ONLINE APPENDIX
Time-Varying Effects of Oil Supply Shocks on the U.S. Economy

Christiane Baumeister
Bank of Canada
cbaumeister@bankofcanada.ca

Gert Peersman
Ghent University
gert.peersman@ugent.be

February 2013

Here we describe the setup and the results of a Monte Carlo simulation exercise that shows that our econometric model is able to capture abrupt changes in a satisfactory manner. We report on the significance of time variation in the responses of the four endogenous variables. We present evidence that the main findings of the paper are robust to changes in the variables included in the model, to alternative model specifications, and to different identification assumptions.

1 A Monte Carlo study

To explore whether our econometric model with smooth transitions is well suited to capture abrupt changes in the data, we carry out a Monte Carlo exercise based on simulated data where the underlying data-generating process is characterized by a one-time break. Given that our benchmark model is too complex to be amenable to a Monte Carlo study, we assess the performance of its main building blocks by conducting two separate experiments. We generate data from (1) an AR(1) model with one exogenous regressor that features a one-off regime shift in its coefficients, and (2) a bivariate version of our benchmark VAR(4) model with an abrupt break in the variance. These two simpler models provide a parsimonious way to assess the appropriateness of modeling structural change in a smoothly evolving way as opposed to a regime switch.

1.1 A regression model with a break in the coefficients

To illustrate the effects of incorrectly assuming a smooth process for the evolution of the coefficients, we simulate data from the following stationary AR(1) model with one exogenous regressor, written in demeaned form:

$$y_t = \alpha_i y_{t-1} + \beta_i x_{t-1} + \epsilon_t \qquad\qquad \epsilon_t \sim N(0,1) \tag{1}$$

where we set $\alpha_1 = 0.2$ and $\beta_1 = 0.5$ for the first half of the sample and $\alpha_2 = 0.6$ and $\beta_2 = 1.5$ for the second half of the sample. For each sample generated with this parameterization, we estimate a model that postulates that the coefficient vector $\delta = [\alpha \quad \beta]'$ evolves smoothly according to a driftless random walk process:

$$\delta_t = \delta_{t-1} + \eta_t \qquad\qquad \eta_t \sim N(0,Q) \tag{2}$$

We estimate the state-space model in equations (1) and (2) by Bayesian methods described in Kim and Nelson (1999). The unrestricted prior for the initial state is Gaussian:

$$\delta_0 \sim N(\widehat{\delta}_{OLS}, 4 \cdot \widehat{V}(\widehat{\delta}_{OLS})) \tag{3}$$

where $\widehat{\delta}_{OLS}$ and $\widehat{V}(\widehat{\delta}_{OLS})$ are the OLS point estimate and asymptotic variance based on a training sample as in our benchmark model. For the variance σ^2 in the observation equation, we postulate an inverse-gamma distribution:

$$\sigma^2 \sim IG\left(\frac{\lambda}{2}, \frac{\nu}{2}\right) \tag{4}$$

with scale parameter $\lambda = 0.01$ and degrees-of-freedom parameter $\nu = 2$. The prior for Q is assumed to be inverse Wishart:

$$Q \sim IW\left(\overline{Q}^{-1}, \xi\right) \tag{5}$$

where $\overline{Q} = 0.01 \cdot \xi$ and $\xi = 3$. The starting values for the coefficients are set to the OLS estimates, $\sigma_0 = 1$, and $Q_0 = 0.1 \cdot I_2$ where I_2 is a 2×2 identity matrix. The time-varying coefficients are drawn using the Carter and Kohn (1994) algorithm outlined in Appendix B of the paper. We constrain α_t to be less than one in absolute value at all dates t. The first $2,000$ draws in the Gibbs simulation process are discarded to ensure convergence. The posterior mean of $\widehat{\delta}_t$ is computed based on the remaining $1,000$ generated values. To evaluate how well this model can pick up the break imposed in the data-generating process, we also obtain an estimate of $\widehat{\delta}_t^{dummy}$ from a model that includes a dummy variable that takes a value of 0 before the break and 1 thereafter. In this way, we can construct error bands that capture the parameter uncertainty in estimating the true model.

We construct sample sizes of $T_1 = 200$ and $T_2 = 600$ after discarding the first $1,000$ periods to remove the influence of initial values. A sample size of 200 can be considered the equivalent of the typical sample length for quarterly time series available for the post-WWII period and 600 is representative of such a dataset at monthly frequency. There are $T/2$ data points on each side of the break date. We carry out $1,000$ Monte Carlo replications for each model and sample size.

Figure 2A reports the mean of the estimates for the exogenous coefficient and for the AR coefficient for the smooth-transition model and for the discrete-break model together with the 68% and 90% posterior credible sets for the two sample sizes. The estimation results show that the drifiting coefficient model locates the break in a satisfactory manner and moves relatively swiftly to the new regime.

1.2 A bivariate VAR model with a break in the variance

In the second experiment, the data-generating process is a bivariate VAR(4) model similar to equation (1) in the main text:

$$y_t = X_t'\theta + \varepsilon_t \tag{6}$$

where y_t denotes a vector of variables; X_t is a matrix including four lags of y_t and a constant; θ is a coefficient matrix, and $\varepsilon_t \sim N(0, \Omega_i)$, $i = 1, 2$ with the following variance-covariance matrices for two subperiods:

$$\Omega_1 = \begin{bmatrix} 20 & 5 \\ 5 & 30 \end{bmatrix} \qquad \Omega_2 = \begin{bmatrix} 1.5 & -4 \\ -4 & 300 \end{bmatrix} \tag{7}$$

To obtain a realistic parameterization for Ω_1 and Ω_2, we take guidance from the estimation of a bivariate VAR model for oil production and the real price of oil over two subsamples. The data generated from this model mimic a specific feature of the observed oil production and oil price series, namely a considerable decrease in oil production volatility and an increase in oil price volatility after the break in the variance. Figure 3A, panel A illustrates this behavior in one such random sample. The length of each sample generated from this model is 750, and the initial 500 periods are removed to yield a sample similar in size to that used in the empirical analysis.

For each sample, we estimate the time-varying VAR model with stochastic volatility presented in Appendix B of the paper. We retain the same priors as in the benchmark model and obtain initial values from the estimation of a constant-coefficient VAR(4) over a twenty–five-year training sample. This leaves us with 150 observations for the actual estimation, and the regime switch occurs at $t = 68$. Given the greater complexity of this model, we can only perform 250 Monte Carlo replications, and the results should consequently be viewed as suggestive. It should, however, be sufficient to examine the speed of transition from one regime to the other which is the main feature of interest.

