JOURNAL OF APPLIED ECONOMETRICS
J. Appl. Econ. **23**: 173–191 (2008)
Published online 1 November 2007 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/jae.978

BAYESIAN COUNTERFACTUAL ANALYSIS OF THE SOURCES OF THE GREAT MODERATION

KIM CHANG-JIN,[a] JAMES MORLEY[b]* AND JEREMY PIGER[c]

[a] *Deparment of Economics, Korea University, Seoul, Korea;*
Department of Economics, University of Washington, Seattle, Washington, USA
[b] *Department of Economics, Washington University, St Louis, Missouri, USA*
[c] *Department of Economics, University of Oregon, Eugene, Oregon, USA*

SUMMARY

We use counterfactual experiments to investigate the sources of the large volatility reduction in US real GDP growth in the 1980s. Contrary to an existing literature that conducts counterfactual experiments based on classical estimation and point estimates, we consider Bayesian analysis that provides a straightforward measure of estimation uncertainty for the counterfactual quantity of interest. Using Blanchard and Quah's (1989) structural VAR model of output growth and the unemployment rate, we find strong statistical support for the idea that a counterfactual change in the size of structural shocks alone, with no corresponding change in the propagation of these shocks, would have produced the same overall volatility reduction as what actually occurred. Looking deeper, we find evidence that a counterfactual change in the size of aggregate supply shocks alone would have generated a larger volatility reduction than a counterfactual change in the size of aggregate demand shocks alone. We show that these results are consistent with a standard monetary VAR, for which counterfactual analysis also suggests the importance of shocks in generating the volatility reduction, but with the counterfactual change in monetary shocks alone generating a small reduction in volatility.

1. INTRODUCTION

A striking stylized fact of the US macroeconomy is the large reduction in its volatility since the early 1980s, a feature that has been called the 'Great Moderation' (Bernanke, 2004). This volatility reduction, first documented by Niemira and Klein (1994), Kim and Nelson (1999) and McConnell and Perez-Quiros (2000), has spawned a large literature attempting to isolate its source.[1] The literature has focused on three explanations: (1) improved macroeconomic policy; (2) changes in the private sector's behavior; and (3) good luck.

A popular approach to distinguishing between these competing explanations is through the use of counterfactual experiments (e.g., Stock and Watson, 2002; Boivin and Giannoni, 2003; Ahmed *et al*., 2004). In these counterfactual experiments, a macroeconomic model is estimated over a pre-volatility reduction sample, denoted Period 1, and a post-volatility reduction sample, denoted Period 2, and the counterfactual variance of a variable of interest is calculated in which the parameter estimates from Periods 1 and 2 are intermingled. For example, a counterfactual variance

* Correspondence to: James Morley, Department of Economics, Washington University in St. Louis, Box 1208, One Brookings Drive, St. Louis, MO 63130-4899, USA. E-mail: morley@wustl.edu
Contract/grant sponsor: Weidenbaum Center on the Economy, Government, and Public Policy.
[1] See Simon (2000), Blanchard and Simon (2001), Chauvet and Potter (2001), Kahn *et al*. (2002), Stock and Watson (2002), Boivin and Giannoni (2003), Ahmed *et al*. (2004), Kim *et al*. (2004), Sensier and van Dijk (2004), Herrera and Pesavento (2005), Choi and Jung (forthcoming), and Gordon (2006).

could be calculated in which Period 2 estimates of model parameters representing monetary policy are mixed with the non-policy parameter estimates of Period 1. If the counterfactual variance is considerably lower than the sample variance in Period 1, improved monetary policy would be a strong candidate for the source of the volatility reduction.[2]

In this paper we also conduct counterfactual experiments to isolate sources of the volatility reduction in US real GDP. However, unlike the previous literature that has focused on classical estimation and point estimates, we consider Bayesian analysis in which posterior densities for the counterfactual variances are constructed. This represents an improvement over the previous analysis because it provides a sense of statistical precision about the counterfactual quantities of interest. In particular, Bayesian estimation procedures make it straightforward to capture the implications of both parameter uncertainty and uncertainty regarding timing of structural change on inferences about the counterfactual variance. In the classical context, there is no known analytical measure of estimation uncertainty for counterfactual variances, as they are complicated nonlinear functions of the underlying model parameters, while a measure based on linear approximation (i.e., the 'delta method') is likely to be highly inaccurate, and it is unclear how to incorporate uncertainty about the timing of structural change into such a measure. Furthermore, while classical inference is difficult even with just two sources of uncertainty, it would be relatively straightforward to incorporate additional sources of uncertainty, such as that about lag length or more general model specification assumptions, in Bayesian analysis.

