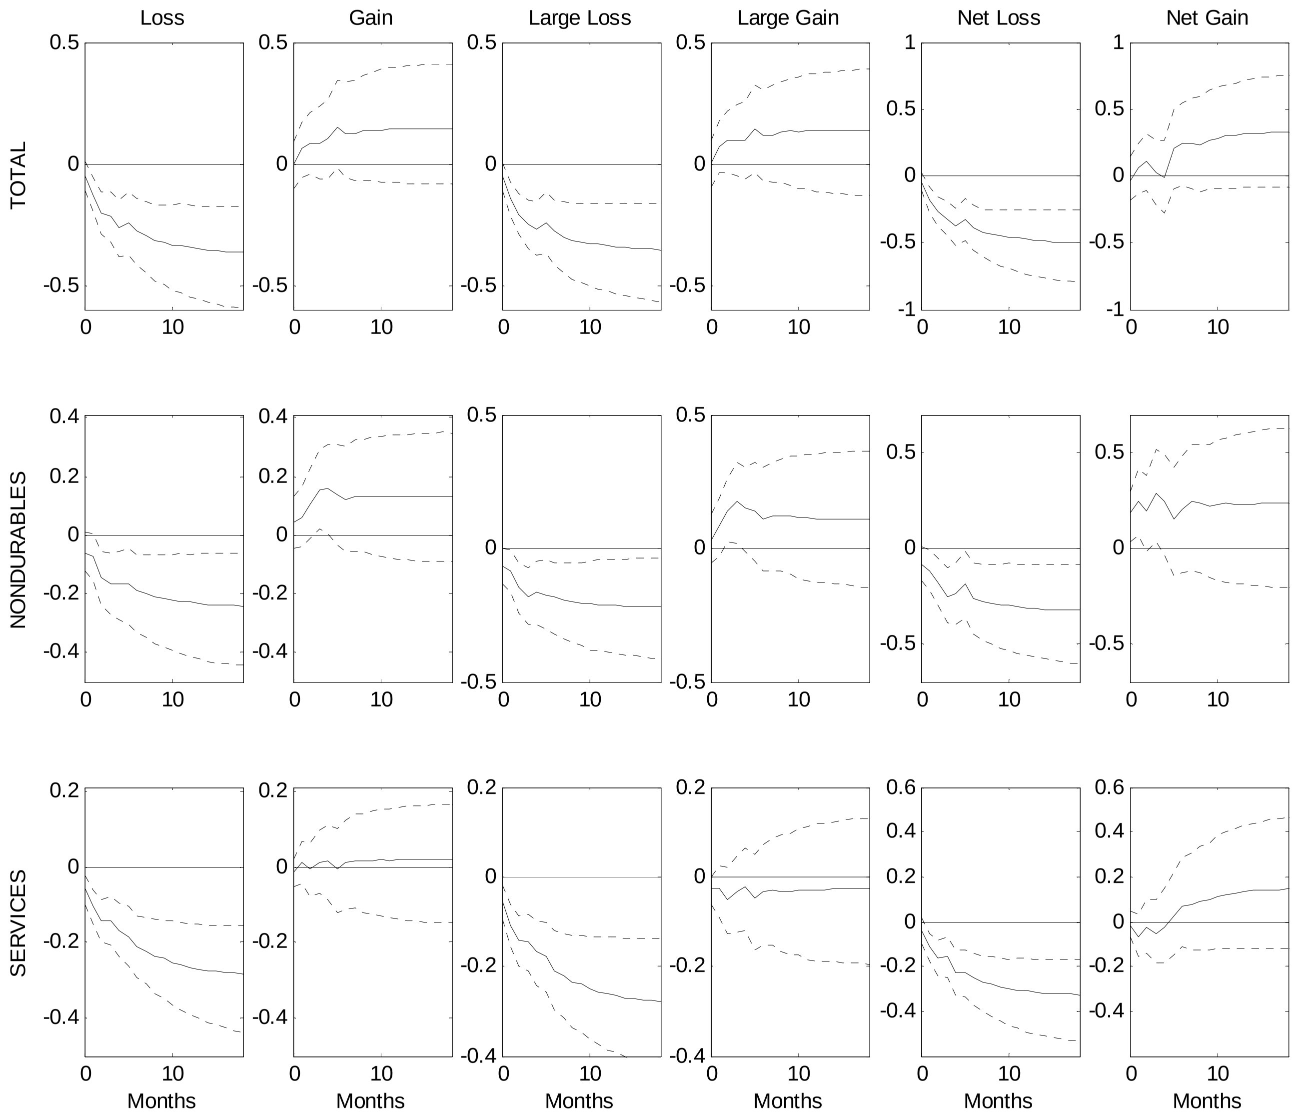

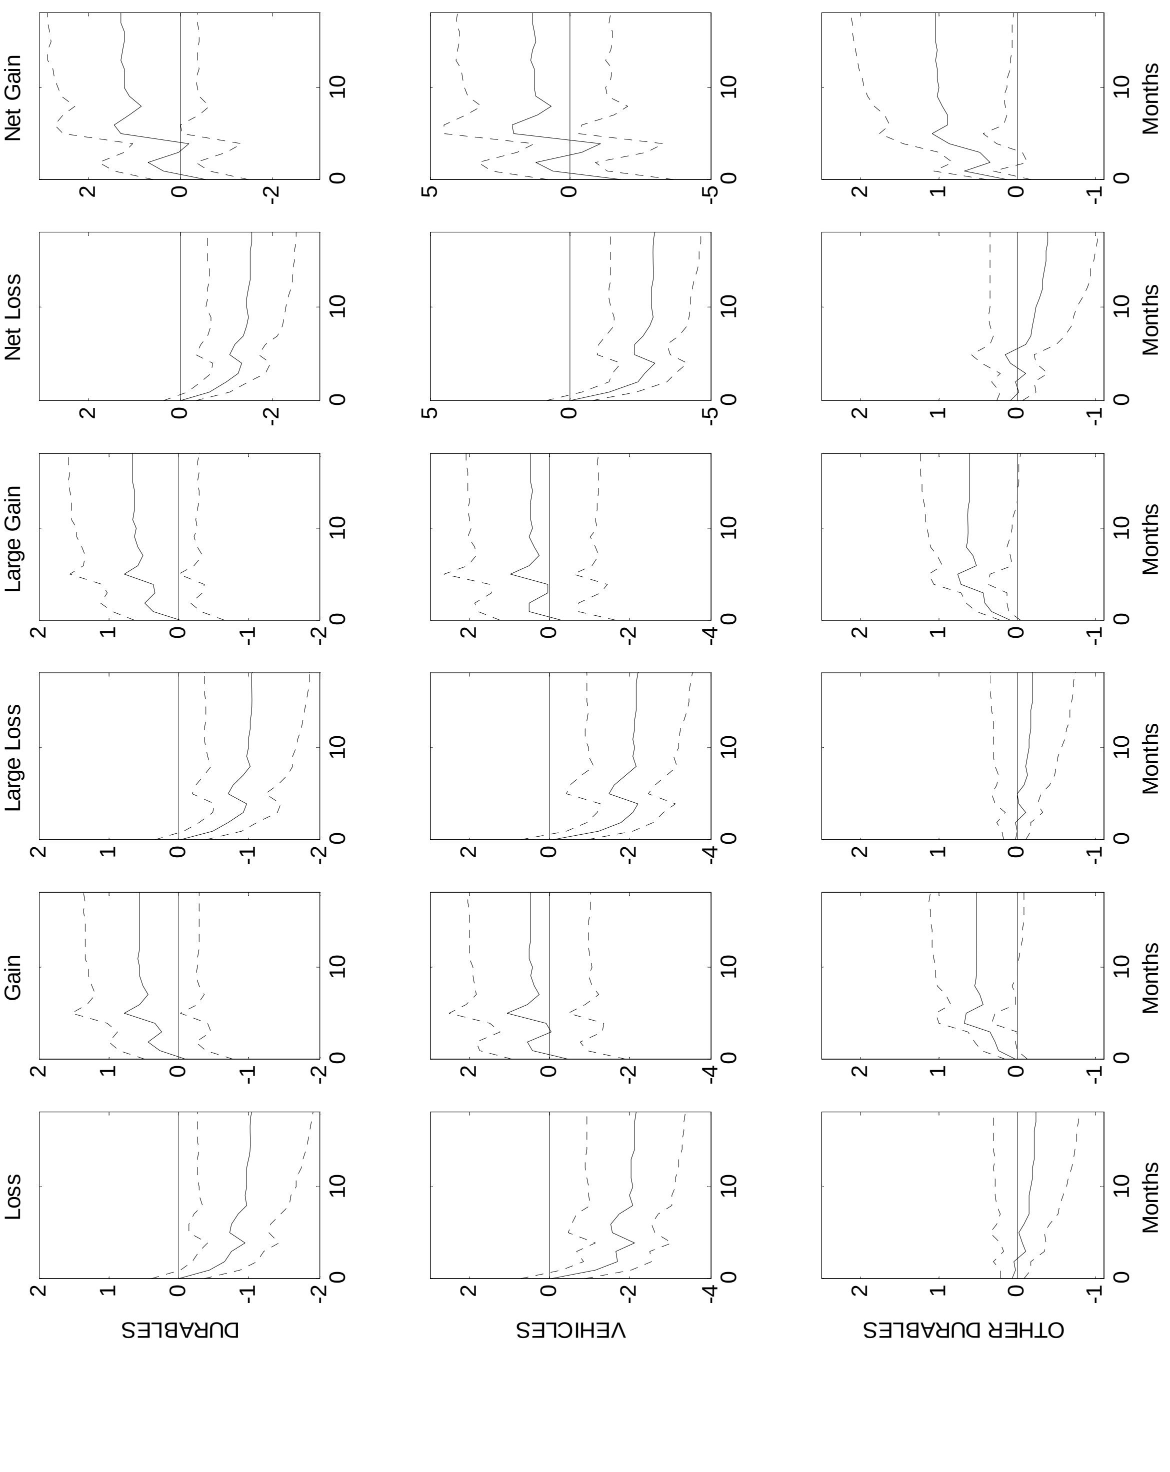
Loss
Gain
Large Loss
Large Gain
Net Loss
Net Gain
DURABLES
VEHICLES
OTHER DURABLES
Months

Figure 3.6: Response of Unemployment Rate to Purchasing Power Shocks: 1970.2-2006.7

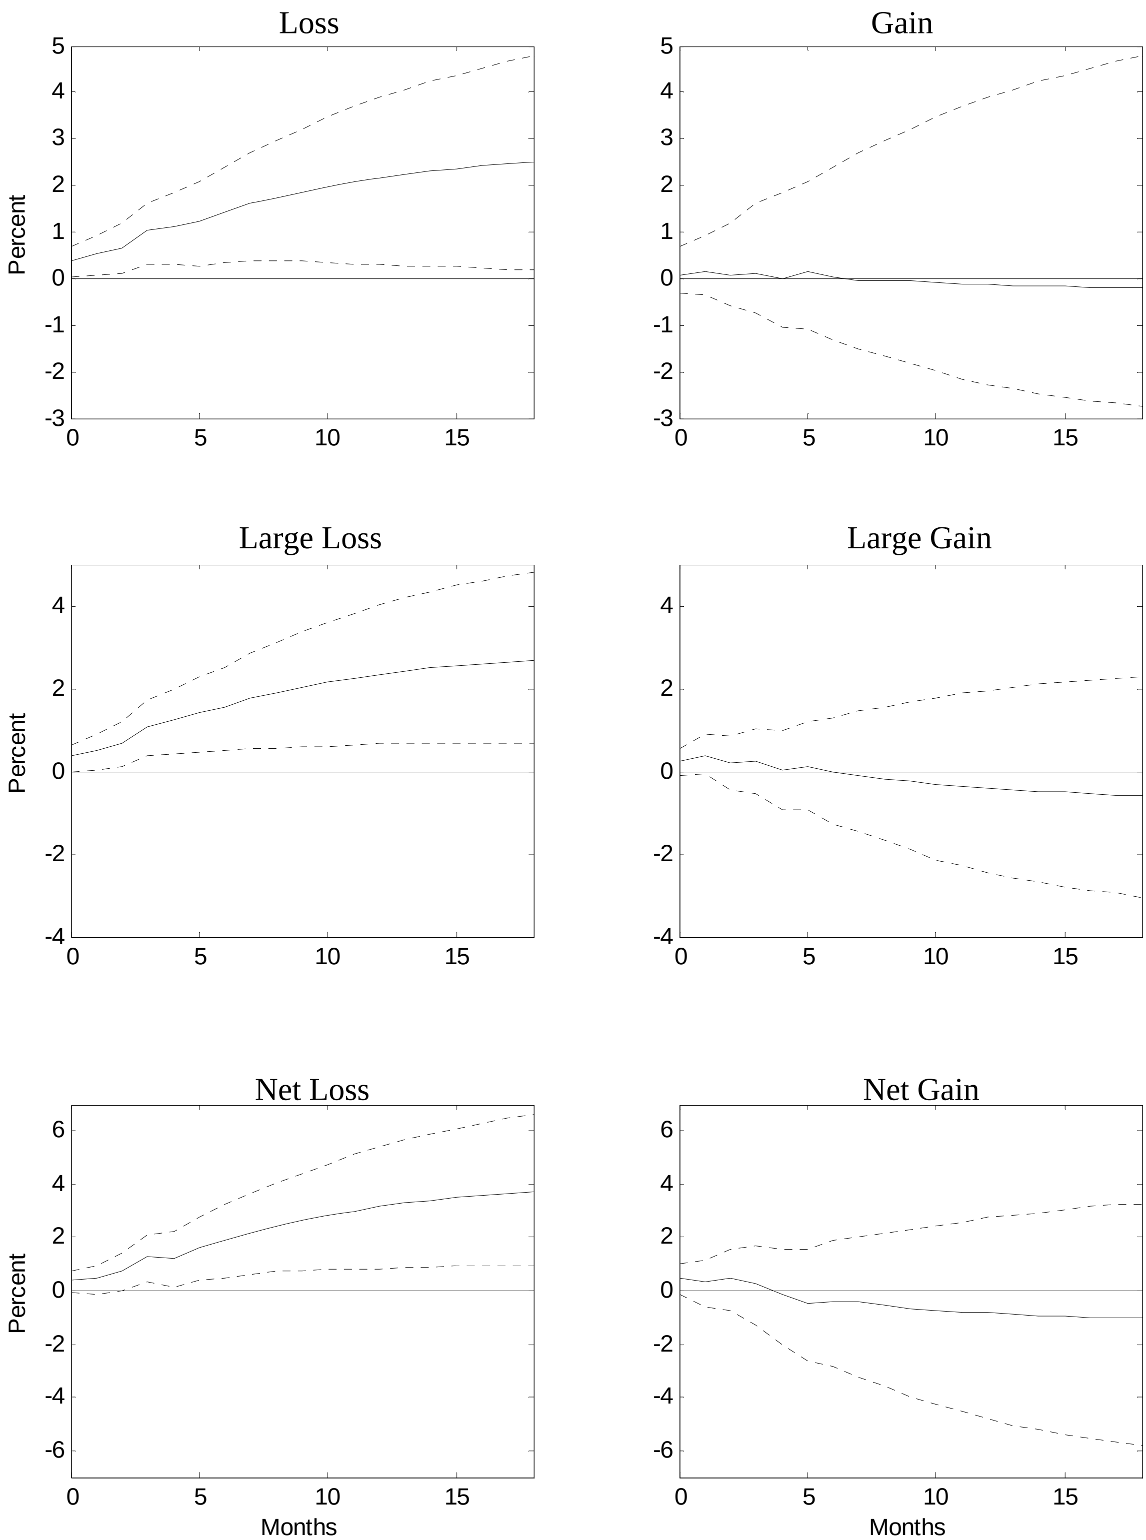

Figure 3.7: Response of Consumer Expectations to Purchasing Power Shocks: 1978.1 -2006.5

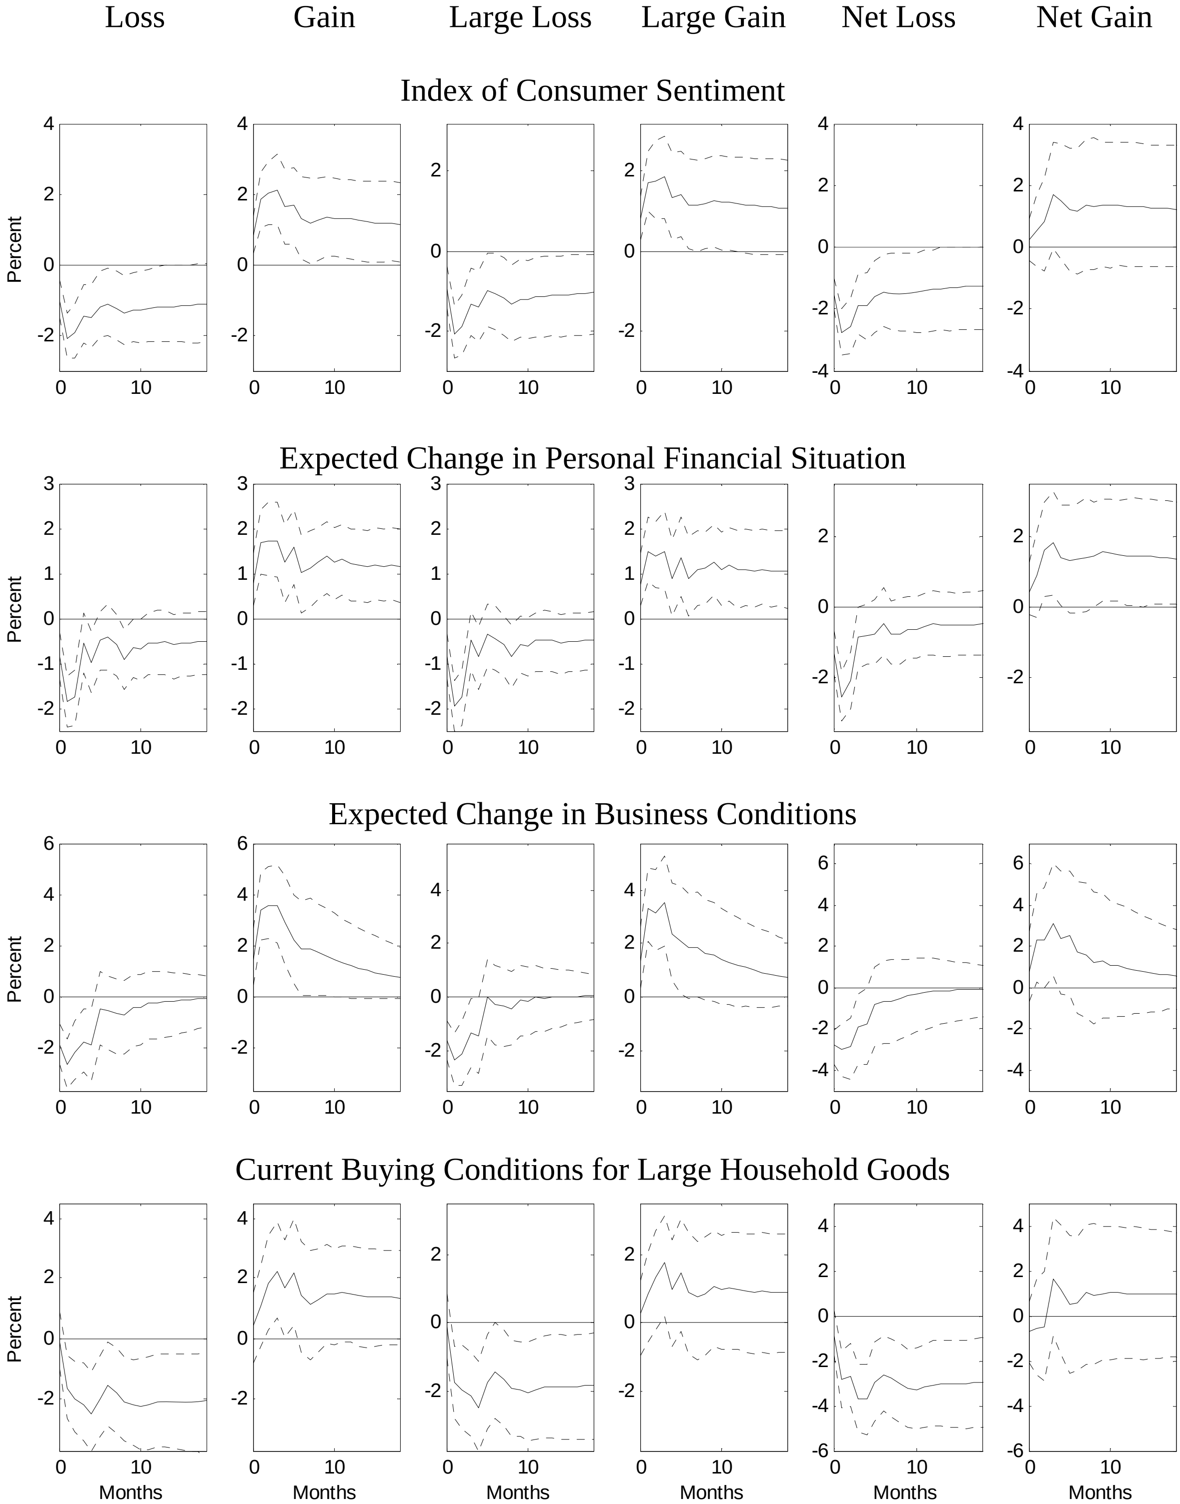

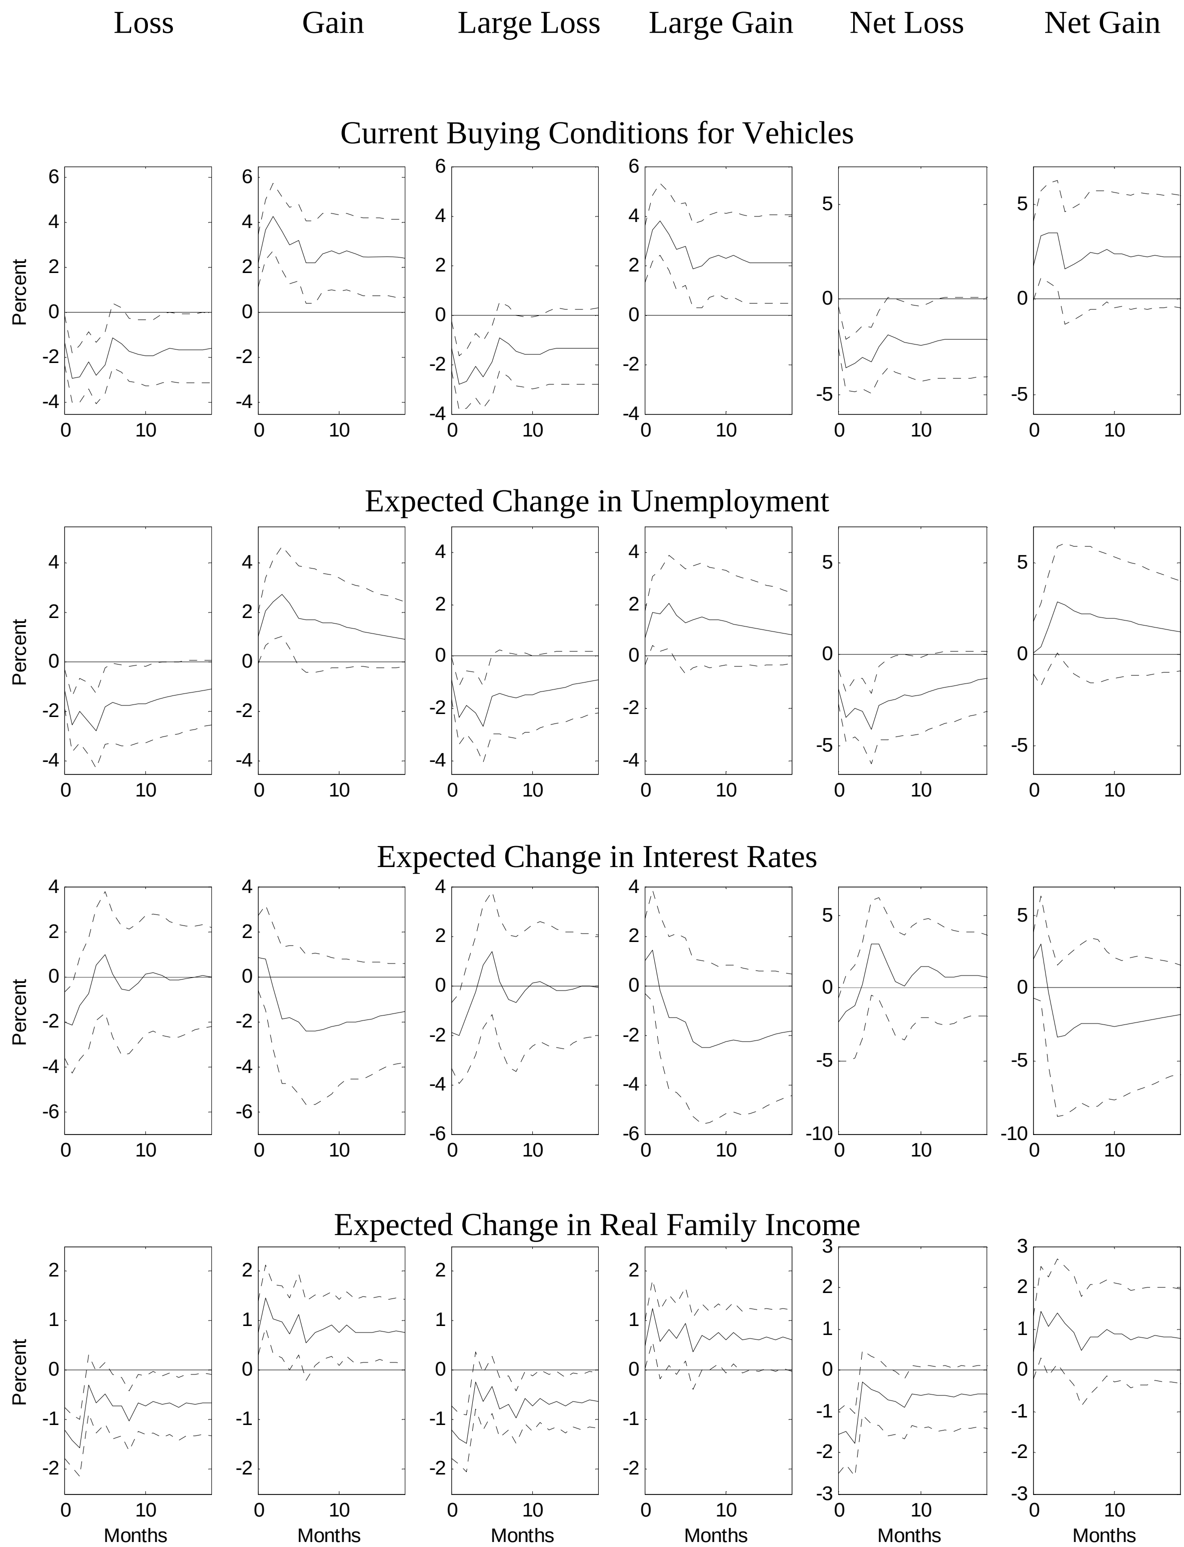

Loss
Gain
Large Loss
Large Gain
Net Loss
Net Gain
Current Buying Conditions for Vehicles
Expected Change in Unemployment
Expected Change in Interest Rates
Expected Change in Real Family Income
Percent
Months

Figure 3.8: Response of Purchasing Power to an Unanticipated Purchasing Power Shock: 1970.2-2006.7

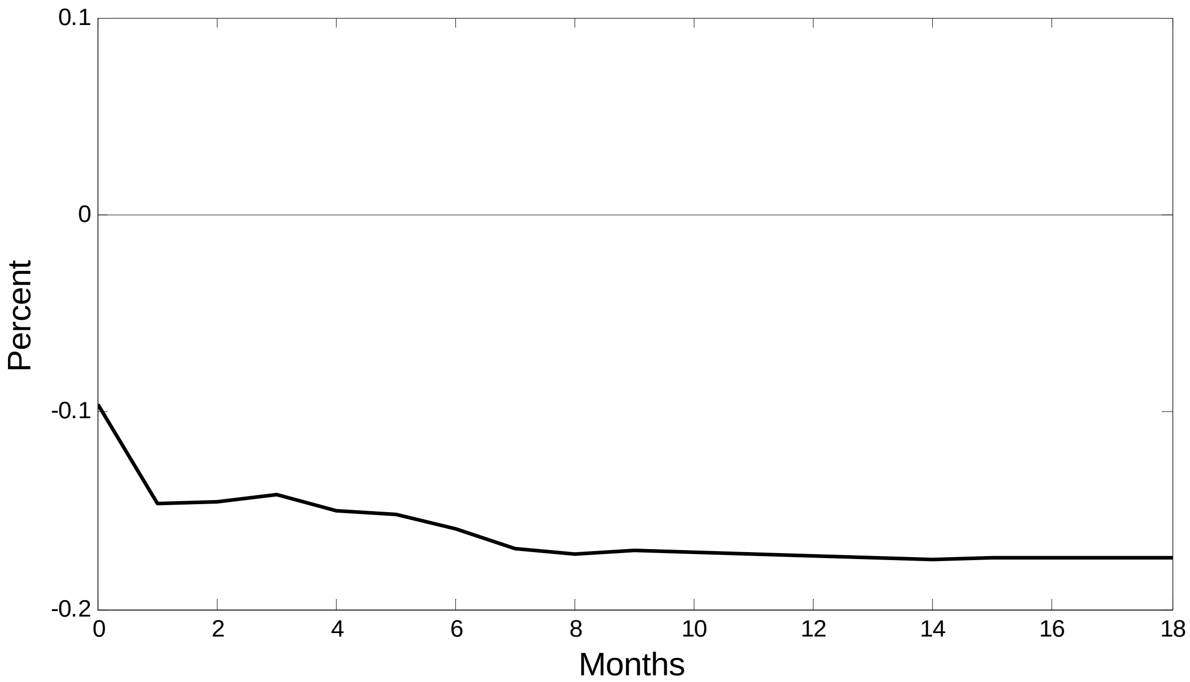

Figure 3.9: Responses of Real Consumption to a Purchasing Power Shock: 1970.2-2006.7

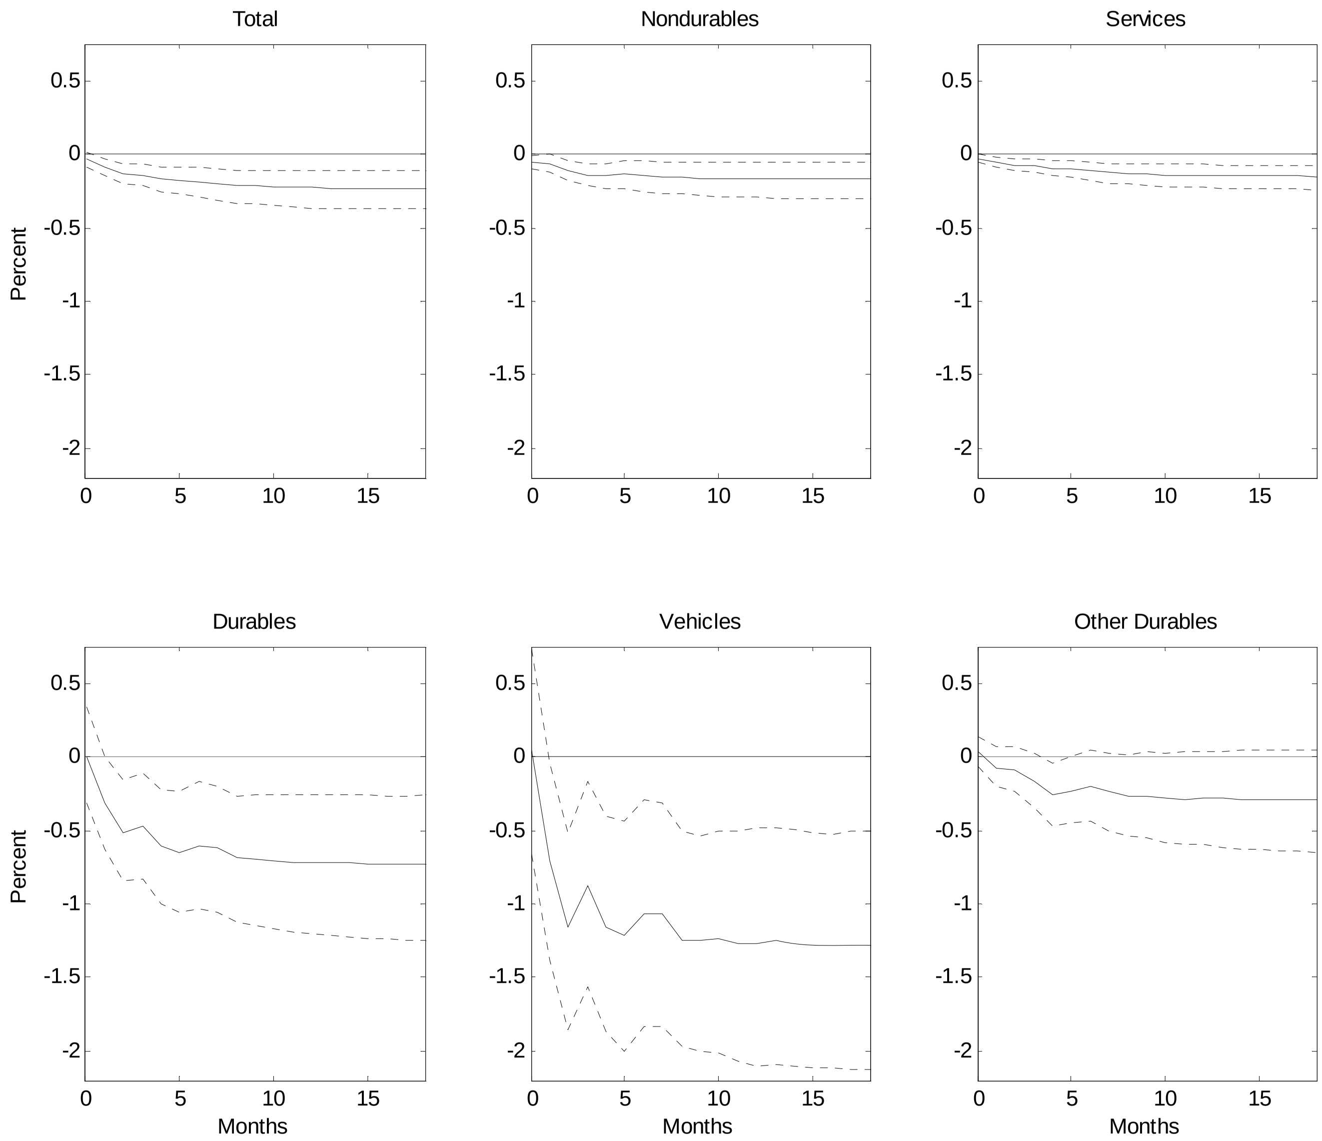

Figure 3.10: Response of Real Energy Consumption to a Purchasing Power Shock: 1970.2-2006.7