Figure 3A, panel B displays the time profile of the average of the variance estimates over the Monte Carlo simulations together with both the 16^{th} and 84^{th} and the 5^{th} and 95^{th} percentiles of the posterior distribution. The results indicate that our approach has the power to detect the regime shift to a satisfactory degree even in a relatively short sample.

2 Further analysis

2.1 Evidence for time variation

In assessing the relative importance of time variation over the sample, we consider the joint posterior distribution of impulse responses across selected pairs of oil market episodes presented in a scatterplot along the lines of Cogley, Primiceri, and Sargent (2010). Shifts of this distribution away from the 45-degree line are indicative of a systematic change over time. Figure 4A reports the joint posterior distributions of the impact response of oil production and the cumulated responses of real GDP and CPI four quarters after an oil supply shock normalized such that it raises the real price of oil by 10% on impact for pairs of representative dates. Figure 5A reports the joint posterior distributions for the case when the oil supply shock corresponds to a 1% shortfall in oil production on impact for the same combinations of dates. Values for the earlier date are always plotted on the x-axis and those for the later date on the y-axis so that the location of the joint distribution of positive (negative) responses above the 45-degree line indicates an increase (decrease) over time and below the 45-degree line a decrease (increase) over time.

The evidence for time variation is most compelling for world oil production. The joint posterior draws for almost all combinations of dates are clustered above the dividing line, pointing towards a systematic decrease in the responses to oil supply shocks as time progresses. The joint posterior distributions for the real price of oil indicate a systematic increase in the magnitude of the impact of an oil supply shock. The exception is the pair $1979Q3 : 1986Q1$ for which the pairwise draws are more scattered suggesting no significant difference.

For real GDP, the points of the joint distribution are almost equally spread out around the 45-degree line for pairs of dates that are not too far apart from each other for both normalizations. This suggests the absence of a significant change in the reaction of real GDP to oil supply shocks. There is, however, some evidence for differences between the magnitudes of output responses during episodes that are more distant in time. In particular, an oil supply shock normalized on oil production is more contractionary in the more recent past compared to early periods since a considerable fraction of the pairwise draws lies below the threshold. The opposite results emerges for the normalization based on the real price of oil. While the dispersion of pairwise posterior draws for consumer prices implies similar responses for pairs of dates in the early part of the sample, there appears to be a systematic increase in price responses in the last decade with the majority of draws lying above the dividing line for both normalizations. Overall, this provides strong evidence for sizeable changes in the responses of oil market variables and some support for changes in the responses of the aggregate economy over time.

2.2 Sensitivity analysis

We summarize the results for the alternative model specifications discussed in section 3.3 in the main text and conduct two additional robustness checks.

To enable the reader to compare our results with previous studies that have investigated the macroeconomic consequences of a fixed oil price increase across time, Figure 6A reports the median responses of the four endogenous variables after an oil supply shock identified with sign restrictions and normalized to a 10% increase in the real price of oil in a constant-coefficient VAR model estimated over the two subperiods $1974Q1 - 1985Q4$ and $1986Q1 - 2011Q1$, together with the 16^{th} and 84^{th} percentiles. We observe that the decrease in world oil production is considerably smaller in the second subsample confirming once again that the short-run price elasticity of oil demand has declined over time. The evidence for time variation in the responses of real GDP and consumer prices across subsamples is weak given that the posterior intervals overlap.

In what follows, an oil supply shock is normalized to correspond to a 1% decrease in world oil production. Figure 7A, panel A displays the time-varying median responses of the unemployment rate (left) and the implicit GDP deflator (right) which yield the same pattern of time variation as real GDP and CPI. For ease of comparison, the dotted lines in panels B and C and in Figure 7A depict the median estimates obtained with the baseline model. Figure 7A, panel B, shows that the time-varying responses obtained with different oil price measures are essentially identical. Panel C presents the evolution of the responses for the model augmented by the federal funds rate which exhibit a pattern remarkably similar to the baseline case. Figure 8A, panel A displays the time-varying responses when sign restrictions are imposed from $t = 4$ to $t = 8$. As before, the time-varying responses closely track the dotted lines that represent the benchmark model, demonstrating that our main conclusions are not sensitive to the modified identification assumptions.

As pointed out by Fry and Pagan (2011), sign restrictions impose only weak information. Building on this insight, Kilian and Murphy (2012) argue that the sign identification strategy needs to be complemented with additional information in order to derive economically meaningful results in the context of oil market models. They propose the use of empirically plausible boundary restrictions on the magnitudes of the implied short-run price elasticities of oil demand and oil supply as auxiliary identification criteria to eliminate those structural models that are associated with implausibly high elasticities. To verify the robustness of our findings, we follow Kilian and Murphy (2013) in imposing that the short-run price elasticity for oil demand cannot exceed its long-run counterpart which may be inferred to be about -0.8 using cross-sectional evidence from U.S. household surveys (see e.g. Hausman and Newey 1995). Figure 8A, panel B shows that our baseline results are not affected by this additional identifying restriction. Even more stringent bounds on the impact price elasticity of oil demand have little effect on the time profile of the impulse responses.

To explore the sensitivity of our results with regard to the data frequency, we estimate a monthly VAR(12) model that includes the growth rates of world oil production, the real refiners' acquisition cost of crude oil imports, U.S. industrial production, and U.S. CPI over two subsamples, $1974M1 - 1985M12$ and $1986M1 - 2011M3$.[1] In line with the quarterly model, the sign restrictions are postulated to hold over a horizon of 12 months after the

[1] Estimating the benchmark time-varying VAR model with monthly data is not an option since this would result in a proliferation of free parameters given the number of lags that are required to allow for sufficient dynamics.

shock. Figure 9A displays the median responses of the four variables to an oil supply shock together with the 16^{th} and 84^{th} percentiles of the posterior distribution. The results for the monthly specification paint much the same picture of time variation as in the quarterly split sample model.

References

[1] **Carter, Christopher K., and Robert Kohn.** 1994. "On Gibbs Sampling for State Space Models." *Biometrika*, 81(3): 541-53.

[2] **Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent.** 2010. "Inflation-Gap Persistence in the US." *American Economic Journal: Macroeconomics*, 2(1): 43-69.

[3] **Fry, Renée, and Adrian Pagan.** 2011. "Sign Restrictions in Structural Vector Autoregressions: A Critical Review." *Journal of Economic Literature*, 49(4): 938-60.

[4] **Hausman, Jerry A., and Whitney K. Newey**. 1995. "Nonparametric Estimation of Exact Consumers' Surplus and Deadweight Loss." *Econometrica*, 63: 1445-76.

[5] **Kilian, Lutz, and Daniel P. Murphy**. 2012. "Why Agnostic Sign Restrictions Are Not Enough: Understanding the Dynamics of Oil Market VAR Models." *Journal of the European Economic Association*, 10(5): 1166-88.

