Asymmetric Responses of Consumer Spending to Energy Prices: A Threshold VAR Approach

Edward S. Knotek II and Saeed Zaman

We document asymmetric responses of consumer spending to energy price shocks: Using a multiple-regime threshold vector autoregressive model estimated with Bayesian methods on US data, we find that positive energy price shocks have a larger negative effect on consumption compared with the increase in consumption in response to negative energy price shocks. For large shocks, the cumulative consumption responses are three to five times larger for positive than for negative shocks. Digging into disaggregated spending, we find that the estimated asymmetric responses are strongest for durable goods, but asymmetries are also present in the responses of nondurables and services.

Edward S. Knotek II is at the Cleveland Fed (edward.knotek@clev.frb.org). Saeed Zaman is at the Cleveland Fed and the University of Strathclyde, UK (saeed.zaman@clev.frb.org). The authors thank Christiane Baumeister, Mark Bils, Oli Coibion, Jim Hamilton, Ana Maria Herrera, Kurt Lunsford, Rob Vigfusson, their discussants Nida Çakir Melek and Karel Mertens, and conference participants at the 5th IAAE Annual Conference, the CEBRA Workshop for Commodities and Macroeconomics, the Federal Reserve Bank of Kansas City Day-Ahead System Energy Meeting, and the 25th Annual SNDE Conference for helpful comments, discussions, and suggestions.

I   Introduction

The nominal price per barrel of crude oil fell by 65% from June 2014 to December 2015, and overall, the prices of consumer energy goods and services as measured by the personal consumption expenditures price index fell by 23% over this period. A conventional view of the relationship between energy prices and consumer spending suggests that increases in oil prices slow consumer spending, akin to tax increases; in linear versions of this conventional view, a large decline in energy prices would be similar to a large tax cut, which would provide a strong boost to consumer spending. Thus, it is notable that—in real time—analysts, forecasters, and energy experts were generally surprised by the lack of a sizable response of consumer spending to the declines in energy prices; see, e.g., Hamilton (2015).[1]

There is a vast literature examining the relationship between the U.S. macro economy and oil price fluctuations (see, e.g., Hamilton 2005 and Kilian 2008 for surveys), and a key question in this literature has been whether the empirical evidence favors asymmetries and nonlinearities in this relationship. Surprisingly, while most of the impact of energy price shocks on economic growth is thought to flow through adjustments in consumer spending (e.g., Bernanke 2006), relatively little research has focused on potential asymmetries in the response of consumer spending to oil and energy prices. Using measures of positive oil price changes (following Mork 1989) or net oil price increases (following Hamilton 1996, 2003), Mehra and Petersen (2005) present empirical evidence suggesting that consumer spending responds

---

[1] More generally, Baumeister and Kilian (2016b) note that movements in oil prices continue to surprise economists, policymakers, consumers, and financial market participants.

asymmetrically to oil price increases and decreases.[2] More recently, Alsalman and Karaki (2019) use the nonlinear simultaneous equation model of Kilian and Vigfusson (2011a) to investigate the nonlinear relationship between consumer spending and oil prices and find strong evidence of asymmetries and nonlinearities in the response of consumption to oil price shocks. By contrast, Edelstein and Kilian (2009) argue that one cannot reject the null hypothesis of symmetry in the impulse response functions of consumer spending to energy price shocks; thus, their analysis focuses on symmetric linear models.

This paper provides empirical evidence of asymmetries in the response of consumer spending to energy price increases and decreases in the context of a multivariate, multiple-regime threshold vector autoregression (TVAR). TVARs provide a parsimonious and flexible methodology to investigate potential asymmetric and nonlinear relationships, as regimes endogenously evolve in response to changing conditions. Because the existence of multiple regimes raises the potential for overfitting due to parameter proliferation, we estimate the TVAR with Bayesian methods. When we use monthly U.S. data on real consumption expenditures, energy inflation, and inflation excluding food and energy prices over an estimation sample spanning 1959-2014, the Bayesian measure of model fit—marginal likelihood—favors the two-regime model over a single regime, thus allowing for the possibility of nonlinearities and asymmetric responses of consumption to energy price shocks. The threshold variable separating our regimes is based on movements in real energy prices averaged over the past nine months.[3]

---

[2] Wang (2013) uses a panel smooth transition regression model to examine the relationship between oil price changes and consumer spending in the G7 countries and documents evidence in support of nonlinear and asymmetric relationships.

[3] Koop and Potter (2010) document support for a regime-switching threshold model among the many potential candidates in their examination of the nature of the relationship between real GDP growth and oil prices. In that study, the regime-switching threshold model that was most favored by the data was based on a threshold defined as oil price inflation averaged over the past four quarters.

Within the context of our model, this assumption implies that a sustained period of relatively high or low real energy inflation—or a large enough positive or negative shock—would lead to different consumption dynamics. With two regimes, our posterior mean estimates put the threshold at 0.36%. Such a threshold estimate is highly intuitive: effectively, consumers behave in two different fashions depending on whether real energy inflation is positive or negative. In response to small energy price shocks, the economy remains near the threshold separating the two regimes and the cumulative responses of consumption to positive and negative energy price shocks show modest asymmetries. In response to larger shocks, the asymmetries are large: at the posterior median, the decline in consumption after one year in response to a 5% increase in energy prices is up to three times as large in absolute terms as consumption's response to a 5% decrease in energy prices, while the ratio increases to five in response to 10% increases and decreases in energy prices.[4]

Digging into the components of consumption, we find marked asymmetric responses of durable goods spending; asymmetric responses to nondurable goods spending and services spending are present but less pronounced.

The energy literature has noted a number of channels that could generate asymmetric economic responses, including an uncertainty channel that can cause consumers to postpone purchases of durables in response to heightened uncertainty from falling energy prices (Bernanke 1983) and a reallocation channel through which declines in energy prices have smaller impacts than increases in energy prices because the reallocation of resources in response to declines in energy prices offsets their economic benefits (Hamilton 1988).

---

[4] For GDP, Kilian and Vigfusson (2011b) find no evidence of asymmetries for small shocks and statistically significant evidence of asymmetries for large shocks; see also Hamilton (2011).

The paper is structured as follows. The next section describes the related literature. Section III describes the data and the empirical models. Section IV discusses the empirical results. Section V reports the results for the components of consumption. Section VI concludes.

II   Related Literature

The vast majority of the literature that has investigated the relationship between energy prices (or oil prices) and aggregate macroeconomic variables has focused on industrial production, real GDP, real stock returns, unemployment, and payroll employment. With the exceptions of Mehra and Petersen (2005), Edelstein and Kilian (2009), Wang (2013), and Alsalman and Karaki (2019), which are the works most closely related to this paper, much less attention has focused on investigating the direct empirical relationship between aggregate consumption (including its disaggregates) and energy prices.[5] This is surprising given that theoretical justifications of the channels contributing to asymmetries and nonlinearities in the relationship between energy prices and the macro economy primarily stem from the dynamics of consumer spending along with firms' spending behavior (e.g., Hamilton 1988; Hamilton 2005).

There are multiple theoretical reasons why economic activity in the U.S. should respond asymmetrically to energy price increases and decreases. Arguably, the two main theoretical explanations for such asymmetries include the uncertainty effect and the reallocation effect; see Edelstein and Kilian (2009), Kilian (2008), or Kilian and Vigfusson (2011b) for more detailed discussions on the channels through which energy prices may impact consumer spending. The

---

[5] Baumeister and Kilian (2016a) provide an extended discussion of the response of consumer spending to the decline in energy prices over 2014-2015. Gelman et al. (2016) use micro data to study consumption responses in the context of the 2014-2015 decline in energy prices. These analyses do not focus on potential asymmetries.

uncertainty effect posits that unexpected changes in energy prices may help lead to uncertainty about the future trajectory of energy prices, which, in turn, may cause consumers to postpone purchases of consumer durables; see Bernanke (1983), Pindyck (1991), and Hamilton (2009). Edelstein and Kilian (2009) contrast this type of wait-and-see behavior in response to an increase in uncertainty with precautionary savings behavior, in which any change in energy prices generates an increase in uncertainty over future employment prospects that weighs on consumer spending; this precautionary savings behavior may be more plausible for big shocks than for small shocks.[6] The reallocation effect refers to allocative disturbances that lead to intersectoral and intrasectoral shifts that cause capital and labor to be unemployed and thus reduce consumption, as in Hamilton (1988). For example, reduced spending on consumer durables will lead to capital and labor shifting away from the auto sector (or changing production within the auto sector to different types of vehicles), but given the presence of frictions in both the capital and labor markets it will temporarily cause unemployment resulting in a further pull-back in consumption. Leduc et al. (2016) posit that consumption responses would depend on consumers' expectations regarding the persistence of the oil price changes: consumers respond to oil price shocks that they perceive to be persistent compared with temporary movements. Using a macro model of learning, they show that consumers in late 2014 perceived the sharp oil price decline—at least initially—as temporary.[7]

---

[6] This precautionary savings behavior would contribute to asymmetric consumption responses, because the precautionary behavior would dampen the positive impacts on spending from energy price declines and augment the negative impacts on spending from energy price increases.

[7] Aladangady and Sahm (2015) show that the households in the University of Michigan's Survey of Consumers who expected the sharp decline in gasoline prices observed in the second half of 2014 to reverse were less enthusiastic about the survey question on whether it is a good time to buy durable goods such as automobiles.

Despite these theoretical arguments for asymmetric effects, the empirical evidence on the existence of such a relationship is far from settled. Starting in the late 1980s, a considerable body of research documented empirical evidence in favor of asymmetries and nonlinearities in the relationship between oil prices and the U.S. macro economy; see, e.g., Lougani (1986); Mork (1989); Lee et al. (1995); Hooker (1996a, 1996b); Davis and Haltiwanger (2001); Balke et al. (2002); Hamilton (1996, 2003); Huang et al. (2005); and Herrera et al. (2011).[8] Kilian and Vigfusson (2011a, 2011b) provide evidence to question the results of those earlier studies, with the main critique targeting the econometric methodology and diagnostic tests used to assess the degree of asymmetry.[9] This has led to reexaminations of the previous studies using alternative econometric methodologies, which have collectively found no statistically significant evidence of asymmetric relationships between oil prices and the macro economy; see, e.g., Kilian (2014); Herrera et al. (2011, 2015); Alsalman and Herrera (2015); and Herrera and Karaki (2015).[10] Koop and Potter (2010) use an extremely flexible econometric framework to examine the relationship between real GDP growth and oil price inflation. Their framework uncovers strong support for a regime-switching threshold model, suggesting strong evidence in favor of a nonlinear relationship between oil prices and aggregate economic activity.