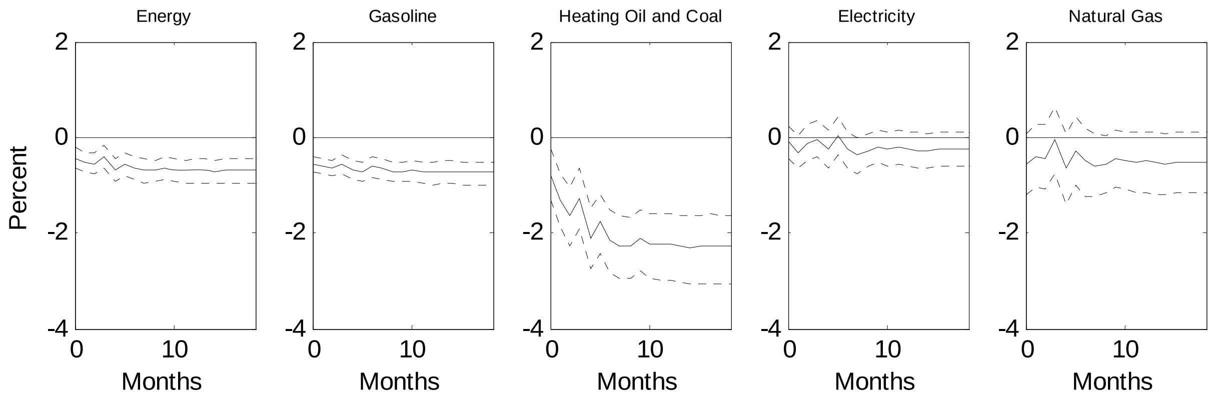

Figure 3.11: Response of Unemployment Rate to a Purchasing Power Shock: 1970.2-2006.7

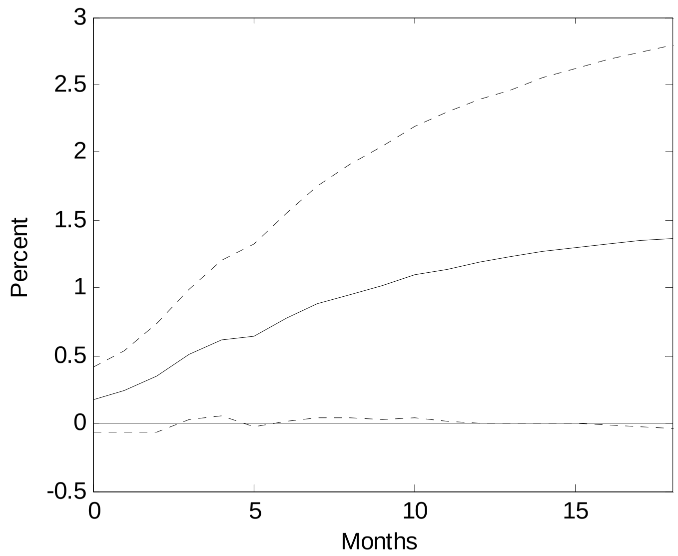

Figure 3.12: Response of Consumer Expectations to a Purchasing Power Shock: 1978.1-2006.5

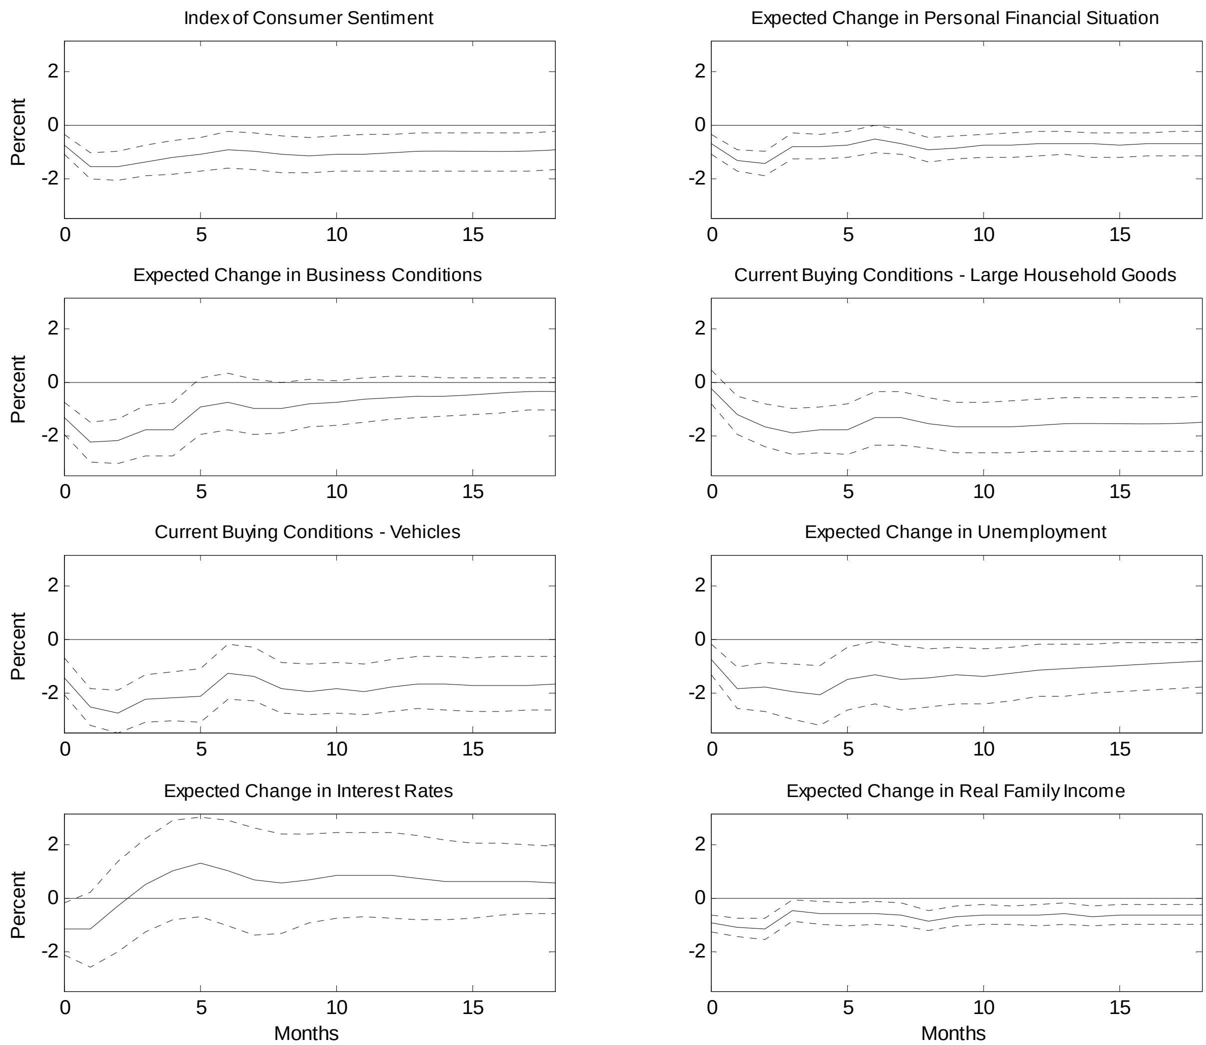

Figure 3.13: Response of Real Consumption by Expenditure Item: 1970.2-2006.7

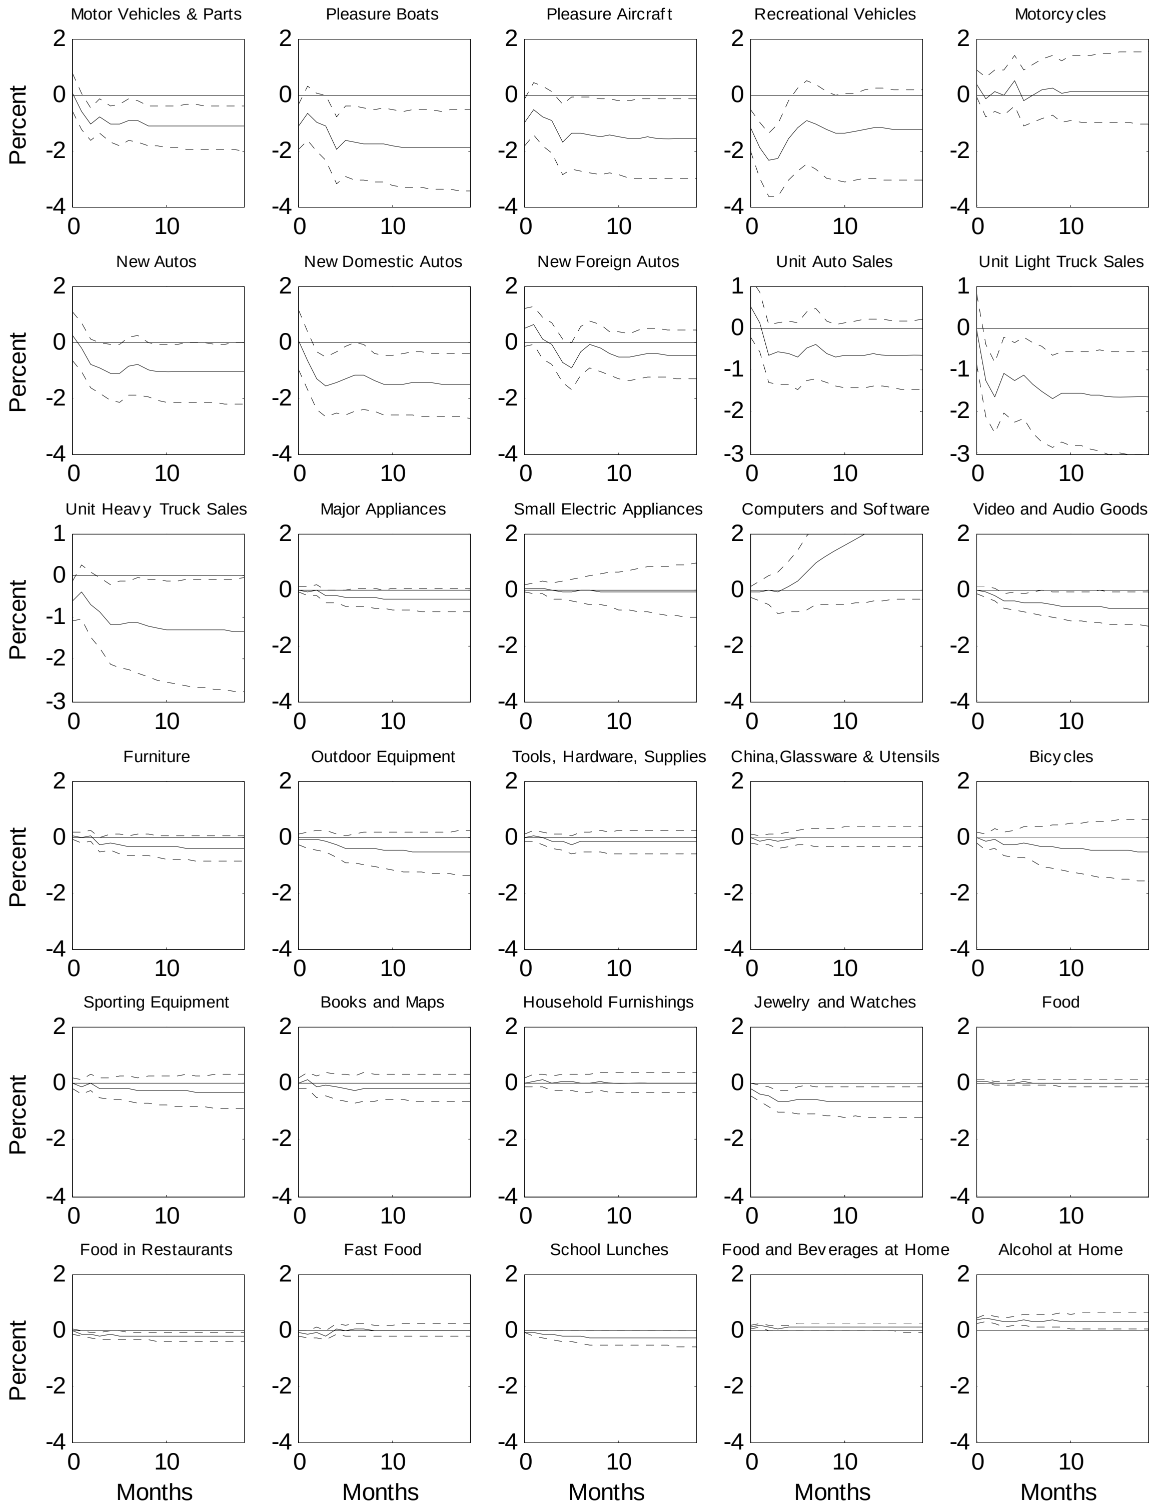

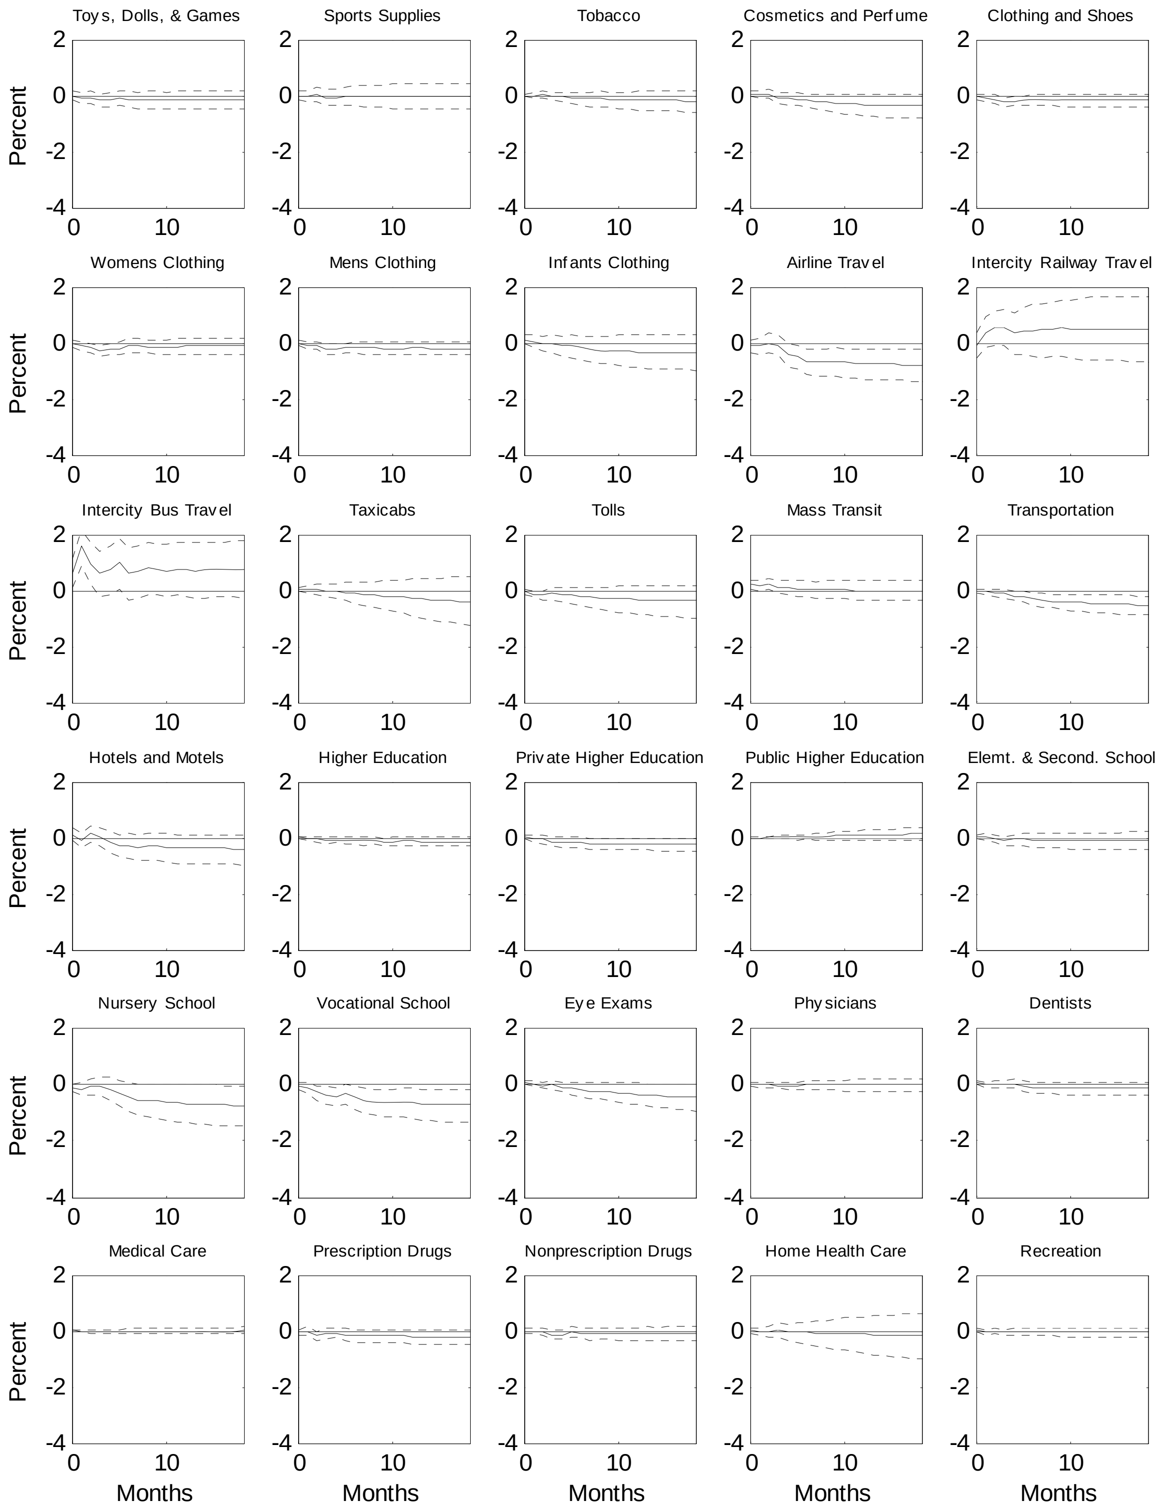
Toys, Dolls, & Games
Sports Supplies
Tobacco
Cosmetics and Perfume
Clothing and Shoes
Womens Clothing
Mens Clothing
Infants Clothing
Airline Travel
Intercity Railway Travel
Intercity Bus Travel
Taxicabs
Tolls
Mass Transit
Transportation
Hotels and Motels
Higher Education
Private Higher Education
Public Higher Education
Elemt. & Second. School
Nursery School
Vocational School
Eye Exams
Physicians
Dentists
Medical Care
Prescription Drugs
Nonprescription Drugs
Home Health Care
Recreation
Percent
Months
2
0
-2
-4
0
10

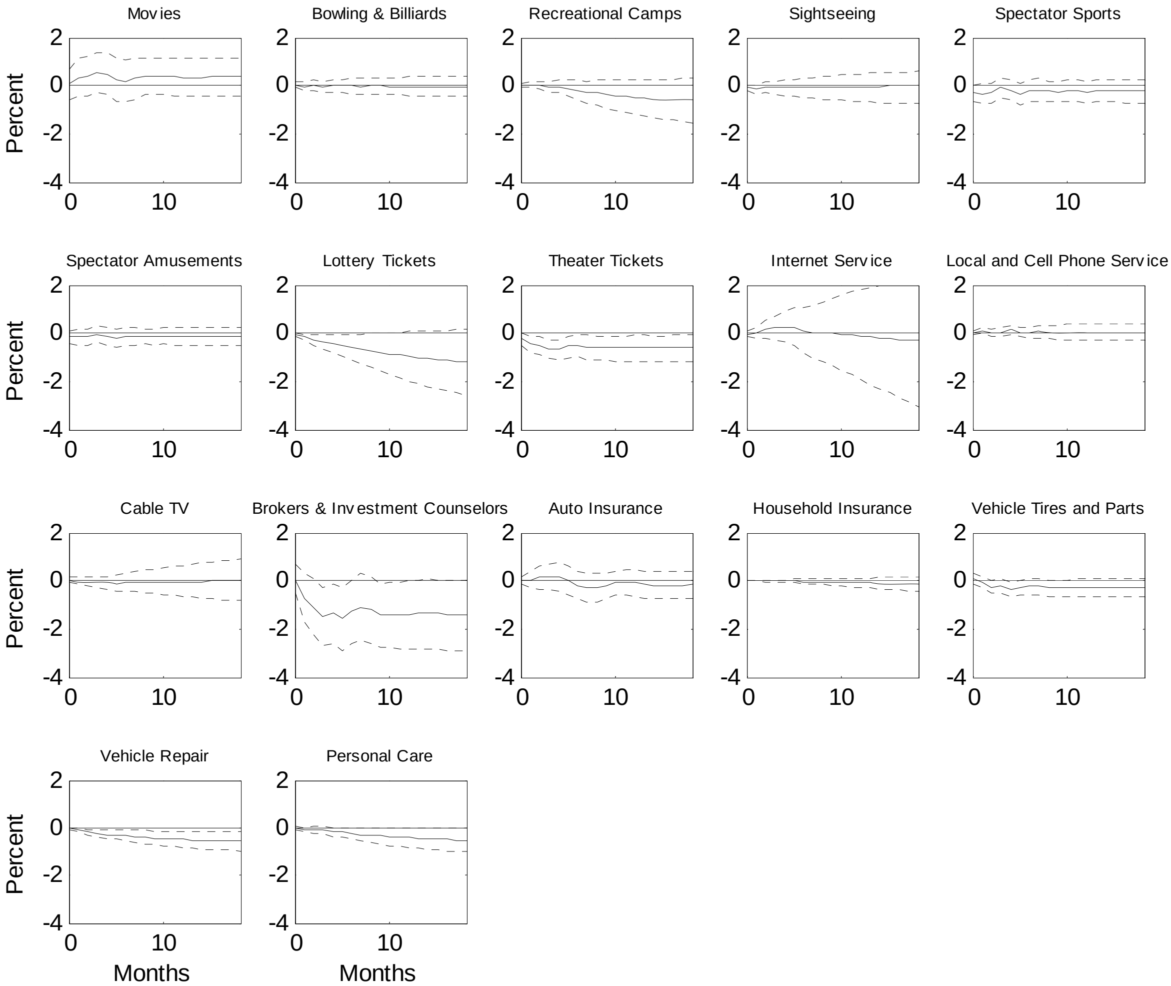
Movies
Bowling & Billiards
Recreational Camps
Sightseeing
Spectator Sports
Spectator Amusements
Lottery Tickets
Theater Tickets
Internet Service
Local and Cell Phone Service
Cable TV
Brokers & Investment Counselors
Auto Insurance
Household Insurance
Vehicle Tires and Parts
Vehicle Repair
Personal Care
Percent
Months
2
0
-2
-4
0
10