[6] **Kilian, Lutz, and Daniel P. Murphy**. 2013. "The Role of Inventories and Speculative Trading in the Global Market for Crude Oil." *Journal of Applied Econometrics*, forthcoming.

[7] **Kim, Chang-Jin, and Charles R. Nelson.** 1999. *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications.* Cambridge, MA: MIT Press.

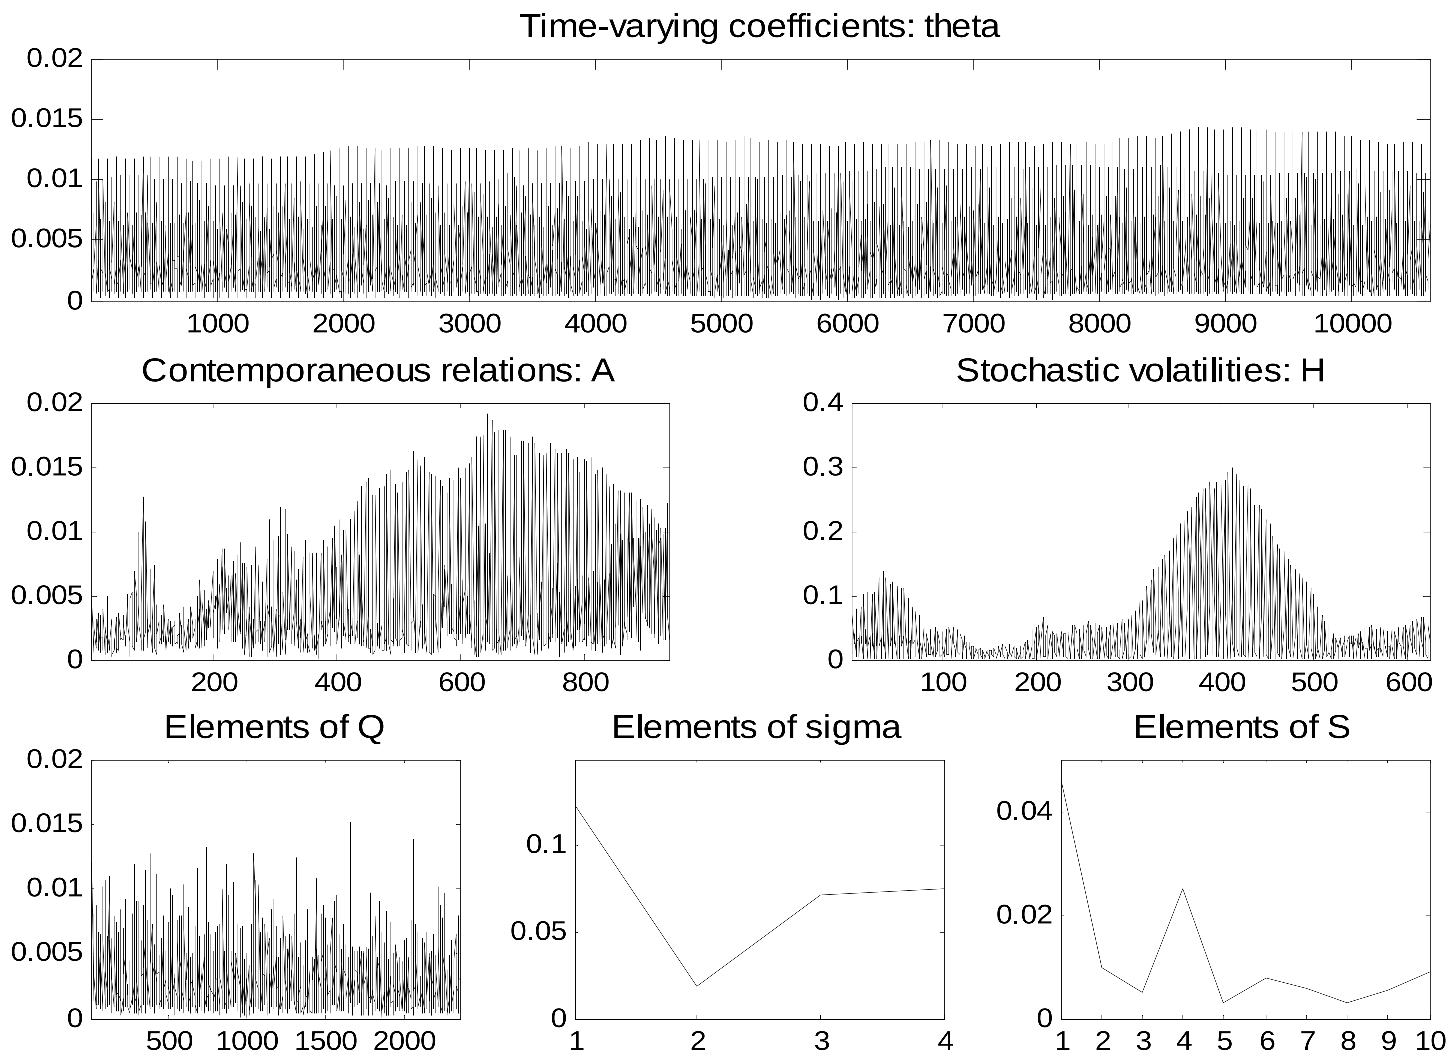

Figure 1A: Assessing the convergence of the Markov chain: inefficiency factors for the draws from the ergodic distribution for the states and hyperparameters.