TVARs have been widely used in other applications, some of which have involved oil or energy prices. Univariate versions of threshold time series models have existed at least going

---

[8] International evidence in favor of a nonlinear relationship between oil prices and real GDP has been documented by Cunado and Perez de Gracia (2003) for a host of European countries; Jimenez-Rodrigueza and Sanchez (2005) for Japan, Norway and the UK; Kim (2012) for six industrialized countries; Engemann et al. (2011) for OECD countries; and Donayre and Wilmot (2016) for Canada. Herrera et al. (2015) find weak evidence of asymmetry in industrial production across OECD countries.

[9] Hamilton (2011) suggests that the relationship continues to be nonlinear for large shocks, pointing to the results of statistically significant nonlinearity for 2 standard deviation shocks in Kilian and Vigfusson (2011b).

[10] Herrera et al. (2011) find mixed evidence: using a longer sample that starts prior to 1973, they find evidence of asymmetry and nonlinearity, but that evidence becomes much weaker with data starting post-1973.

back to Tong (1978), but multivariate models have increased in popularity since Balke (2000), who investigated the role of credit in contributing to nonlinear dynamics in a model of aggregate output and credit conditions. Huang et al. (2005) use a TVAR with two regimes to study the relationship between real stock returns (and industrial production) and oil prices; the threshold variable is based on movements in the nominal price of energy. Van Robays (2016) uses macroeconomic uncertainty, which is strongly associated with oil price volatility, as the threshold variable in a TVAR to study the relationship between oil prices and macroeconomic uncertainty.

III   Empirical Models and Data

Linear vector autoregressions (VARs) typically provide a useful starting point for assessing the economic impacts of energy price shocks; e.g., Edelstein and Kilian (2009) use a bivariate VAR in consumer spending and purchasing power changes related to energy price fluctuations. As a baseline, we use trivariate VARs where the vector Y_t contains energy inflation, π_t^{energy}, inflation excluding food and energy prices, π_t^{core}, and real consumption growth, $\Delta \ln c_t$.[11] A general representation of a linear VAR(p) model is:

$$Y_t = A_c + A_1 Y_{t-1} + ... + A_p Y_{t-p} + u_t \tag{1}$$

where $t=1,\ldots,T$, A_c is a vector of constants, $A_1,\ldots,A_p$ are matrices of coefficients, and the error vector u_t is normally distributed with mean zero and variance-covariance matrix $\Sigma = E u_t u_t'$.

---

[11] Clark and Terry (2010) use a four-variable system of real GDP, energy inflation, core inflation, and the federal funds rate to study time variation in the pass-through of energy shocks to core inflation. Our results are unchanged by the addition of the federal funds rate.

VARs are susceptible to overfitting, and so to deal with this curse of dimensionality we use Bayesian estimation methods as discussed in Banbura et al. (2010), Koop (2013), and Carriero et al. (2015). With a stationary model, the priors on the coefficients are set to zero for all variables, and we impose otherwise standard Minnesota and sum-of-coefficient priors in the estimation.

To compute impulse response functions, we identify the energy price shock through Cholesky factorization, in which energy inflation is ordered first followed by core inflation and then real consumption growth. This recursive identification scheme assumes that shocks to energy prices affect other variables contemporaneously, but other shocks affect energy inflation only with a lag. We work with monthly data, and this exogeneity assumption over this timeframe is consistent with the findings of Kilian and Vega (2011) and a long line of empirical work; e.g., Edelstein and Kilian (2009) make a similar assumption. In our main results, we identify the "average" energy price shock in our VAR, abstracting from energy supply and energy demand shocks as in Kilian (2009); we consider robustness below. Given this linear model, computation of the impulse response functions is straightforward and standard.

The linear VAR model in Equation (1) imposes restrictions that may prove to be inappropriate for our investigation. Crucially, the structural impulse response functions from the linear VAR model impose symmetry in the responses of the system to a structural shock, both on impact and thereafter: the responses to a positive structural shock are mirror images of the responses to a negative structural shock. Second, the responses to a shock necessarily follow the principle of proportionality; the impact of a 10% structural shock is exactly twice the response to a 5% structural shock. Third, the responses of the linear VAR to a structural shock do not depend on when the structural shock occurs—i.e., the responses are state invariant.

9

A large literature beginning with Hamilton (1996) has argued that consumers (and firms) may respond more to big movements in energy prices than to small movements in energy prices, suggesting that a constant linear VAR model as in Equation (1) may not be an appropriate framework in this context.[12] For additional flexibility, a natural extension is to allow for the possibility that VAR dynamics may be determined by different regimes that vary over time with some endogenous state. In particular, consider the nonlinear threshold VAR (TVAR) that allows for two regimes:

$$Y_t = \left[A_{c,1} + A_{1,1}Y_{t-1} + ... + A_{p,1}Y_{t-p} + u_{t,1} \right] D_t + \left[A_{c,2} + A_{1,2}Y_{t-1} + ... + A_{p,2}Y_{t-p} + u_{t,2} \right](1 - D_t) \quad (2)$$

where

$$D_t = \begin{cases} 1 & \text{if } Y_{t-d}^* \leq \gamma^* \\ 0 & \text{if } Y_{t-d}^* > \gamma^* \end{cases} \quad (3)$$

and where the errors $u_{1,t}$ and $u_{2,t}$ have variance-covariance matrices Σ_1 and Σ_2, respectively. In this TVAR, the regimes are determined by the endogenous state D_t: when $Y_{t-d}^* \leq \gamma^*$, the VAR dynamics are (potentially) governed by a different set of coefficients and shock variances compared with when $Y_{t-d}^* > \gamma^*$. The unobserved threshold value is γ^*, and the threshold variable is denoted Y_{t-d}^*. In our application, we define the threshold variable in terms of a moving average of lagged real energy inflation—i.e., energy inflation minus core inflation.

$$Y_{t-d}^* = \frac{1}{K} \sum_{j=1}^{K} \left(\pi_{t-j}^{energy} - \pi_{t-j}^{core} \right). \quad (4)$$

---

[12] Researchers have dealt with this issue by using a nonlinear transformation of the oil price in a constant linear VAR, commonly denoted as a censored linear VAR. This nonlinear transformation is supposed to remove oil price declines and only measure oil price increases relative to some reference period (e.g., 1, 2, or 3-year peaks) and has usually taken the following form (e.g., Hamilton 2003): $x_t = \max\{0, X_t - \max\{X_{t-1},...,X_{t-k}\}\}$ where X is the price of oil and x is the nonlinear transformation of the oil price X.

Our results are similar if we use energy inflation instead of real energy inflation as the threshold variable. As a baseline, we use a nine-month moving average to calculate the threshold variable, $K=9$.[13] We set the parameter delay lag $d=1$.[14] It is worth noting that with a delay lag $d=1$, in response to a shock the regime can only switch after at least a one-month delay. For practical purposes, this means that, on impact, positive and negative shocks have symmetric effects until the regime can switch, suggesting that asymmetries found with this approach are likely to actually be conservative, a point we return to below.[15]

Estimating the TVAR with Bayesian methods and producing impulse responses are computationally intensive. A priori, we assume the regimes are identical and we impose the same Minnesota and sum-of-coefficient priors used earlier in the linear VAR case, and our prior for the threshold value is the sample mean for real energy inflation. We estimate the TVAR model by following the Gibbs sampling procedure detailed in Alessandri and Mumtaz (2017) and Chib and Greenberg (1995). For each simulation of the Gibbs sampling, their procedure includes a Metropolis-Hastings step to sample the threshold value. Using the estimation results along with the data sample, we compute generalized impulse responses in the spirit of Koop et al. (1996), but with our identified structural shocks, as in Kilian and Vigfusson (2011a). Unlike

---

[13] The use of a moving average process for the threshold variable is relatively common to have more persistence in the switching variable and to overcome month-to-month volatility in the underlying raw series. A rough grid search over the number of lags p and the number of months K to include in the moving average process found that $p=9$ and $K=9$ maximized the marginal likelihood of the sample.

[14] While it is possible in principle to estimate d, we fix it to reduce the computational burden.

[15] Kilian and Lütkepohl (2016, p. 596) note that the rationale for a "hard" threshold (which produces an abrupt transition) is more compelling at the microeconomic level than at the macro level. Because the thresholds governing consumer behavior would likely vary across individuals, in the aggregate a smooth threshold may therefore be more plausible, such as a regime-switching smooth transition regression based on a logistic transition function (STRL). The empirical studies using STRL (e.g., Granger et al. 1993; Weise 1999; Shintani et al. 2013) on the macro aggregates have estimated large values for the smoothness parameter of the transition function—where a larger value implies a less smooth or more abrupt transition—lending support to our use of a hard threshold. Hence, we view the use of a TVAR for our application as a reasonable approach.

in the linear VAR case, in computing responses from this nonlinear TVAR model it is important to account for the initial state of the system when the shock hits, the sign of the shock, and the size of the shock, because all three of these factors can affect the endogenous evolution of the regimes; e.g., a starting position close to the threshold value is more likely to be crossed than a starting position far from it. Thus, the impulse responses we show are broadly representative of the different potential starting points in the different regimes. Their construction requires simulation techniques; the technical appendix provides detailed estimation steps and the algorithm for constructing the impulse responses.

To take the linear VAR and multiple-regime TVAR to the U.S. data, our baseline results use monthly data beginning in January 1959 and ending in June 2014. We end our primary sample in June 2014 so that we do not double-count the most recent decline in energy prices, which we used to motivate this study based on expressed concerns about the lack of a strong boost in consumer spending from the declines in energy prices from mid-2014 through the end of 2015. Our results are little changed by extending the sample through later dates. We measure real consumer spending using monthly growth in real personal consumption expenditures (PCE). Our measure of energy inflation is monthly inflation in the PCE price index for energy goods and services. Finally, our measure of core inflation is monthly inflation in the PCE price index excluding food and energy. All three measures come from the Bureau of Economic Analysis.[16] With monthly data, we use $p=9$ lags in the linear VAR and the TVAR.

---

[16] In all cases, monthly growth or inflation rates are computed as 100 times the natural log first difference.

IV   Empirical Results

To compare and contrast the linear VAR and TVAR models, we proceed by examining model fit and the plausibility of the estimated threshold, and then we move on to analyzing the impulse response functions.