Figure 3.14: Contribution of Purchasing Power Shocks to Real Consumption Growth

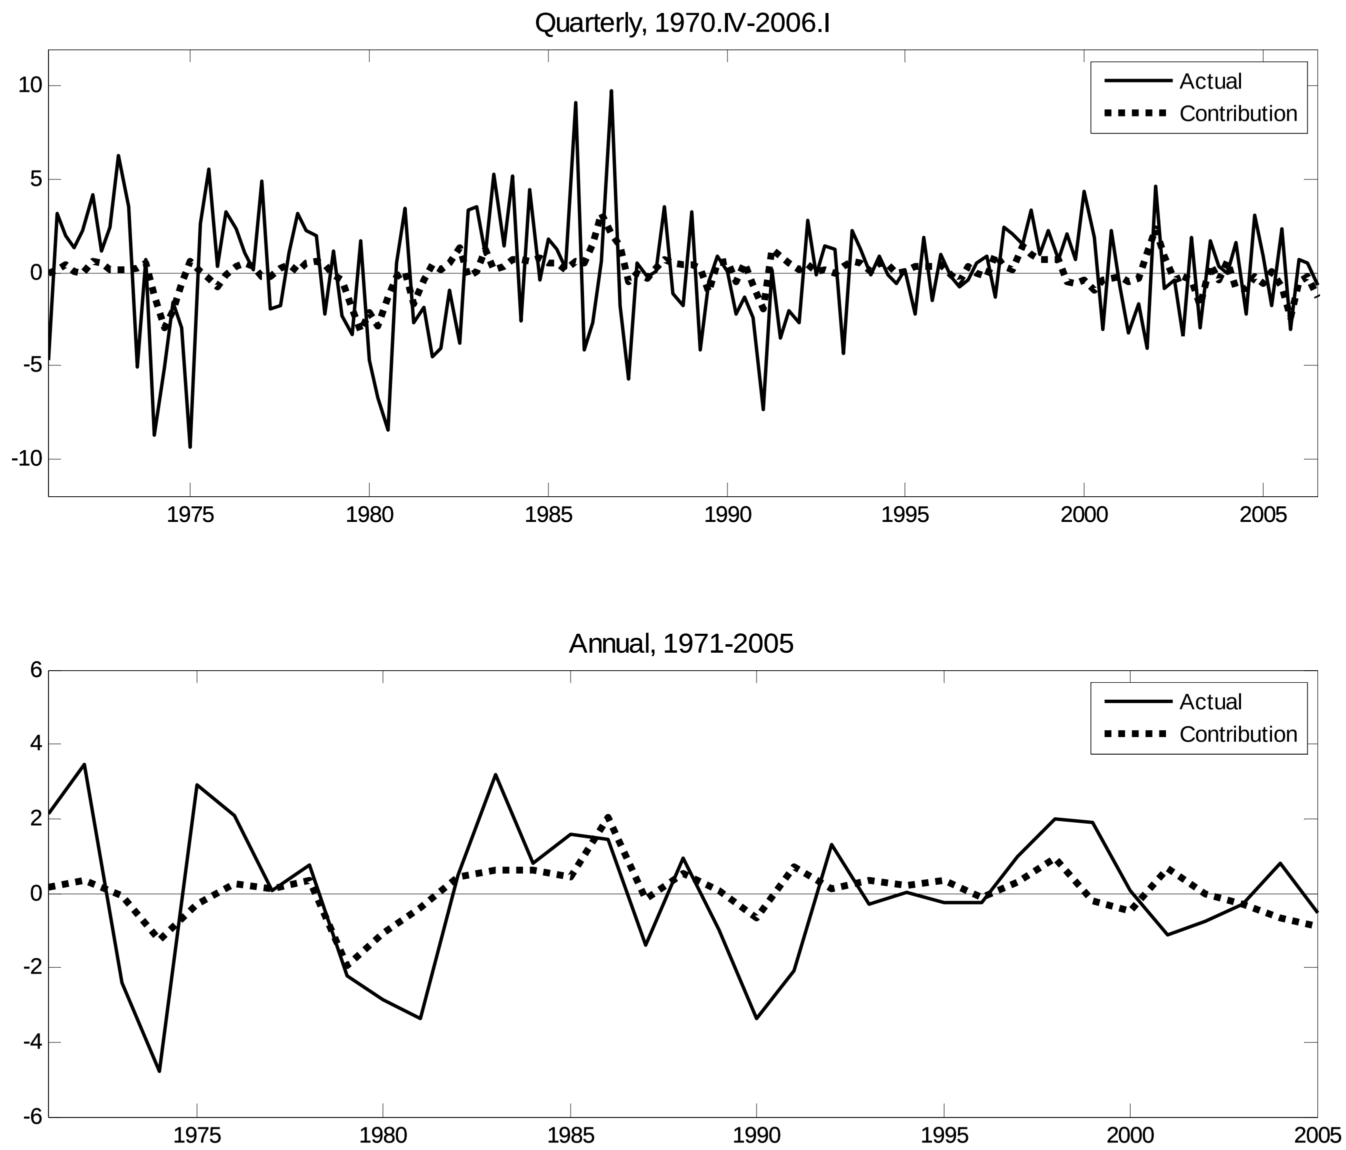

Figure 3.15: Selected Responses by Sample Period

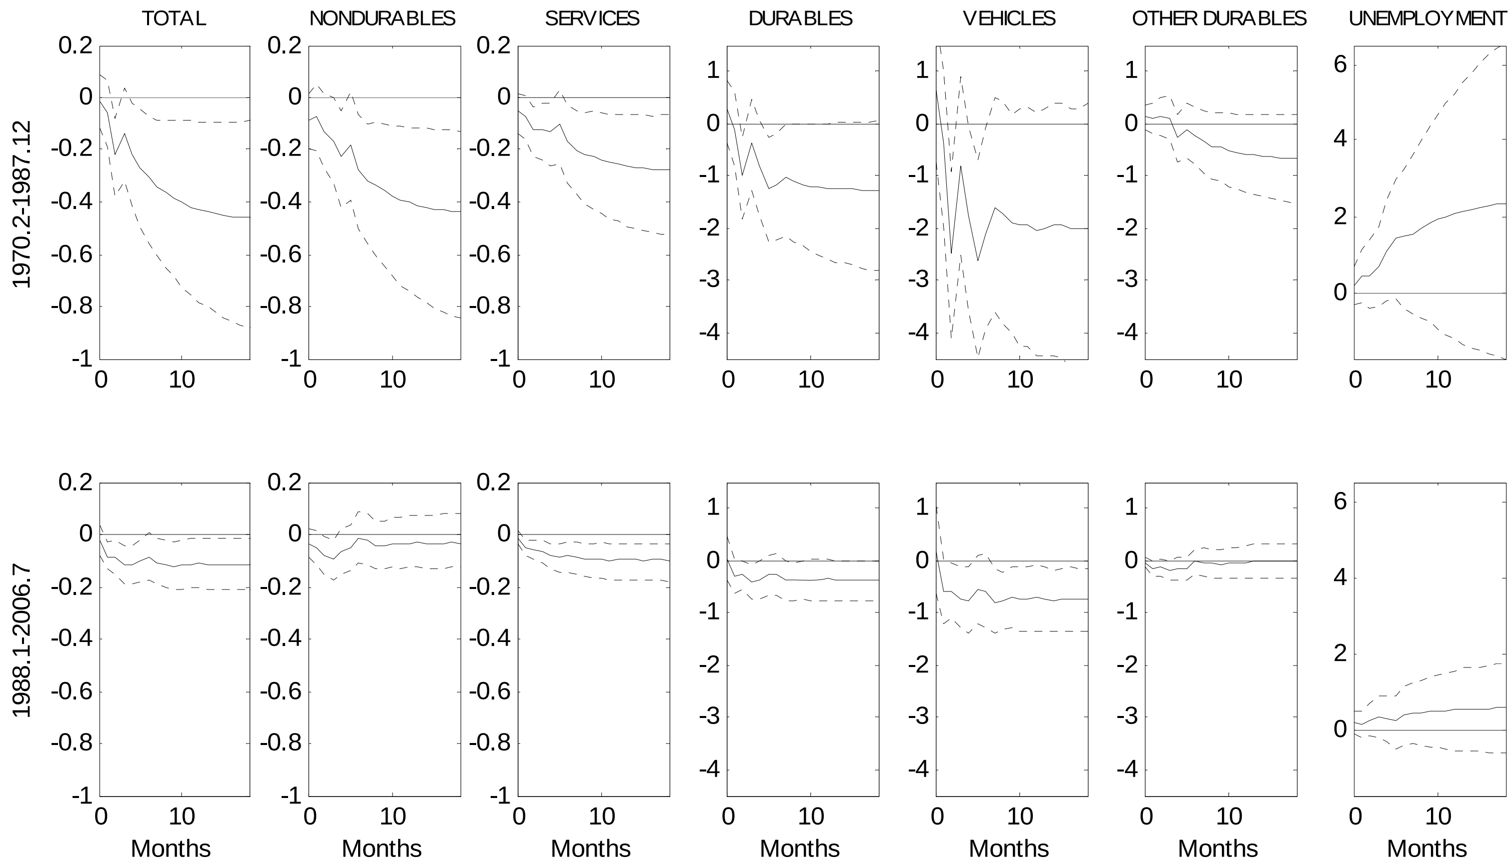

Figure 3.16: Response of New Auto Consumption to a Purchasing Power Shock

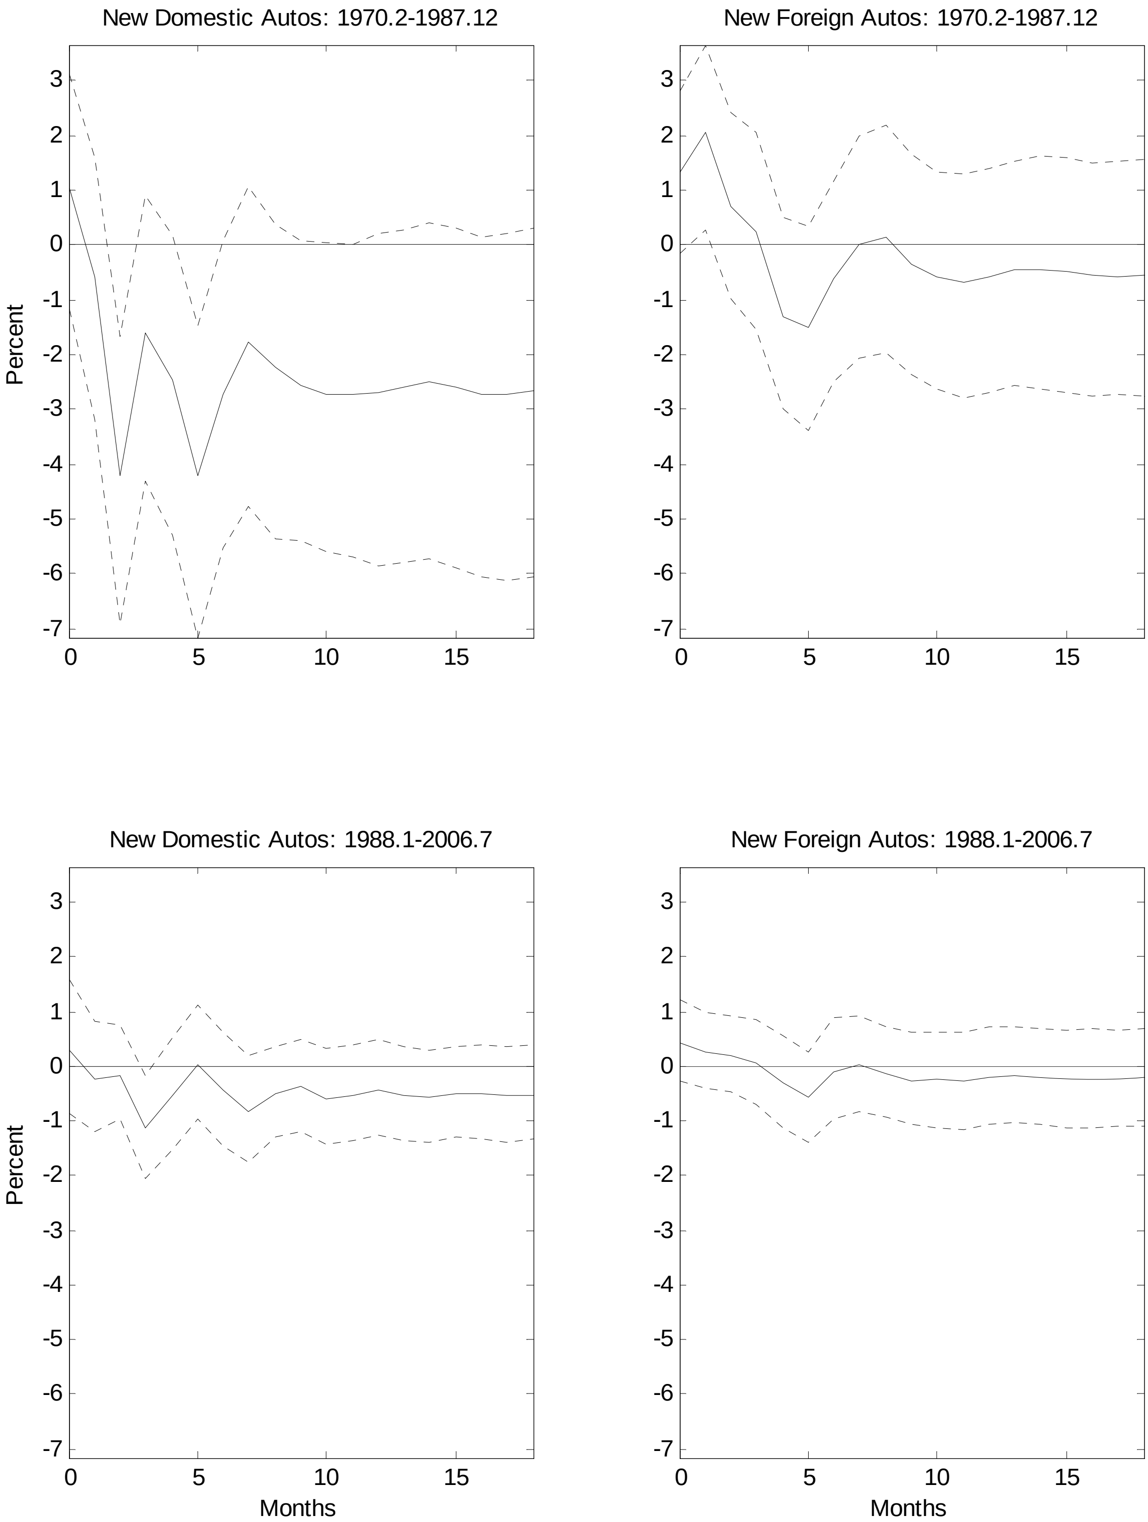

BIBLIOGRAPHY

[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] Behravesh, 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] Greenwood, T. (2006), "Tourism Likely Down This Year in Michigan, Say Analysts," *The Detroit News* , October 30.

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

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

[27] Hamilton, J.D. (2005), "Oil and the Macroeconomy," forthcoming: *New Palgrave Dictionary of Economics.*

[28] Herrera, A.-M. (2006), "Inventories, Oil Shocks and Macroeconomic Behavior," mimeo, Department of Economics, Michigan State University.

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

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

[31] 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.

[32] House, C., and M.D. Shapiro (2006), "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," mimeo, Department of Economics, University of Michigan.

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

[34] Ip, G. (2006), "Falling Oil Prices May Spell Relief for Consumers," *The Wall Street Journal*, September 13.

[35] 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), Article 2.

[36] Jorgenson, D.W., and K.-Y. Yun (1991). *Tax Reform and the Cost of Capital.* Oxford University Press.

[37] Kilian, L. (1998), "Small-Sample Confidence Intervals for Impulse Response Functions," *Review of Economics and Statistics*, 80(2), 218-230.

[38] Kilian, L. (2001), "Impulse Response Analysis in Vector Autoregressions with Unknown Lag Order," *Journal of Forecasting*, 20(3), 161-179.

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

[40] 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.

[41] 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.

[42] Mehra, Y.P., and J.D. Peterson (2005), "Oil Prices and Consumer Spending," *Federal Reserve Bank of Richmond Economic Quarterly*, 91, 53-72.

[43] 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.

[44] Peters, J.W. (2006), "Growth Slowing Across the Country, Fed Regional Reports Say," *The New York Times*, July 27.

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

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

[47] "Retail Sales, Expected to Rise, Instead Fell in June," *The New York Times*, July 15, 2006.

[48] U.S. Bureau of Labor Statistics (2005), "Consumer Expenditures in 2003," Report 986. Available at: http://www.bls.gov/cex/csxann03.

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

CHAPTER IV

The Response of Business Fixed Investment to Changes in Energy Prices: A Test of Some Hypotheses About the Transmission of Energy Price Shocks

4.1 Introduction

It is widely accepted that in the absence of a major disruption in consumers' and firms' expenditures, the effects of energy price shocks on the U.S. economy are bound to be small (see Hamilton 2005). In related work, Edelstein and Kilian (2007) have investigated the response of consumers' expenditures to changes in energy prices, including the response of residential fixed investment. Little is known, however, about the response of firms' investment expenditures with the notable exception of Herrera's (2006) work on inventory adjustments in response to energy price shocks. This paper will investigate in detail how nonresidential fixed investment in structures and in equipment responds to changes in energy prices.

There are two main channels by which energy price shocks may affect nonresidential fixed investment. One channel is the increase in the marginal cost of production associated with an increase in energy prices. This cost channel depends on the cost share of energy. A second channel operates through reduced demand for the firm's output, as consumer expenditures fall in response to rising energy prices (see Edelstein and Kilian 2007).

The response of nonresidential fixed investment need not be symmetric in energy price changes. For example, changes in energy prices are thought to create uncertainty about future energy prices, causing firms to postpone irreversible investment decisions (see Bernanke 1983 and Pindyck 1991). This uncertainty effect has implications for both supply-side and demand-side accounts of the transmission of energy price shocks. Specifically, firms may respond to uncertainty about future production costs or to uncertainty about future sales and revenue. In either case, when energy prices rise, the uncertainty effect will reinforce the decline in firms' investment expenditures due to reduced consumer demand and higher energy costs. When energy prices fall, in contrast, the uncertainty effect counteracts the increase in investment expenditures driven by lower costs and increased consumer demand, dampening the increase in investment spending.

Models of asymmetric transmission mechanisms such as the uncertainty effect have been widely used in the empirical literature on oil prices to explain the apparent breakdown of the linear relationship between real GDP growth and oil prices in the mid-1980s (see, e.g., Mork 1989, Hooker 1996). This breakdown became apparent after the collapse of OPEC in late 1985, when crude oil prices plummeted, yet the U.S. economy failed to boom. As Edelstein and Kilian (2007) show, the weak economic growth of 1986 was not caused by weak growth in household consumption or residential fixed investment. Rather it can be traced to an unprecedented drop in the growth rate of firms' fixed investment. Thus, it is natural to investigate the possibility that nonresidential fixed investment may respond asymmetrically to energy price shocks.

In this paper, we allow for a variety of different measures of energy price shocks including percent changes in energy prices, large percent changes and net percent changes. We perform formal statistical tests for the presence of asymmetry in the response of nonresidential fixed investment to energy price shocks of different sign, but the same magnitude. Symmetry in this context means that the sum of the impulse response function to energy price increases and of the impulse response function to energy price decreases is jointly equal to zero at all horizons. The existing literature has not assessed the evidence for asymmetries by formal statistical tests of the symmetry assumption. One common approach is to include oil price increases and decreases as separate variables in a single-equation model for output growth, and to perform a Wald test for the equality of the coefficients on the lags of these variables (see, e.g., Mork 1989, Dotsey and Reid 1992, Hooker 1996, and Hooker 2002). A drawback of this approach is that this test only alerts us to differences in the slope coefficients, whereas we are really interested in whether the impulse responses to positive and negative energy price shocks are different. Another common approach has been simply to inspect the point estimates of the impulse response functions without formal testing (see, e.g., Davis and Haltiwanger 2001). That comparison, however, tells us nothing about the statistical significance of the difference. Nor do tests for pointwise statistical significance of the differences in the impulse response function constitute a formal test of the symmetry assumption.

The question of whether nonresidential fixed investment responds symmetrically to energy prices increases and decreases has not been addressed in the literature. Our study is the first to provide a detailed examination of the components of nonresidential fixed investment in structures and equipment. Using the BEA's National Income and Product Account data, we study the responses of nonresidential investment in structures (including commercial and health care, manufacturing, power and communication, mining, and other structures) and in equipment (including information processing, industrial,transportation, mining and oil field machinery, and all other equipment). Special attention is given to mining and oil field machinery and to mining structures devoted to petroleum, coal and natural gas exploration and extraction.

We show that the apparent asymmetry in the estimated responses of business fixed investment in equipment and structures cannot be reconciled with standard theoretical explanations of asymmetric responses. Rather this evidence is an artifact (1) of the aggregation of mining-related expenditures by the petroleum industry and all other expenditures, and (2) of ignoring an exogenous shift in investment mainly caused by the 1986 Tax Reform Act. After controlling for these factors, formal statistical tests are unable to reject the hypothesis of symmetric responses to energy price shocks for all components of investment in structures. For equipment there is weak statistical evidence of asymmetries in some components, but not in total investment in equipment. Thus, there is no empirical support for theoretical models of the uncertainty effect on business fixed investment expenditures such as Bernanke (1983) or Pindyck (1991) nor is there support for the reallocation effect emphasized in Hamilton (1988).

Once symmetry is imposed and petroleum-industry related expenditures are excluded, the estimated response of business fixed investment in equipment and structures is small and statistically insignificant. Our conclusions are in most cases robust to defining energy price shocks in terms of percent changes, large percent changes or net percent changes in energy prices; they are also robust to using alternative measures of energy prices and to weighting energy prices by the energy share in value added. Historical decompositions show that for most forms of investment in structures and equipment, the contribution of energy price shocks historically has been negligible. For mining and oil field machinery the cumulative contribution of energy price shocks has been somewhat larger. An even larger cumulative effect is obtained for mining structures (which consist mainly of structures devoted to the exploration and extraction of petroleum, natural gas and coal).

The remainder of the paper is organized as follows. In section 2, we propose a measure of intermediate energy prices that forms the basis of our analysis. We further document the evolution of the share of energy inputs in U.S. value added, and we summarize the results of various sensitivity analyses we conducted. In section 3 we investigate the evidence for asymmetric responses of business fixed investment to energy price shocks. Building on the results of sections 3, in section 4 we impose symmetry in estimating the response of nonresidential fixed investment in structures and in equipment. That same model is also used to quantify the cumulative contribution of energy price shocks to fluctuations in real nonresidential fixed investment. We conclude in section 5.

4.2 Specification of the Energy Price Series

Since there is no evidence for an asymmetric effect of energy price changes on consumption, as shown in Edelstein and Kilian (2007), if business investment does respond asymmetrically to changes in energy prices, then it must be due to uncertainty about future energy costs. This point has implications for the specification of energy price shocks in the context of this paper. There is reason to believe that crude oil prices do not reflect the energy costs faced by firms. Few firms use crude oil as an input to production. Rather firms tend to purchase refined petroleum products such as gasoline or heating oil, or they rely on electricity and natural gas. Hence, we focus on a broad measure of energy prices in the form of the BLS producer price of processed fuels and lubricants. The choice of this index is not innocuous. Indeed, it is plausible that a given change in energy prices may affect firms differently from households. Whereas the latter's purchasing power depends primarily on the price of motor fuels such as gasoline, as documented in Edelstein and Kilian (2007), firms depend to a much greater extent on electric power and natural gas. Our approach differs from much of the earlier literature which has focused on shocks to crude oil prices (see, e.g., Davis and Haltiwanger 2001; Hamilton 2003; Lee and Ni 2002). By focusing on crude oil prices, these studies potentially omit important price shocks caused in the process of refining crude oil. For example, following Hurricanes Rita and Katrina in late 2005 crude oil prices declined, whereas gasoline prices skyrocketed, as refineries in the Gulf area were shut down. They also may dramatically overstate the magnitude of the energy price shocks faced by firms, as illustrated below.

4.2.1 The Baseline Energy Price Series

As the energy input price, we use the Bureau of Labor Statistic's Processed Fuels and Lubricants producer price index (PPI) which is available on a monthly basis for the 1973.1 - 2006.12 period. This price index includes the prices of all energy goods purchased by firms as inputs in production. Table 1 lists the components of this price series and each component's weight in the construction of the index in 2002. The largest component is electric power (40.3%), followed by natural gas (14.5%) and unleaded gasoline (14%). Given the lack of suitable data prior to January 1973, we extrapolate this price series back to January 1970 using the Fuels and Related Products and Power PPI published by the BLS. The latter PPI includes crude energy and residential energy prices in addition to intermediate energy inputs prices. Given the high correlation of the two series after 1973, the error is likely to be small. We obtain a measure of real energy prices by deflating the resulting index by the PPI for all commodities.

The upper panel of Figure 1 plots the real price of intermediate energy inputs for 1970.1 - 2006.12. It shows that the real price of energy in late 2005 briefly exceeded its 1981 peak and at the end of the sample is similar to its level in 1985, before the collapse of OPEC. Thus, the real energy prices faced by firms are considerably higher by historical standards than the corresponding real price of crude oil (see Figure 1 in Kilian 2007). Moreover, the magnitude of real energy price shocks faced by firms is much smaller in general than the corresponding shocks to crude oil prices, owing to the large share of electric power available at stable prices. For example, in 1974, crude oil prices rose twice as much as intermediate energy prices. Even more strikingly, in 1990, crude oil prices rose by 83%, whereas intermediate energy prices only rose by 12%. This evidence confirms the importance of distinguishing between crude oil prices and the energy prices faced by firms.

The panel below shows the quarterly growth rate in real energy prices faced by firms for 1970.II - 2006.IV. The latter series will form the basis of the empirical analysis in the sections below. As expected, the series shows marked spikes in 1974, 1979/1980, 1981, 1986, 1990/91, and 2003, during times of major disturbances in crude oil markets, and in 2005, when U.S. refining capacity was sharply reduced in the wake of Hurricanes Rita and Katrina. In the sections below, we will treat innovations to the (suitably defined) percent change in real intermediate energy prices as predetermined with respect to nonresidential fixed investment. This assumption will allow us to estimate the response of investment expenditures to energy price shocks, while controlling for reverse causality.

4.2.2 To Weight or Not to Weight

It is a commonly held view that the share of energy in value added has declined over time, making changes in energy prices less important for the U.S. economy (see, e.g., Bernanke 2006). Such changes in the economy-wide energy share are important for empirical work since they may invalidate standard regression analysis based on unweighted energy price innovations. While estimates of this share have been constructed, for example, by Rotemberg and Woodford (1996), the evolution of this share over time does not appear to have been documented. Following Rotemberg and Woodford (1996), the share of energy inputs in GDP is taken to be the sum of nominal value added in oil and gas extraction and imports of petroleum and petroleum products, divided by nominal GDP. This definition is natural in our context since we are interested in economy-wide expenditures on fixed investment and are unable to break down the aggregates by industry.

Notwithstanding the common view that the energy share has declined since the first oil shock of 1973, there do not appear to be data available on this energy share prior to 1977.[1] The annual energy share in value added for 1977-2005 is plotted in the lower panel of Figure 1. The average share over our sample period is 2.3%. Starting at 3.3% in 1977, the energy share rises through the late 1970s, briefly reaching a maximum of 5% in 1981. The share then falls throughout the 1980s and 1990s, reaching a minimum of 1% in 1998 following the Asian crisis. Since this time, it has risen back to its 1977 level of 3.3% in 2005. Thus, the common view that the share of energy in GDP has fallen since the 1970s must be partially revised.

An important question is whether these changes in the energy share affect the main conclusions of this paper. In sensitivity analysis (not reported in this paper) we verified that

[1] Nominal value added in oil and gas extraction is reported by the BEA on an annual basis from 1977 through 2005. Nominal imports of petroleum and petroleum products are also available from the BEA for 1978-2005. We extrapolate the latter series back to 1977 using the percent change in the dollar value of imported barrels of crude oil (defined as the number of barrels of crude oil imported into the U.S. times the refiner acquisition cost for imported crude oil per barrel of crude oil). Since year-over-year changes in nominal imports of petroleum and petroleum products and in nominal imports of crude oil are highly correlated over the 1977-2005 period, this is likely to be a good approximation.

this is not the case for the 1977-2005 period. In fact, weighting the change in intermediate energy prices by the energy share makes little difference.[2] As the weighted and unweighted quarterly energy price changes have a correlation of 99%, the estimated responses of different types of structures and equipment investment to weighted price shocks (properly scaled) are virtually identical to the responses to unweighted price shocks. We therefore abstract from the weighting issue altogether, which allows us to extend the sample back to 1970.II. This increase in sample size not only improves the accuracy and power of the statistical tests at the center of our analysis, but it also allows us to assess the role of energy during the two oil price shocks of the 1970s, which would not be possible using weighted energy price changes.

4.2.3 Other Energy Price Measures

While there are good reasons to treat intermediate energy prices as the relevant measure of energy prices for firms, as discussed above, other choices are possible and different measures of energy prices might produce different results. Firms may adjust their investment spending in response to changes in energy costs as well as to changes in revenues. If firms react primarily to production costs, then a measure of intermediate energy prices is most appropriate. On the other hand, if firms respond primarily to changes in consumer demand for their goods and services, then a measure of consumer energy prices is more appropriate. While all empirical results reported in the sections below are based on real energy input prices, we verified that our main conclusions are robust to alternative measures of energy prices. Notably, we experimented with the BEA's PCE price index for energy goods and services, which includes gasoline (and other motor fuels), natural gas, electricity, and all other energy goods (including heating oil, coal, and oil lubricants). This price index was deflated by the overall PCE price index. The primary difference between the real consumer energy price and the real intermediate energy inputs price is the weight devoted to different components of each index. The correlation between the two indices is quite

[2]Since the data used in computing the energy share are only available at annual frequency, we use linear interpolation to generate quarterly weights for 1977.IV-2005.IV.

high ($\rho = 0.97$). In sensitivity analysis (not reported in this paper), we found that our qualitative conclusions are unaffected whether we use the baseline measure of intermediate energy price changes, unweighted real consumer energy price changes or consumer price changes weighted by the energy share in consumption (see Edelstein and Kilian 2007).

4.3 Does Nonresidential Fixed Investment Respond Asymmetrically to Energy Price Increases and Decreases?

In this section, we study impulse response estimates based on a series of bivariate VAR models. Each VAR model includes a suitably defined transformation of the quarterly change in energy prices defined in section 2. It also includes the quarterly percent change of the component of real nonresidential fixed investment that is of interest, as reported by the BEA. The sample period is 1970.II - 2006.IV. The VAR models are identified recursively with energy prices ordered first, indicating that energy price innovations are not affected contemporaneously by innovations to real nonresidential fixed investment.[3] The advantage of using a VAR model is that it isolates the linearly unpredictable component of energy price changes and allows for reverse causality.[4]

Throughout this paper, we impose a VAR lag order of 4 quarters. This lag order tends to be larger than the lag order selected by the Akaike Information Criterion conditional on an upper bound of 12 quarterly lags, which in most cases selected an implausibly low lag order. We present point estimates of the cumulative response of real investment to a one-time, one standard deviation energy price shock. The maximum horizon of the impulse response function is six quarters. The 95% bootstrap confidence intervals are based on the bias-corrected method of Kilian (1998).

We investigate the evidence for asymmetries as follows: Our baseline approach is to split percent changes in the real price of energy into positive and negative changes. We then run

[3]Under the assumption that 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.

[4]Our 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 2007). Identifying these shocks is not straightforward when dealing with refined products. 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.

one bivariate model in energy price increases and the percent change in investment, and a second model in energy price decreases and the percent change in investment. Then we compare the response of investment while controlling for the magnitude of the shock. This allows us to characterize the nature of the asymmetry, if any, as described further below. In addition, we conduct a formal Wald test for symmetry in the response functions. This test is 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 preserves the contemporaneous correlations in the data across models.

The baseline approach assumes that firms respond proportionately to the magnitude of the shock. Alternatively, it is possible that firms only respond to large shocks. We allow for that possibility by alternatively defining energy price shocks as increases or decreases that exceed 4.04%, which corresponds to one standard deviation of the percent change in energy prices. A third possibility is that firms respond only to changes in energy prices that are unprecedented in recent history (see Hamilton 1996, 2003). We address this possibility by using a measure of the net energy price increase, defined as the difference between the current price and the maximum price over the previous year if the current price is larger than the previous maximum, and zero otherwise. Edelstein and Kilian (2007) extend this measure to include net energy price declines that are constructed in a similar fashion.[5] All three measures of energy price shocks are shown in Figure 2. For each specification, increases in prices are represented by positive values and decreases by negative values, allowing both series to be shown in the same plot. Since increases and decreases are mutually exclusive, the respective series for increases and for decreases are shown in the same plot.

A natural starting point for the analysis is the full-sample impulse response estimates for 1970.II-2006.IV. Full-sample estimates are potentially misleading, however. As discussed in Edelstein and Kilian (2007), a complicating factor in testing for asymmetric responses

[5]Our results are qualitatively robust to considering the net change relative to the preceding three years.

is the existence of exogenous shifts of nonresidential fixed investment spending during this sample period. The most prominent of these shifts is associated with the 1986 Tax Reform Act, which had a negative effect on investment spending at a time of rapidly falling energy prices precipitated by the collapse of the OPEC cartel in late 1985. Real nonresidential investment in equipment fell by 4.65% that year, while real investment in structures fell by 16.35% (relative to their respective means), as the investment tax credit was repealed in the first quarter of 1986 and real estate tax shelters were eliminated. Failing to account for this exogenous shift at a time of sharply declining energy prices will tend to bias the test results in favor of asymmetry.

We attempt to account for the effects of the 1986 Tax Reform Act in two alternative ways. The first approach is to split the sample and to focus on the 1988.I-2006.IV period. The break date of 1987.IV was chosen to exclude all short-run effects of the 1986 Tax Reform Act from the second part of the sample.[6] A drawback to this approach is that splitting the sample in half may undermine the power of the symmetry tests. On the other hand, this approach does not constrain any of the model parameters and, by eliminating all confounding influences that are associated with the 1970s and early 1980s, may provide the most accurate representation of the effects of energy price shocks on investment at the end of the sample. The second approach involves regressing the growth rate of real investment on a constant and a dummy for each of the four quarters of 1986 and using the residuals from this regression as the second variable in the recursively identified VAR model. This approach implicitly assumes that the 1986 Tax Reform Act only affected the mean growth rate of real investment and that the effects of the 1986 Tax Reform Act on investment were concentrated in 1986. This approach allows us to make use of the full sample, increasing the power of statistical tests of symmetry, but it abstracts from other forms of structural change.

[6]The results are similar if the sample is split at the 1987.I quarter.

4.3.1 Testing for Asymmetries in the Response of Nonresidential Fixed Investment

Tables 2a-2e summarize the pattern of impulse response estimates for various components of nonresidential fixed investment for the 1970.II-2006.IV period, the 1988.I-2006.IV period, the 1970.II-1987.IV period, the 1970.II-1987.IV period (including 1986 dummies), and the 1970.II-2006.IV period (including 1986 dummies), respectively. We show separate results for three specifications of energy price shocks: percent changes in energy prices (C), large percent changes in energy prices (LC) and net percent changes in energy prices relative to the preceding year (NC), as defined earlier. To conserve space the plots of the underlying impulse response estimates are collected in the Appendix.[7] Rather than present each of the 168 impulse response functions and the corresponding 168 Wald tests of the symmetry hypothesis, we summarize our findings succinctly in tables. We begin by classifying the response patterns. For a typical investment expenditure item such as manufacturing structures, one would expect one of the following response patterns:

No response: The responses of investment to positive and negative energy price shocks are not statistically significant at the 95% confidence level.

Symmetry: The negative response to positive energy price shocks and the positive response to negative price shocks are approximately equal in magnitude and at least one of the responses is statistically significant.

Classical asymmetry: There is a statistically significant negative response of investment to positive energy price shocks and a response of investment to negative energy price shocks that is bounded above by the absolute magnitude of the response to a positive shock.

Perverse asymmetry: Any asymmetry in the responses that departs from the classical asymmetry predicted by commonly used economic models of the transmission of energy price shocks.

[7] Each row in Figures A.1a-A.1e shows the impulse response for a specific component of structures investment to a positive and negative energy price shock under each of the three energy price shock specifications.

As it turns out, these four patterns exhaust the possibilities encountered in our impulse response function estimates with one important exception. That exception is investment in *mining structures* and investment in *mining and oil field equipment*. Since crude oil, coal and natural gas producers account for 92.2% of all capital expenditures of the mining sector, energy price increases tend to stimulate investment in mining activities.[8] Hence, the notion of symmetry (or of classical asymmetry) for these mining-related expenditures is the mirror image of how these terms are defined for all other investment expenditures:

> *No response*: The responses of investment to positive and to negative energy price shocks are not statistically significant at the 95% significance level.
>
> *Symmetry*: The positive response to positive energy price shocks and the negative response to negative energy price shocks are approximately equal in magnitude and at least one of the responses is statistically significant.
>
> *Classical asymmetry*: There is a statistically significant negative response of investment to negative energy price shocks and a response of investment to positive energy price shocks that is bounded above by the absolute magnitude of the response to a negative shock.
>
> *Perverse asymmetry*: Any asymmetry in the responses that departs from the classical asymmetry predicted by commonly used economic models of the transmission of energy price shocks.

We use X's in Tables 2a-2e in the appropriate columns to identify which responses follow which pattern. In addition, we use *, **, and ***, to indicate rejections of the symmetry hypothesis at the 10%, 5% and 1% level, respectively. It is possible for a response to be asymmetric in the sense that the Wald test rejects symmetry, yet be classified as *No response*, if the impulse response coefficients themselves are not statistically different from zero.

[8] See U.S. Census Bureau 2005, Table 1. Industry Statistics: 2002. We added capital expenditures for coal mining (5%) and for oil and gas extraction (78%), and the respective support activities (10%) including drilling.

Under the null hypothesis of symmetry the response of investment to a positive energy price shock at each horizon is exactly the mirror image of the response to a negative energy price shock of the same magnitude. This hypothesis may be expressed as a Wald test statistic with an asymptotic χ^2_7-distribution, given our choice of impulse response horizon. We estimate the variance of the estimator 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. The bootstrap algorithm is designed to preserve the contemporaneous correlations in the data across regression models.[9]

It is useful to analyze separately the responses of nonresidential investment in structures and in equipment, which account for 32% and 68%, respectively, of total nonresidential fixed investment. We begin with nonresidential structures.

4.3.1.1 Asymmetries in Nonresidential Investment in Structures

A natural conjecture is that total nonresidential fixed investment in structures should decline, as energy prices rise, and increase, as energy prices fall. The results for structures are shown in the upper panel of Tables 2a-2e. We further differentiate between *mining structures*, *commercial and health care structures*, *manufacturing structures*, *power and communication structures* and *other structures*.

Full-Sample Analysis for 1970.II-2006.IV: The full-sample results in Table 2a show a perverse asymmetry for total nonresidential investment in structures reflecting the fact that these expenditures do not fall significantly in response to positive energy price shocks, but fall significantly in response to negative energy price shocks. This seemingly paradoxical result holds regardless of the specification of the energy price shocks. While not statistically

[9] Our approach to testing symmetry is more stringent than previous approaches. One common approach in the literature has been to fit a single-equation regression model that includes lags of both increases and decreases in oil prices and to test the equality of the coefficients at a given lag by means of a Wald test (see, e.g., Mork 1989, Dotsey and Reid 1992, Hooker 1996a, Hooker 2002). 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 energy prices. Rather than simply inspecting the point estimates as in Davis and Haltiwanger (2001) or testing for symmetry at selected horizons (ignoring the dependence of response coefficients across horizons), we test for the symmetry of the entire response function.

significant at conventional significance levels, this evidence of asymmetry is clearly at odds with conventional explanations of asymmetries in the response of nonresidential investment expenditures.

It is useful to decompose structures investment further. The second row of Table 2a suggests that mining structures are characterized by highly significant classical asymmetries of the type anticipated based on conventional economic theories. Thus, one possible explanation of the perverse asymmetry in total structures investment is a simple composition effect. Total structures is the sum of investment in industries that use energy as an input (and that depend on consumer demand) and of investment by mining industries that produce energy in the form of crude oil, coal or natural gas. Thus, if the response of mining expenditures exhibits a classical asymmetry, as shown in Table 2a, and if nonresidential structures excluding mining do not, then aggregating these two components will by construction, produce a perverse asymmetry.

This point is illustrated in Figure 3 using a hypothetical example. For expository purposes, we postulated a small symmetric response of nonresidential structures excluding mining (upper panel). The panel below shows a classical asymmetry in the response of expenditures on mining structures. Such an asymmetry could arise for any number of reasons, including an uncertainty effect that dampens the positive response to higher energy prices and strengthens the negative response to falling energy prices. An alternative explanation could be that there were few domestic investment opportunities in response to rising energy prices in the United States during our sample period, making it difficult to invest in energy production domestically, while the option of cutting investment in mining in response to falling energy prices always has been available. For our purposes, the source of this classical asymmetry is immaterial. Assuming equal weights for expositional purposes and summing the symmetric response in the first panel and the hypothesized asymmetric response of mining in the second panel, produces a response for total nonresidential structures that looks perverse from the point of view of conventional economic theories of the transmission of energy price shocks in that total investment rises somewhat in response to positive energy price shocks and falls strongly in response to negative energy price shocks.

This composition effect is quantitatively important in the data despite the relatively small share of mining structures in total structures, which ranges from under 10% to 30%, depending on the sample period. It can be shown that under the C specification, for example, the p-value of the symmetry test after excluding mining rises sharply from 0.20 to 0.48, indicating weakening evidence of perverse asymmetries. Similar results hold for the other two specifications of energy price shocks. Nevertheless, Table 2a suggests that the composition effect is not the only source of the perverse asymmetry in the responses of total nonresidential structures. In fact, total nonresidential structures excluding mining also exhibit a perverse asymmetry.

What then is the source of the perverse asymmetry in total structures excluding mining? Table 2a shows that all but one component of structures investment do not respond significantly to energy price shocks. Only for *other structures* is there some statistically insignificant evidence of a classical asymmetry.[10] The latter asymmetry is not the cause of the perverse asymmetry in total structures excluding mining, however, which is unmatched by any of the disaggregates and indeed is not statistically significant. We conclude that in no case is there evidence of statistically significant asymmetries outside the mining industry.