Panel A

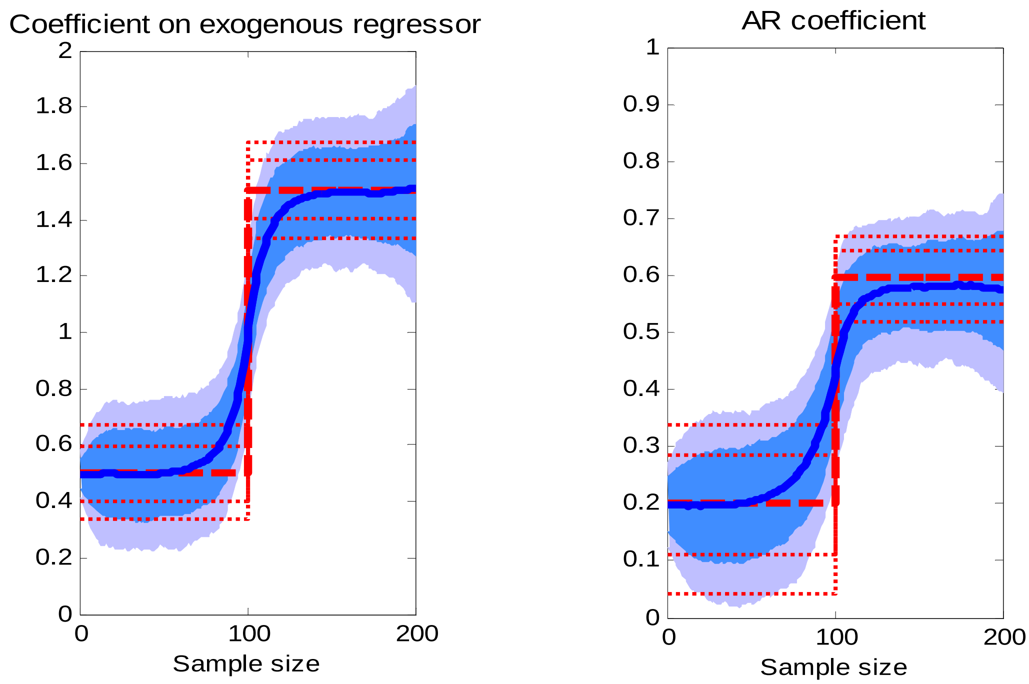

Panel B

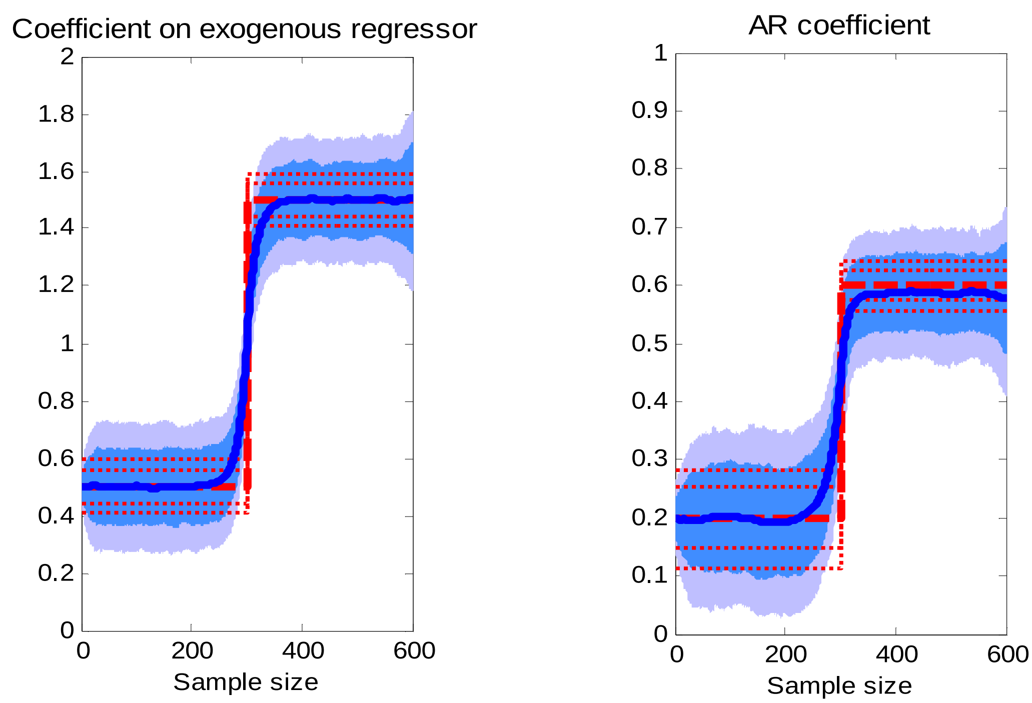

Figure 2A: Mean, 68% and 90% posterior credible sets of coefficient estimates from a smooth-transition model (solid line and shaded areas) and a discrete-break model (dashed and dotted lines) for 1,000 Monte Carlo replications.
Panel A: Sample size *T=200*.
Panel B: Sample size *T=600*.