One advantage of estimating our models with Bayesian methods is the convenience of model comparisons in terms of model fit. Using the Bayesian metric of marginal likelihood (ML), we can rank model specifications in terms of the degree to which they fit the sample data (see the technical appendix for details). In our case, model fit also provides evidence on potential nonlinearities and asymmetries. If the linear VAR has a higher ML value compared to any of our TVAR specifications, then that would provide evidence in favor of a linear relationship; a higher ML value from a TVAR specification provides evidence in favor of nonlinearities. However, it does not determine the shape of the nonlinearities; e.g., the nonlinearities could be in energy inflation rather than the response of consumption to energy price shocks. Impulse response analysis is necessary to examine the nature of the nonlinearities.

Table 1 reports the logarithm of the ML for the specifications of the linear VAR and the baseline TVAR with two regimes (one threshold) and K=9 months for the moving average in the threshold variable. The ML criterion strongly favors having multiple regimes over a single regime. This finding holds for various values of K, the number of terms in the moving average. While not shown, based on the ML criterion, the data actually favor a TVAR model with three regimes over the two-regime specification in the table. The drawback to three regimes is the potential for one regime to have a relatively small number of observations, which turns out to be

13

the case in this preferred specification.[17] To avoid drawing inference on the basis of a small number of observations, we limit attention to the two-regime case.

In a threshold VAR, the estimated threshold value partitions the sample into the desired number of regimes depending on the behavior of the threshold variable. Table 1 also presents an estimate of the threshold value from our baseline TVAR specification. The posterior mean of the threshold is 0.36%. Such a value, when related to the threshold variable, appears plausible and is relatively intuitive: it essentially says that if real energy inflation has recently been zero or negative—i.e., $Y_{t-1}^{*} = (1/9)\sum_{j=1}^{9}(\pi_{t-j}^{energy} - \pi_{t-j}^{core}) \leq 0.36$ at the posterior mean—then the economy is in one regime, which we label generically regime 1; if it has recently been positive, then the economy is in regime 2. Large negative shocks are therefore likely to either push the economy into regime 1 or keep it there; large positive shocks are likely to do the same for regime 2.

Figure 1 shows a time plot of the threshold variable along with the posterior mean estimated threshold from the TVAR model. Prior to the 2000s, most of the time in the sample was spent in regime 1, with the exception of notable spikes in oil prices in the 1970s and the late 1980s/early 1990s associated with OPEC actions and the first Gulf War. Overall, 73% of observations lie below the estimated posterior mean threshold value, consistent with high real energy inflation being somewhat rare based on historical experience—although more common later in the sample.

To assess the behavior of consumer spending in response to energy price shocks, we turn to impulse response exercises. In the TVAR model specification, impulse responses can differ qualitatively and quantitatively based on the size of the shock, the sign of the shock, and the

---

[17] In the three regime TVAR estimation, the two estimated posterior mean thresholds remain near zero: one is −0.47% and the other is 0.35%. But only 10 percent of observations fall into regime 1.

starting position of the economy within one regime or another because the regime—and hence the model's dynamics—can endogenously evolve over time. As a result, the figures corresponding to impulse responses for TVARs show responses to both positive and negative shocks and to smaller and larger shocks. Specifically, to get a visual sense of the evidence of asymmetry and nonlinearity we report results for responses of consumption to ±1%, ±5%, and ±10% shocks to energy prices. While monthly shocks in the range of ±10% are rare, they are more plausible when viewed over the span of a few months.[18]

Figure 2 shows the cumulative consumption responses to ±1% energy price shocks in the linear VAR and conditional on the shock occurring in either regime 1 (low/negative real energy inflation) or regime 2 (positive real energy inflation); we report cumulative consumption responses at selected time periods in Table 2. In all cases, the responses in the linear VAR to the positive and negative shocks are simply mirror images of each other, because the model imposes symmetry by assumption. As noted above, on impact, the impulse responses within a given regime in the TVAR exercise are symmetric for positive and negative shocks because the regime is fixed within that period; it is only in the period after the shock occurs that the regime is able to change, which may generate nonlinearities and asymmetries. Our threshold approach may be underestimating cumulative consumption asymmetries to the extent that differential responses occur on impact to positive and negative shocks.

That said, in response to ±1% energy price shocks, evidence of asymmetries in the consumption response to energy price shocks is limited. In regime 1, a positive 1% energy price shock reduces consumption by a cumulative 0.15% after 12 months; by contrast, a negative 1%

---

[18] For example, the cumulative decline based on the monthly growth rate in PCE energy prices from October 2014 to March 2015 was -17.28%, and from October 2014 to December 2015 it was -23.24%.

energy price shock that occurs in regime 1 boosts cumulative consumption by 0.11% after 12 months. The differences are smaller in regime 2—the positive real energy inflation regime. Nevertheless, the TVAR model produces notable differences vis-à-vis the canonical linear VAR. The cumulative posterior median estimates for the linear VAR are consistently smaller, in absolute terms, than the cumulative posterior median estimates from the TVAR: the linear VAR tends to underestimate the sensitivity of consumer spending to relatively small energy price shocks, both positive and negative.

Larger energy shocks generate noticeable asymmetries. Figure 3 plots the impulse responses to ±5% energy price shocks. While the linear model's impulse response to a −5% shock produces an increase in cumulative consumer spending that is quantitatively similar to those from the TVAR model regardless of the starting regime, the estimated effects from a +5% shock in the linear model are much smaller compared with those from the TVAR. Starting from regime 1, a +5% shock to energy prices eventually results in a 0.93% decrease in cumulative consumer spending after 12 months in the TVAR specification—approximately three times as large as the linear model would predict. Clear evidence of regime dependence also emerges in response to larger shocks. Asymmetries are more pronounced in regime 1 between positive and negative shocks. After 12 months, a +5% shock generates a 0.93% decrease in cumulative consumer spending, while a −5% shock generates a 0.33% increase in cumulative consumer spending; the consumption response to the positive energy price shock is nearly three times as large in absolute terms as the response to the negative energy shock. In regime 2, the consumption response to a +5% energy price shock is less than twice as large in absolute terms as the response to a −5% energy price shock.

16

Qualitatively similar results are obtained from ±10% shocks in Figure 4. The posterior median from the linear VAR tends to understate the consumption response to a positive shock of this size while overstating the response to the large negative shock. A +10% energy price shock in regime 1 tends to have a large depressing effect on consumption, while a −10% shock provides a relatively small boost—such that the 70% probability bands for the cumulative consumption response include zero throughout the response horizon. As shown in Table 2, starting from regime 1, the consumption response to a +10% shock is five times larger in absolute value after 12 months than the response to a −10% shock.

Overall, these exercises suggest several conclusions. First, the data tend to favor the TVAR model based on the Bayesian metric of model fit, marginal likelihood. Second, the impulse responses coming from linear VARs understate the responsiveness of consumption to small energy price shocks; for larger energy price shocks, the linear VAR continues to understate the sensitivity of consumption to energy price increases and to overstate its sensitivity to energy price decreases. Third, while evidence of asymmetries in consumption responses is visible following small energy price shocks, the asymmetries are more prominent in response to relatively large energy price shocks.

V   Responses of Consumption Components and Robustness

The estimated asymmetric and nonlinear behavior between aggregate consumption and energy prices may not necessarily hold for each of the components of aggregate consumption. To get a sense of which of the underlying components of the consumption basket may be contributing to the observed asymmetric behavior of aggregate consumption, we repeat the above exercises by

17

examining the responses of the three main consumption components—consumer spending on durables, nondurables, and services—to energy price shocks.

In each case, we find that the marginal likelihood is maximized by the TVAR model estimated using Bayesian methods rather than the linear VAR model, and the posterior mean estimated thresholds are all incredibly similar—either 0.36% or 0.37%, compared with 0.36% in our baseline estimation with total aggregate consumption. The online appendix provides further details and a complete set of figures for the impulse responses. We compactly summarize these results by looking at the responses of real spending on durables in Table 3, real spending on nondurables in Table 4, and real spending on services in Table 5.

Qualitatively, the patterns of asymmetries are identical to those for total consumption across the three major consumption components. In all cases, the linear model understates the responsiveness of consumer spending on durables, nondurables, and services to small energy shocks, and it markedly understates the responsiveness of consumption at the component level in response to large positive energy shocks. We also document sizeable asymmetries in response to the same-sized positive and negative shocks: consumer spending on durables, nondurables, and services is more responsive to positive energy shocks than to negative energy shocks, especially in response to relatively larger shocks.

Quantitatively, the extent of the asymmetries differs somewhat at the component level. After 12 months, the decline in durable consumption in response to a +5% energy shock is 3.3 times larger than the response to a −5% shock when starting in regime 1 (−2.29% v. +0.70%). For nondurable consumption and services consumption, the comparable ratio is 2.7 times (−0.79% v. +0.29% for nondurables, and −0.70% v. +0.26% for services). As with the results

18

for aggregate consumer spending, the asymmetries are smaller when starting in regime 2; the comparable ratios are 1.7 for durables and 1.5 for both nondurables and services.

In other robustness exercises, we examine asymmetries in the relationship between consumer spending and energy prices in the period since 1985. As reported in the online appendix, we find that the asymmetries documented above are still present but are generally smaller.[19] In results not reported, we find that extending the sample to include recent periods of declining energy prices or augmenting the specification of the model to include inflation expectations and financial conditions does not affect our results, and we obtain similar results when replacing our energy price series with the purchasing power series in Edelstein and Kilian (2009). We also investigated asymmetries in the consumption response to energy supply and energy demand shocks, based on Kilian (2009). Using our TVAR framework estimated using Bayesian methods, we find that energy-specific demand shocks produce the strongest asymmetric responses of consumption, complementing the findings of Alsalman and Karaki (2019) using the Kilian and Vigfusson (2011a) nonlinear simultaneous modeling framework.

VI   Conclusion

This paper reconsiders the issue of potential asymmetries and nonlinearities in the relationship between consumer spending and energy prices. Using a multivariate threshold vector autoregressive model estimated with Bayesian methods, we document that the historical U.S.

---

[19] Using industrial production data, Herrera et al. (2011) find stronger evidence of asymmetries and nonlinearities using data before 1973 compared with more recent data. Hooker (1996a, 1996b), Mehra and Petersen (2005), and Edelstein and Kilian (2009) document that the relationship between oil prices and many macroeconomic variables has weakened over time.

data prefer multiple regimes. We document notable asymmetries of consumer spending in response to large energy shocks: consumer spending declines more in response to large energy price increases than it increases in response to the same-sized energy price decreases. By contrast, evidence of asymmetries is more limited in response to small shocks. Notably, we provide evidence that the linear model not only misses the asymmetries, but it also underestimates the sensitivity of consumption to small shocks (both negative and positive) as well as large positive shocks. Finally, we find qualitatively similar results for spending on consumer durables, nondurables, and services.