We apply the Bayesian counterfactual analysis to Blanchard and Quah's (1989) structural VAR model of output growth and the unemployment rate. The model has three structural components: aggregate supply shocks, aggregate demand shocks, and a propagation mechanism for the structural shocks. The structural components are identified using a long-run restriction in which aggregate demand shocks have no long-run effect on the level of real GDP. Estimation results suggest a large reduction in output volatility in the 1980s, with clear evidence of a reduction in the size of the structural shocks. We find strong statistical support for the idea that a counterfactual change in the size of structural shocks alone would have produced the same overall volatility reduction as actually occurred. By contrast, a counterfactual change in propagation alone would have produced little reduction in volatility, although these inferences are somewhat less precise. Meanwhile, a change in the size of aggregate supply shocks alone would have generated a larger volatility reduction than a change in the size of aggregate demand shocks alone. Thus, consistent with the previous literature, these results point towards the role of good luck rather than a change in private sector behavior or macroeconomic policy.

Because of concerns about identifying aggregate demand shocks using a long-run restriction, and to explicitly consider the role of monetary policy, we check the robustness of our results using a standard monetary VAR similar to that considered in Stock and Watson (2002), Boivin and Giannoni (2003), and Ahmed *et al.* (2004). We find consistent results in that shocks are much more successful than propagation in generating a large volatility reduction. Further, we find that a change in monetary shocks alone is only able to generate a small fraction of the total reduction in volatility. Thus, the primary role for shocks appears to be arising from the non-monetary shocks in the system.

[2] It is worth noting that counterfactual analysis is subject to the Lucas (1976) critique, in that the experiments proceed by changing sets of model parameters while holding others constant. However, a number of papers have documented limited empirical relevance of the Lucas critique for vector autoregressions estimated to US post-war data (e.g., Rudebusch, 2005; Leeper and Zha, 2003; Sims and Zha, 2006), which is the application considered in this paper.

The Bayesian approach we take to counterfactual analysis is closely related to that developed in Sims and Zha (2006). They employ Bayesian techniques to estimate and compare a number of multiple regime structural VAR models of the US economy, including models with regime switches in propagation and/or shock parameters. They also simulate counterfactual histories of the level of inflation under alternative policy regimes to ascertain the extent to which changes in monetary policy parameters might account for the rise and subsequent fall in inflation observed in the 1970s and 1980s.[3] While the Bayesian approach is similar, we consider a different counterfactual quantity, namely the unconditional variance of real GDP growth under the assumption of a single regime change.

In the next section we discuss the principles underlying counterfactual experiments in the context of structural change. Section 3 describes the Bayesian approach taken in this paper. Section 4 presents details of the structural VAR and counterfactual experiments that we use to investigate the sources of the volatility reduction in US real GDP growth. Section 5 reports the results of the Bayesian estimation and counterfactual experiments. Section 6 examines the robustness of the results to the alternative monetary VAR model. Section 7 concludes.

2. COUNTERFACTUAL EXPERIMENTS

The principles behind counterfactual experiments in the context of structural change are most easily illustrated using the example of a stationary AR(1) process:

$$x_t = \phi x_{t-1} + e_t, e_t \sim N(0, \sigma^2) \tag{1}$$

where $|\phi| < 1$. Suppose x_t undergoes a structural break corresponding to a reduction in its variance. There are two possible sources of the variance reduction: 'shocks' and 'propagation'. To see this, note that the variance of x_t is

$$\text{var}(x_t) \equiv \gamma_0 = \frac{\sigma^2}{1 - \phi^2} \tag{2}$$

where σ^2 corresponds to the variance of the shocks and ϕ^2 corresponds to the propagation. A reduction in either σ^2 or ϕ^2 will reduce the size of γ_0. In practice, a variance reduction could reflect a change in both shocks and propagation.

Counterfactual experiments consider the hypothetical changes that would have occurred if either only shocks or only propagation had changed. To illustrate, let $i, j = 1, 2$ index the structural regime for shocks and propagation, respectively:

$$\gamma_0^{(i,j)} = \frac{(\sigma^{(i)})^2}{1 - (\phi^{(j)})^2} \tag{3}$$

Then $i = j$ produces the actual variances for the two structural regimes, while $i \neq j$ produces the counterfactual variances based on changes in shocks only or propagation only.

[3] Another example of a counterfactual quantity constructed using Bayesian techniques can be found in Dueker and Nelson (2003), who use Bayesian analysis to simulate counterfactual histories of several macroeconomic variables given alternative realizations of a latent business cycle indicator.