Panel A

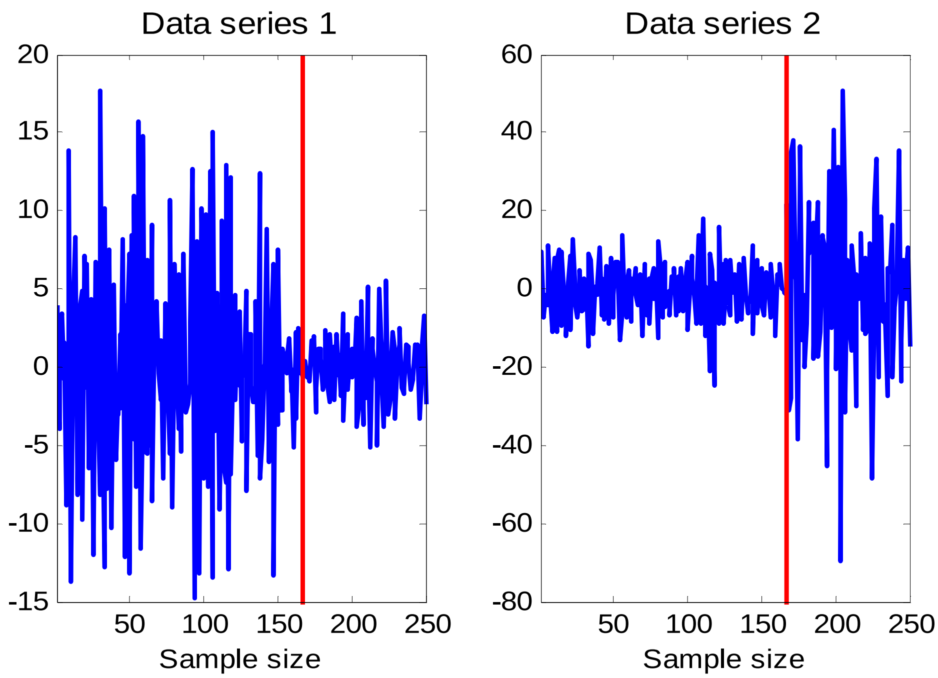

Panel B

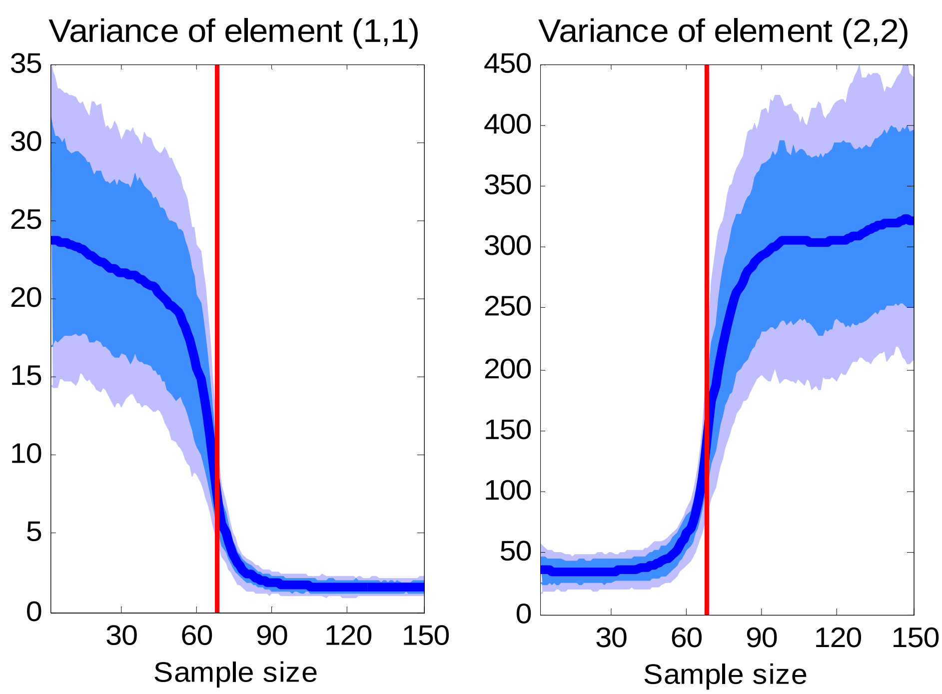

Figure 3A: Panel A: Random sample generated from bivariate VAR(4) model with break in variance at *t=168* (vertical line).

Panel B: Mean, 68% and 90% posterior credible sets of variance estimates for 250 Monte Carlo replications.

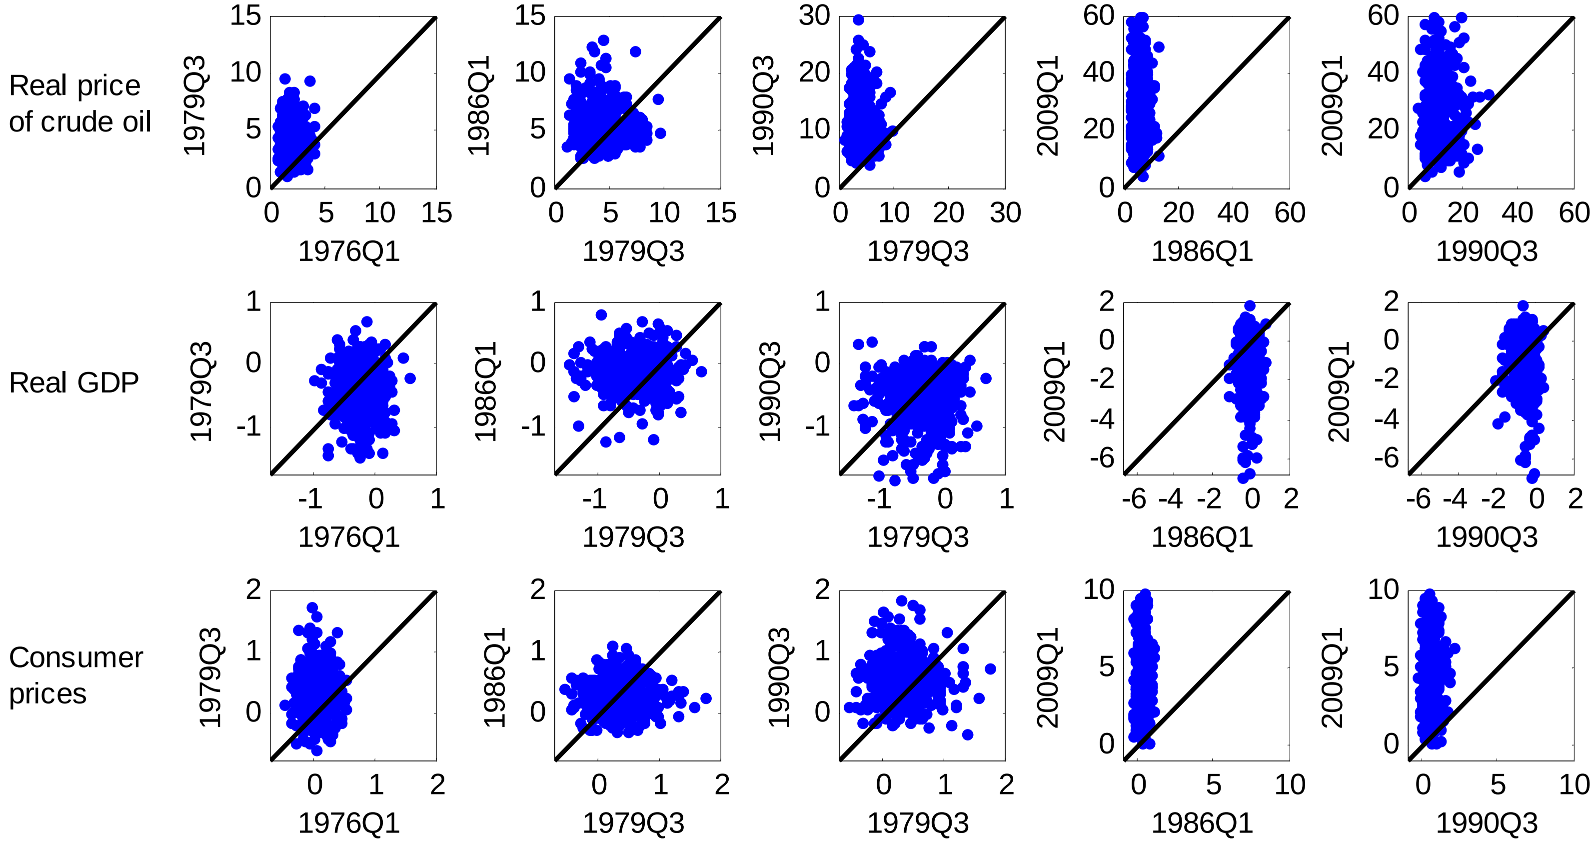

Figure 4A: Joint posterior distribution of the responses of the real price of oil, real GDP and consumer prices to an oil supply shock normalized such that it decreases world oil production by 1% on impact for selected pairs of dates.