VII   Technical Appendix

Linear VAR Estimation

   In the linear VAR model of Eq. (1), coefficient estimates in $A_1, ..., A_p$ and Σ are shrunk to their prior means. The prior beliefs for the mean and variances of the coefficient matrices are as follows:

$$E[A_k^{(i,j)}] = \begin{cases} \delta_i, & if \; i = j, k = 1 \\ 0, & otherwise \end{cases}$$

$$Var[A_k^{(i,j)}] = \lambda^2 \frac{1}{l^2} \frac{\sigma_i^2}{\sigma_j^2}, \qquad l = 1, ..., p$$

(A.1)

Since all variables are transformed to growth rates (i.e., they are stationary) prior to estimation, we set $\delta_i = 0$ for all variables. The scale factor $1/l^2$ helps impose the prior belief that recent lags play a more influential role compared to the distant ones by proportionally shrinking the variances on the more distant lags (centered on the prior mean of zero). The prior parameter σ_i

equals the standard deviation of the residuals obtained from regressing the variable y_i on its own

p lags and a constant over the sample period up to time t. The hyperparameter λ governs the

tightness of our prior beliefs. As $\lambda \to 0$ the prior dominates and so the posterior equals the prior,

i.e., the data have no say; as $\lambda \to \infty$, the prior has no influence and so the posterior estimates

converge to OLS estimates. We set λ to 0.2, a value often used in the literature.[20]

Imposing sum-of-coefficient priors on the equations of the VAR can be helpful in

improving forecast accuracy. The parameter μ governs the tightness of this prior and is set to 1,

indicating a very loose prior. Prior specification for each equation is symmetric in its treatment

of own lags of the dependent variable and lags of other variables. As a result, we have a prior

that is a natural conjugate (normal-inverted Wishart prior), which proves to be convenient when

solving for the model as these priors can be implemented easily by augmenting the data matrices

with fictitious data, permitting OLS to estimate the model.

For notational convenience, we re-write the linear VAR model of Eq. (1) in a compact

matrix notation,

$$Y = XA + U \qquad (A.2)$$

where $Y = [y_1, \dots, y_T]'$ is a $T \times N$ matrix, and $X = [x_1, \dots, x_T]'$ is a $T \times k$ matrix.

$A = [A_c \ A_1 \ \dots \ A_p]'$ is a $k \times N$ matrix of VAR coefficients. $U = [u_1 \ u_2 \ \dots \ u_T]'$ is a

$T \times N$ matrix of error terms.

Let $\alpha = vec(A)$ of size $kN \times 1$.

Then, the natural conjugate prior or normal inverse-Wishart prior (N-IW) is specified as,

---

[20] In principle, the value of these prior parameters could be obtained through a grid search that maximizes the marginal likelihood. The value we are using happens to be close to values that other researchers have obtained through a grid search optimization.

$$\alpha|\Sigma \sim N\left(\underline{\alpha}, \Sigma \otimes \underline{\Omega}\right) \qquad (A.3)$$

and $$\Sigma \sim iW\left(\underline{S}, \underline{v}\right) \qquad (A.4)$$

where $\underline{\alpha} = vec(\underline{A})$, $\underline{\Omega}$, $\underline{S}$, and $\underline{v}$ are prior hyperparameters whose values are set using the prior moment conditions specified in Eq. (A.1).

The resulting joint posterior density, $p(\alpha, \Sigma | Y)$ also has the form N-IW.

$$\alpha|\Sigma, Y \sim N\left(\overline{\alpha}, \Sigma \otimes \overline{\Omega}\right) \qquad (A.5)$$

and $$\Sigma|Y \sim iW\left(\overline{S}, \overline{v}\right) \qquad (A.6)$$

where $\overline{\alpha} = vec(\overline{A})$

$$\overline{A} = \left(\underline{\Omega}^{-1} + X'X\right)^{-1}\left(\underline{\Omega}^{-1}\underline{A} + XY'\right) \qquad (A.7)$$

$$\overline{\Omega} = \left(\underline{\Omega}^{-1} + X'X\right)^{-1} \qquad (A.8)$$

$$\overline{v} = \underline{v} + T \qquad (A.9)$$

$$\overline{S} = \underline{A} + \hat{u}'\,\hat{u} + \hat{A}'X'X\hat{A} + \underline{A}'\,\underline{\Omega}^{-1}\underline{A} - \hat{A}'\overline{\Omega}^{-1}\hat{A} \qquad (A.10)$$

where $\hat{A}$ is the OLS estimate and $\hat{u}$ are OLS residuals defined as $\hat{A} = (X'X)^{-1}X'Y$ and $\hat{u} = Y - X\hat{A}$ respectively.

Computation of the Marginal Likelihood in a Linear VAR

The marginal likelihood is,

$$p(Y) = \int p(Y|\theta)p(\theta)d\theta$$

where θ collects all the VAR model parameters. An advantage of working with an N-IW prior is that a closed-form solution exists for $p(Y)$ and it is,

$$p(Y) = \left(\frac{1}{\pi}\right)^{\frac{NT}{2}} \times \left|(I + X\underline{\Omega}X')^{-1}\right|^{\frac{N}{2}} \times |\underline{S}|^{\frac{v}{2}} \times \frac{\Gamma_N\left(\frac{v+T}{2}\right)}{\Gamma_N\left(\frac{v}{2}\right)}$$

$$\times \left|\underline{S} + (Y - X\underline{A})'(I + X\underline{\Omega}X')^{-1}(Y - X\underline{A})\right|^{-\frac{v+T}{2}} \qquad (A.11)$$

Impulse Response in a Linear VAR

To compute impulse response functions, we require an identification strategy to solve for the energy price shock. We adopt a Cholesky factorization method to identify the energy price shock, in which energy inflation is ordered first, followed by core inflation and then real consumption growth. This recursive identification scheme assumes that shocks to energy prices affect other variables contemporaneously, but other shocks affect energy inflation only with a lag.

Formally, the impulse responses are constructed by a simple Monte Carlo procedure, which sequentially samples from the model's conditional posterior densities defined in Eq. (A.5), and Eq. (A.10). For each posterior draw of parameters, A and Σ, the corresponding structural coefficients are computed using the Cholesky method (which involves triangular decomposition of the Σ). The structural coefficients along with the desired magnitude of the energy shock are used to construct impulse response functions. Cycling through all the posterior draws provides us with an empirical posterior distribution of the impulse response functions, which are used to compute the posterior median response and the posterior probability intervals.

TVAR Estimation

To permit Bayesian inference in the TVAR model, a posterior simulator using an MCMC algorithm is developed. Specifically, we use the Metropolis-within-Gibbs algorithm identical to the one used by Alessandri and Mumtaz (2017), which in turn is based on the Gibbs sampler of Chen and Lee (1995).

<u>A</u>. Prior specification: Define priors for all the TVAR model parameters

    a.   For the TVAR parameters, $A = \{A_{r1}, A_{r2}\}$ and $\Sigma = \{\Sigma_{r1}, \Sigma_{r2}\}$, the prior assumptions are identical to the one specified for the parameters, A and Σ in the linear VAR model, i.e., Eq. (A.3) and Eq. (A.4). Note, A_{r1} and Σ_{r1} corresponds to the VAR parameters and error variance matrix in regime 1, and similarly, A_{r2} and Σ_{r2} correspond to the VAR parameters and error variance matrix in regime 2.

    b.   For the threshold parameter, γ^* , a normal prior is assumed: $\gamma^* \sim N(\underline{\gamma}, \underline{v})$, where $\underline{\gamma}$ is the mean of the threshold variable Y_t^*, and $\underline{v} = 10$. Set the starting value, $\gamma^{*\,(0)} = \underline{\gamma}$.

<u>B</u>: Develop a posterior simulator: In order to perform posterior inference in the TVAR model in Eq. (2), we require sampling from the joint posterior density, $p(A_{r1}, \Sigma_{r1}, A_{r2}, \Sigma_{r2}, \gamma^*|Y)$. It can be seen that conditional on γ^* and with the assumption of conjugate N-IW priors, the TVAR model specified in Eq. (2) collapses into two separate linear VAR models with conjugate N-IW priors. Hence, using the rules of probability, the joint posterior density can be written as

$$p(A_{r1}, \Sigma_{r1}, A_{r2}, \Sigma_{r2}, \gamma^*|Y) = p(A_{r1}, \Sigma_{r1}|Y_{r1}, \gamma^*) \; p(A_{r2}, \Sigma_{r2}|Y_{r2}, \gamma^*) \, p(\gamma^*|Y).$$

The first two components on the right-hand side are the joint posterior densities of VAR parameters in regime 1 and regime 2, respectively, and the third component is the conditional

posterior density (or the marginal posterior density) for the threshold parameter. Drawing from the joint posterior density of VAR parameters corresponding to regime 1, $p(A_{r1}, \Sigma_{r1}|Y_{r1}, \gamma^*)$ is equivalent to drawing first from the conditional posterior, $p(A_{r1}|Y_{r1}, \Sigma_{r1}, \gamma^*)$, which is normal, and then drawing from the conditional posterior, $p(\Sigma_{r1}|Y_{r1}, A_{r1}, \gamma^*)$, which is inverse-Wishart. Therefore, drawing from the joint posterior density, $p(A_{r1}, \Sigma_{r1}|Y_{r1}, \gamma^*)$ is straightforward as it takes a convenient form of N-IW density. Similarly, drawing from joint posterior $p(A_{r2}, \Sigma_{r2}|Y_{r2}, \gamma^*)$ is straightforward. However, the conditional posterior density $p(\gamma^*|Y)$ does not take the form of any standard density suggesting the need for a Metropolis-Hastings step to draw from this conditional posterior. Accordingly, to simulate posterior density of the TVAR parameters, the posterior draws are obtained via an MCMC algorithm that sequentially cycles through the following four steps, for g=1 to G draws.

Step 1: Split the estimation sample.

Conditional on the draw of γ^*, split the estimation sample into $r1, r2$ regimes

Regime 1 (denoted $r1$): $Y_{t-1}^* \leq \gamma^*$

$$Y_{r1} = X_{r1} A_{r1} + U_{r1} \qquad (A.12)$$

Set $\alpha_{r1} = vec(A_{r1})$ and Σ_{r1} is the error covariance matrix in regime 1

Regime 2 (denoted $r2$): $Y_{t-1}^* > \gamma^*$

$$Y_{r2} = X_{r2} A_{r2} + U_{r2} \qquad (A.13)$$

Set $\alpha_{r2} = vec(A_{r2})$ and Σ_{r2} is the error covariance matrix in regime 2

Step 2: Draw from $p(\alpha_{r1}|\Sigma_{r1}, Y_{r1}, \gamma^*)$ and $p(\Sigma_{r1}|Y_{r1}, A_{r1}, \gamma^*)$ details below.