DOI: 10.1002/jae

It should be noted that counterfactual experiments do not provide a formal decomposition of what caused the variance of x_t to change. In particular, it is easily seen from (3) that shocks and propagation interact in a nonlinear fashion to determine the variance. As a result, the changes in variance implied by the counterfactual experiments do not necessarily sum to the actual change in variance. Thus, some caution should be employed in interpreting counterfactual experiments. In essence, they represent hypothetical scenarios only.

Nevertheless, a hypothetical may be very revealing. For example, suppose the counterfactual variance corresponding to a change in shocks but not propagation ($\gamma_0^{(2,1)}$) is of a similar magnitude to the variance after the structural break ($\gamma_0^{(2,2)}$), while the counterfactual variance corresponding to a change in propagation but not shocks ($\gamma_0^{(1,2)}$) is of a similar magnitude to the variance before the structural break ($\gamma_0^{(1,1)}$). Then, the findings would be highly suggestive of a large role for shocks in producing the variance reduction. Of course, in practice, one might be concerned about the statistical precision of these counterfactual inferences. The Bayesian approach taken in this paper is designed to address this concern.

3. THE BAYESIAN APPROACH

In Bayesian analysis, beliefs about model parameters are described using probability distributions. If data provide a lot of information about the values of parameters, posterior densities will be relatively tight. If the data are uninformative, posterior densities will be more spread out and largely reflect the specification of prior beliefs, rather than sample information.

Given posterior distributions for model parameters, it is possible to simulate from posterior distributions for functions of those parameters, including counterfactual variances. Continuing with the simple example of a structural break in the variance of the AR(1) process in the previous section, we can repeatedly draw realizations of $(\phi^{(1)}, \phi^{(2)}, \sigma^{(1)}, \sigma^{(2)})$ from their joint posterior distribution and use the unconditional variance formula in (3) to construct implied realizations and, therefore, distributions for $\gamma_0^{(i,j)}$ for $i, j = 1, 2$. For $i = j$, we would have posterior distributions for the subsample variances of x_t. For $i \neq j$, we would have posterior distributions for the counterfactual variances.[4]

Bayesian analysis also allows us to make inferences about the relative sizes of different variances. In particular, we can calculate ratios of functions of the variances and evaluate the probabilities that they are less than a specified fixed value. A simple example would be the

[4] Note that the counterfactual variances will reflect the weaker dependence between shock and propagation parameters across subsamples than within subsamples. For example, suppose that shock and propagation parameters are negatively related within subsamples such that the posteriors for the actual unconditional variances are much tighter than would be implied by independent draws from the marginal distributions for the parameters. Then, given weak dependence or independence (as would be the case under a noninformative prior) between parameters across subsamples, the posteriors for the counterfactual variances will be much less precise than for the actual variances. However, if the dependence between shock and propagation parameters reflects the relative size of the shock variance, then it is possible to use this dependence when constructing counterfactual variances in order to get more precise inferences. To capture this dependence, we consider sorting the realized propagation parameters by the realized shock variances and assigning them a rank, which can be matched to the quantile of the realized shock variance in the other subsample. In practice, we find that sorting makes counterfactual posteriors only slightly more precise. Thus, for simplicity of presentation, we report the basic results without sorting.

comparison of an actual and counterfactual variance:

$$\Pr\left[\frac{\gamma_0^{(i,j)}}{\gamma_0^{(i,i)}} < k\right]$$

That is, we can evaluate the probability that the counterfactual variance when $i \neq j$ is less than some percentage, denoted by the fraction k, of the actual variance in sample period i. The important benchmark is a probability of 0.5, which, given $k = 1$, suggests 'even odds' that the one variance is bigger than the other. Meanwhile, as we vary k from one towards zero, probabilities much higher or lower than 50% provide strong statistical support for one variance being significantly (in an economic sense) larger than the other.

A particularly attractive feature of the Bayesian framework is that it allows us to account for uncertainty about the exact timing of structural change in a relatively straightforward manner. In particular, estimation of an unknown breakdate can be added to estimation of the parameters and the actual and counterfactual variances. Accounting for uncertainty about the timing of structural change represents an improvement over the existing literature, which has conducted counterfactual analysis conditional on a given breakdate.

4. COUNTERFACTUAL ANALYSIS OF THE VOLATILITY REDUCTION IN US REAL GDP GROWTH

In this section we present the design of our Bayesian counterfactual analysis of the recent volatility reduction in US real GDP growth. For our model of the US macroeconomy, we consider Blanchard and Quah's (1989) long-run structural VAR model of output growth and the unemployment rate. Beyond the basic distinction between shocks and propagation, the model allows us to consider aggregate supply and aggregate demand shocks as separate possible sources of the volatility reduction. In terms of propagation, the model also allows us to identify contemporaneous structural propagation of shocks across series, in addition to the dynamic propagation considered in counterfactual experiments based on reduced-form VARs (e.g., Stock and Watson, 2002).