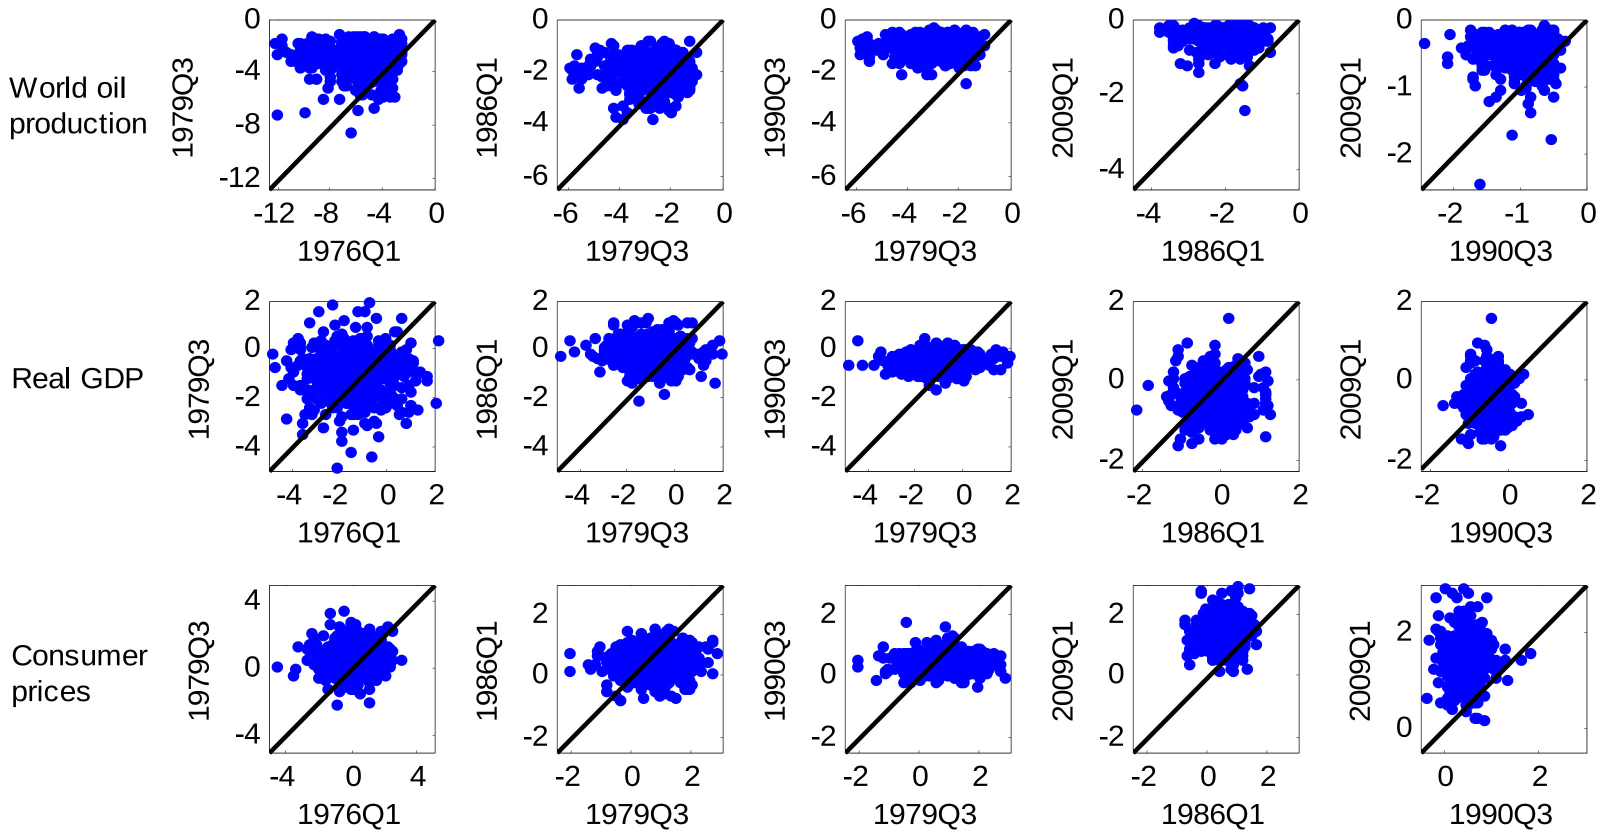

Figure 5A: Joint posterior distribution of the responses of world oil production, real GDP and consumer prices to an oil supply shock normalized such that it raises the real price of oil by 10% on impact for selected pairs of dates.