Step 3: Draw from $p(\alpha_{r2}|\Sigma_{r2}, Y_{r2}, \gamma^*)$ and $p(\Sigma_{r2}|Y_{r2}, A_{r2}, \gamma^*)$ details below.

For notational sake, let's denote the regimes $r = \{r1, r2\}$, then

25

$$\alpha_r | \Sigma_r, Y_r, \ \gamma^* \sim N\left(\overline{\alpha}_r, \ \Sigma_r \otimes (X^{*'} X^*)^{-1}\right) \qquad (A.14)$$

where $X^* = \left[X_r; \ X_{dummy}\right]$ and $\overline{\alpha}_r = vec(\overline{A}_r)$

$$\Sigma_r | Y_r, \gamma^* \ \sim iW\left(\overline{S}_r, \overline{v}_r\right) \qquad (A.15)$$

where $\overline{v}_r = \ \underline{v} + T_r \ ; T_r$ represents the number of observations in regime r

$$\overline{S}_r = \underline{A}_r + \hat{u}_r{}' \hat{u}_r + \hat{A}_r{}' X' X \hat{A}_r + \underline{A}_r{}' \underline{\Omega}_r^{-1} \underline{A}_r \ - \ \hat{A}_r{}' \overline{\Omega}_r^{-1} \hat{A}_r$$

where $\underline{\Omega} = \ \underline{\Omega}_r$ and $\underline{A} = \ \underline{A}_r$; $\hat{u}_r = Y_r - \ X_r \hat{A}_r$

Note: Symbols with a bar on top represent the posterior mean estimates; those with a bar below represent prior mean estimates; and those with a hat on top represent OLS estimates.

<u>Step 4</u>: Draw from $\ p(\gamma^* | Y)$

Conditional on the draw of parameters sampled in steps 2 and 3 above, a candidate draw for the threshold parameter γ^* is obtained using the random walk Metropolis-Hastings step, which generates candidate draws as,

$$\gamma^*_{new} = \ \gamma^*_{old} + \ \Psi^{1/2} v, \quad v \sim N(0,1) \qquad (A.16)$$

where Ψ is the scaling factor whose value is chosen to keep the acceptance probability in the range of 20 to 40 percent.

Acceptance probability, $\ \varrho(\gamma^*_{old}, \gamma^*_{new}) = min\left[\frac{p(\gamma^* = \gamma^*_{new} \mid Y, \Theta)}{p(\gamma^* = \gamma^*_{old} \mid Y, \Theta)}, 1\right] \qquad (A.17)$

where p(.) represents the posterior density of γ^*; Θ denotes all other TVAR model parameters; the *min* operator ensures that the acceptance probability does not exceed 1.

Deriving the posterior density of $\ \gamma^*$:

The Bayes theorem implies

$$p(\gamma^* | Y) \ \propto p(Y | \gamma^*) p(\gamma^*) \qquad (A.18)$$

26

and
$$p(Y|\gamma^*) = p(Y_{r1}|\gamma^*)\, p(Y_{r2}|\gamma^*) \qquad (A.19)$$

where $p(Y_{r1}|\gamma^*)$ is the marginal likelihood for the linear VAR process of regime 1, and

$p(Y_{r2}|\gamma^*)$ is the marginal likelihood for the linear VAR process of regime 2.

The above relationship implies

$$p(\gamma^* = \gamma^*_{new}\ |Y)\ \propto p(Y|\gamma^* = \gamma^*_{new})p(\gamma^*) \qquad (A.20)$$

Similarly,

$$p(\gamma^* = \gamma^*_{old}\ |Y)\ \propto p(Y|\gamma^* = \gamma^*_{old})p(\gamma^*) \qquad (A.21)$$

Accordingly,

$$\frac{p(\gamma^* = \gamma^*_{new}\ |Y)}{p(\gamma^* = \gamma^*_{old}\ |Y)}\ \propto\ \frac{p(Y|\gamma^* = \gamma^*_{new})}{p(Y|\gamma^* = \gamma^*_{old})} \qquad (A.23)$$

where $p(Y|\gamma^* = \gamma^*_{new})$ is the marginal likelihood of the TVAR model evaluated at $\gamma^* = \gamma^*_{new}$

and $p(Y|\gamma^* = \gamma^*_{new})$ is the marginal likelihood of the TVAR model evaluated at $\gamma^* = \gamma^*_{old}$

Note: $p(\gamma^* = \gamma^*_{new}\ |Y) = p(\gamma^* = \gamma^*_{new}\ |Y,\Theta)$ and $p(\gamma^* = \gamma^*_{old}\ |Y) =$

$p(\gamma^* = \gamma^*_{old}\ |Y,\Theta)$

Therefore, the acceptance probability is the ratio of the marginal likelihoods of the TVAR

model evaluated at thresholds γ^*_{new} and γ^*_{old}.

We set G=11,000, and use the first 10,000 posterior draws as a burn-in sample (i.e.,

discard draws) and store the remaining 1,000 posterior draws for posterior inference.

Marginal Likelihood in a Threshold VAR

The marginal likelihood of the TVAR conditional on threshold γ^* is

$$p(Y|\gamma^*) = p(Y_{r1}|\gamma^*)\, p(Y_{r2}|\gamma^*) \qquad (A.24)$$

$$p(Y_{r1}|\gamma^*) = \left(\frac{1}{\pi}\right)^{\frac{NT_{r1}}{2}} \times \left|(I + X_{r1}\underline{\Omega}X_{r1}')^{-1}\right|^{\frac{N}{2}} \times \left|\underline{S}\right|^{\frac{v}{2}} \times \frac{\Gamma_N\left(\frac{v + T_{r1}}{2}\right)}{\Gamma_N\left(\frac{v}{2}\right)}$$

$$\times \left|\underline{S} + (Y_{r1} - X_{r1}\underline{A})'(I + X_{r1}\underline{\Omega}X_{r1}')^{-1}(Y_{r1} - X_{r1}\underline{A})\right|^{-\frac{v + T_{r1}}{2}} \qquad (A.25)$$

$$p(Y_{r2}|\gamma^*) = \left(\frac{1}{\pi}\right)^{\frac{NT_{r2}}{2}} \times \left|(I + X_{r2}\underline{\Omega}X_{r2}')^{-1}\right|^{\frac{N}{2}} \times \left|\underline{S}\right|^{\frac{v}{2}} \times \frac{\Gamma_N\left(\frac{v + T_{r2}}{2}\right)}{\Gamma_N\left(\frac{v}{2}\right)}$$

$$\times \left|\underline{S} + (Y_{r2} - X_{r2}\underline{A})'(I + X_{r2}\underline{\Omega}X_{r2}')^{-1}(Y_{r2} - X_{r2}\underline{A})\right|^{-\frac{v + T_{r2}}{2}} \qquad (A.26)$$

To obtain the marginal likelihood $p(Y)$ we need to evaluate the $p(Y|\gamma^*)$ for every possible threshold and then average them

$$p(Y) = \int p(Y|\gamma^*)p(\gamma^*)d\gamma^* \qquad (A.27)$$

TVAR Impulse Response Construction

To construct the impulse response functions from the TVAR specifications, we closely follow the steps detailed in Alessandri and Mumtaz (2017), and Chiu and Hoke (2016) which in turn are based on Koop et al. (1996).[21] We quantify the response of the system to an energy price shock of size(s) φ. In particular, we choose φ from the set $\{-10\%, -5\%, -1\%, 1\%, 5\%, 10\%\}$. Analyzing the responses of the system to shocks of various magnitudes and signs helps assess the degree of asymmetry and nonlinearity in the relationship between energy prices and other variables of the model. Lets denote the *regimes* by $\{1, ..., R\}$, and let H be the number of

---

[21] Koop et al. (1996) developed the algorithm for generating the nonlinear VAR impulse response functions, and that procedure has been referred to as generalized impulse response functions. As mentioned in Kilian and Lütkepohl (2017), the focus of Koop et al. (1996) was on reduced-form impulse response functions, not structural impulse response functions. We modify the procedure of Koop et al. (1996) along the lines of Kilian and Vigfusson (2011a) to generate structural impulse response functions, using Cholesky ordering to identify the shocks.

months of responses that are of interest after the initial period t when the shock is realized. We then use the following algorithm to construct the impulse responses.

For each of the energy price shock, φ from $\{-10\%, -5\%, -1\%, 1\%, 5\%, 10\%\}$ do

  For each of the regimes, $r = 1,...,R$ do

    For each of the *retained* draws of Gibbs sampling, $i = 1,...,G$ do

      For each of the history corresponding to the data sample falling under the regime r, *state=1,..,T* do

1. Pick a block of 'p' consecutive actual values of all the variables. This block constitutes the history (or the initial state: lagged values) at time t, denoted as the His_{t-1}

2. For *reps=1:Reps* do

3. Simulate two forecast paths (unconditional and conditional on the energy price shock) of the variables of the system for *horizon=0,..,H* months

4. The first simulated forecast path is constructed by randomly drawing forecast errors for each variable of the system and for each *horizon* from an independent Gaussian distribution with mean zero and variance covariance matrix, $\Sigma^{r,i}$. Along with these error draws, and VAR parameter draw i, the evolution of $Y_{t+horizon}$ over the $H+1$ periods is constructed. This resulting path of $Y_{t+horizon}$ is the unconditional forecast and is denoted as $Y_{t+horizon}(Hist_{t-1})$.

5. The second simulated forecast path is constructed by first applying Cholesky decomposition to the error draw obtained in the previous step 4 to orthogonalize the shocks which allows us to solve for the energy price shock. The value of the energy price shock is replaced with the energy price shock of interest φ and the entire forecast error matrix is reconstructed (i.e., responses are rescaled and normalized). Along with this updated forecast error matrix, and VAR parameter draw i, the evolution of $Y_{t+horizon}$ over the $H+1$ periods is constructed. This resulting path of $Y_{t+horizon}$ is the forecast conditional on a one-period energy price shock and is denoted as $Y_{t+horizon}(\varphi_t , Hist_{t-1})$

6. The draw of the generalized impulse responses of interest is the difference between:

   $GIRF_Y{}^{r,i,state,reps}(\varphi_t) = Y_{t+horizon}(\varphi_t , Hist_{t-1}) - Y_{t+horizon}(Hist_{t-1})$

   End reps loop

29

7. Averaging over all the *reps* gives the consistent estimate of the GIRF corresponding to Gibbs draw *i,* shock φ_t and conditional on a particular history of regime *r*: $GIRF_Y^{\,r,i,state}(\varphi_t)$

End state loop

8. Average over all the histories (i.e., initial states) corresponding to regime *r*, which gives the average GIRF corresponding to regime *r* and Gibbs draw *i* for an energy price shock φ_t : $GIRF_Y^{\,r,i}(\varphi_t)$

End *i* loop

9. We end up with *G* generalized impulse responses (i.e., a distribution of GIRFs). Averaging over all these *G* draws of the generalized impulse responses obtained in step 8 gives us the average generalized impulse response (posterior response) for an energy price shock φ_t conditional on being in regime *r*: $GIRF_Y^{\,r}(\varphi_t)$. Similarly, we use the distribution of the GIRF to construct the 70% probability bands. We show both the posterior GIRF median response and the 70% probability bands in the figures shown in the paper.