Split-sample Analysis for 1970.II-1987.IV and 1988.I-2006.IV: There has been much discussion of the proposition that the effect of energy price shocks has declined since the 1970s and early 1980s. It is therefore instructive to split our sample in half. We choose to break the sample in 1987.IV, for reasons discussed below. Table 2b focuses on the second half of the sample: 1988.I-2006.IV. Compared to the full sample, we see that none of the asymmetries remain statistically significant, and that the qualitative patterns of the re-

[10] *Other structures* include religious structures, educational and vocational structures, lodging, amusement and recreation structures, farm structures, transportation structures, and other infrastructure.

sponses change. Except under specification C, there is no evidence at all of an asymmetry in the response of *total structures*. The perverse asymmetry encountered under the first specification can be traced to the weak evidence of classical asymmetries in the response of *mining*. Excluding *mining*, there is no longer any evidence of perverse asymmetries in nonresidential structures investment. Under the C specification, for example, the p-value of the symmetry test rises slightly from 0.85 to 0.88 after excluding *mining*, consistent with the weak evidence of classical asymmetries in *mining*. None of the components of structures exhibits statistically significant asymmetry in the second half of the sample. Thus, for 1988.I-2006.II, there seems to be little support for the classical asymmetries discussed in the literature.

In contrast, the results for the first half of the sample in Table 2c are far less clear-cut. Table 2c shows a clear pattern of classical asymmetries in mining structures. The only difference from the full sample is the lack of statistical significance. Using the percent change specification of energy price shocks (C), there is some evidence of a perverse asymmetry in total nonresidential structures (whether *mining* is excluded or not), but there is no such evidence for the other two specifications. Excluding *mining* sharply raises the p-value of the symmetry test from 0.51 to 0.84, again under the C specification. There also is systematic, if statistically insignificant, evidence of classical asymmetries in *other structures* and under the LC specification for *total structures*.

1970.II-1987.IV and 1988.I-2006.IV Revisited: Why are the results for the two subsamples so different? As discussed in Edelstein and Kilian (2007), there is reason to believe that the results for the first half of the sample are distorted. The distortion arises because we do not control for exogenous shifts in investment that occurred at the same time as the sharp decline in energy prices in the mid-1980s. Perhaps the most important candidate for such an exogenous shift in nonresidential investment is the 1986 Tax Reform Act. Investment in commercial structures including office buildings often served as a tax shelter prior to 1986. By removing the real estate tax shelter, the 1986 Tax Reform Act caused a sharp exogenous drop in investment in commercial and office space in 1986 that may easily be mistaken for an asymmetry (see *Survey of Current Business* 1987, p.4). The fact that the only evidence for perverse asymmetries in Table 2c arises for commercial and health care structures is consistent with this view.[11]

In addition, Edelstein and Kilian (2007) have suggested that the seemingly asymmetric response of investment in the mining sector in 1986 may have been driven by the mining industry's reaction to the collapse of OPEC in late 1985. That response far transcended the response that one would have expected based on the decline in energy prices in 1986 alone, suggesting that the sudden collapse of OPEC was viewed as an exogenous shock by the mining industry over and above the decline in energy prices.

In Table 2d we control for both of these effects by specifying four intercept dummies in the VAR model, one for each quarter of 1986. Table 2d shows that this modification eliminates all evidence of perverse asymmetries in structures for the 1970.II-1987.IV period. With the exception of *other structures* which appear to exhibit a classical asymmetry, none of the components of structures nor *total structures* show a statistically significant response to energy price shocks. These results are the same across all three specifications.

One drawback of the evidence in Tables 2b and 2d is that the reduction in sample size associated with sample-splitting may render some results statistically insignificant merely because of lower power. There is also a greater chance of obtaining an apparent asymmetric pattern merely by chance. Table 2e therefore revisits the full-sample analysis with 1986 dummies.

The results are stronger than for the second half of the sample in that there is no evidence of asymmetries in *total structures*. The results in Tables 2b, 2d, and 2e also agree in that there is no significant asymmetry in total nonresidential *structures excluding mining*. Unlike in Table 2b, there is no evidence of asymmetry in the response of *mining structures*

[11] Data limitations prevent us from disentangling commerical and office space from health care structures.

in Table 2e. Unlike in Table 2d there is evidence of a symmetric response of *mining structures*. The classical asymmetry in *other structures* in Table 2e is not statistically significant. No evidence at all remains of the perverse asymmetries that plagued the earlier analysis, which raises our confidence in the results. This result illustrates the importance of controlling for exogenous shifts in nonresidential investment in 1986.

We conclude that nonresidential investment in structures does not respond asymmetrically to energy price shocks. That result is robust to alternative specifications of energy price shocks. There is no statistically significant evidence of the asymmetric patterns suggested by various economic models of the transmission of energy price shocks (or of any other asymmetric patterns).

4.3.1.2 Asymmetries in Nonresidential Investment in Equipment

Having found no compelling statistical evidence of asymmetries in structures investment, we now turn to the evidence for nonresidential investment in equipment. The results for equipment are shown in the bottom panels of Tables 2a-2e. We distinguish between *mining and oil field machinery*, *equipment for information processing*, *industrial equipment*, *transportation equipment*, and all *other equipment*.

Table 2a shows that in the full sample, there is no evidence of perverse asymmetries in equipment, but a robust pattern of classical asymmetries with the exception of *information processing*. Only after excluding *mining and oil field machinery* is that pattern marginally statistically significant, and even in that case the rejection occurs only under one of three specifications. If we focus on the results for 1988.I-2006.IV only (Table 2b), much of the evidence of classical asymmetries vanishes. An exception is *mining and oil field machinery*. There is also some evidence of a perverse, but statistically insignificant asymmetry in the response of *industrial equipment*.

In contrast, the results for 1970.II-1987.IV in Table 2c are closer to those for the full sample; the only difference being some evidence of perverse asymmetries in the response of *mining and oil field machinery*. That evidence weakens considerably in Table 2d, once we include the 1986 dummies. The 1986 Tax Reform Act 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.[12] After controlling for these effects, all of the perverse responses for *mining equipment* become statistically insignificant. There is now a clear pattern of classical asymmetries in most components of equipment, but none that are statistically significant.

Moving to the full-sample results for the model with 1986 dummies in Table 2e, all of the perverse responses vanish and a very clear picture emerges. With the exception of *information processing* all forms of nonresidential investment in equipment appear to exhibit classical asymmetries. The strongest evidence against symmetry is obtained for the *LC* specification, including a 10% rejection of symmetry for *total equipment.* In contrast, under the *NC* specification, symmetry is never rejected, and under the *C* specification symmetry is rejected only for two disaggregates.

While it is difficult to compare formally the *NC* specification with the other two, we can test the support for the *LC* specification against the *C* specification, since one specification is nested in the other. This fact allows us to test formally the hypothesis that firms only respond to large percent changes in energy prices as opposed to all percent changes in energy prices. Wald tests for the equality of the impulse responses to large percent changes and to all percent changes in energy prices generate p -values of 0.63 for total equipment and of 0.67 for total equipment excluding mining and oil field machinery. Thus, there is no reason to reject the *C* specification which makes efficient use of all data on energy price changes in favor of the more limited *LC* specification. Whereas we cannot reject symmetry for the response of *total nonresidential equipment* or for *mining and oil field machinery* under the

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

C specification, we marginally reject symmetry at the 10% level for *total nonresidential equipment excluding mining and oil field machinery.* The latter rejection appears to be driven by *industrial equipment.* Thus, there is weak evidence for classical asymmetries in some forms of equipment investment under the C specification, but none under the NC specification. There is no evidence against symmetry in *total investment in equipment* for either specification, suggesting that there is not enough evidence to abandon the standard linear VAR model framework for investment.

4.3.1.3 Revisiting the Specification of the Energy Price Series

The results in Tables 2a-2e illustrate the importance of distinguishing investment in mining-related structures and equipment from other nonresidential investment. This distinction is all the more important, as the rationale for using intermediate energy prices for expenditures related to mining activities is questionable. A case can be made that investment expenditures on mining structures and on mining and oil field machinery should depend on the output price of primary energy goods more than on the input price of refined energy goods. We therefore computed an alternative set of results for mining structures and equipment based on the producer price index for crude oil (as a proxy for all primary energy goods). The results are similar to those in Table 2a in that for 1970.II-2006.IV there is statistically significant evidence of classical asymmetries in mining structures. Qualitatively similar results are obtained for the subsamples, but the statistical significance is lost in all cases. As in Table 2d, once we control for the 1986 dummies, the evidence for asymmetries in the first half of the sample weakens. For the full sample with 1986 dummies, we find impulse response patterns consistent with classical asymmetries, but none that are statistically significant.

For mining and oil field machinery, the results are more sensitive to the choice of the energy price series and more erratic under the crude oil price specification than the intermediate energy goods specification. There is weak evidence of asymmetries under the LC specification for the second part of the sample, but the C specification cannot be statistically rejected in favor of the *LC* specification. Likewise, the qualitative results for the full sample with 1986 dummies remain unchanged in that there is no evidence of statistically significant asymmetries. Thus, the choice of the energy price series does not affect our conclusions.

4.4 Quantifying the Effect of Energy Price Shocks on Real Nonresidential Fixed Investment

Given the lack of evidence for asymmetries in the responses of structures as well as equipment to energy price shocks, in this section our objective is to quantify these responses based on models that impose symmetry. Apart from the imposition of symmetry, our estimation methodology is the same as in the previous section. A positive one standard deviation shock amounts to a 3.9% quarterly increase in energy prices. We present results only for percent changes in energy prices. The estimated responses to large price changes and net price changes are qualitatively similar. We focus on the results for 1970.II-2006.IV (including 1986 dummies), 1970.II-1987.IV (including 1986 dummies) and 1988.I-2006.IV, so that we can assess the extent to which the estimates have declined over time. The 95% bootstrap confidence intervals are again based on the bias-corrected method of Kilian (1998).

4.4.1 The Response of Nonresidential Investment in Structures to Energy Price Shocks

Figure 4a focuses on the results for nonresidential investment in structures. The response of *total structures* is flat, if not positive, and never statistically significant, whether in the full sample or in the subsamples. Investment in *mining structures* increases by 5.1% after six quarters. The response is highly statistically significant. Interestingly, the response estimate for the second half of the sample is even more statistically significant, but that for the first half is statistically insignificant. Excluding *mining structures* does not alter the response of total structures investment much; they remain statistically insignificant and close to zero. The same is true for the disaggregates. With the exception of *other structures*, the disaggregates tend to have flat responses. The response of *manufacturing structures* is associated with noticeably larger sampling uncertainty than the other responses.

4.4.2 The Response of Nonresidential Investment in Equipment to Energy Price Shocks

Figure 4b shows an insignificant decline in *total equipment* investment. In the full sample, a one-standard deviation shock lowers *total equipment* investment by -1.61%. The point estimate is dominated by the data for the first half of the sample. In the second half, the estimated response is virtually flat. In contrast, *mining and oil field machinery* shows a highly statistically significant increase in the full sample. The point estimate is 7.4% after six quarters and quite similar to the split-sample results, except that the split-sample results are not statistically significant. Excluding *mining and oil field machinery* from total investment in equipment makes little difference for the response estimates or their statistical significance.

Of the remaining forms of equipment investment, only the responses of *transportation* and of *other equipment excluding mining* are statistically significant. They also show the largest responses. Investment in *transportation equipment* drops by -3.7% after six quarters. The point estimate for 1970.II-1987.IV is -5.8% (yet only partially significant), whereas the point estimate for 1988.I-2006.IV is only -1.7% (and statistically insignificant). These results indicate that the response of investment in *transportation equipment* (unlike that of *total equipment excluding mining and oil field machinery*) has become less responsive to energy price shocks, a result that mirrors similar findings for motor vehicle consumption in Edelstein and Kilian (2007). *Other equipment excluding mining* drops by -3.1% in the full sample, and by -5.8% and -0.9% in the respective subsamples.

4.4.3 Quantifying Historical Fluctuations in Nonresidential Fixed Investment due to Energy Price Shocks

It is useful to put the impulse response estimates in perspective. The magnitude of a one standard deviation shock to energy prices is 3.9%. The biggest shock in our sample occurred in 2003.I and amounted to a 15.7% increase in energy prices. A shock of this size would leave *total nonresidential investment in structures* virtually unchanged after six quarters. It would raise investment in *mining structures* by 20.5% and lower all other nonresidential investment in structures by -0.7%. If the same shock were applied to investment in equipment, we would observe a decline of -6.5% in *total equipment excluding mining and oil field machinery*, but a sharp increase in investment in *mining and oil field machinery* of 29.7%, resulting in a -6.5% decline in *total nonresidential equipment investment.* Of course, such large shocks are rare and this prediction presumes that no other shocks occurred subsequently, whereas in reality positive shocks are often followed by negative shocks, as was indeed the case after the 2003 spike (see Figure 2).

While the estimated impulse response functions provide a measure of the response of real business fixed investment to a hypothetical permanent shock to energy prices, an equally insightful question is how important energy prices have been overall in driving fluctuations in nonresidential fixed investment during the 1970.II-2006.IV period. The answer to this question may be obtained from historical decompositions of the data on the basis of the full-sample estimates of the bivariate VAR model of the previous section (including the 1986 dummies).

The uppermost panel of Figure 5 shows the actual (demeaned) growth rates for total real structures investment excluding mining structures. It also shows the growth rates predicted based on the cumulative effect of the energy price shocks alone. The difference between the two series measures the extent to which investment growth is not explained by energy price shocks. The panel below shows the corresponding results for mining structures.

It is evident that only a small part of the fluctuations in non-mining structures investment is accounted for by the cumulative effect of energy price shocks. There is some evidence of a decline in investment growth in non-mining structures in the second half of the 1970s and in the early 1980s, for example, that is driven by energy prices. There is also some evidence of increased investment in the mid- and late 1980s and the late 1990s, but all these effects are dwarfed by the large positive spikes in investment in non-mining structures in 1971-1973, 1976-1979, 1984-85, and the large negative spikes in 1975, 1980, and 1982/83. This evidence is important because it supports the view that not only the booms in investment of the 1970s and 1980s, but also the major declines in investment in 1975, 1980 and 1982-83, were not primarily driven by energy price shocks (see, e.g., Barsky and Kilian 2002, 2004).

Energy price shocks do a somewhat better job at tracking the growth in investment dedicated to mining structures, but there remain important exceptions. For example, the sharp increase in domestic investment in mining structures in 1978/79 (during the time of the Iranian Revolution) is not explained by the cumulative energy price shocks. Similarly, the drop in investment in 1985 and in 1986 (after the collapse of OPEC) far exceeds what the model predicts.

The remaining two panels deal with mining and oil field machinery and with non-mining equipment investment. The results for the latter aggregate are very similar to those for non-mining structures. Energy price shocks do a somewhat better job at explaining fluctuations in investment in mining and oil field machinery than in non-mining equipment investment. Interestingly, the decline in investment in mining and oil field machinery in 1986 is much less pronounced than for mining structures and to some extent is explained by energy prices. On the other hand, there are some major spikes in 1992/93 (and to a lesser extent in 2002/2003) that are not explained by energy price shocks.

We conclude that, at least outside the mining sector, energy prices historically have had only a negligible effect on the growth rate of real nonresidential fixed investment. Only for mining structures and to a lesser extent for mining and oil field machinery, have energy prices played a somewhat more important role.

4.5 Conclusion

This paper investigated one of the two main channels for the transmission of energy prices to the U.S. economy that have been discussed in the literature on oil and the macroeconomy (see, e.g., Hamilton 2005). There is a widespread perception that nonresidential expenditures on structures and equipment are sensitive to energy price shocks. Little is known about how sensitive these expenditures are and whether the responses are symmetric or not.

The absence of an economic expansion in 1986 following the rapid fall in energy prices as a result of the OPEC collapse in 1985, has led a number of researchers to incorporate asymmetries into models of the transmission of energy price shocks to the economy such as the uncertainty effect on irreversible investment decisions described in Bernanke (1983) and Pindyck (1991). A common view is that positive energy price shocks will lower nonresidential fixed investment, while negative energy price shocks of the same magnitude will not increase investment to the same extent. We referred to this pattern as a *classical asymmetry*. In contrast, a more conventional symmetric response would involve a negative response to positive shocks and a positive response to negative shocks of the same magnitude that sum to zero. We referred to this pattern as *symmetry*.

Despite the prominence of classical asymmetries in policy discussions and in economic modeling, it remains an unresolved empirical question whether nonresidential fixed investment responds asymmetrically to energy price shocks. Clearly, the absence of the expected asymmetric pattern in the data would have immediate implications for the credibility of economic models of the transmission of energy price shocks that appeal to classical asymmetries. Our analysis addressed this question by comparing the responses of nonresidential fixed investment to positive and to negative energy price shocks of the same magnitude under the assumption that energy price shocks may be treated as predetermined with respect to nonresidential investment at the quarterly frequency. We showed that uncritical application of the VAR methodology will result in spurious evidence of asymmetries in the responses of nonresidential fixed investment and in particular, in evidence of *perverse asymmetries*, in the sense that these asymmetries are at odds with predictions of commonly used economic models of the transmission of energy price shocks.

We traced this evidence to two problems in particular. One problem is that the aggregation of mining-related investment by domestic producers of coal, natural gas and crude oil with other forms of investment expenditures may generate perverse asymmetries (in the sense that the observed asymmetries are seemingly at odds with the predictions of economic models). We showed that this problem can be quantitatively important. The other problem is that impulse response estimates tend to be highly sensitive in small samples to an exogenous shift in nonresidential investment that occurred in 1986 during a major decline in energy prices. This shift was related in large part to the 1986 Tax Reform Act. We controlled for this shift by introducing quarterly intercept dummies for 1986 and/or by splitting the sample in 1987.IV. Excluding investment in mining-related activities and including the 1986 dummies removed much of the evidence of asymmetries in the 1970.II-1987.IV and 1988.I-2006.IV subsamples. The little evidence that remained became statistically insignificant.

In addition, we illustrated the importance of relying on long samples when estimating the effect of energy price shocks on real nonresidential investment. We showed that using the full 1970.II-2006.IV period in conjunction with the 1986 dummies helps eliminate all evidence of asymmetries that are economically implausible. Results for this final specification are free of all perverse asymmetries and in general quite robust to alternative specifications of energy price shocks.

We concluded that there is no compelling evidence of asymmetries in the responses of aggregate nonresidential investment in structures and no compelling evidence of asymmetries in the response of aggregate nonresidential equipment investment. Only in two cases did we find a marginal rejection of symmetry at the 10% level for components of equipment investment. There is no significant rejection of symmetry at all for the components of structures investment. These results are qualitatively consistent with recent findings by Edelstein and Kilian (2007) for the response of consumer durables and residential fixed investment. Our results are robust to weighting energy price changes by the share of energy in value added (the evolution of which we documented in the paper) and robust to alternative specifications of the energy price series.

Theoretical models that incorporate asymmetric effects of energy price changes have played an important role in the recent literature. While there are several potential sources of asymmetry in the responses of nonresidential investment, our evidence casts doubt, in particular, on the view that investment is subject to an uncertainty effect of the type described in Bernanke (1983) and Pindyck (1991). If there is such an effect, it seems too small or too imprecisely estimated to be detected by our statistical tools. Similarly, our results are inconsistent with the existence of a reallocation effect of the type emphasized in Hamilton (1988).

Thus, there is no compelling reason to abandon the assumption of symmetric responses to energy price shocks. Our analysis in the remainder of the paper quantified the effect of energy price shocks on nonresidential investment based on linear models that impose symmetry. We showed that, in response to a one percent increase in energy prices, *real nonresidential investment in equipment* falls by a statistically insignificant -0.41% after six quarters, whether or not *mining and oil field machinery* is included. The largest responses are found in *transportation equipment* with a statistically significant -0.95% decline, and in *mining and oil field machinery* which rises by a statistically significant 1.90%. Other forms of real equipment investment are largely unresponsive to energy price shocks. Real investment in *total nonresidential structures* increases by a statistically insignificant 0.03%, while real investment in *mining structures* (used for petroleum, coal and natural gas exploration and extraction) increases by a statistically significant 1.31%. Other forms of real structures investment are unresponsive to energy price shocks. Excluding *mining structures*, the response of nonresidential structures is -0.04%, but remains statistically insignificant.

We found no compelling evidence that nonresidential fixed investment in structures has become less responsive to energy price shocks since the mid-1980s with the exception of *other structures* (including religious, educational, transportation, and farm structures, among others). We found some evidence of a decline in the responsiveness of nonresidential fixed investment in equipment, mainly in *transportation equipment* and in *other equipment*. The declining response of investment in *transportation equipment* is consistent with evidence of a decline in the response of real consumption of vehicles in Edelstein and Kilian (2007).

We also addressed the closely related question of how important energy price shocks have been overall for U.S. nonresidential fixed investment. Historical decompositions show only a negligible cumulative effect on investment in equipment and structures. A partial exception is investment in mining structures and to a lesser extent investment in mining and oil field machinery.

Table 4.1: Components of Processed Fuels and Lubricants PPI

Component	Percent Share
Liquefied petroleum gas	2.40
Electric power	40.32
Natural gas	14.50
Unleaded gasoline	13.98
Aviation gasoline	0.57
Kerosene	0.16
Jet Fuel	9.73
Home heating oil	2.06
#2 diesel fuel	11.36
Residual Fuels	2.42
Lubricating grease	0.14
Lubricating and similar oils	1.29
Petroleum and coal products, n.e.c.	1.10

Notes: Based on 2002 weights in processed fuels and lubricants. Source: U.S. Bureau of Labor Statistics

Table 4.2a: Classification of Responses to Innovations in Real Energy Prices: 1970.II-2006.IV

	No Response			Symmetric Response			Classical Asymmetry			Perverse Asymmetry		
	C	LC	NC	C	LC	NC	C	LC	NC	C	LC	NC
Nonresidential Structures												
Total										X	X	X
Mining							X**	X*	X***			
Total Excluding Mining										X	X	X
Commercial and Healthcare	X	X	X									
Manufacturing	X	X	X									
Power and Communication	X	X*	X*									
Other	X							X	X			
Nonresidential Equipment												
Total							X	X	X			
Mining and Oil Field Machinery							X	X	X			
Total Excluding Mining and Oil Field Machinery							X	X*	X			
Information Processing	X	X	X									
Industrial							X	X	X			
Transportation							X	X	X			
Other Excluding Mining and Oil Field Machinery							X	X	X			

Note: C: Change in real energy price. LC: Large change in real energy price. NC: Net change in real energy price. An *X* in a given column indicates that the estimated responses follow the corresponding pattern. A formal definition of these patterns can be found in the text.

* Symmetry rejected at 10% level. ** Symmetry rejected at 5% level. *** Symmetry rejected at 1% level.

Table 4.2b: Classification of Responses to Innovations in Real Energy Prices: 1988.I-2006.IV

	No Response			Symmetric Response			Classical Asymmetry			Perverse Asymmetry		
	C	LC	NC	C	LC	NC	C	LC	NC	C	LC	NC
Nonresidential Structures												
Total		X	X							X		
Mining							X	X	X			
Total Excluding Mining	X	X	X									
Commercial and Healthcare	X	X	X									
Manufacturing	X	X	X									
Power and Communication	X	X	X									
Other		X	X							X		
Nonresidential Equipment												
Total	X	X	X									
Mining and Oil Field Machinery							X	X	X			
Total Excluding Mining and Oil Field Machinery	X	X	X									
Information Processing	X	X	X									
Industrial	X										X	X
Transportation	X	X	X									
Other Excluding Mining and Oil Field Machinery	X	X	X									

Note: C: Change in real energy price. LC: Large change in real energy price. NC: Net change in real energy price. An *X* in a given column indicates that the estimated responses follow the corresponding pattern. A formal definition of these patterns can be found in the text.

* Symmetry rejected at 10% level. ** Symmetry rejected at 5% level. *** Symmetry rejected at 1% level.

Table 4.2c: Classification of Responses to Innovations in Real Energy Prices: 1970.II-1987.IV

	No Response			Symmetric Response			Classical Asymmetry			Perverse Asymmetry		
	C	LC	NC	C	LC	NC	C	LC	NC	C	LC	NC
Nonresidential Structures												
Total			X					X		X		
Mining							X	X	X			
Total Excluding Mining	X	X	X									
Commercial and Healthcare		X	X							X		
Manufacturing	X	X	X									
Power and Communication	X	X	X									
Other							X	X	X			
Nonresidential Equipment												
Total							X	X	X			
Mining and Oil Field Machinery		X**								X		X**
Total Excluding Mining and Oil Field Machinery							X	X	X			
Information Processing	X	X	X									
Industrial							X	X	X			
Transportation							X	X	X			
Other Excluding Mining and Oil Field Machinery							X	X	X			

Note: C: Change in real energy price. LC: Large change in real energy price. NC: Net change in real energy price. An *X* in a given column indicates that the estimated responses follow the corresponding pattern. A formal definition of these patterns can be found in the text.

* Symmetry rejected at 10% level. ** Symmetry rejected at 5% level. *** Symmetry rejected at 1% level.

Table 4.2d: Classification of Responses to Innovations in Real Energy Prices: 1970.II-1987.IV with 1986 Dummies

	No Response			Symmetric Response			Classical Asymmetry			Perverse Asymmetry		
	C	LC	NC	C	LC	NC	C	LC	NC	C	LC	NC
Nonresidential Structures												
Total	X	X	X									
Mining	X	X	X									
Total Excluding Mining	X	X	X									
Commercial and Healthcare	X	X	X									
Manufacturing	X	X	X									
Power and Communication	X	X	X									
Other							X	X	X			
Nonresidential Equipment												
Total							X	X	X			
Mining and Oil Field Machinery		X								X		X
Total Excluding Mining and Oil Field Machinery							X	X	X			
Information Processing	X	X	X									
Industrial							X	X	X			
Transportation							X	X	X			
Other Excluding Mining and Oil Field Machinery							X	X	X			

Note: C: Change in real energy price. LC: Large change in real energy price. NC: Net change in real energy price. An *X* in a given column indicates that the estimated responses follow the corresponding pattern. A formal definition of these patterns can be found in the text.

* Symmetry rejected at 10% level. ** Symmetry rejected at 5% level. *** Symmetry rejected at 1% level.

Table 4.2e: Classification of Responses to Innovations in Real Energy Prices: 1970.II-2006.IV with 1986 Dummies

	No Response			Symmetric Response			Classical Asymmetry			Perverse Asymmetry		
	C	LC	NC	C	LC	NC	C	LC	NC	C	LC	NC
Nonresidential Structures												
Total	X	X	X									
Mining				X	X	X						
Total Excluding Mining	X	X	X									
Commercial and Healthcare	X	X	X									
Manufacturing	X	X	X									
Power and Communication	X	X	X									
Other							X	X	X			
Nonresidential Equipment												
Total							X	X*	X			
Mining and Oil Field Machinery							X	X	X			
Total Excluding Mining and Oil Field Machinery							X*	X**	X			
Information Processing	X	X*	X									
Industrial							X*	X	X			
Transportation							X	X	X			
Other Excluding Mining and Oil Field Machinery							X	X	X			

Note: C: Change in real energy price. LC: Large change in real energy price. NC: Net change in real energy price. An *X* in a given column indicates that the estimated responses follow the corresponding pattern. A formal definition of these patterns can be found in the text.

* Symmetry rejected at 10% level. ** Symmetry rejected at 5% level. *** Symmetry rejected at 1% level.