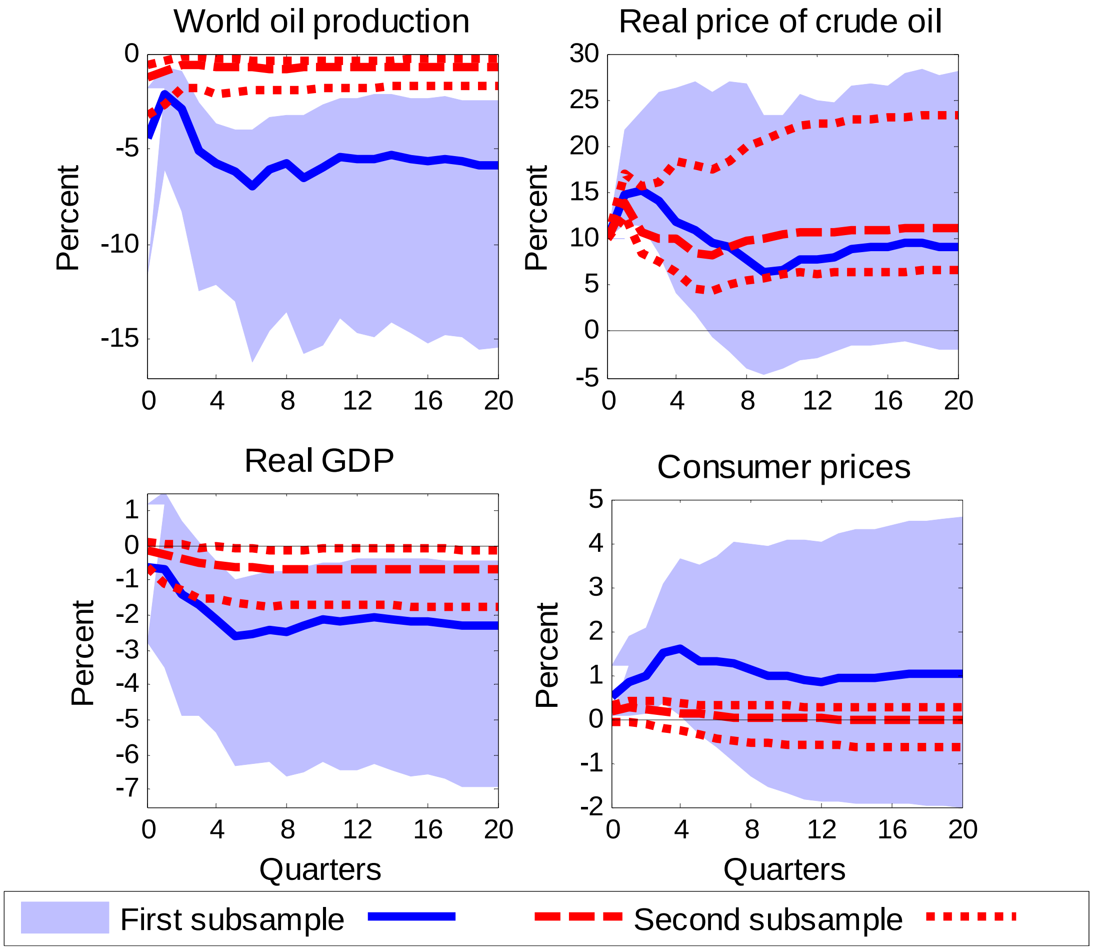

Figure 6A: Posterior median responses with 16th and 84th percentiles after an oil supply shock identified with sign restrictions and normalized to a 10% increase in the real price of crude oil in a constant-coefficient VAR estimated over two subsamples, 1974*Q1*–1985*Q4* and 1986*Q1*–2011*Q1*.

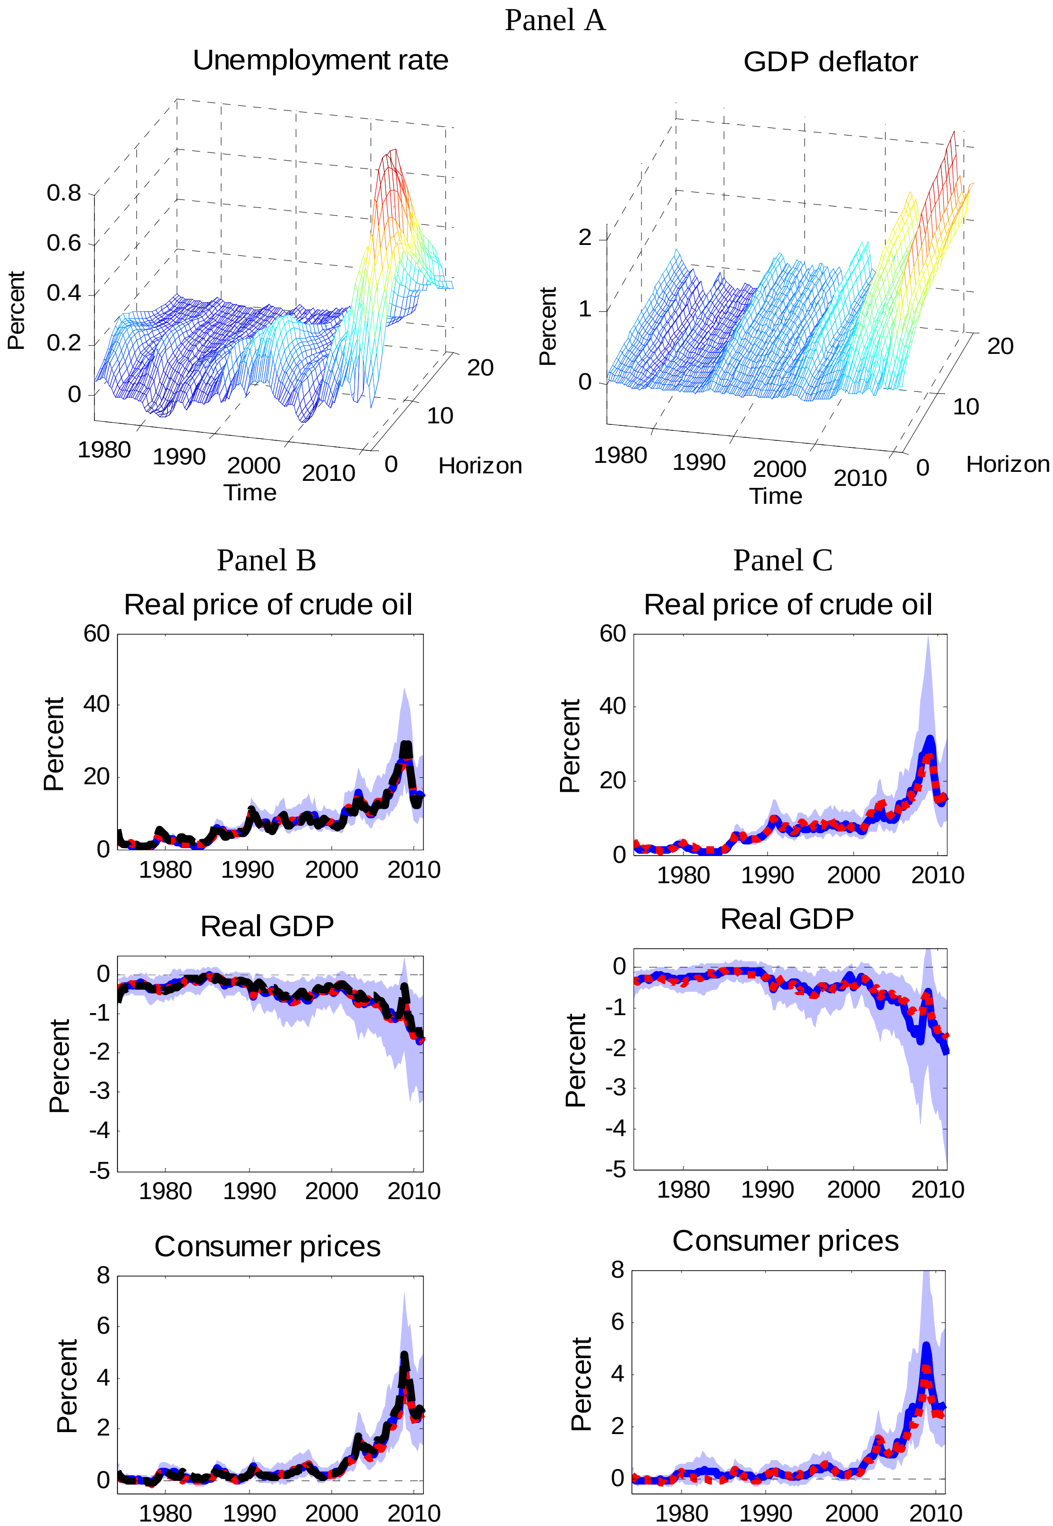

Figure 7A: Panel A: Median responses of U.S. unemployment (left) and GDP deflator (right).
Panel B: Median responses for model with WTI (dashed line) and composite refiners' acquisition cost (solid line) with 68% posterior credible set (shaded area).
Panel C: Median responses for specification with federal funds rate.