End *r* loop

End φ loop

References

Aladangady, Aditya, and Claudia R. Sahm (2015) "Do Lower Gasoline Prices Boost Confidence?" *FEDS Notes* 2015 (1501). https://doi.org/10.17016/2380-7172.1501.

Alessandri, Piergiorgio, and Haroon Mumtaz (2017) "Financial Conditions and Density Forecasts for US Output and Inflation." *Review of Economic Dynamics* 24 (March): 66–78. https://doi.org/10.1016/j.red.2017.01.003.

Alsalman, Zeina N., and Ana María Herrera (2015) "Oil Price Shocks and the U.S. Stock Market: Do Sign and Size Matter?" *The Energy Journal* 36 (3): 171–88. https://doi.org/10.5547/01956574.36.3.7.

Alsalman, Zeina N., and Mohamad B. Karaki (2019) "Oil Prices and Personal Consumption Expenditures: Does the Source of the Shock Matter?" *Oxford Bulletin of Economics and Statistics* 81 (2): 250–70. https://doi.org/10.1111/obes.12276.

Balke, Nathan S. (2000) "Credit and Economic Activity: Credit Regimes and Nonlinear Propagation of Shocks." *Review of Economics and Statistics* 82 (2): 344–49. https://doi.org/10.1162/rest.2000.82.2.344.

Balke, Nathan S., Stephen P.A. Brown, and Mine K. Yucel (2002) "Oil Price Shocks and the U.S. Economy: Where Does the Asymmetry Originate?" *The Energy Journal* 23 (3): 27–52. https://doi.org/10.5547/ISSN0195-6574-EJ-Vol23-No3-2.

Bańbura, Marta, Domenico Giannone, and Lucrezia Reichlin (2010) "Large Bayesian Vector Auto Regressions." *Journal of Applied Econometrics* 25 (1): 71–92. https://doi.org/10.1002/jae.1137.

Baumeister, Christiane, and Lutz Kilian (2016a) "Lower Oil Prices and the U.S. Economy: Is This Time Different?" *Brookings Papers on Economic Activity* 2016 (2): 287–357. https://doi.org/10.1353/eca.2016.0029.

Baumeister, Christiane, and Lutz Kilian (2016b) "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *Journal of Economic Perspectives* 30 (1): 139–60. https://doi.org/10.1257/jep.30.1.139.

Bernanke, Ben S. (1983) "Irreversibility, Uncertainty, and Cyclical Investment." *The Quarterly Journal of Economics* 98 (1): 85–106. https://doi.org/10.2307/1885568.

Bernanke, Ben S. (2006) "Energy and the Economy" Speech. Board of Governors of the Federal Reserve System (U.S.). https://www.federalreserve.gov/newsevents/speech/bernanke20060615a.htm.

Carriero, Andrea, Todd E. Clark, and Massimiliano Marcellino (2015) "Bayesian VARs: Specification Choices and Forecast Accuracy: BAYESIAN VARS." *Journal of Applied Econometrics* 30 (1): 46–73. https://doi.org/10.1002/jae.2315.

Chen, Cathy W. S., and Jack C. Lee (1995) "Bayesian Inference of Threshold Autoregressive Models." *Journal of Time Series Analysis* 16 (5): 483–92. https://doi.org/10.1111/j.1467-9892.1995.tb00248.x.

Chib, Siddhartha, and Edward Greenberg (1995) "Understanding the Metropolis-Hastings Algorithm." *The American Statistician* 49 (4): 327–35. https://doi.org/10.1080/00031305.1995.10476177.

Chiu, Ching-Wai (Jeremy), and Sinem Hacioglu Hoke (2016) "Macroeconomic Tail Events with Non-Linear Bayesian VARs." Bank of England working paper 611. Bank of England. https://econpapers.repec.org/paper/boeboeewp/0611.htm.

Clark, Todd E., and Stephen J. Terry (2010) "Time Variation in the Inflation Passthrough of Energy Prices." *Journal of Money, Credit and Banking* 42 (7): 1419–33. https://doi.org/10.1111/j.1538-4616.2010.00347.x.

Cuñado, Juncal, and Fernando Pérez de Gracia (2003) "Do Oil Price Shocks Matter? Evidence for Some European Countries." *Energy Economics* 25 (2): 137–54. https://doi.org/10.1016/S0140-9883(02)00099-3.

Davis, Steven J., and John Haltiwanger (2001) "Sectoral Job Creation and Destruction Responses to Oil Price Changes." *Journal of Monetary Economics* 48 (3): 465–512. https://doi.org/10.1016/S0304-3932(01)00086-1.

Donayre, Luiggi, and Neil A. Wilmot (2016) "The Asymmetric Effects of Oil Price Shocks on the Canadian Economy." *International Journal of Energy Economics and Policy* 6 (2): 167–82.

Edelstein, Paul, and Lutz Kilian (2009) "How Sensitive Are Consumer Expenditures to Retail Energy Prices?" *Journal of Monetary Economics* 56 (6): 766–79. https://doi.org/10.1016/j.jmoneco.2009.06.001.

Engemann, Kristie M., Kevin L. Kliesen, and Michael T. Owyang (2011) "Do Oil Shocks Drive Business Cycles? Some U.S. and International Evidence." *Macroeconomic Dynamics* 15 (S3): 498–517. https://doi.org/10.1017/S1365100511000216.

Gelman, Michael, Yuriy Gorodnichenko, Shachar Kariv, Dmitri Koustas, Matthew Shapiro, Dan Silverman, and Steven Tadelis (2016) "The Response of Consumer Spending to Changes in Gasoline Prices." w22969. Cambridge, MA: National Bureau of Economic Research. https://doi.org/10.3386/w22969.

Granger, Clive W., Timo Terasvirta, and Heather M. Anderson (1993) "Modeling Nonlinearity over the Business Cycle." In *Business Cycles, Indicators, and Forecasting*, edited by James H. Stock and Mark W. Watson, 311–26. Chicago: University of Chicago Press. https://ideas.repec.org/h/nbr/nberch/7196.html.

Hamilton, James D. (1988) "A Neoclassical Model of Unemployment and the Business Cycle." *Journal of Political Economy* 96 (3): 593–617. https://doi.org/10.1086/261553.

Hamilton, James D. (1996) "This Is What Happened to the Oil Price-Macroeconomy Relationship." *Journal of Monetary Economics* 38 (2): 215–20. https://doi.org/10.1016/S0304-3932(96)01282-2.

Hamilton, James D. (2003) "What Is an Oil Shock?" *Journal of Econometrics* 113 (April): 363–98. https://doi.org/10.1016/S0304-4076(02)00207-5.

Hamilton, James D. (2008) "Oil and the Macroeconomy." In *The New Palgrave Dictionary of Economics*, edited by Steven N. Durlauf and Lawrence E. Blume, 2nd ed. Basingstoke: Palgrave Macmillan. https://doi.org/10.1057/9781137336583.1307.

Hamilton, James D. (2009) "Causes and Consequences of the Oil Shock of 2007-08." *Brookings Papers on Economic Activity* 2009: 215–83. https://doi.org/10.1353/eca.0.0047.

Hamilton, James D. (2011) "Nonlinearities and the Macroeconomic Effects of Oil Prices." *Macroeconomic Dynamics* 15 (S3): 364–78. https://doi.org/10.1017/S1365100511000307.

Hamilton, James D. (2015) "Energy Prices and Consumer Spending." *Econbrowser* (blog). May 12, 2015. http://econbrowser.com/archives/2015/05/energy-prices-and-consumer-spending.

Herrera, Ana María, and Mohamad B. Karaki (2015) "The Effects of Oil Price Shocks on Job Reallocation." *Journal of Economic Dynamics and Control* 61 (December): 95–113. https://doi.org/10.1016/j.jedc.2015.08.006.

Herrera, Ana María, Latika Gupta Lagalo, and Tatsuma Wada (2011) "Oil Price Shocks and Industrial Production: Is the Relationship Linear?" *Macroeconomic Dynamics* 15 (S3): 472–97. https://doi.org/10.1017/S1365100511000290.

Herrera, Ana María, Latika Gupta Lagalo, and Tatsuma Wada (2015) "Asymmetries in the Response of Economic Activity to Oil Price Increases and Decreases?" *Journal of International Money and Finance* 50 (February): 108–33. https://doi.org/10.1016/j.jimonfin.2014.09.004.

Hooker, Mark A. (1996a) "This Is What Happened to the Oil Price-Macroeconomy Relationship: Reply." *Journal of Monetary Economics* 38 (2): 221–22. https://doi.org/10.1016/S0304-3932(96)01283-4.

Hooker, Mark A. (1996b) "What Happened to the Oil Price-Macroeconomy Relationship?" *Journal of Monetary Economics* 38 (2): 195–213. https://doi.org/10.1016/S0304-3932(96)01281-0.

Huang, Bwo-Nung, M.J. Hwang, and Hsiao-Ping Peng (2005) "The Asymmetry of the Impact of Oil Price Shocks on Economic Activities: An Application of the Multivariate Threshold Model." *Energy Economics* 27 (3): 455–76. https://doi.org/10.1016/j.eneco.2005.03.001.

Jiménez-Rodríguez, Rebeca, and Marcelo Sánchez (2005) "Oil Price Shocks and Real GDP Growth: Empirical Evidence for Some OECD Countries." *Applied Economics* 37 (2): 201–28. https://doi.org/10.1080/0003684042000281561.

Kilian, Lutz (2008) "The Economic Effects of Energy Price Shocks." *Journal of Economic Literature* 46 (4): 871–909. https://doi.org/10.1257/jel.46.4.871.

Kilian, Lutz (2009) "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review* 99 (3): 1053–69. https://doi.org/10.1257/aer.99.3.1053.