Figure 4.1: Intermediate Energy Prices and Energy Share in Value Added

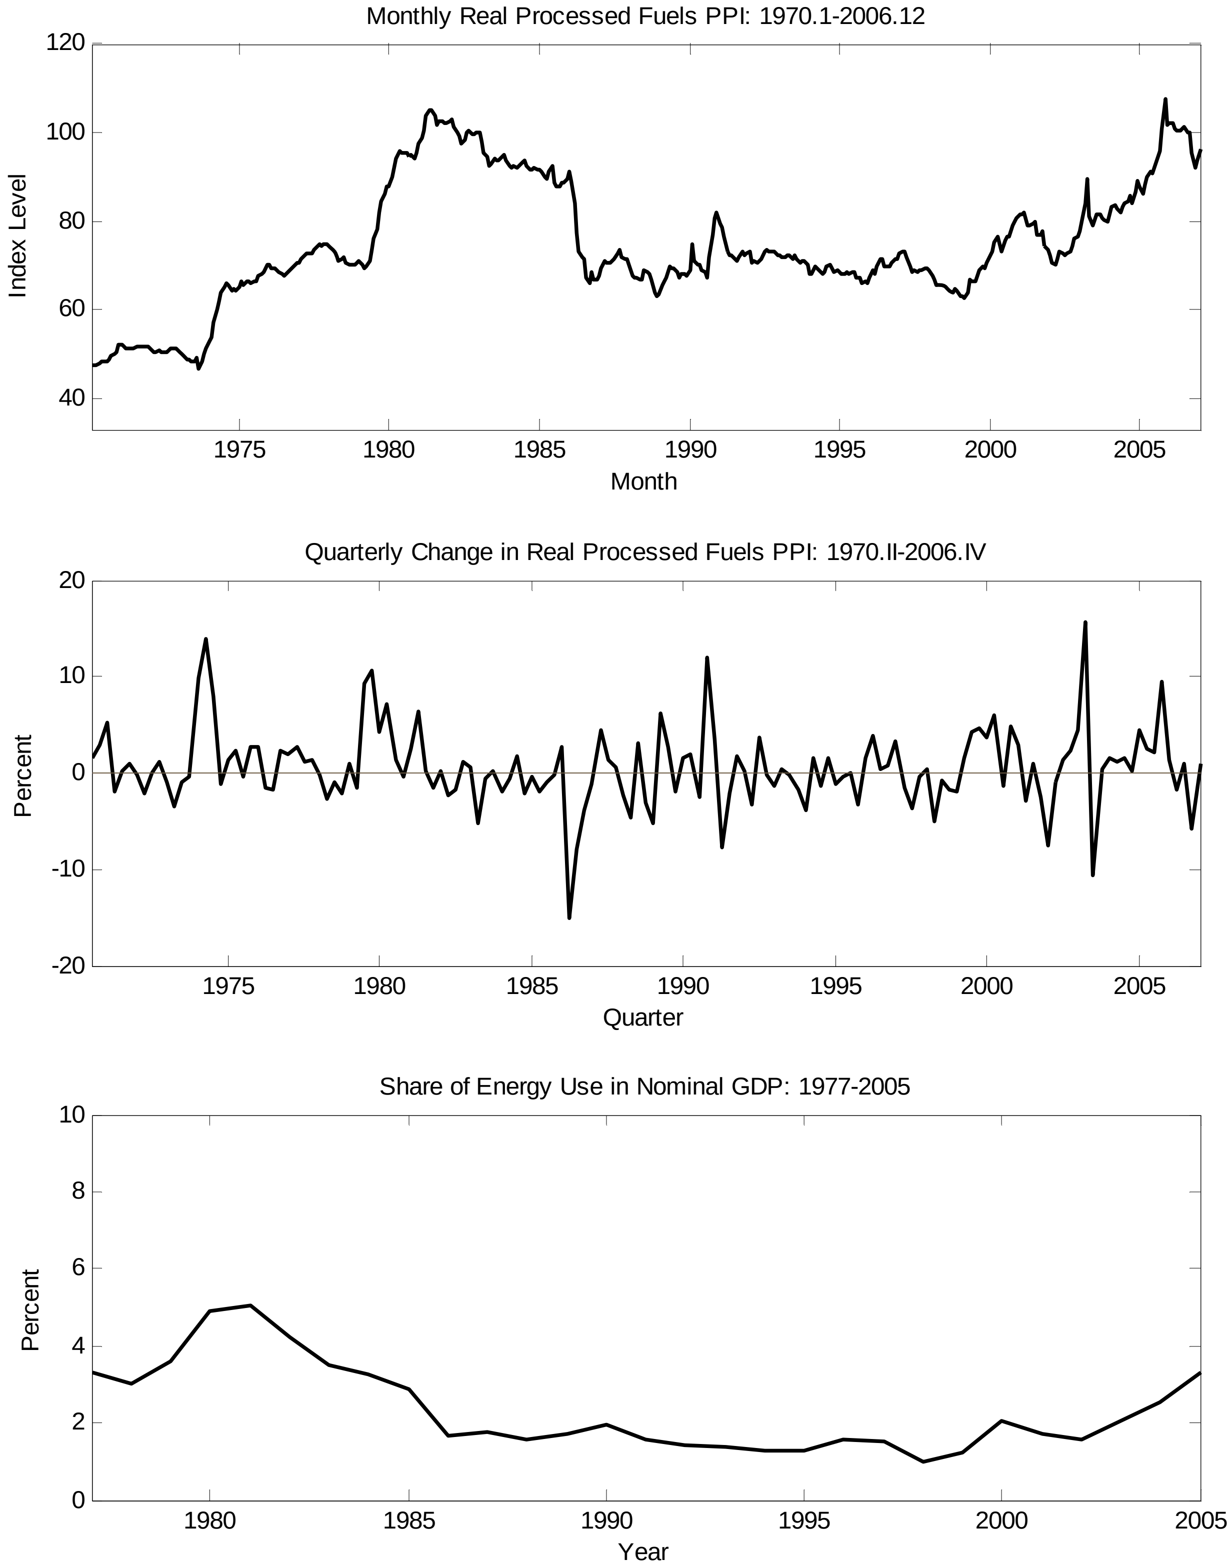

Figure 4.2: Alternative Measures of Energy Price Shocks: 1970.II-2006.IV

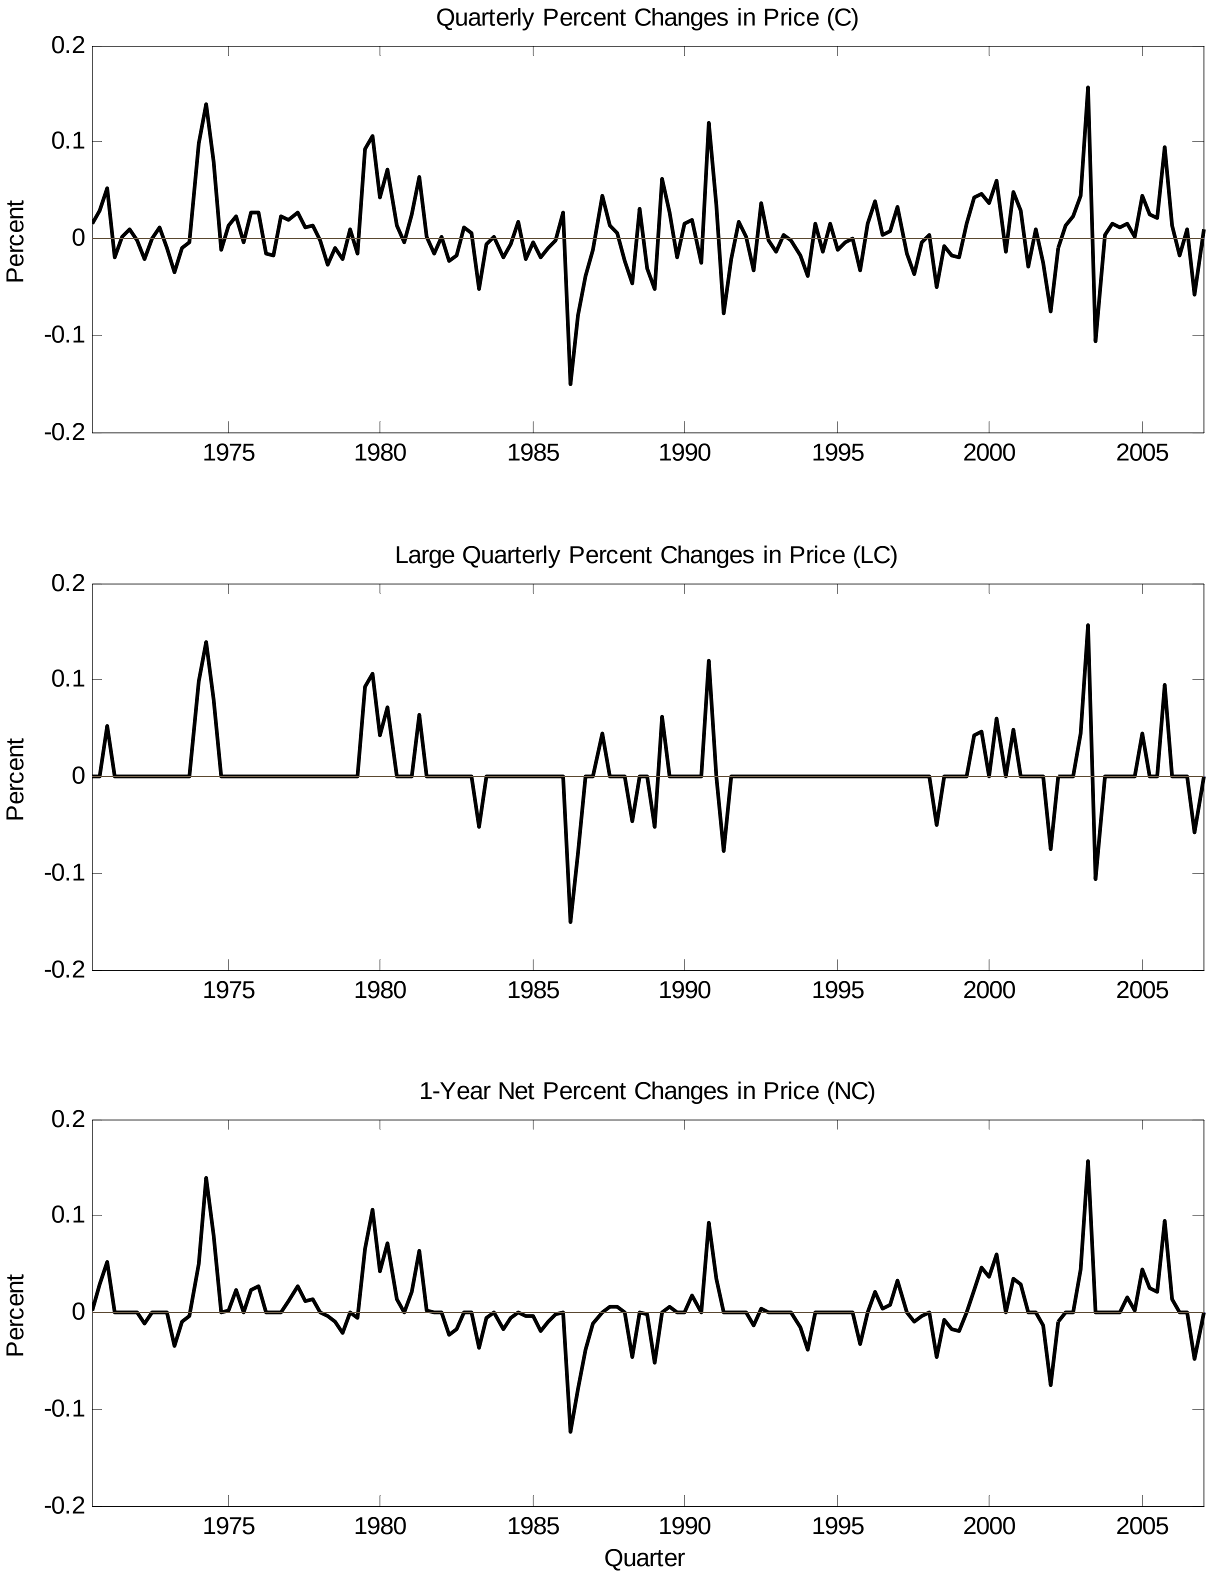

Figure 4.3: Illustration of the Composition Effect

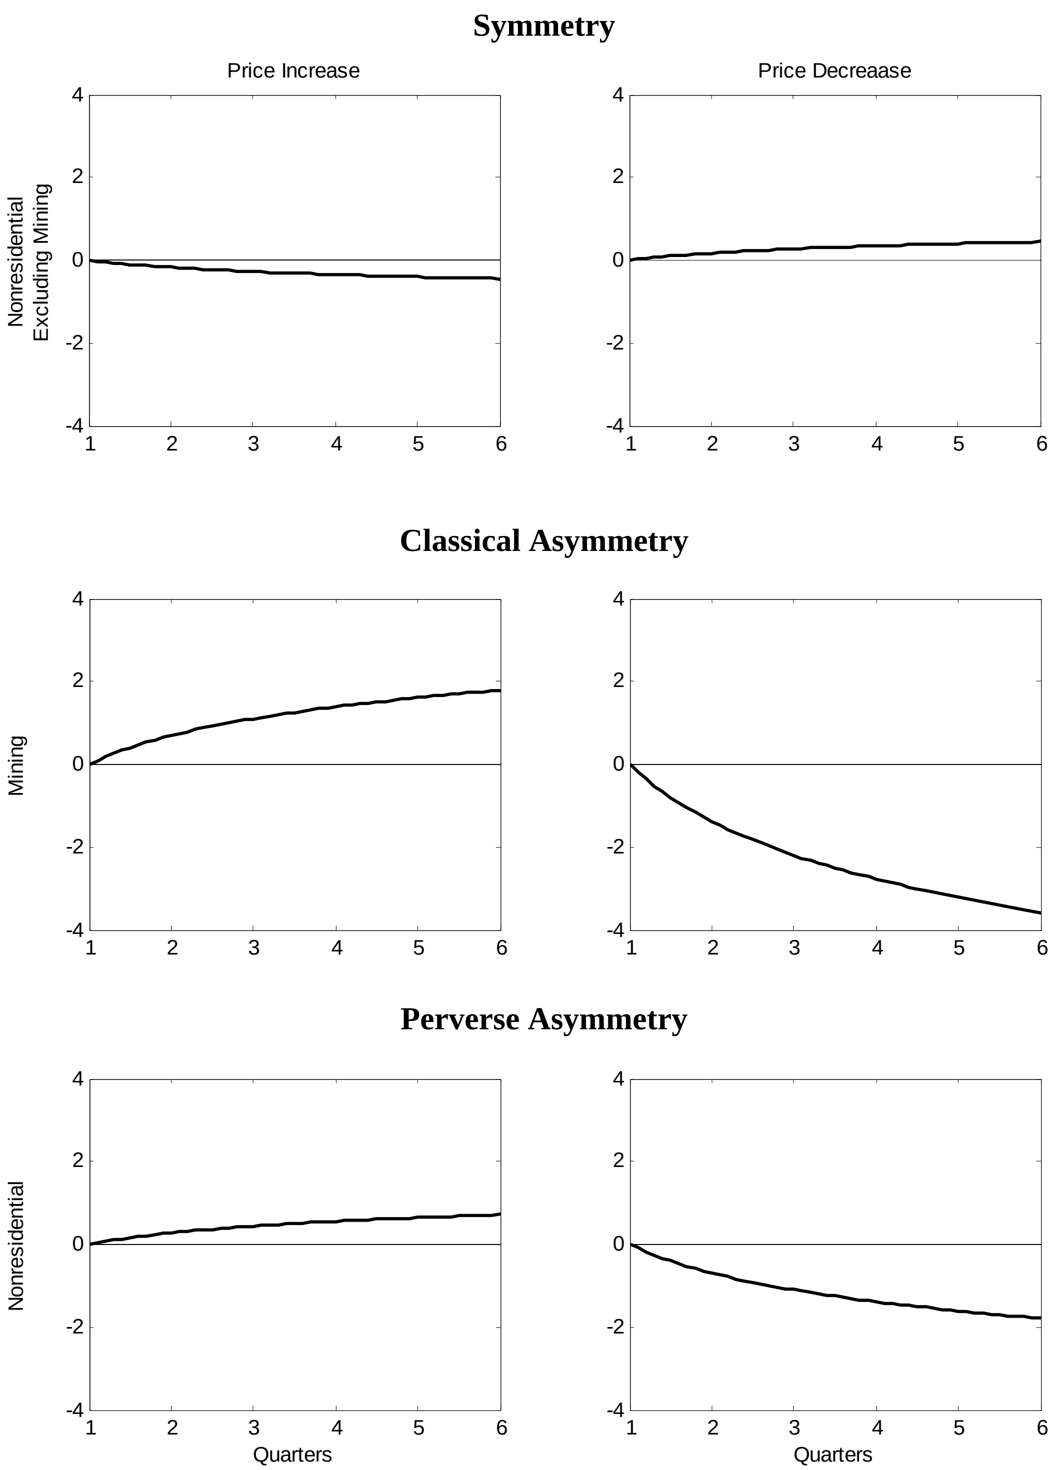

Figure 4.4a: Response of Real Structures Investment to Energy Price Shock

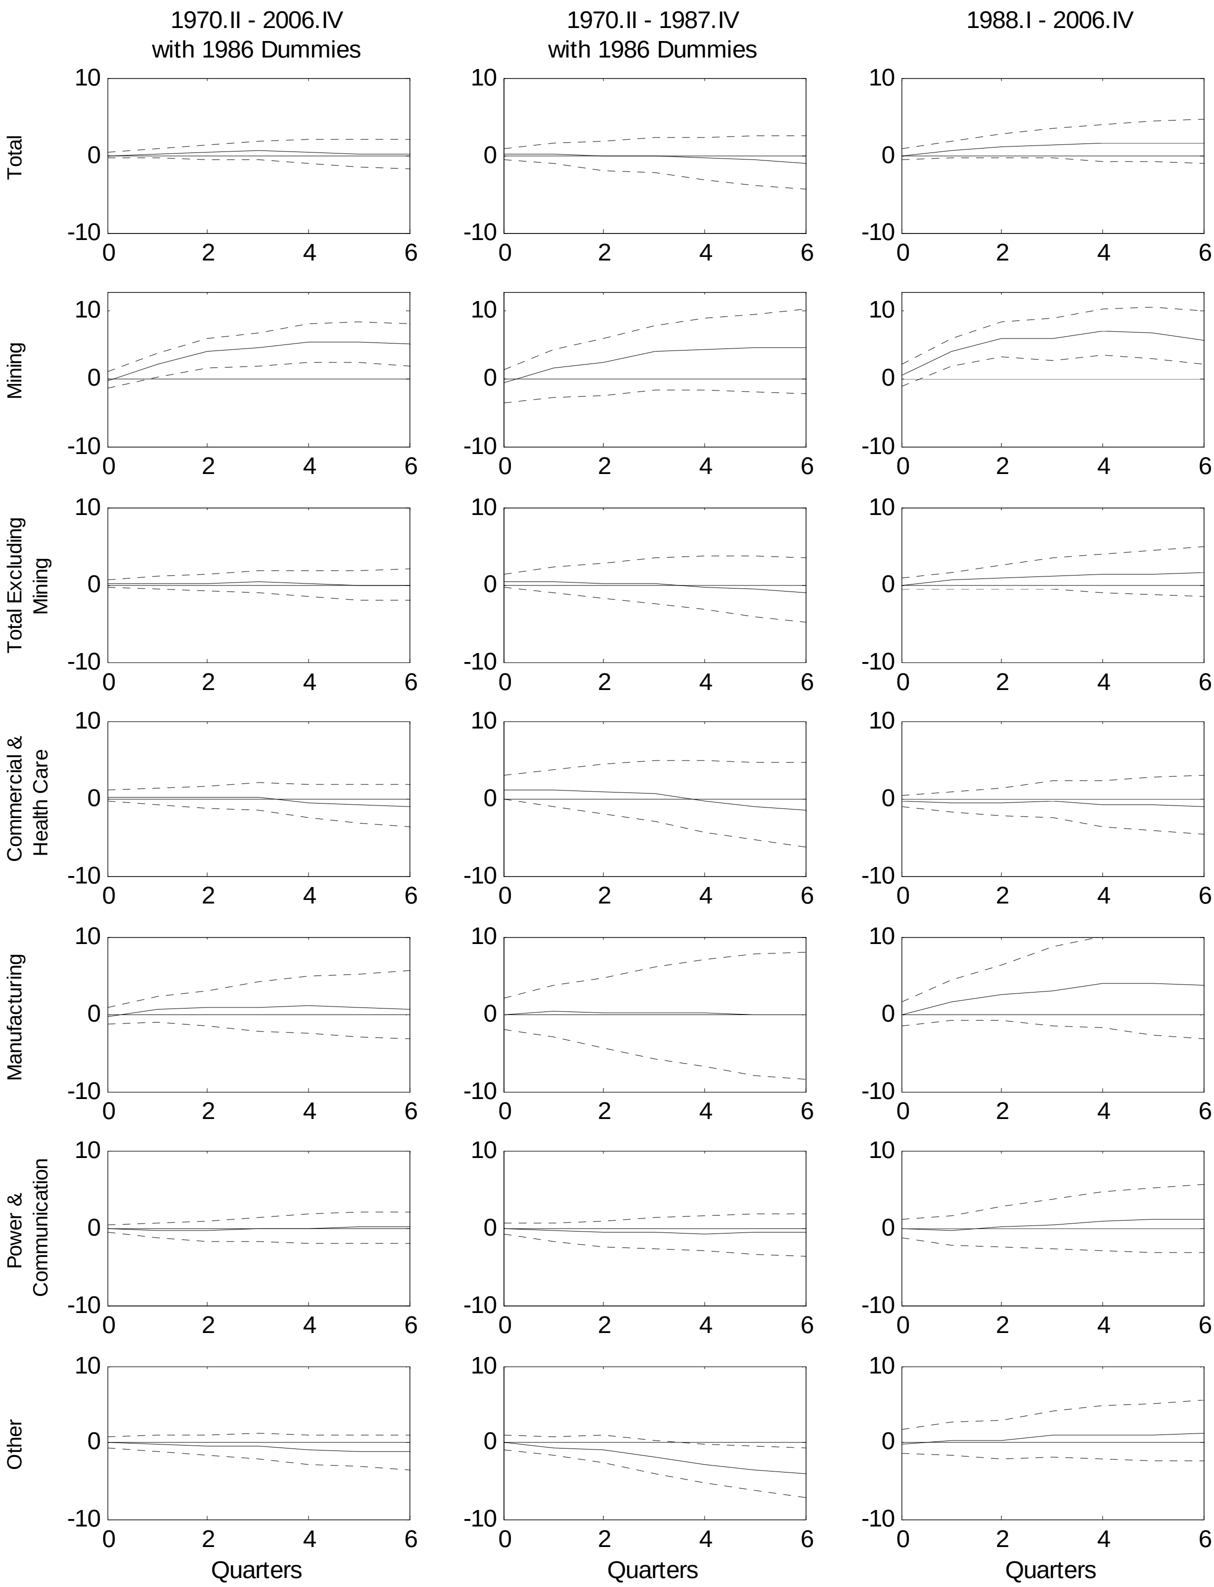

Figure 4.4b: Response of Real Equipment Investment to Energy Price Shock

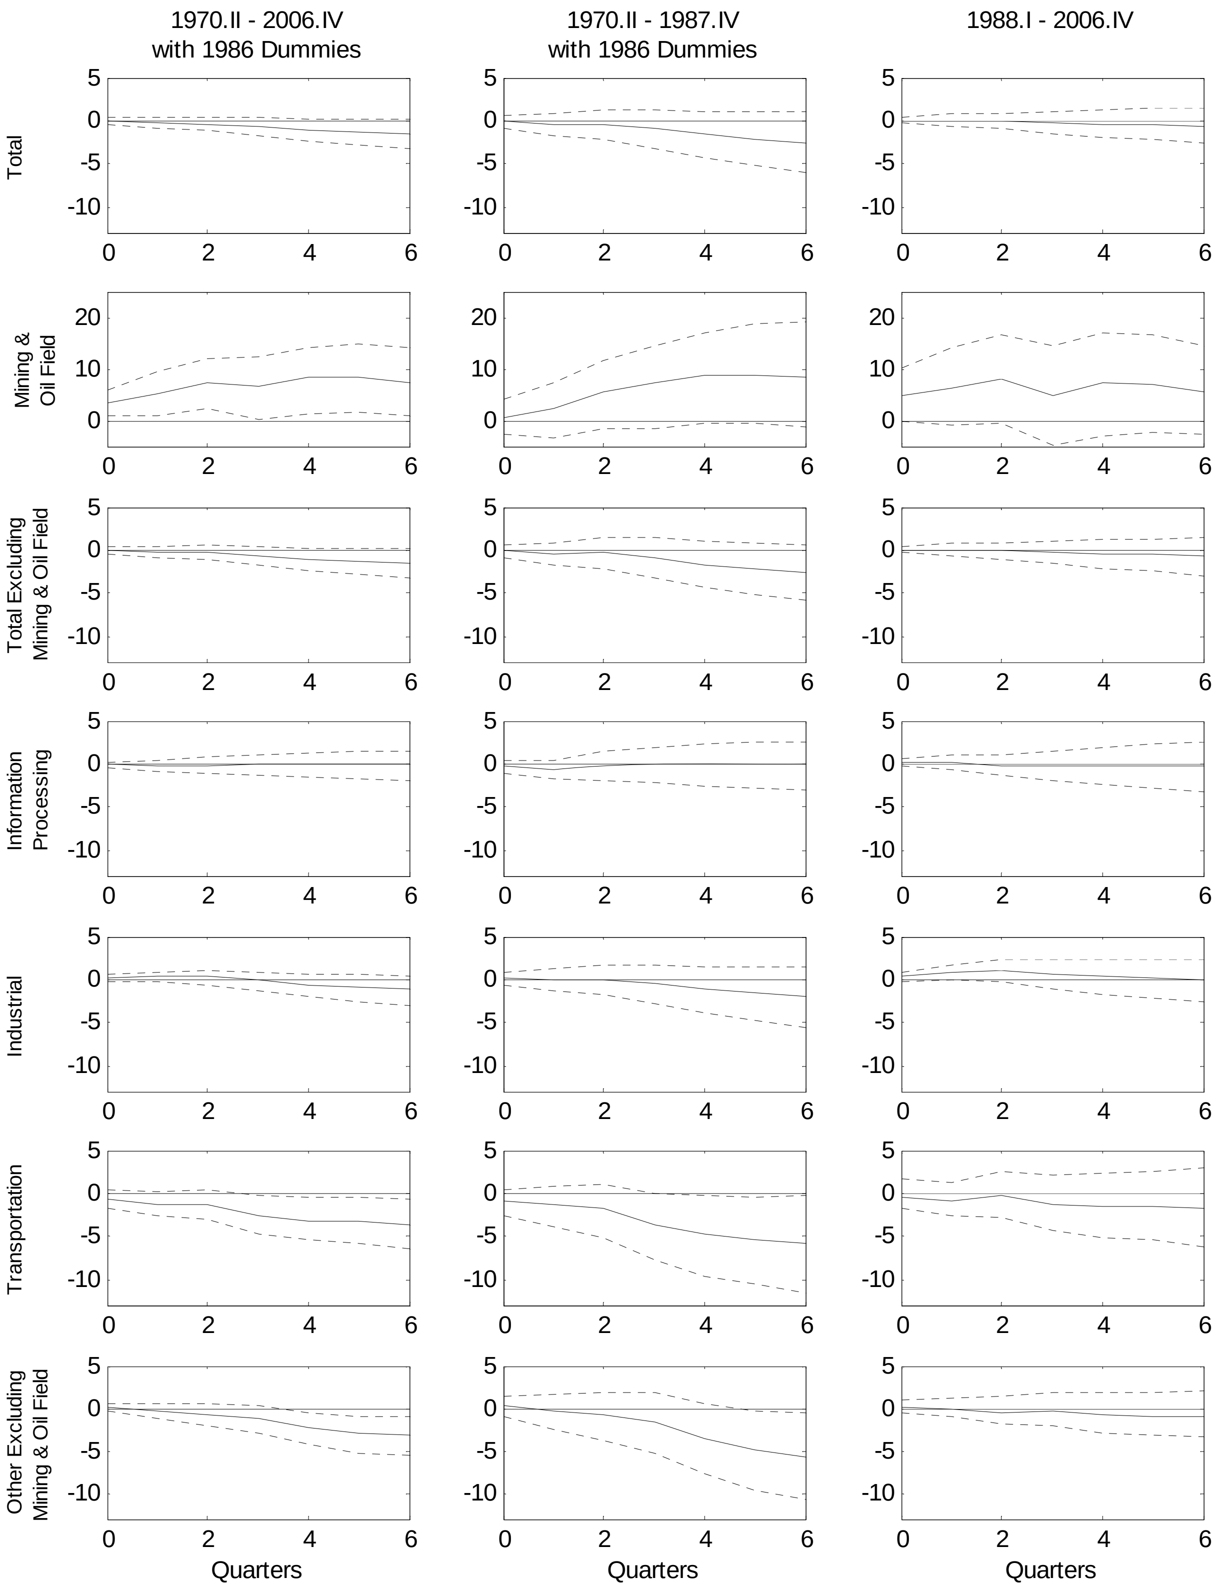

Figure 4.5: Contribution of Energy Price Shocks to Quarterly Change in Real Investment: 1970.II - 2006.IV

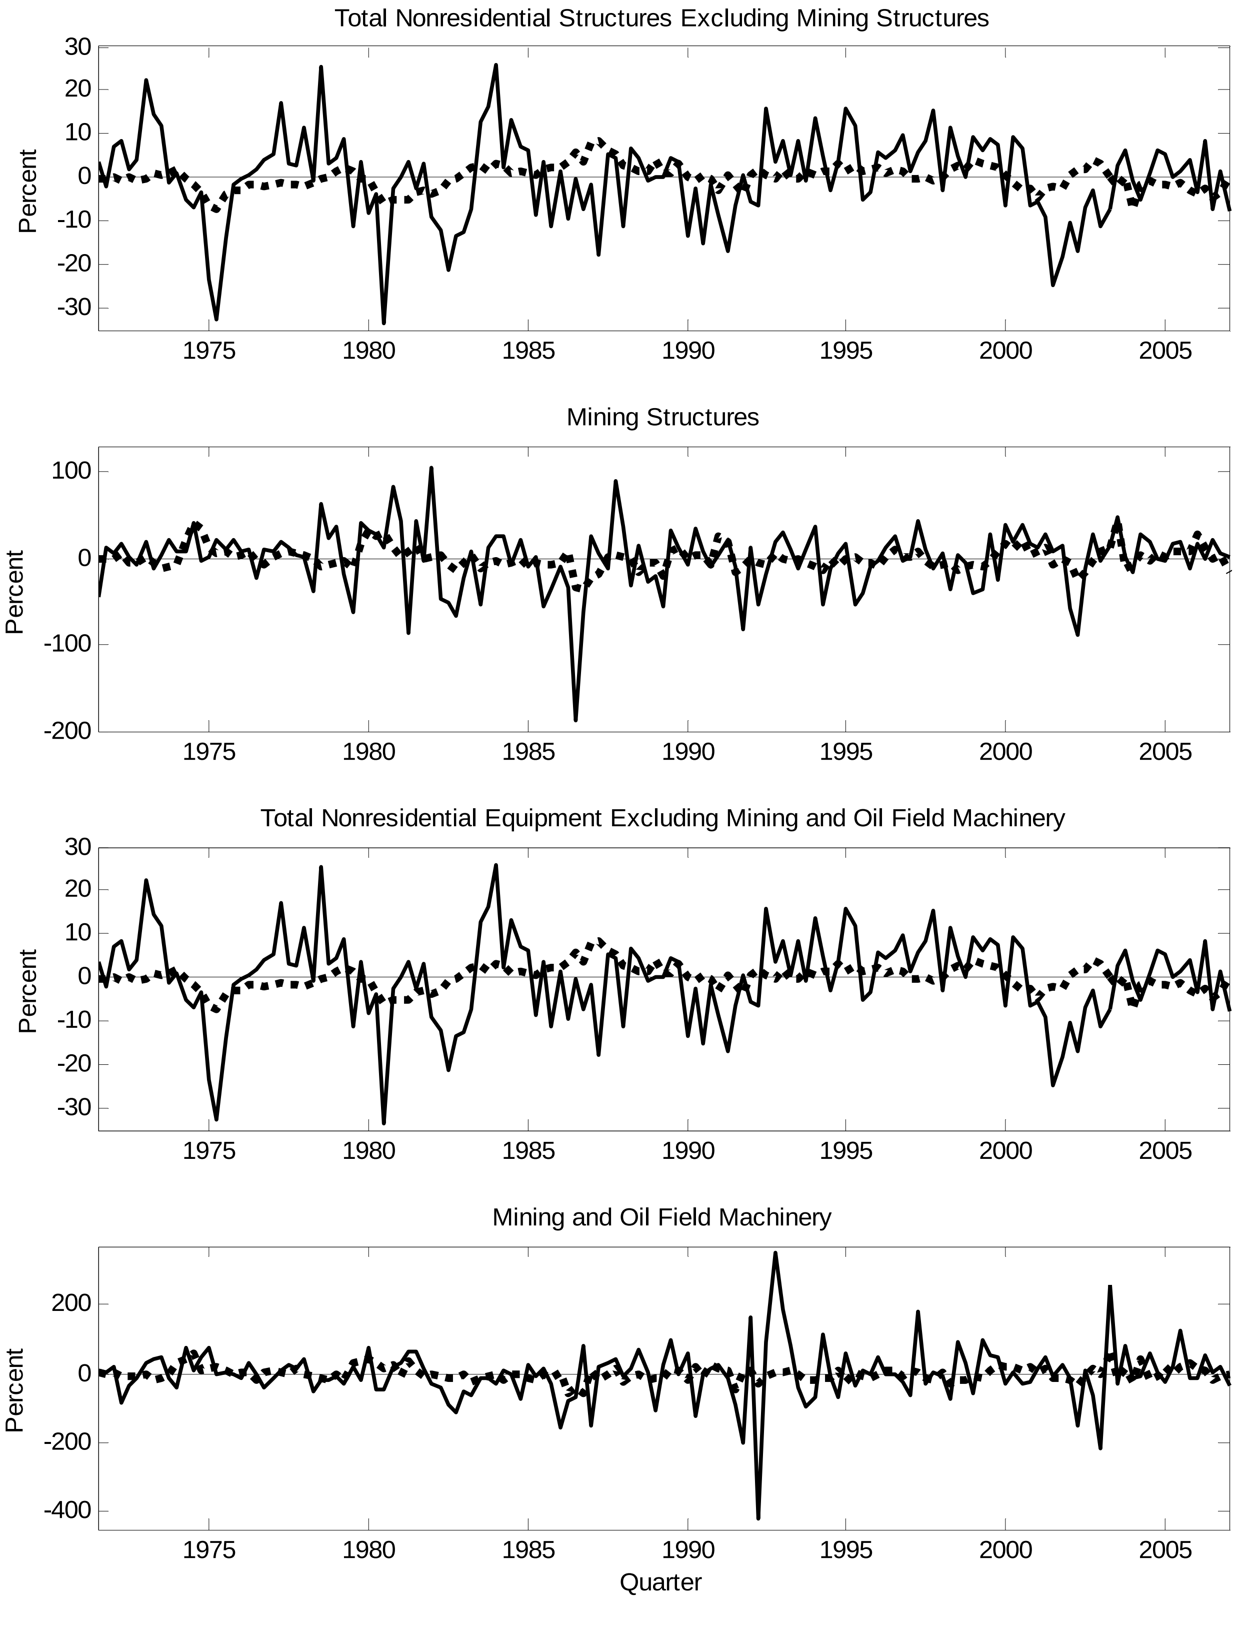

Appendix

The following results are summarized in Tables 4.2a-4.2e. Tables 4.A.1, 4.A.2, and 4.A.3 of this appendix include the p-values for tests of symmetry in the response of real investment to positive and negative energy price shocks under the *percent changes*, *large percent changes*, and *net percent changes* specifications. Figures 4.A.1a-4.A.1e show the impulse response estimates (with 95% confidence bands) for real structures investment to positive and negative energy price shocks under the three specifications. Figures 4.A.2a-4.A.2e show the corresponding results for real equipment investment.

Table 4.A.1: Symmetry Tests - C Specification

	1970.II – 2006.IV	1988.I – 2006.IV	1970.II – 1987.IV	1970.II – 1987.IV with ’86 Dummies	1970.II – 2006.IV with ’86 Dummies
Structures					
Total	0.199	0.852	0.507	0.853	0.591
Mining	0.039	0.898	0.113	0.963	0.944
Total Excluding Mining	0.480	0.877	0.839	0.806	0.625
Commercial and Healthcare	0.690	0.957	0.783	0.669	0.570
Manufacturing	0.895	0.957	0.969	0.859	0.873
Power and Communication	0.115	0.774	0.309	0.783	0.426
Other	0.428	0.361	0.713	0.855	0.620
Equipment					
Total	0.161	0.949	0.535	0.457	0.113
Mining and Oil Field	0.425	0.125	0.124	0.675	0.737
Total Excluding Mining and Oil Field	0.132	0.947	0.484	0.419	0.093
Information Processing	0.188	0.215	0.729	0.662	0.137
Industrial	0.116	0.158	0.172	0.338	0.090
Transportation	0.139	0.781	0.566	0.613	0.164
Other Excluding Mining and Oil Field	0.480	0.708	0.645	0.607	0.431

C: Percent Changes in Real Energy Prices.

Table 4.A.2: Symmetry Tests - LC Specification

	1970.II – 2006.IV	1988.I – 2006.IV	1970.II – 1987.IV	1970.II – 1987.IV with ’86 Dummies	1970.II – 2006.IV with ’86 Dummies
Structures					
Total	0.341	0.900	0.441	0.958	0.786
Mining	0.060	0.954	0.132	0.949	0.817
Total Excluding Mining	0.666	0.906	0.933	0.985	0.824
Commercial and Healthcare	0.817	0.956	0.838	0.861	0.789
Manufacturing	0.961	1.000	0.966	0.954	0.935
Power and Communication	0.098	0.523	0.910	0.980	0.473
Other	0.493	0.912	0.848	0.992	0.771
Equipment					
Total	0.100	0.908	0.760	0.694	0.057
Mining and Oilfield	0.689	0.316	0.038	0.769	0.863
Total Excluding Mining and Oil Field	0.087	0.904	0.706	0.652	0.050
Information Processing	0.114	0.274	0.783	0.697	0.067
Industrial	0.159	0.201	0.443	0.694	0.142
Transportation	0.123	0.800	0.885	0.933	0.131
Other Excluding Mining and Oil Field	0.510	0.812	0.923	0.935	0.435

LC: Large Percent Changes in Real Energy Prices.

Table 4.A.3: Symmetry Tests - NC Specification

	1970.II – 2006.IV	1988.I – 2006.IV	1970.II – 1987.IV	1970.II – 1987.IV with '86 Dummies	1970.II – 2006.IV with '86 Dummies
Structures					
Total	0.154	0.738	0.530	0.963	0.838
Mining	0.002	0.516	0.184	0.979	0.929
Total Excluding Mining	0.379	0.739	0.902	0.936	0.724
Commercial and Healthcare	0.719	0.871	0.856	0.744	0.729
Manufacturing	0.845	0.880	0.998	0.893	0.944
Power and Communication	0.051	0.118	0.504	0.912	0.720
Other	0.510	0.574	0.825	0.913	0.849
Equipment					
Total	0.476	0.831	0.767	0.607	0.375
Mining and Oilfield	0.218	0.869	0.038	0.581	0.782
Total Excluding Mining and Oil Field	0.458	0.815	0.761	0.602	0.349
Information Processing	0.379	0.761	0.647	0.650	0.483
Industrial	0.281	0.561	0.252	0.455	0.280
Transportation	0.415	0.950	0.763	0.702	0.374
Other Excluding Mining and Oil Field	0.764	0.989	0.768	0.772	0.842

NC: 1-Year Net Percent Changes in Real Energy Prices.