Panel A

Panel B

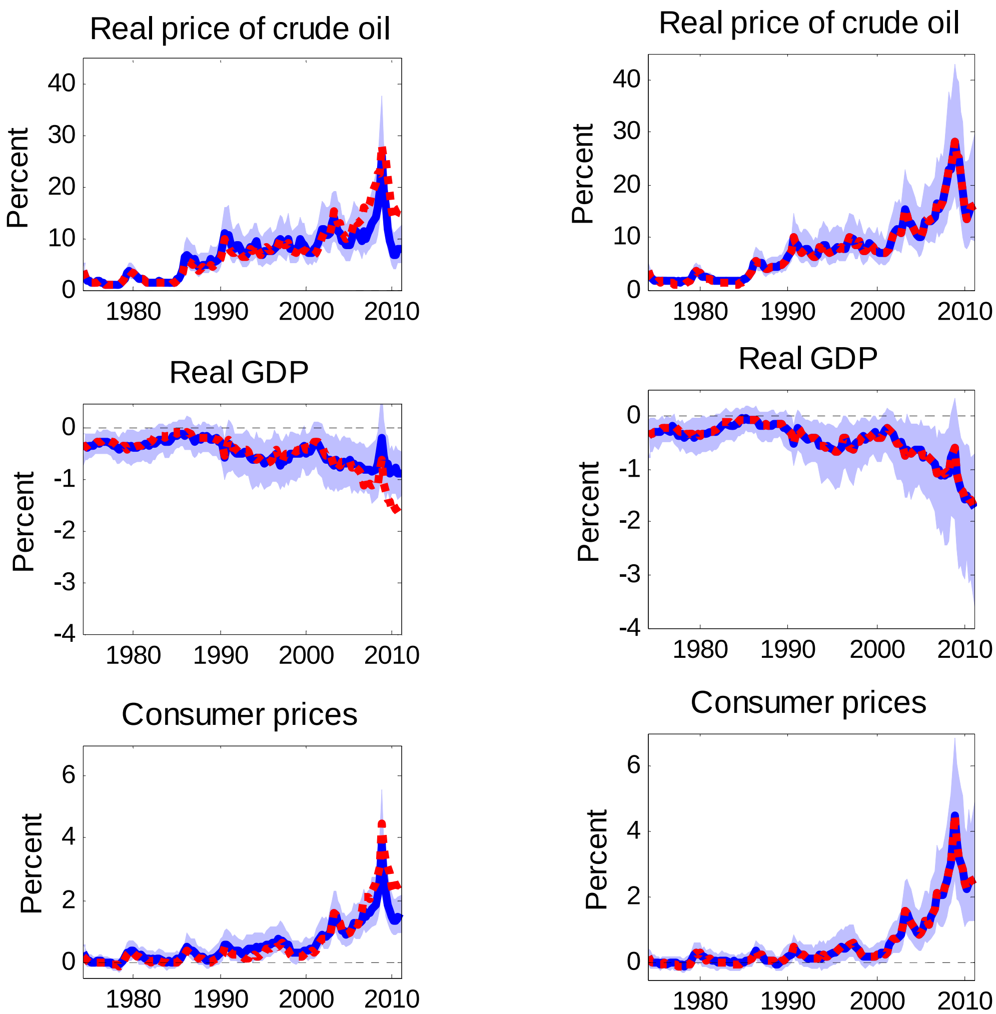

Figure 8A: Panel A: Median responses with 16th and 84th percentiles when sign restrictions are imposed from *t=4* to *t=8*.

Panel B: Median responses with 16th and 84th percentiles when the lower bound for the short-run price elasticity of demand is -0.8

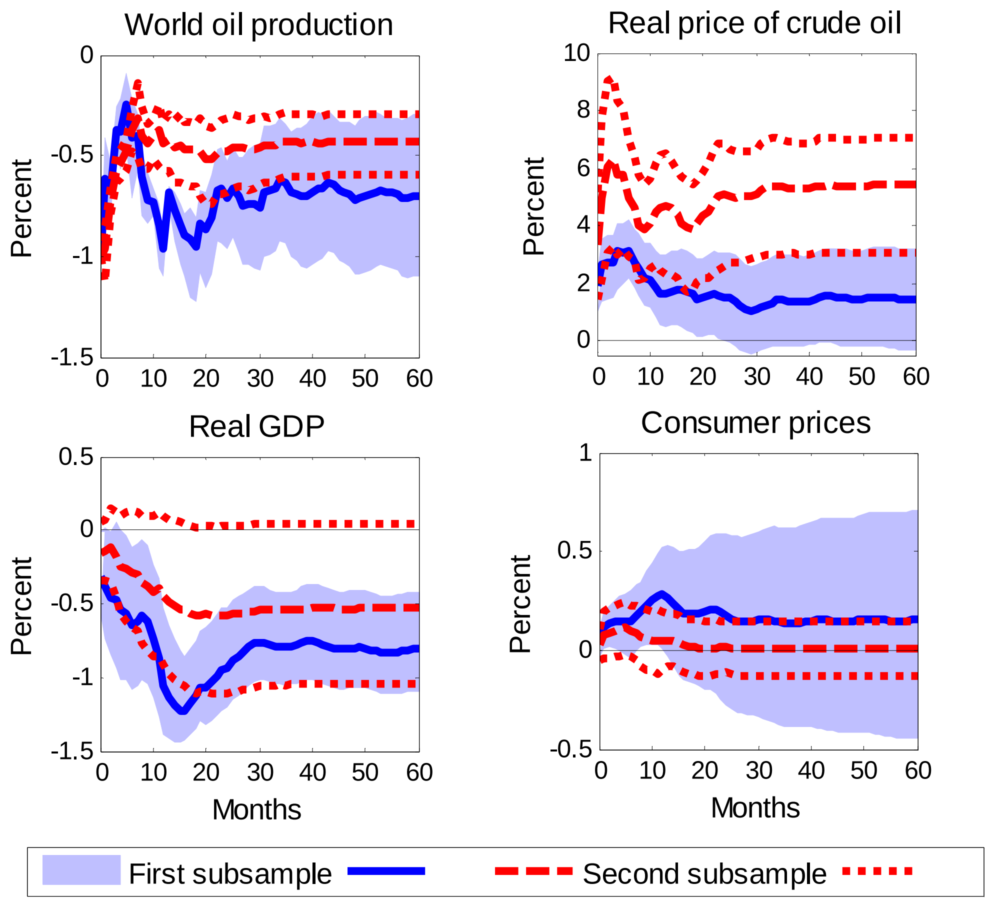

Figure 9A: Posterior median responses with 16th and 84th percentiles after an oil supply shock identified with sign restrictions and normalized to a 1% decrease in oil production in a constant-coefficient VAR estimated over two subsamples, 1974*M1*–1985*M12* and 1986*M1*–2011*M3*.