Kilian, Lutz (2014) "Oil Price Shocks: Causes and Consequences." *Annual Review of Resource Economics* 6 (1): 133–54. https://doi.org/10.1146/annurev-resource-083013-114701.

Kilian, Lutz, and Helmut Lütkepohl (2017) *Structural Vector Autoregressive Analysis:* 1st ed. Cambridge University Press. https://doi.org/10.1017/9781108164818.

Kilian, Lutz, and Clara Vega (2011) "Do Energy Prices Respond to U.S. Macroeconomic News? A Test of the Hypothesis of Predetermined Energy Prices." *Review of Economics and Statistics* 93 (2): 660–71. https://doi.org/10.1162/REST_a_00086.

Kilian, Lutz, and Robert J. Vigfusson (2011a) "Are the Responses of the U.S. Economy Asymmetric in Energy Price Increases and Decreases?: Are Responses of the U.S. Economy Asymmetric?" *Quantitative Economics* 2 (3): 419–53. https://doi.org/10.3982/QE99.

Kilian, Lutz, and Robert J. Vigfusson (2011b) "Nonlinearities in the Oil Price–Output Relationship." *Macroeconomic Dynamics* 15 (S3): 337–63. https://doi.org/10.1017/S1365100511000186.

Kim, Dong Heon (2012) "What Is an Oil Shock? Panel Data Evidence." *Empirical Economics* 43 (1): 121–43. https://doi.org/10.1007/s00181-011-0459-y.

Koop, Gary M (2013) "Forecasting with Medium and Large Bayesian VARS." *Journal of Applied Econometrics* 28 (2): 177–203. https://doi.org/10.1002/jae.1270.

Koop, Gary, M. Hashem Pesaran, and Simon M. Potter (1996) "Impulse Response Analysis in Nonlinear Multivariate Models." *Journal of Econometrics* 74 (1): 119–47. https://doi.org/10.1016/0304-4076(95)01753-4.

Koop, Gary, and Simon M. Potter (2010) "A Flexible Approach to Parametric Inference in Nonlinear and Time Varying Time Series Models." *Journal of Econometrics* 159 (1): 134–50. https://doi.org/10.1016/j.jeconom.2010.05.002.

Leduc, Sylvain, Kevin Moran, and Robert J. Vigfusson (2016) "The Elusive Boost from Cheap Oil." *FRBSF Economic Letter* 2016 (13). https://ideas.repec.org/a/fip/fedfel/00091.html.

Lee, Kiseok, Shawn Ni, and Ronald A. Ratti (1995) "Oil Shocks and the Macroeconomy: The Role of Price Variability." *The Energy Journal* 16 (4): 39–56. https://doi.org/10.5547/ISSN0195-6574-EJ-Vol16-No4-2.

Loungani, Prakash (1986) "Oil Price Shocks and the Dispersion Hypothesis." *The Review of Economics and Statistics* 68 (3): 536–39. https://doi.org/10.2307/1926035.

Mehra, Yash P., and Jon D. Petersen (2005) "Oil Prices and Consumer Spending." Federal Reserve Bank of Richmond *Economic Quarterly* 91 (3): 53–72. https://ideas.repec.org/a/fip/fedreq/y2005isump51-70nv.91no.3.html.

Mork, Knut Anton (1989) "Oil and the Macroeconomy When Prices Go up and down: An Extension of Hamilton's Results." *Journal of Political Economy* 97 (3): 740–44. https://doi.org/10.1086/261625.

Pindyck, Robert S. (1991) "Irreversibility, Uncertainty, and Investment." *Journal of Economic Literature* 29 (3): 1110–48. https://www.jstor.org/stable/2727613.

Shintani, Mototsugu, Akiko Terada-Hagiwara, and Tomoyoshi Yabu (2013) "Exchange Rate Pass-through and Inflation: A Nonlinear Time Series Analysis." *Journal of International Money and Finance* 32 (February): 512–27. https://doi.org/10.1016/j.jimonfin.2012.05.024.

Tong, Howell (1978) "On a Threshold Model." In *Pattern Recognition and Signal Processing*, edited by C. H. Chen, 575–86. Alphen aan den Rijn: Sijthoff & Noordhoff.

Van Robays, Ine (2016) "Macroeconomic Uncertainty and Oil Price Volatility." *Oxford Bulletin of Economics and Statistics* 78 (5): 671–93. https://doi.org/10.1111/obes.12124.

Wang, Yu Shan (2013) "Oil Price Effects on Personal Consumption Expenditures." *Energy Economics* 36 (March): 198–204. https://doi.org/10.1016/j.eneco.2012.08.007.

Weise, Charles L. (1999) "The Asymmetric Effects of Monetary Policy: A Nonlinear Vector Autoregression Approach." *Journal of Money, Credit and Banking* 31 (1): 85–108. https://doi.org/10.2307/2601141.

Table 1: Marginal Likelihood Values and Estimated Threshold

	Linear VAR	TVAR
Marginal likelihood	526.7	589.4
Estimated threshold	---	0.36% [0.34, 0.38]

Notes: The estimated threshold is the posterior mean of the retained draws; the numbers in the brackets correspond to the 70% interval. The reported marginal likelihood values correspond to the log of the expected posterior likelihood. The estimation sample is February 1959 to June 2014.

Table 2: Cumulative Response of Aggregate Consumption to Energy Price Shocks

Cumulative Consumption Response after 6 Months						
	Size of Energy Price Shock					
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.05	+0.05	-0.23	+0.23	-0.45	+0.45
TVAR, Regime 1	-0.10	+0.07	-0.69	+0.17	-1.18	+0.14
TVAR, Regime 2	-0.08	+0.07	-0.44	+0.24	-0.86	+0.29

Cumulative Consumption Response after 12 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.07	+0.07	-0.35	+0.35	-0.69	+0.69
TVAR, Regime 1	-0.15	+0.11	-0.93	+0.33	-1.55	+0.34
TVAR, Regime 2	-0.12	+0.11	-0.68	+0.43	-1.23	+0.55

Cumulative Consumption Response after 18 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.07	+0.07	-0.37	+0.37	-0.74	+0.74
TVAR, Regime 1	-0.16	+0.12	-0.97	+0.34	-1.61	+0.35
TVAR, Regime 2	-0.13	+0.12	-0.72	+0.45	-1.31	+0.58

Notes: All cumulative responses are expressed as percentages. The numbers reported reflect the estimated posterior median responses to the energy price shocks given in the table. The estimation sample is February 1959 to June 2014.

Table 3: Cumulative Response of Durable Consumption to Energy Price Shocks

Cumulative Durable Consumption Responses after 6 Months						
	Size of Energy Price Shock					
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.11	+0.11	-0.55	+0.55	-1.10	+1.10
TVAR, Regime 1	-0.22	+0.15	-1.67	+0.24	-3.01	-0.05
TVAR, Regime 2	-0.19	+0.16	-1.21	+0.54	-2.48	+0.53

Cumulative Durable Consumption Responses after 12 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.21	+0.21	-1.05	+1.05	-2.11	+2.11
TVAR, Regime 1	-0.36	+0.27	-2.29	+0.70	-3.94	+0.61
TVAR, Regime 2	-0.33	+0.29	-1.89	+1.09	-3.55	+1.38

Cumulative Durable Consumption Responses after 18 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.23	+0.23	-1.15	+1.15	-2.31	+2.31
TVAR, Regime 1	-0.38	+0.29	-2.42	+0.75	-4.13	+0.67
TVAR, Regime 2	-0.36	+0.31	-2.05	+1.19	-3.82	+1.53

Notes: All cumulative responses are expressed as percentages. The numbers reported reflect the estimated posterior median responses to the energy price shocks given in the table. The estimation sample is February 1959 to June 2014.

Table 4: Cumulative Response of Nondurable Consumption to Energy Price Shocks

Cumulative Nondurable Consumption Responses after 6 Months						
	Size of Energy Price Shock					
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.04	+0.04	-0.20	+0.20	-0.41	+0.41
TVAR, Regime 1	-0.09	+0.07	-0.59	+0.19	-0.98	+0.22
TVAR, Regime 2	-0.07	+0.06	-0.38	+0.25	-0.71	+0.34

Cumulative Nondurable Consumption Responses after 12 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.05	+0.05	-0.25	+0.25	-0.50	+0.50
TVAR, Regime 1	-0.13	+0.10	-0.79	+0.29	-1.29	+0.35
TVAR, Regime 2	-0.11	+0.10	-0.58	+0.39	-1.04	+0.52

Cumulative Nondurable Consumption Responses after 18 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.05	+0.05	-0.27	+0.27	-0.54	+0.54
TVAR, Regime 1	-0.14	+0.11	-0.83	+0.31	-1.36	+0.37
TVAR, Regime 2	-0.12	+0.11	-0.62	+0.41	-1.12	+0.56

Notes: All cumulative responses are expressed as percentages. The numbers reported reflect the estimated posterior median responses to the energy price shocks given in the table. The estimation sample is February 1959 to June 2014.

Table 5: Cumulative Response of Services Consumption to Energy Price Shocks

Cumulative Services Consumption Responses after 6 Months						
	Size of Energy Price Shock					
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.03	+0.03	-0.15	+0.15	-0.29	+0.29
TVAR, Regime 1	-0.08	+0.06	-0.52	+0.15	-0.86	+0.17
TVAR, Regime 2	-0.05	+0.04	-0.28	+0.17	-0.53	+0.19

Cumulative Services Consumption Responses after 12 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.04	+0.04	-0.20	+0.20	-0.41	+0.41
TVAR, Regime 1	-0.11	+0.09	-0.70	+0.26	-1.13	+0.29
TVAR, Regime 2	-0.08	+0.07	-0.43	+0.28	-0.76	+0.35

| Cumulative Services Consumption Responses after 18 Months						
	+1%	-1%	+5%	-5%	+10%	-10%
Linear VAR	-0.04	+0.04	-0.21	+0.21	-0.43	+0.43
TVAR, Regime 1	-0.12	+0.09	-0.73	+0.26	-1.17	+0.29
TVAR, Regime 2	-0.08	+0.08	-0.45	+0.29	-0.79	+0.35

Notes: All cumulative responses are expressed as percentages. The numbers reported reflect the estimated posterior median responses to the energy price shocks given in the table. The estimation sample is February 1959 to June 2014.