Figure 4.A.1a: Response of Real Nonresidential Fixed Investment in Structures: 1970.II-2006.IV

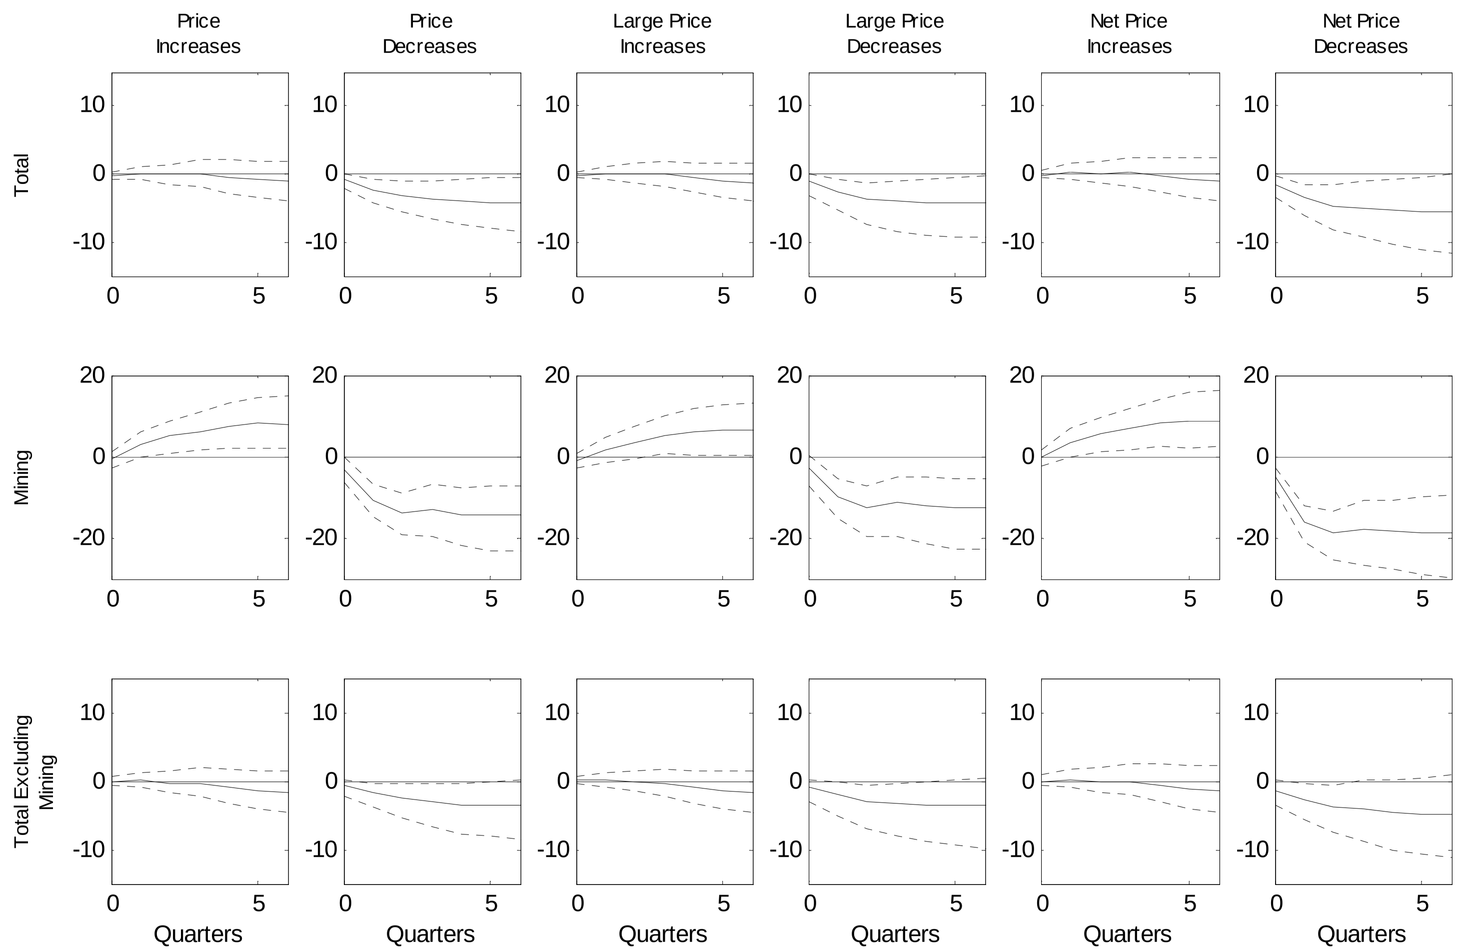

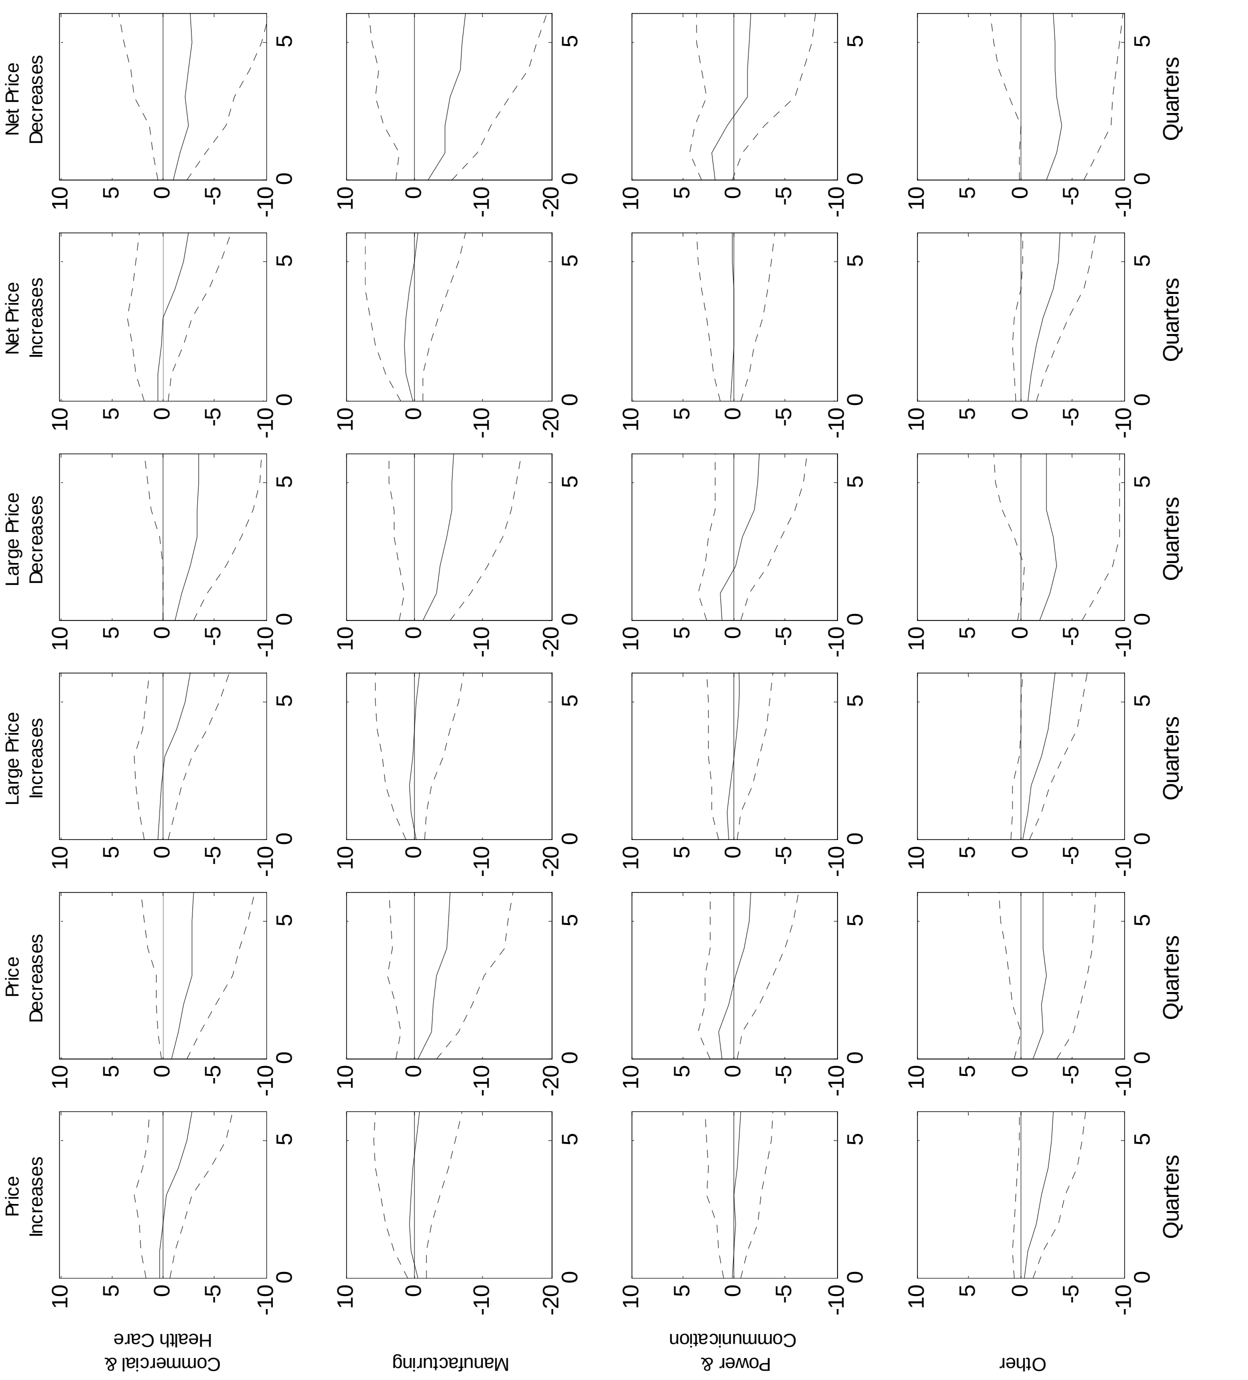
Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Commercial & Health Care
Manufacturing
Power & Communication
Other
Quarters

Figure 4.A.1b: Response of Real Nonresidential Fixed Investment in Structures: 1988.I-2006.IV

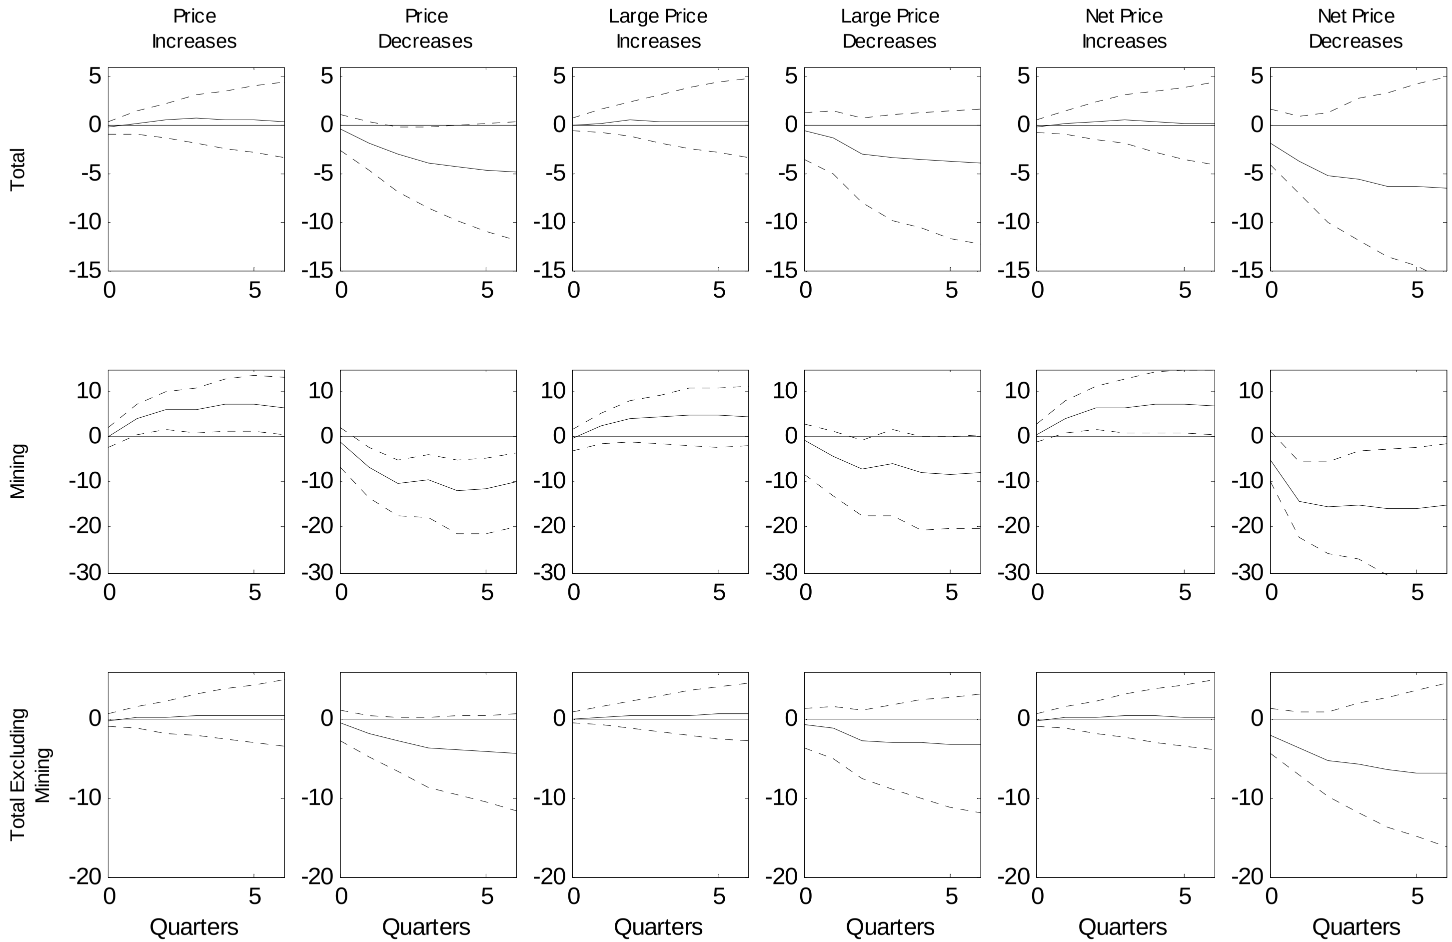

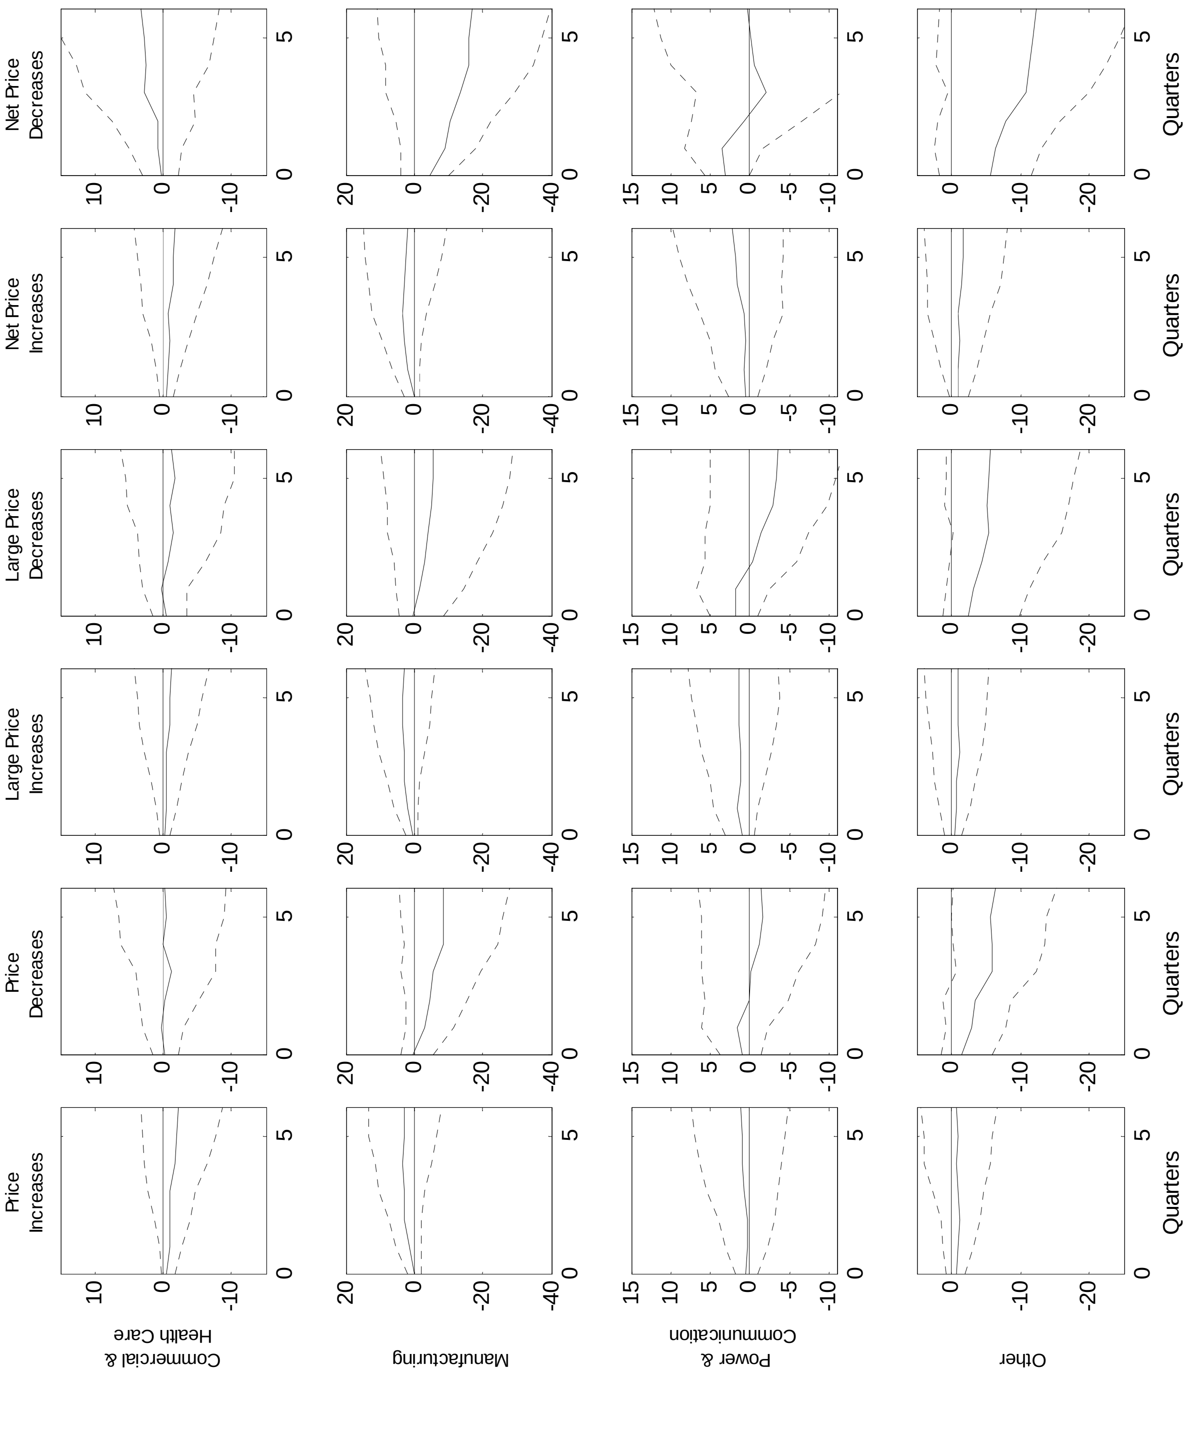

Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Commercial & Health Care
Manufacturing
Power & Communication
Other
Quarters

Figure 4.A.1c: Response of Real Nonresidential Fixed Investment in Structures: 1970.II-1987.IV

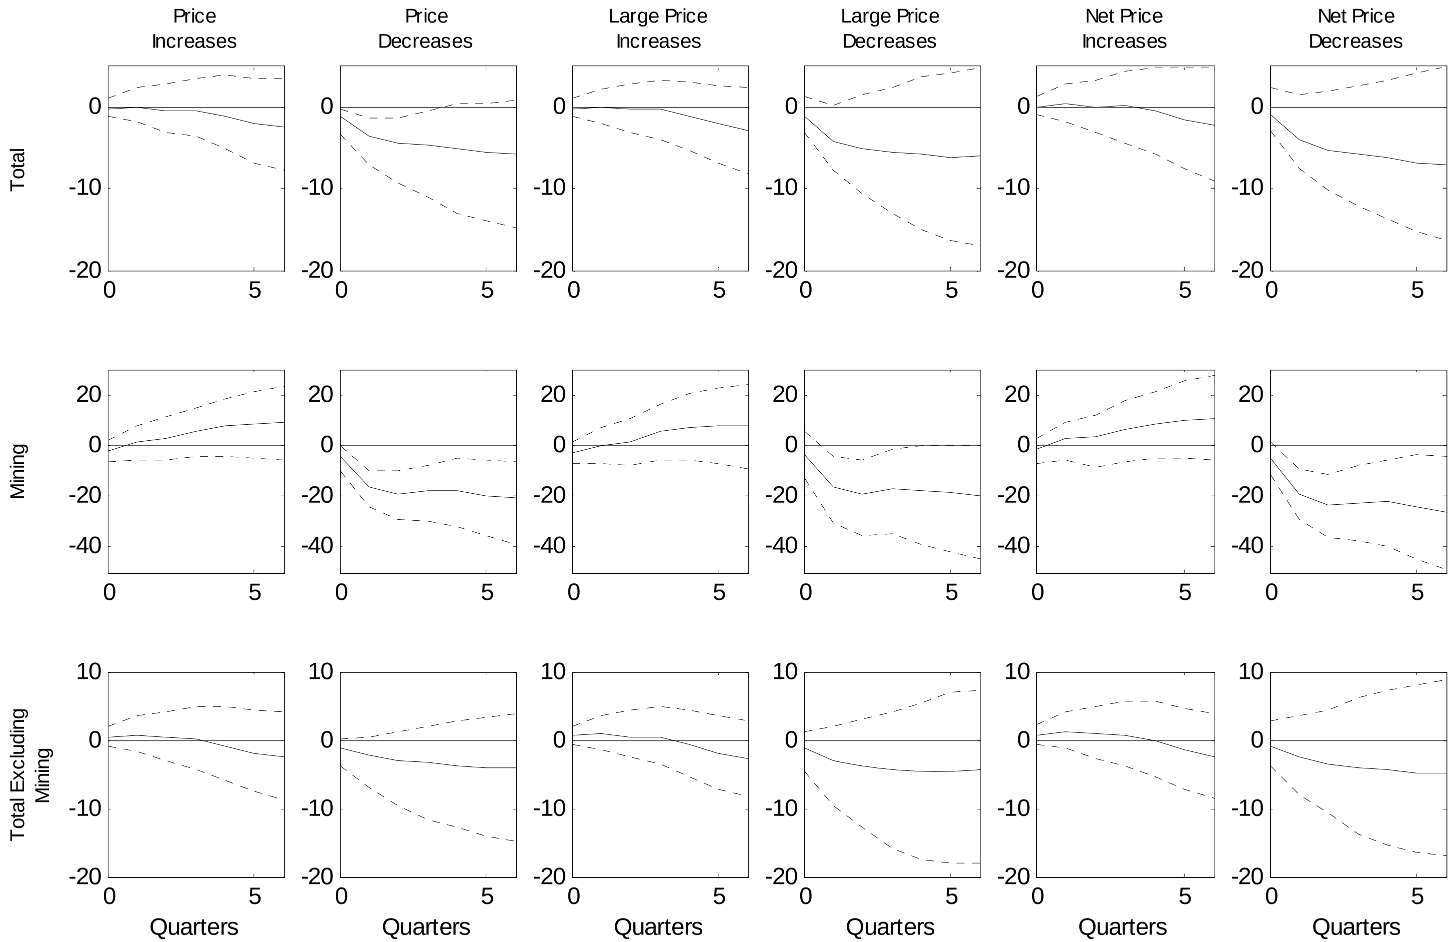

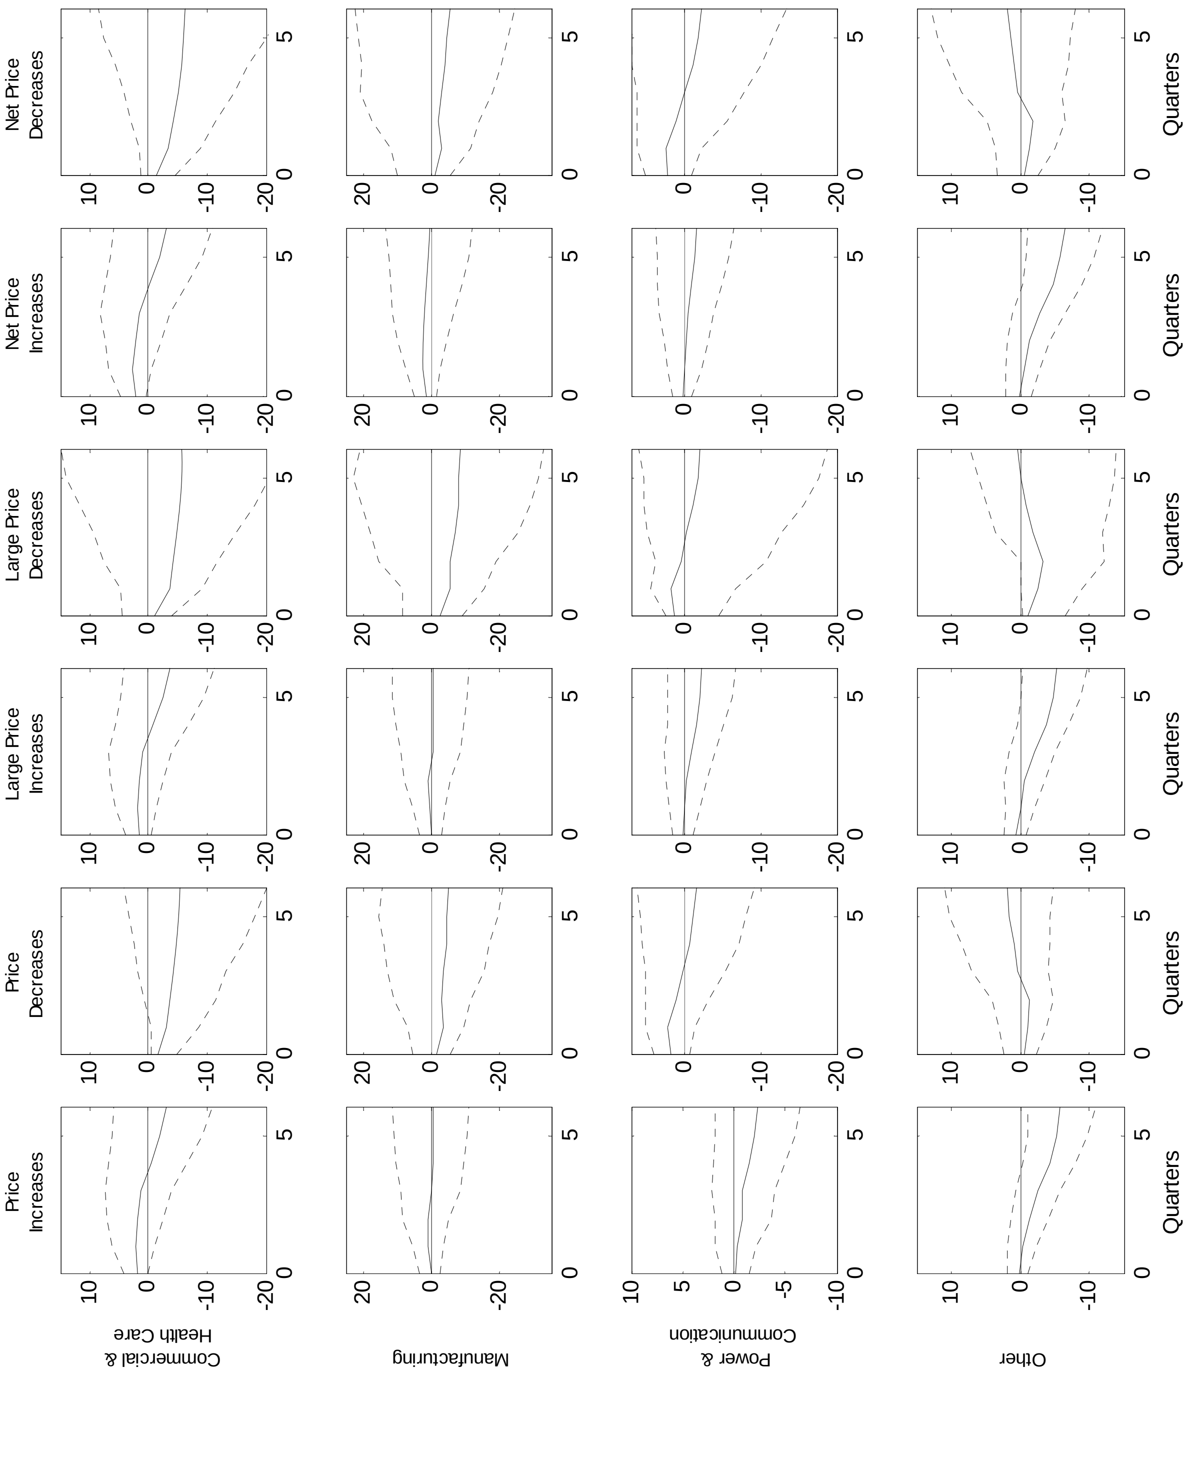
Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Commercial & Health Care
Manufacturing
Power & Communication
Other
Quarters

Figure 4.A.1d: Response of Real Nonresidential Fixed Investment in Structures: 1970.II-1987.IV with 1986 Dummies

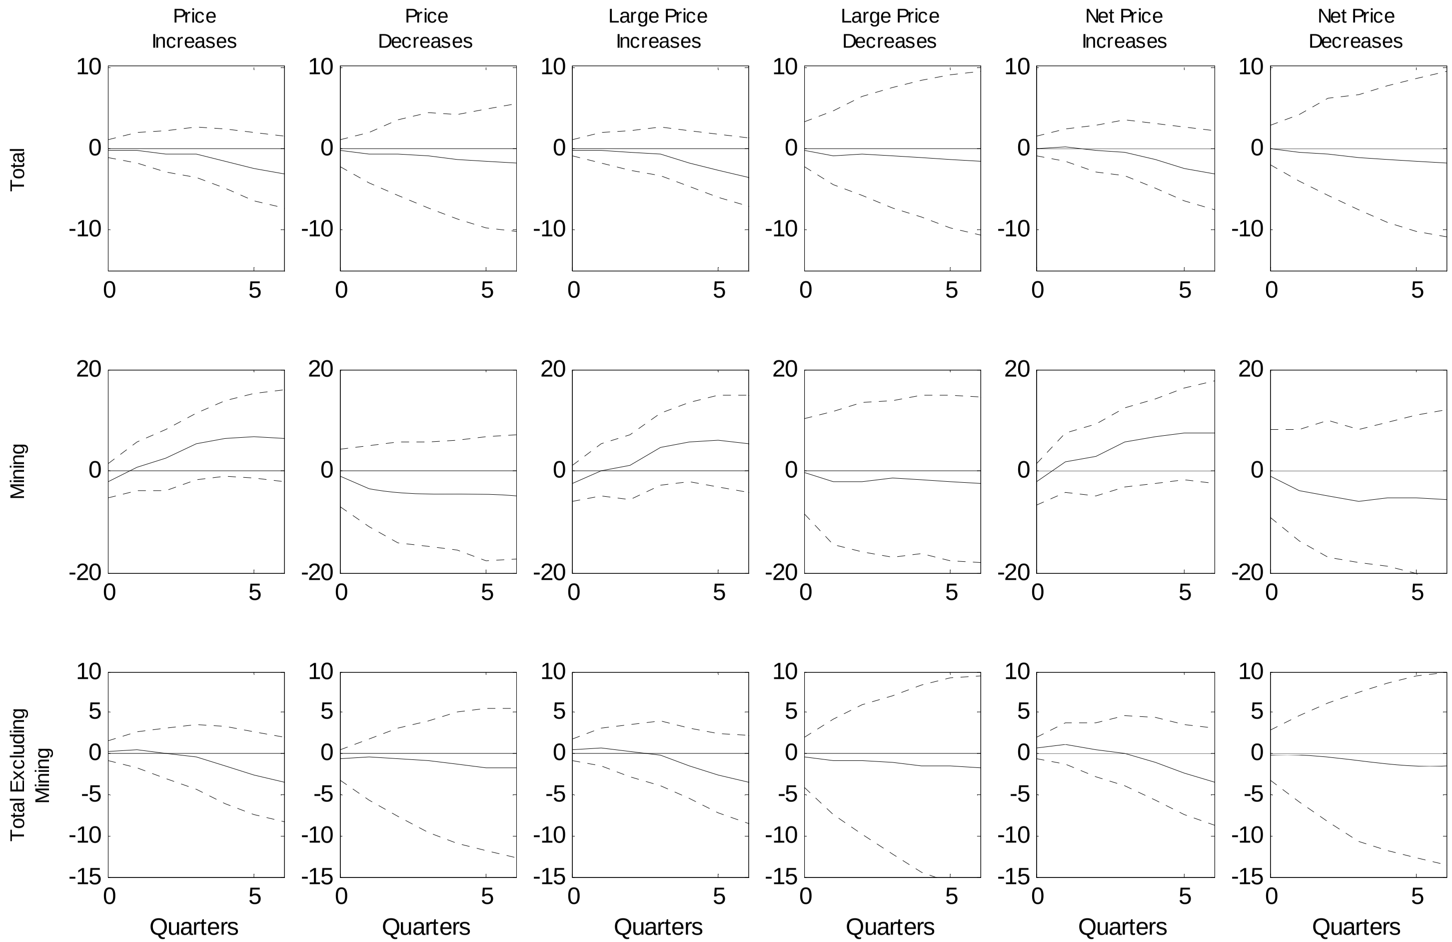

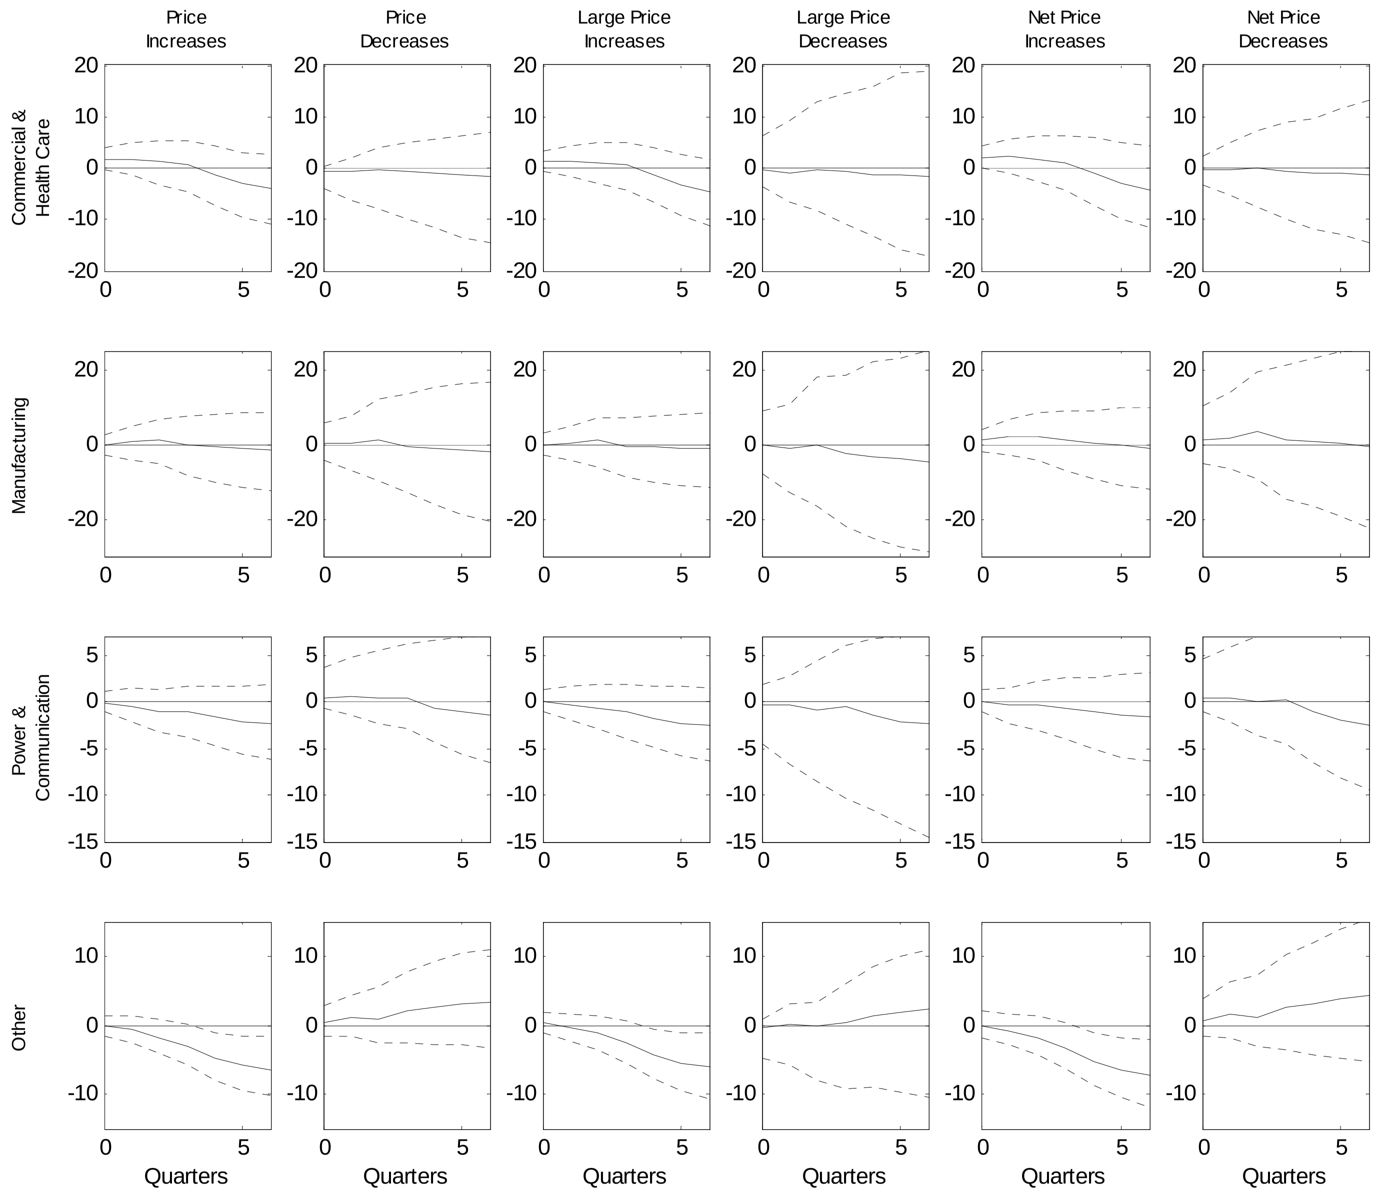

Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Commercial & Health Care
Manufacturing
Power & Communication
Other
Quarters

Figure 4.A.1e: Response of Real Nonresidential Fixed Investment in Structures: 1970.II-2006.IV with 1986 Dummies

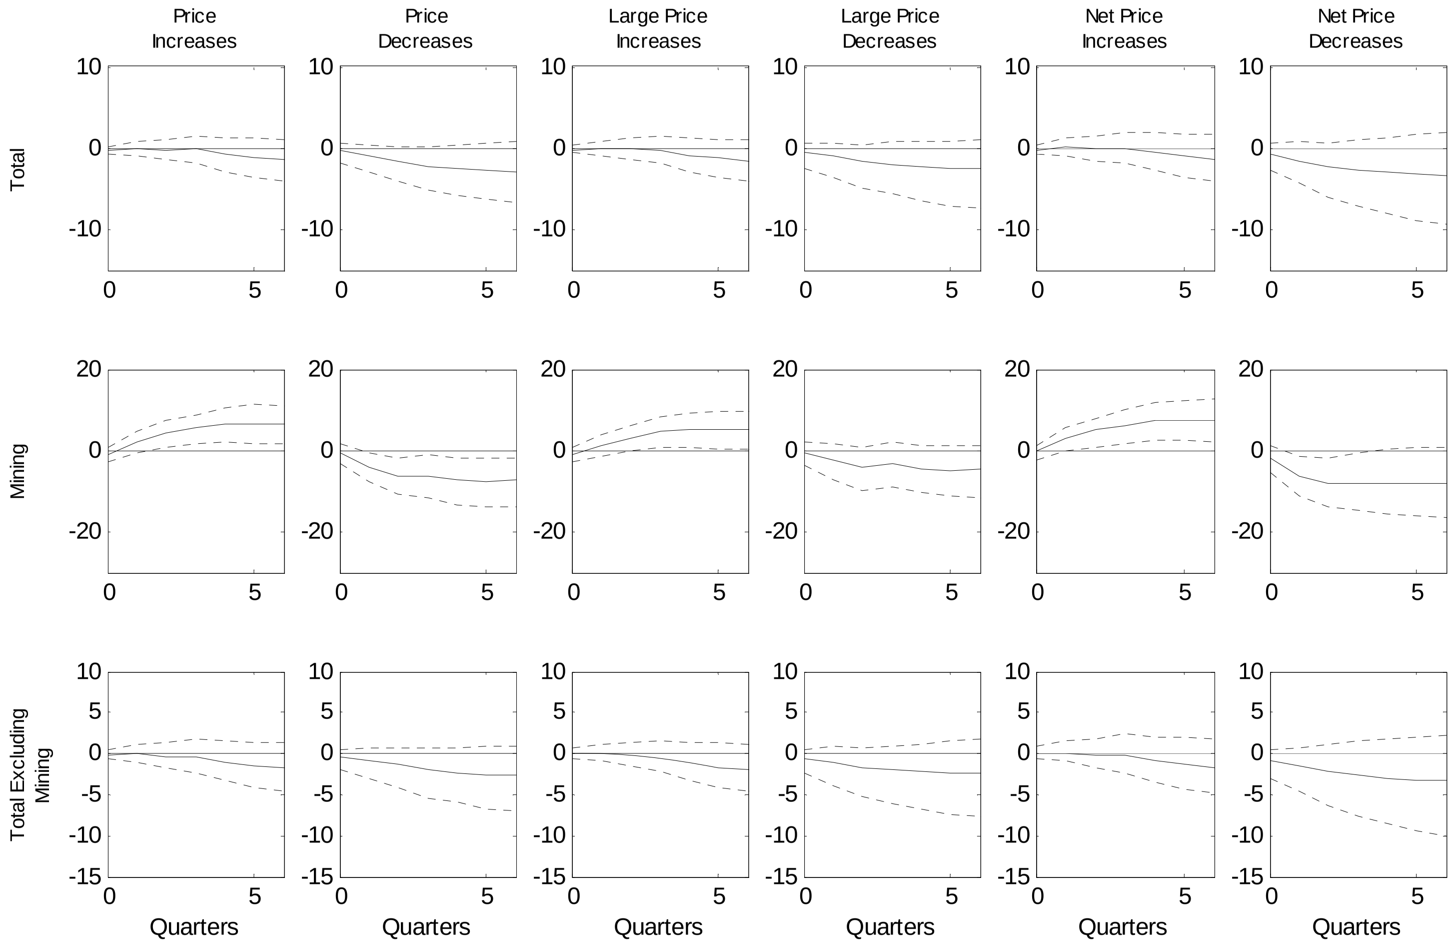

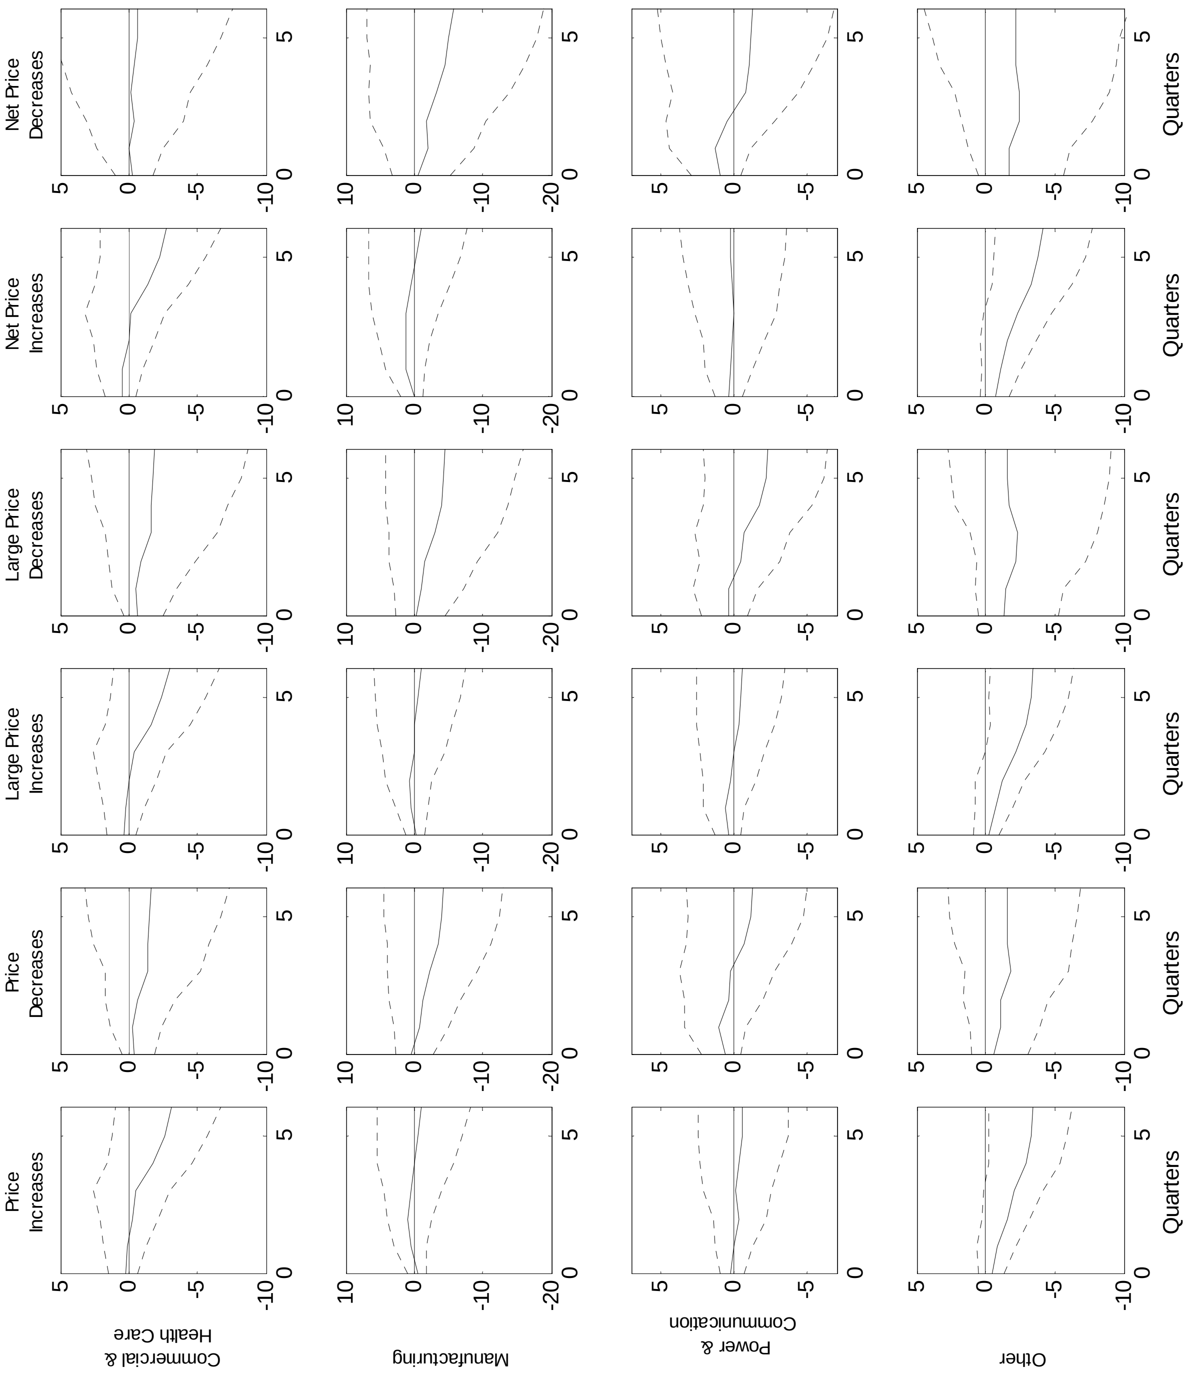
Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Commercial & Health Care
Manufacturing
Power & Communication
Other
Quarters

Figure 4.A.2a: Response of Real Nonresidential Fixed Investment in Equipment: 1970.II-2006.IV

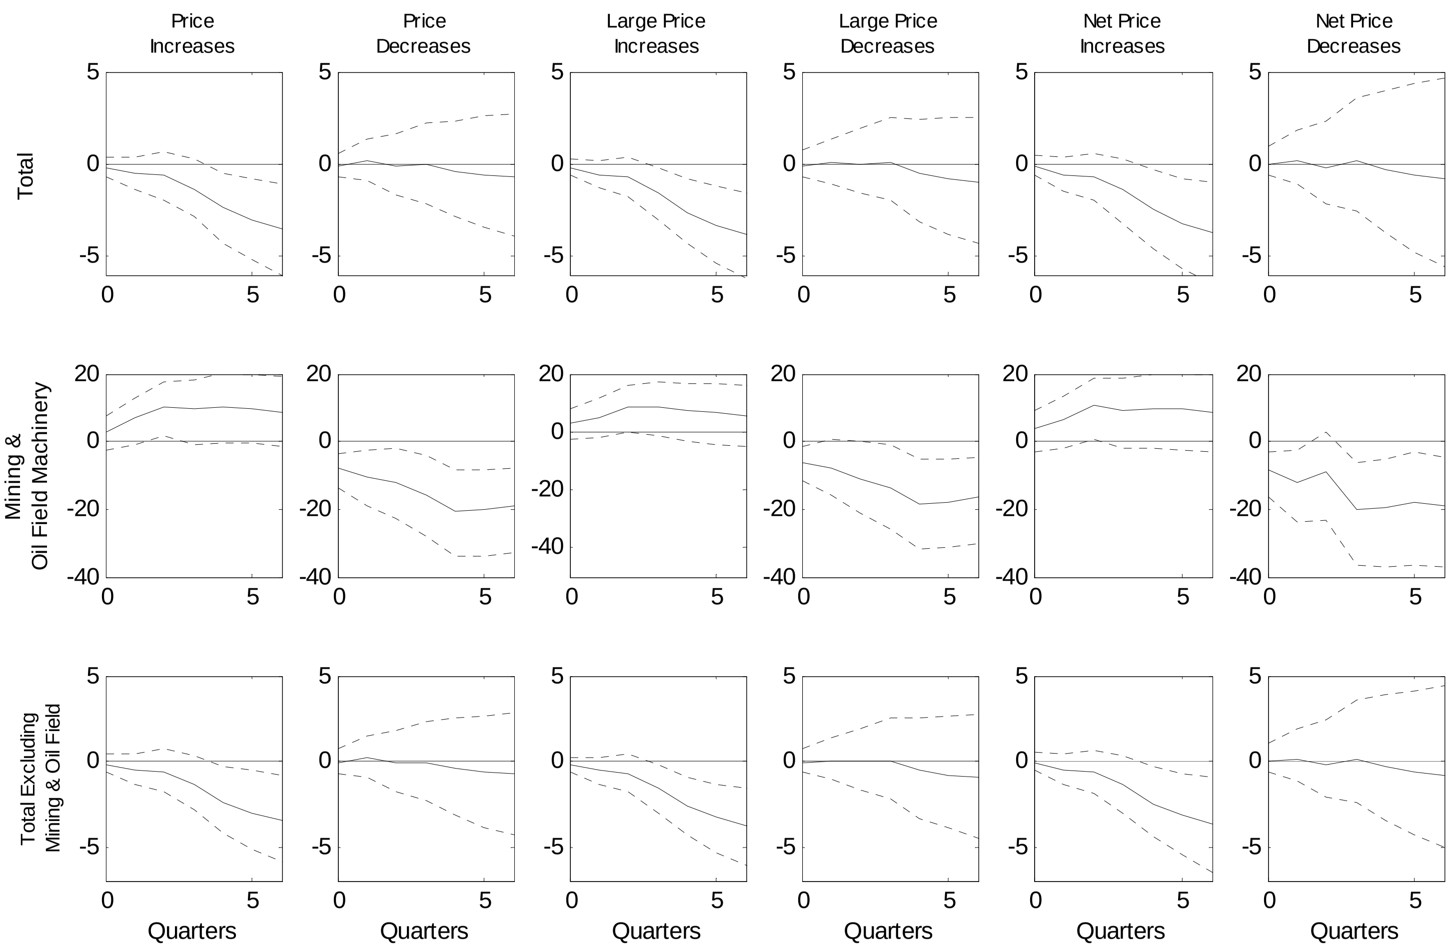

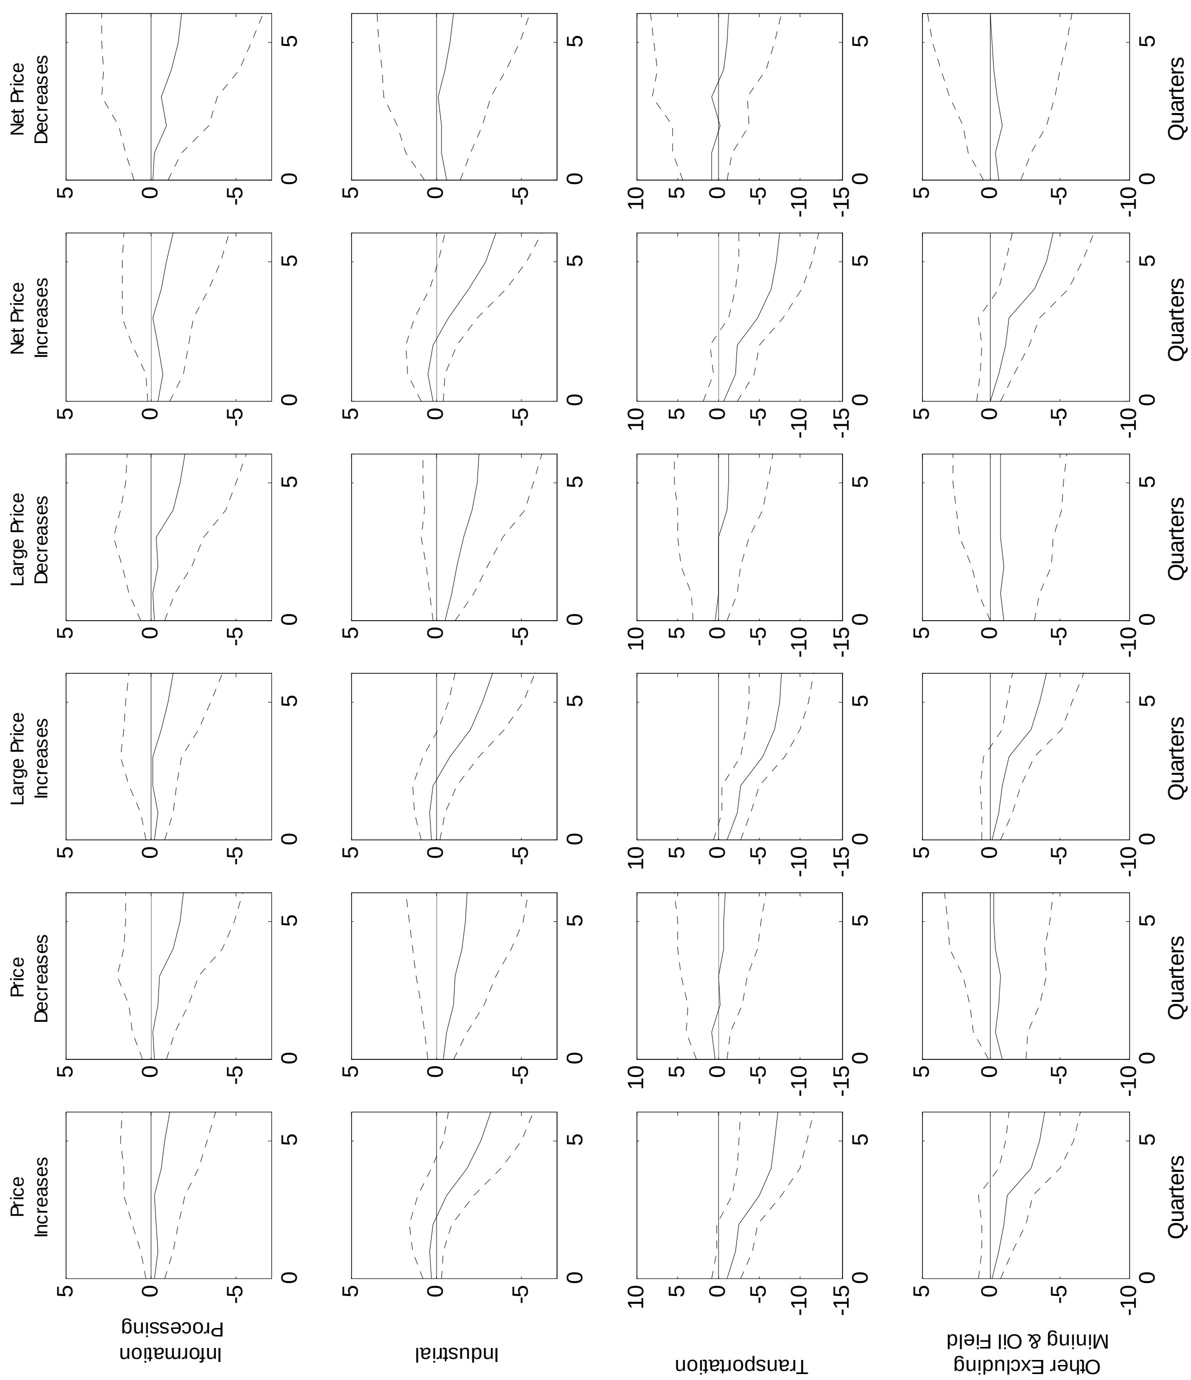

Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Information Processing
Industrial
Transportation
Other Excluding Mining & Oil Field
Quarters

Figure 4.A.2b: Response of Real Nonresidential Fixed Investment in Equipment: 1988.I-2006.IV

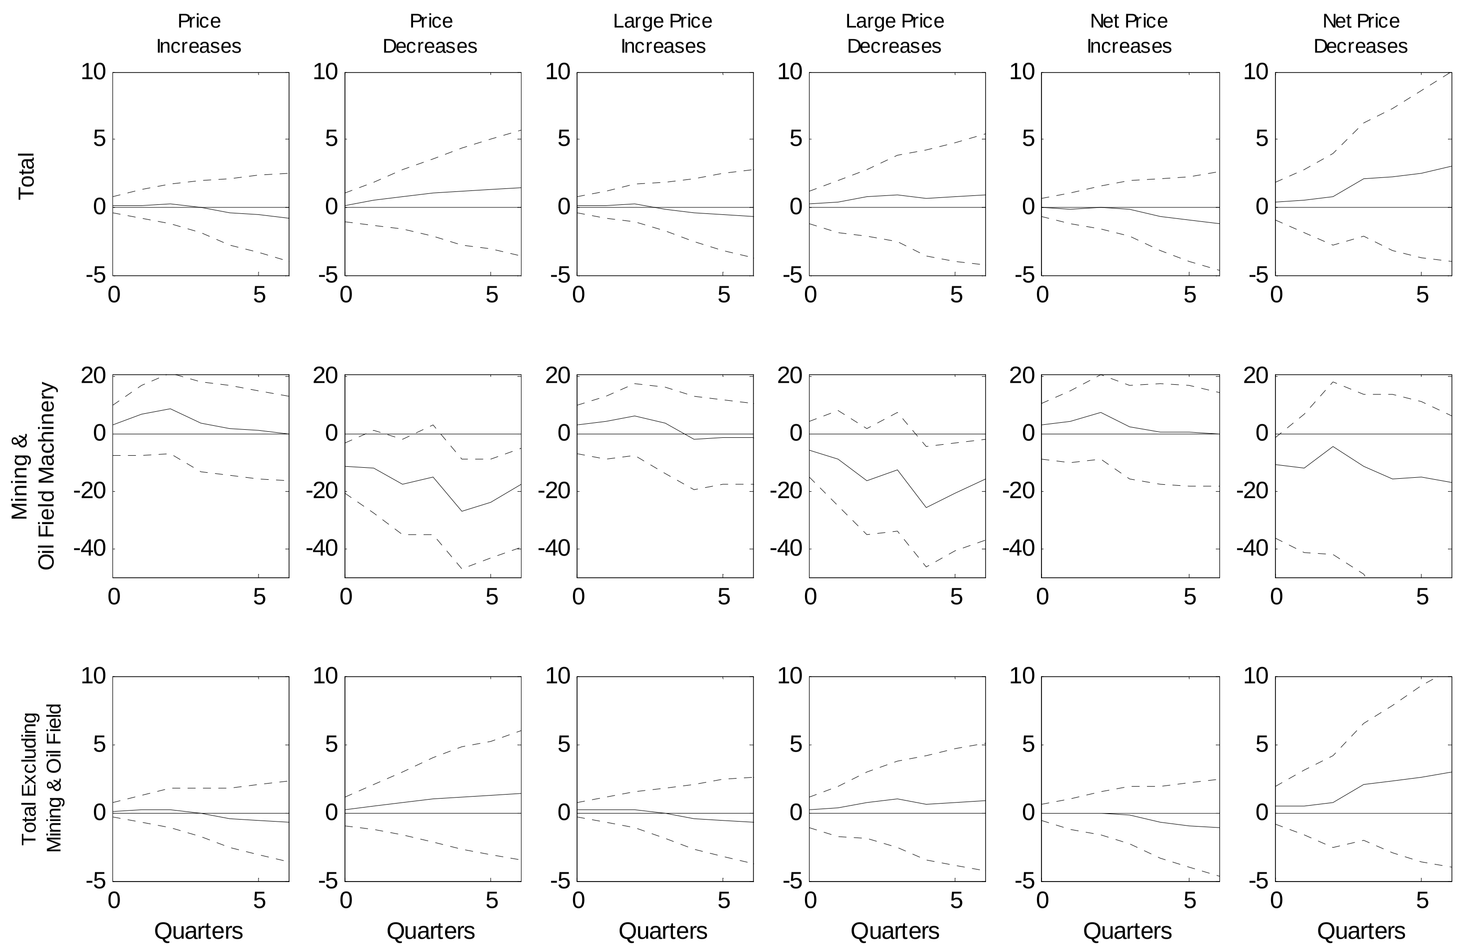

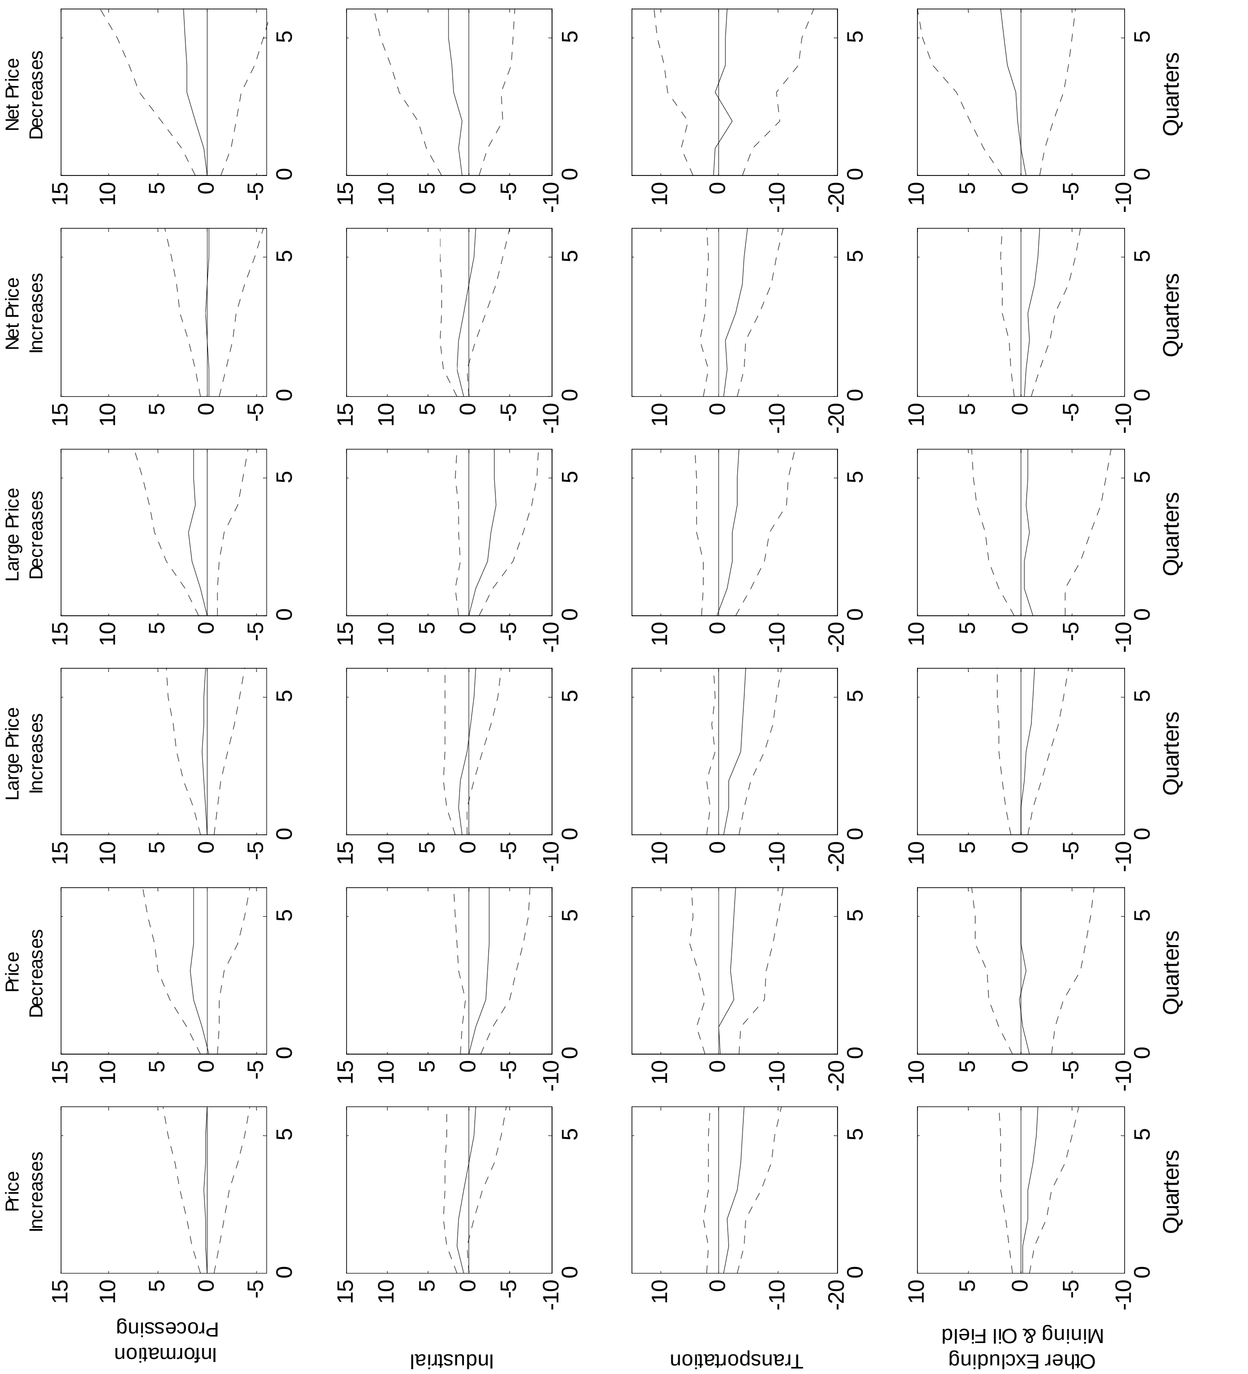
Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Information Processing
Industrial
Transportation
Other Excluding Mining & Oil Field
Quarters