Figure 1: Estimated Thresholds under Two Regimes
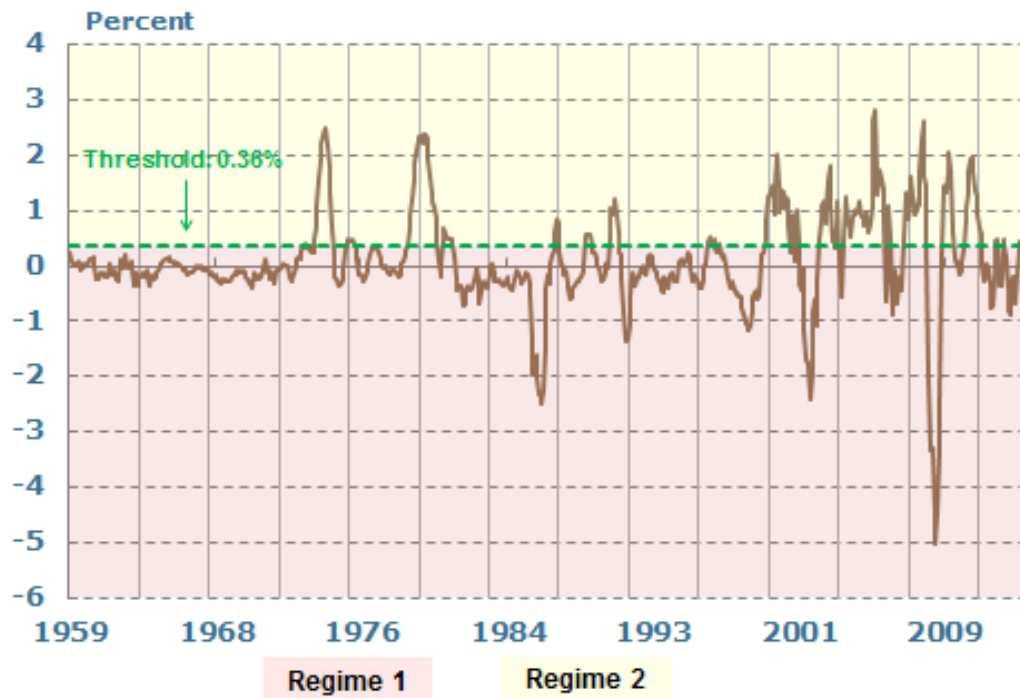
Notes: The figure plots the threshold variable (one-period lagged 9-month moving average of real energy price inflation), estimated threshold (posterior mean), and identified regimes based on the posterior mean. The system is in regime 1 (declining to moderately increasing real energy price inflation) when the threshold variable is less than or equal to 0.36% (indicated by pink shaded area). The system is in regime 2 (increasing real energy price inflation) when the threshold variable exceeds 0.36% (indicated by yellow shaded area). The estimation sample is February 1959 to June 2014.

Figure 2: Cumulative Response of Aggregate Consumption: +1% Energy Shock (upper panel) and -1% Energy Shock (lower panel)
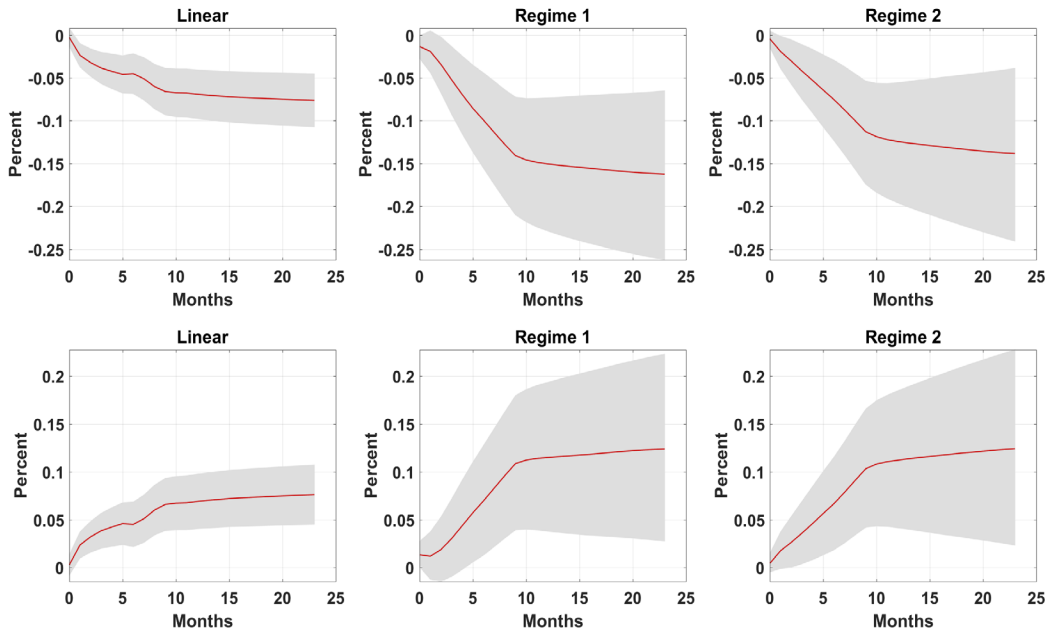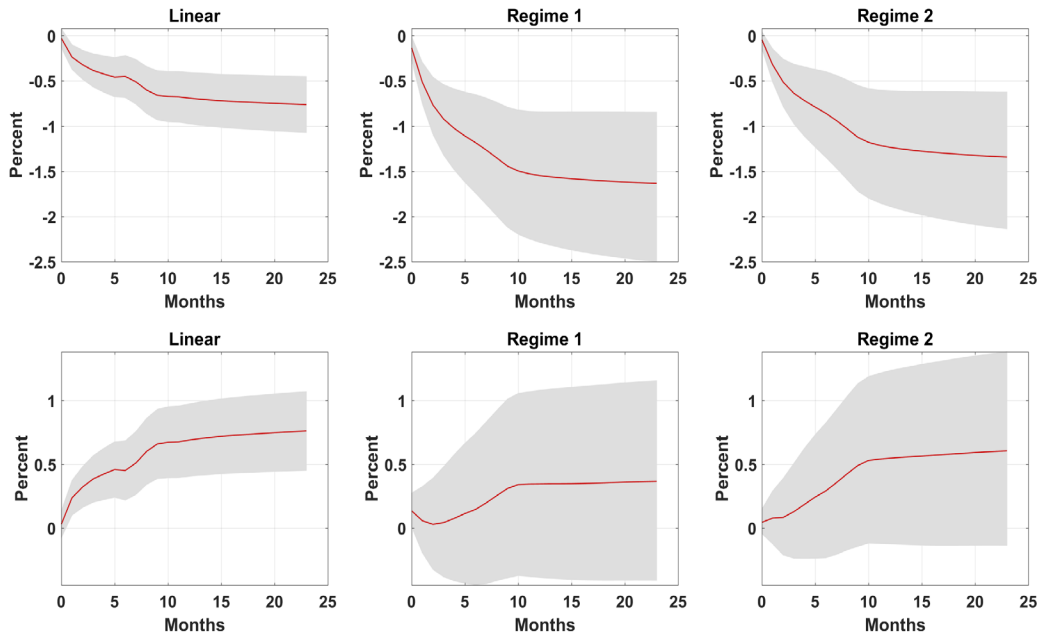

Notes: The figure shows the impulse response functions from the linear VAR and threshold VAR starting in either regime 1 or regime 2 estimated with monthly data over February 1959 to June 2014. The responses represent the cumulative posterior median responses of aggregate consumption to a one-time estimated +1% and -1% shock to energy price inflation. The grey shaded areas are 70% probability bands. Regime 1 corresponds to the low energy inflation regime and regime 2 corresponds to the high energy inflation regime.

Figure 3: Cumulative Response of Aggregate Consumption: +5% Energy Shock (upper panel) and -5% Energy Shock (lower panel)
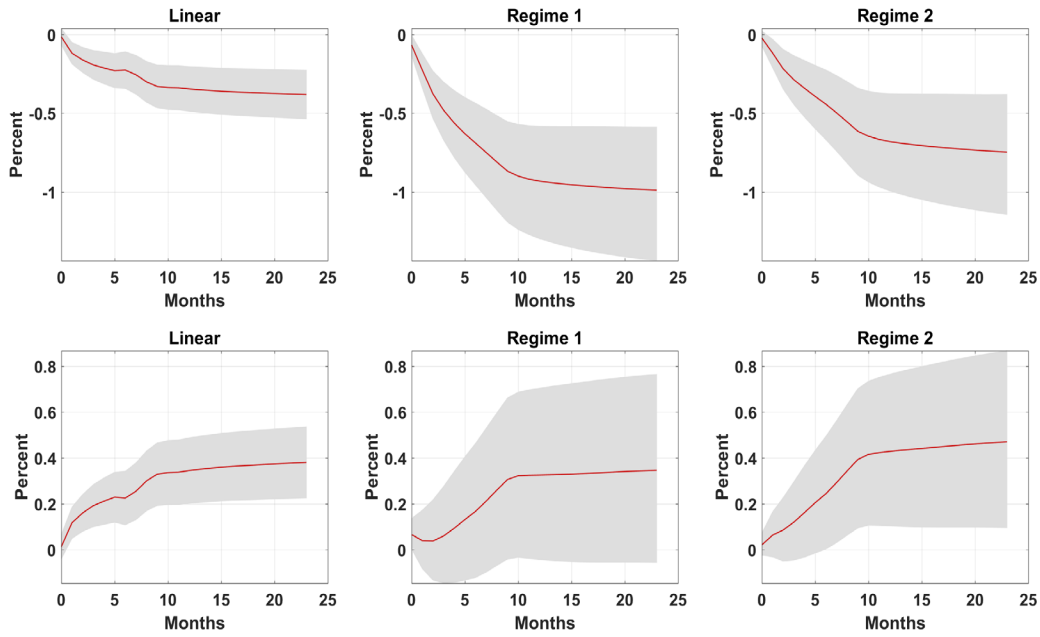
Notes: The figure shows the impulse response functions from the linear VAR and threshold VAR starting in either regime 1 or regime 2 estimated with monthly data over February 1959 to June 2014. The responses represent the cumulative posterior median responses of aggregate consumption to a one-time estimated +5% and -5% shock to energy price inflation. The grey shaded areas are 70% probability bands. Regime 1 corresponds to the low energy inflation regime and regime 2 corresponds to the high energy inflation regime.

Figure 4: Cumulative Response of Aggregate Consumption: +10% Energy Shock (upper panel) and -10% Energy Shock (lower panel)

Notes: The figure shows the impulse response functions from the linear VAR and threshold VAR starting in either regime 1 or regime 2 estimated with monthly data over February 1959 to June 2014. The responses represent the cumulative posterior median responses of aggregate consumption to a one-time estimated +10% and -10% shock to energy price inflation. The grey shaded areas are 70% probability bands. Regime 1 corresponds to the low energy inflation regime and regime 2 corresponds to the high energy inflation regime.