Figure 4.A.2c: Response of Real Nonresidential Fixed Investment in Equipment: 1970.II-1987.IV

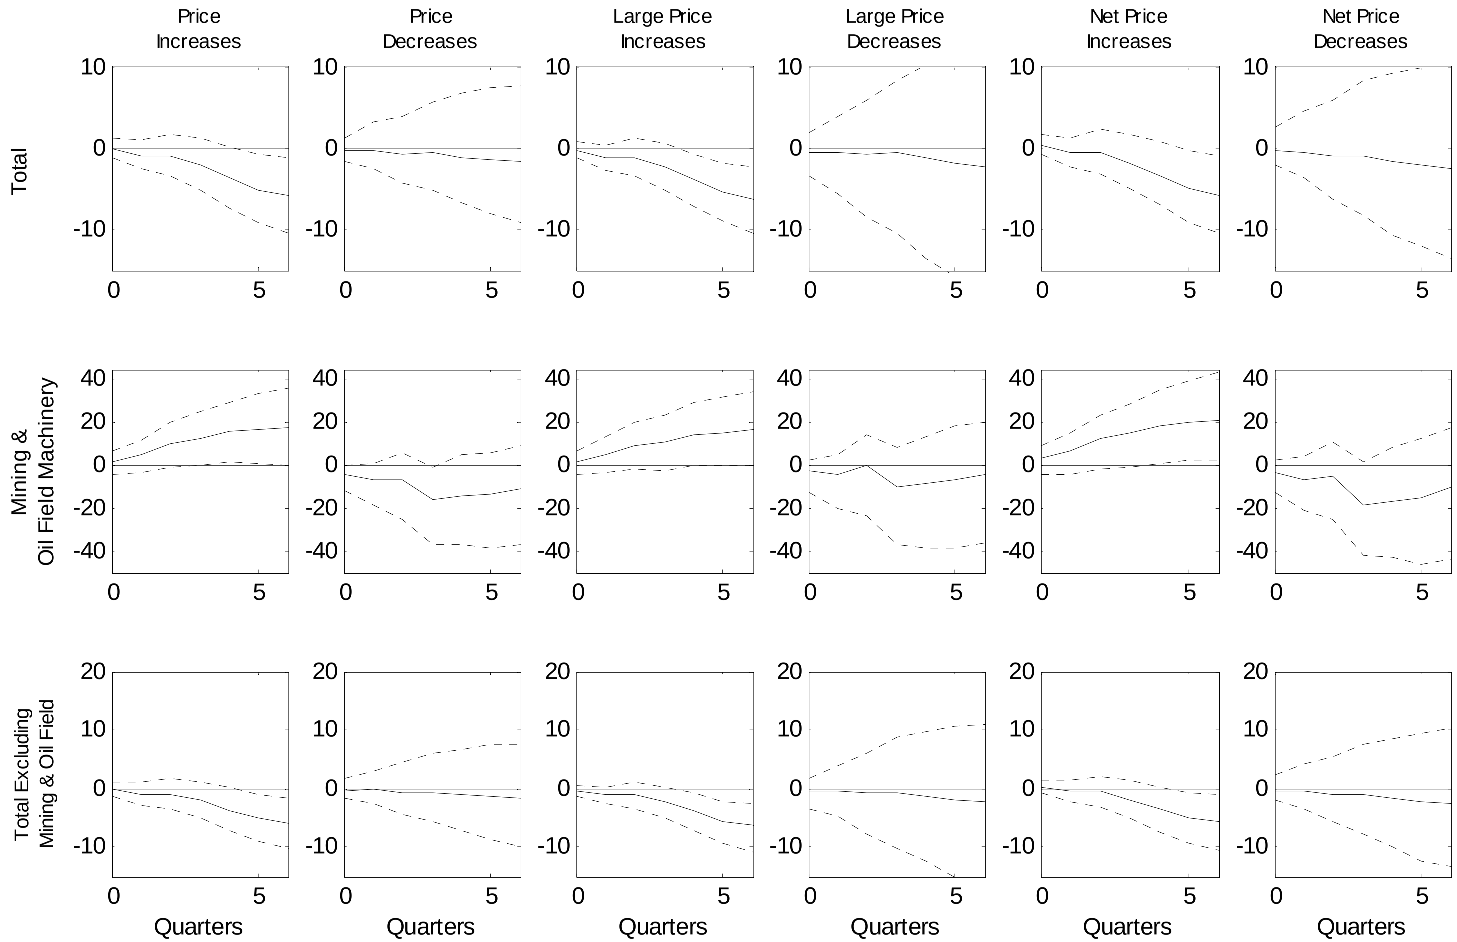

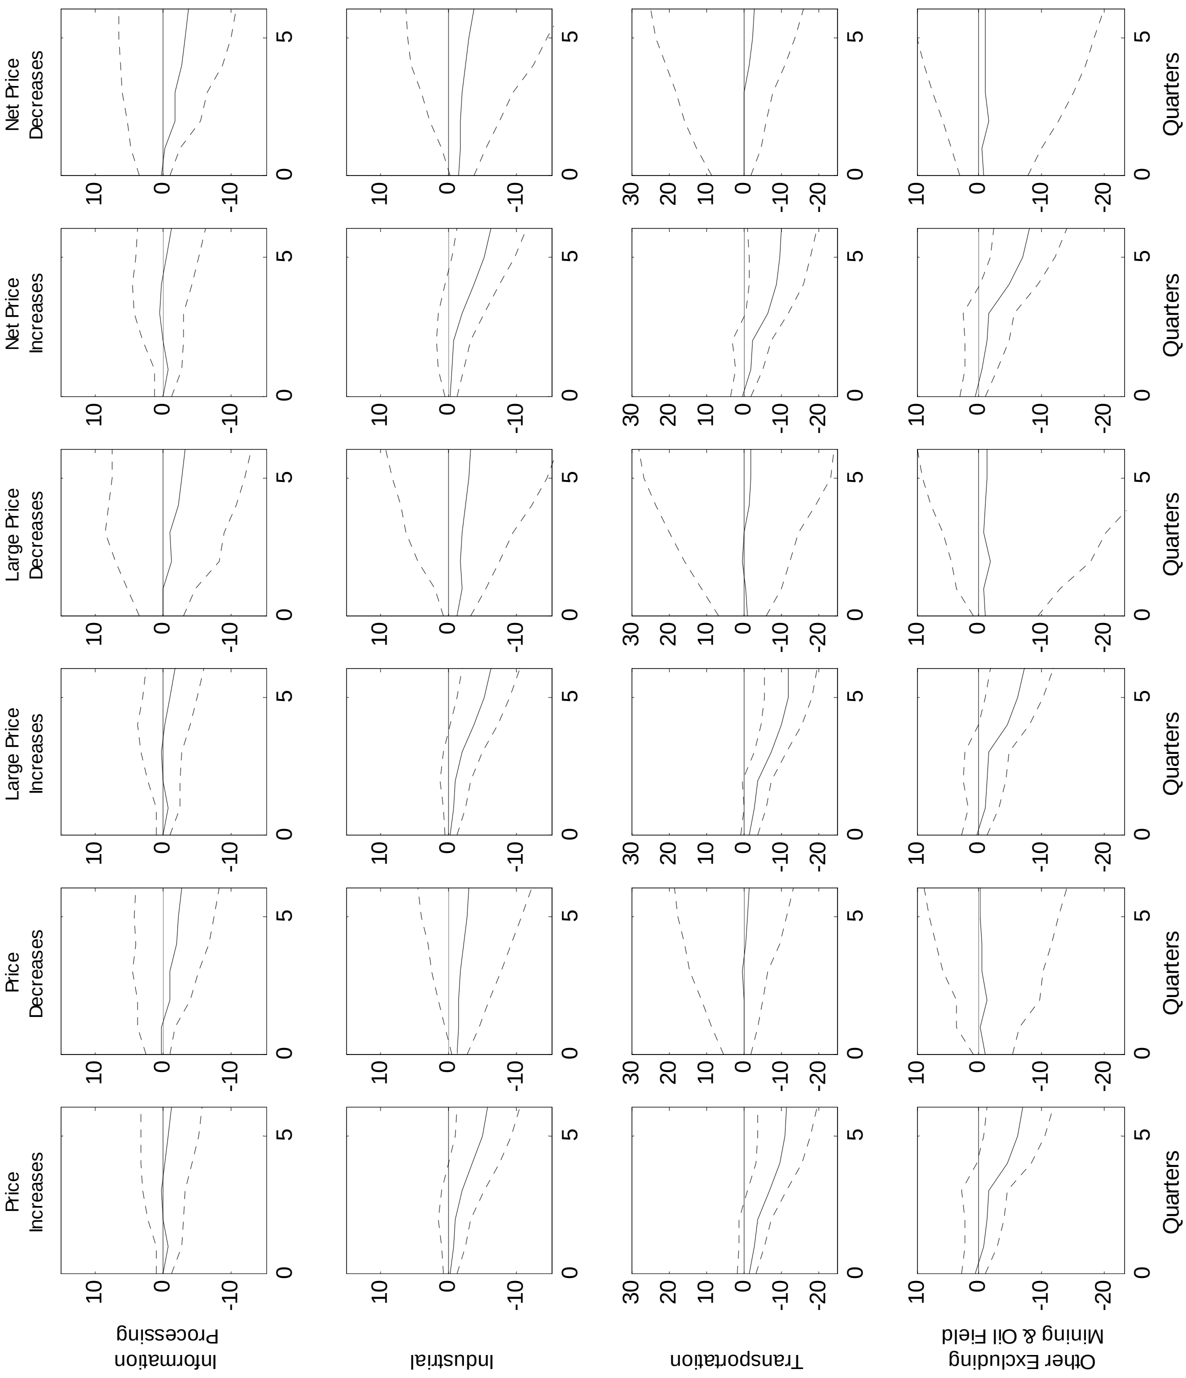

Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Information Processing
Industrial
Transportation
Other Excluding Mining & Oil Field
Quarters

Figure 4.A.2d: Response of Real Nonresidential Fixed Investment in Equipment: 1970.II-1987.IV with 1986 Dummies

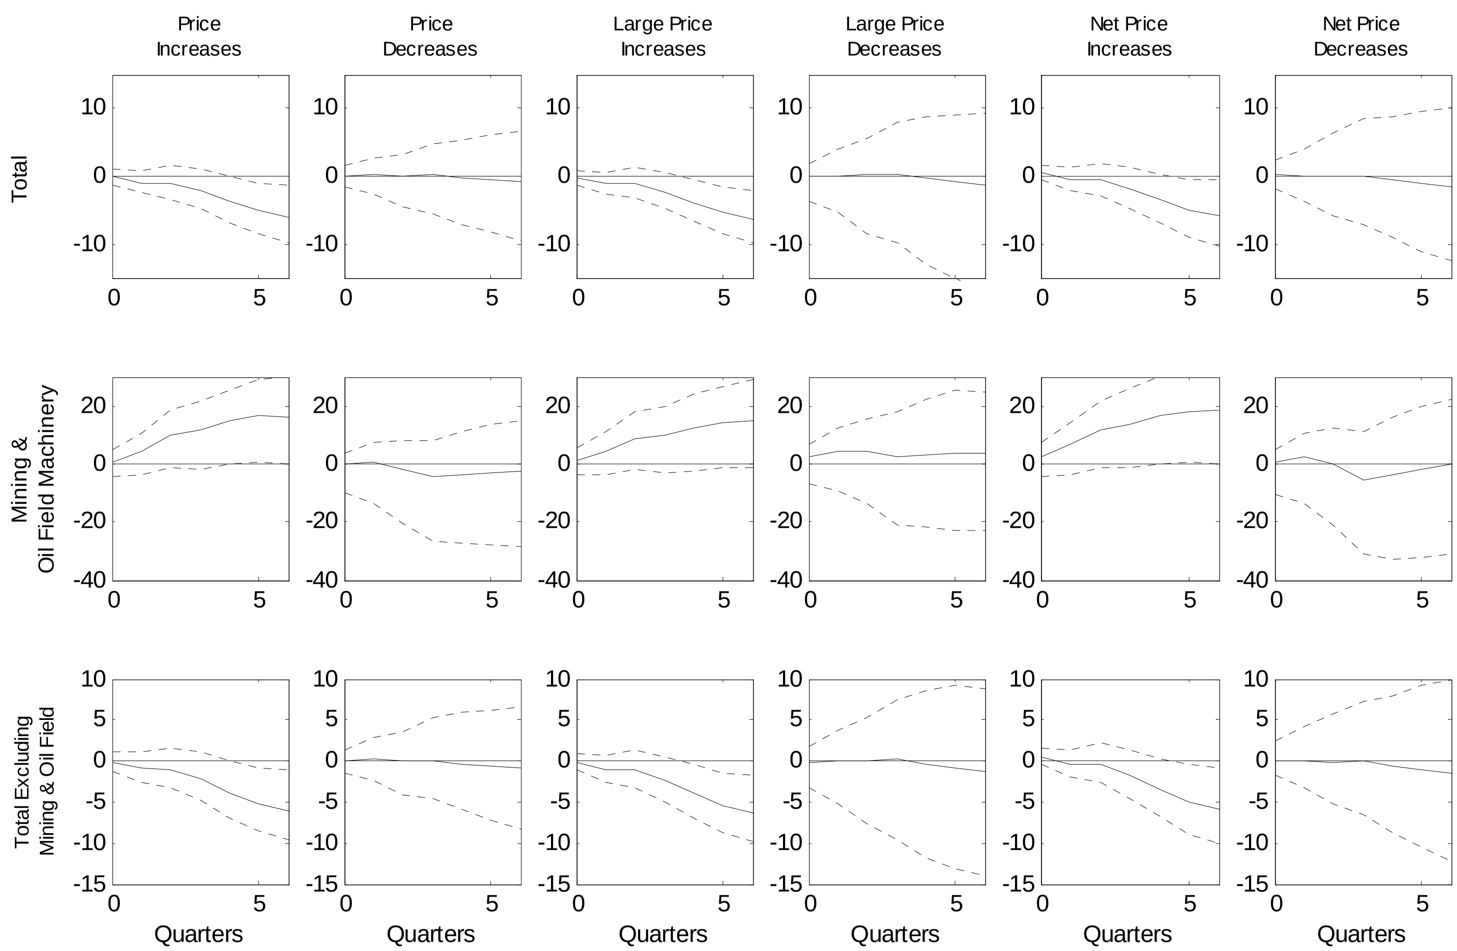

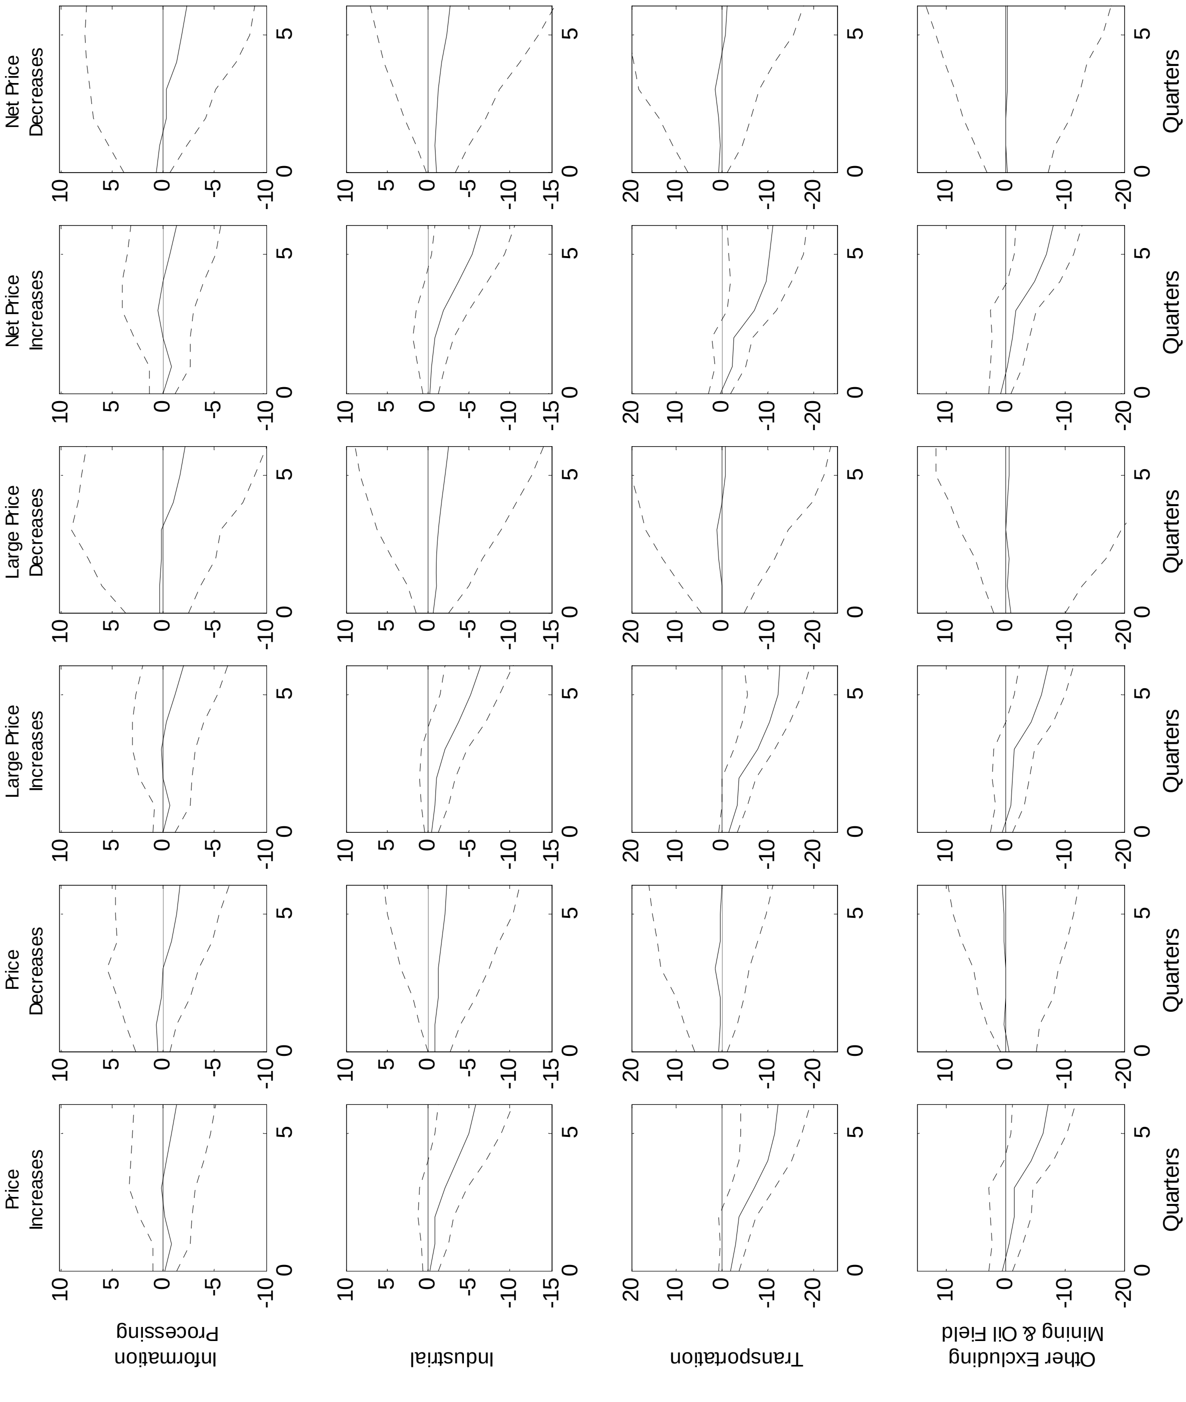

Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Information Processing
Industrial
Transportation
Other Excluding Mining & Oil Field
Quarters

Figure 4.A.2e: Response of Real Nonresidential Fixed Investment in Equipment: 1970.II-2006.IV with 1986 Dummies

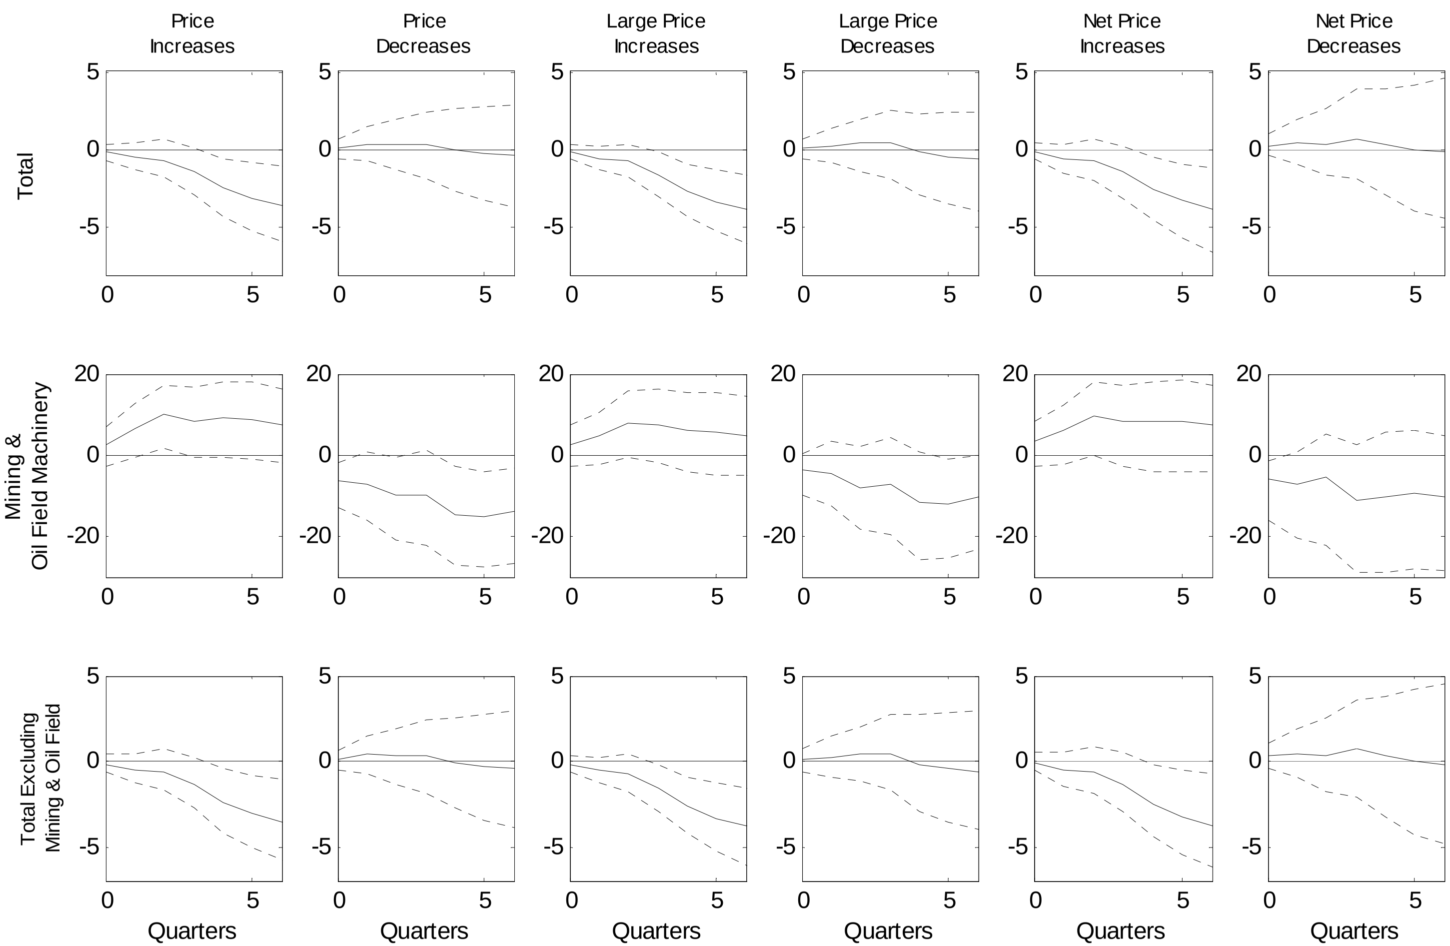

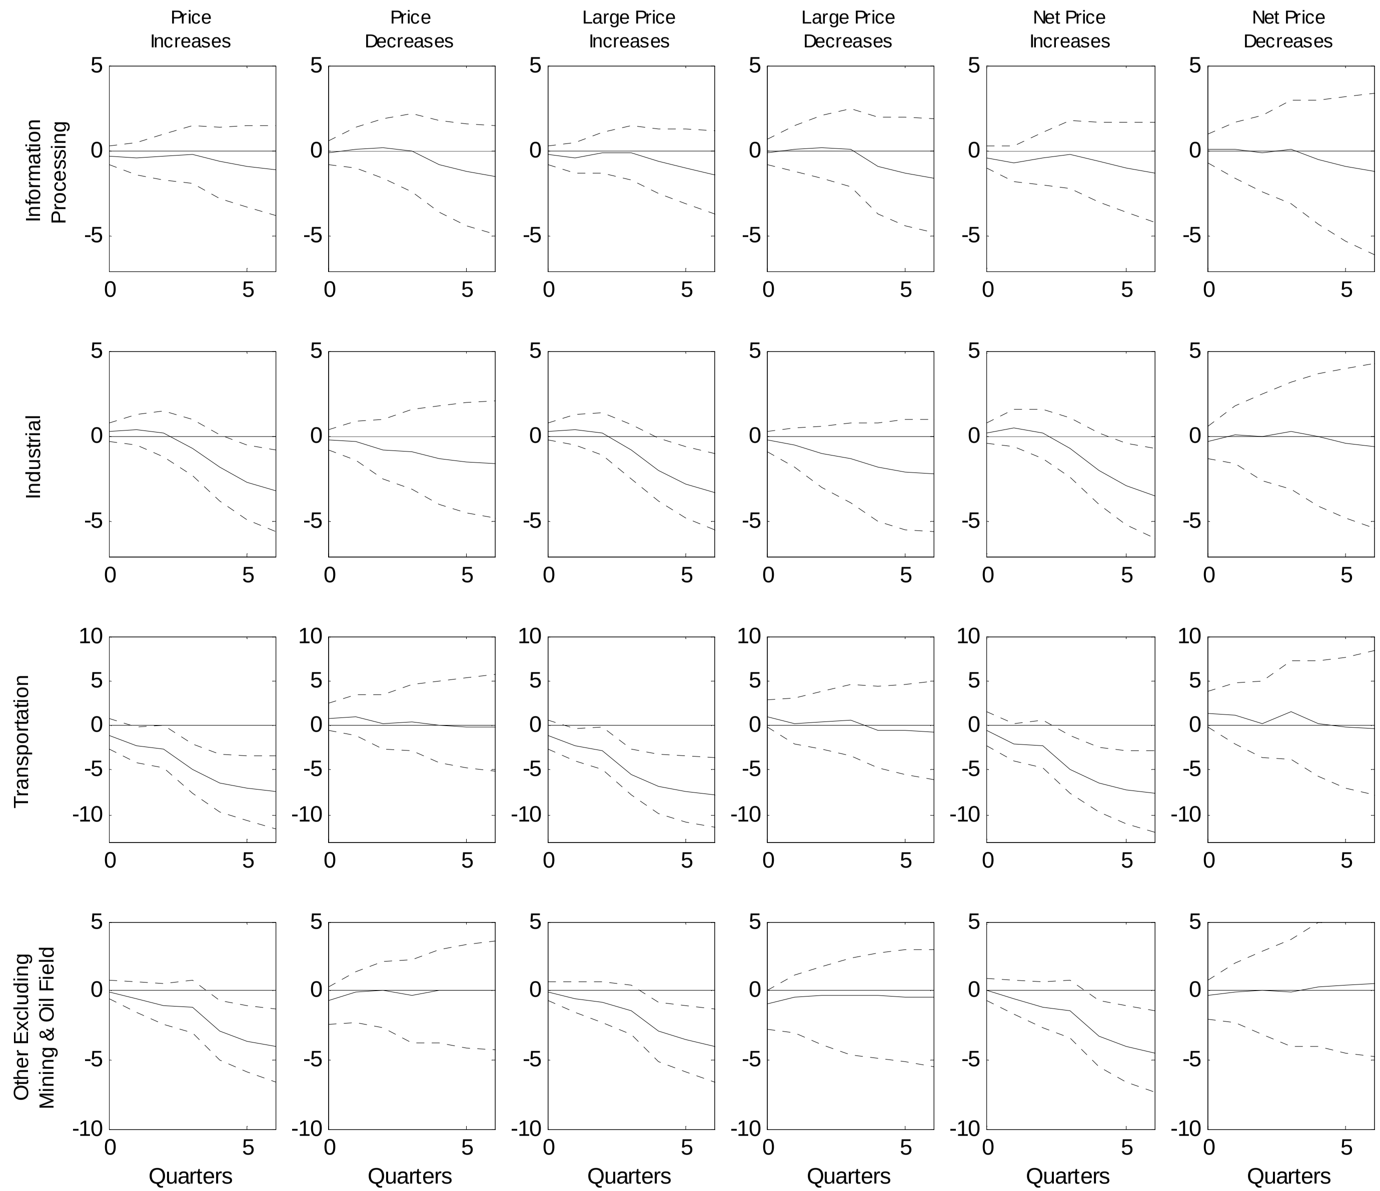
Price Increases
Price Decreases
Large Price Increases
Large Price Decreases
Net Price Increases
Net Price Decreases
Information Processing
Industrial
Transportation
Other Excluding Mining & Oil Field
Quarters

BIBLIOGRAPHY

[1] 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.

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

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

[4] Bernanke, B.S. (2006), "Semiannual Monetary Policy Report to the Congress," testimony to the Committee on Financial Services, U.S. House of Representatives, February.

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

[6] 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.

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

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

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

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

[11] Hamilton, J.D. (2005), "Oil and the Macroeconomy," forthcoming: *New Palgrave Dictionary of Economics.*

[12] Herrera, A.-M. (2006), "Inventories, Oil Shocks and Macroeconomic Behavior," mimeo, Department of Economics, Michigan State University.

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

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

[15] 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.

[16] Kilian, L. (1998), "Small-Sample Confidence Intervals for Impulse Response Functions," *Review of Economics and Statistics*, 80(2), 218-230.

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

[18] 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.

[19] 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.

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

[21] Rotemberg, J.J., and M. Woodford (1996), "Imperfect Competition and the Effects of Energy Price Increases on Economic Activity," *Journal of Money, Credit, and Banking*, 28, 549-577.

[22] U.S. Census Bureau (2005), *2002 Economic Census. Mining. Subject Series*, http://www.census.gov/econ/census02/guide/SUBSUMM.HTM.

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

CHAPTER V

Conclusion

This thesis demonstrates that, contrary to findings in the existing literature, commodity prices help predict inflation even in more recent years. In fact, commodity prices contain information about future inflation not captured by the leading principle components of a broader set of macroeconomic and financial data. These improved inflation forecasts, however, do not provide professional economists with a better understanding of monetary policy.

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.

This thesis also suggests that it is necessary to revise our thinking about the transmission of energy price shocks to the economy. Researchers typically believe that frictions associated with sectoral reallocations are important channels through which energy price shocks are transmitted to the economy. These frictions in turn imply that real consumption, real investment, and aggregate unemployment will respond asymmetrically to energy price increases and decreases. We find no compelling evidence for asymmetries in the response of these variables, suggesting that sectoral reallocations are not a significant transmission mechanism. This does not mean, however, that the effects of energy price shocks are necessarily small. In addition to a discretionary income effect, we find that energy price shocks result in a sizeable precautionary savings effect and a sizeable operating cost effect on energy-using durable goods. The extent to which each of these channels is important depends on the type of expenditure in question. Besides mining-related investment in structures and equipment, the types of expenditure that matter most for the transmission of energy price shocks to the economy include consumer and firm expenditures on motor vehicles and residential fixed investment in housing.