4.1. A Long-Run Structural VAR Model of Output and Unemployment

We start with a reduced-form VAR model for output growth and the unemployment rate:

$$\Phi(L)Y_t = c + e_t, e_t \sim N(0, \Omega) \tag{4}$$

where $Y_t \equiv (\Delta y_t, u_t)'$, $\Phi(L) \equiv I - \sum_{k=1}^{p} \Phi_k L^k$, $c \equiv (c_1, c_2)'$, $e_t \equiv (e_{1t}, e_{2t})'$, and $\Omega = \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix}$. We assume that the lag order, p, is finite and that Y_t is stationary, meaning that we can invert $\Phi(L)$ to solve for the vector Wold form:

$$Y_t = \mu + \Psi(L)e_t \tag{5}$$

where $\mu = \Phi(1)^{-1}c$, and $\Psi(L) = \Phi(L)^{-1}$.

DOI: 10.1002/jae

The key structural assumption is that the reduced-form representation in (5) corresponds to the following linear structural model:

$$Y_t = \mu + \beta(L)\varepsilon_t,\ \varepsilon_t \sim N(0, D) \tag{6}$$

where $\varepsilon_t \equiv (\varepsilon_t^{\mathrm{AS}}, \varepsilon_t^{\mathrm{AD}})'$, $D \equiv \begin{bmatrix} \sigma_{\mathrm{AS}}^2 & 0 \\ 0 & \sigma_{\mathrm{AD}}^2 \end{bmatrix}$, $\beta(L) \equiv \sum_{k=0}^{\infty} \beta_k L^k$, and $\beta_0 = \begin{bmatrix} 1 & 1 \\ \beta_{0,21} & \beta_{0,22} \end{bmatrix}$. That is, output growth and the unemployment rate depend on current and lagged values of aggregate supply shocks ($\varepsilon_t^{\mathrm{AS}}$) and aggregate demand shocks ($\varepsilon_t^{\mathrm{AD}}$). For the purpose of separating out changes in the size of shocks from changes in the contemporaneous propagation of shocks across series, we normalize the size of both structural shocks in terms of their initial impact on output growth.[5] In terms of this model, D reflects the size of shocks and $\beta(L)$ reflects the propagation of the shocks. The matrix β_0 captures the proportional impact of the shocks on each series, rather than the size of shocks. In particular, regardless of the size of shocks or how the shocks are normalized, β_0 only changes given a change in the contemporaneous propagation of shocks across series.

We identify aggregate demand shocks, aggregate supply shocks, and the propagation of the shocks using a long-run restriction, as in Blanchard and Quah (1989). Briefly, long-run identification proceeds by assuming no long-run effect of the aggregate demand shock on the level of output (y_t); that is, $\beta(1)$ is a lower triangular matrix. This restriction is then used to solve for $\beta(L)$ and D. The details of this solution are provided in the Appendix.

4.2. Counterfactual Experiments for the Structural VAR

Given the structural model in (6), the variance of Y_t is a function of the shock variances D and the propagation $\beta(L)$:

$$\Gamma_0 = \sum_{k=0}^{\infty} \beta_k D \beta_k' \tag{7}$$

As with the simple AR(1) example, we can consider a structural break in the variance of Y_t and perform counterfactual analysis based on the structural break. Letting $i, j = 1, 2$ index the structural regime for propagation and shocks, respectively:

$$\Gamma_0^{(i,j)} = \sum_{k=0}^{\infty} \beta_k^{(j)} D^{(i)} \beta_k'^{(j)} \tag{8}$$

Then, $i = j$ produces the actual variance estimates for the two structural regimes and $i \neq j$ produces counterfactual variances.

4.3. Bayesian Estimation for the Structural VAR

The reduced-form VAR in (4) is equivalently written as

$$Y_t = \Pi' X_t + e_t \tag{9}$$

[5] Given the normalization, there is an implicit assumption that both structural shocks have a non-zero initial impact on output growth.

DOI: 10.1002/jae

where $\Pi = [c \quad \Phi_1 \quad . \quad . \quad \Phi_p]'$ and $X_t = [1 \quad Y'_{t-1} \quad . \quad . \quad Y'_{t-p}]'$. We allow for a one-time structural break in the parameters of (9) at the breakdate τ. That is, we have

$$Y_t = \left\{ \begin{array}{ll} \Pi^{(1)'}X_t + e_t, e_t \sim N(0, \Omega^{(1)}); & t < \tau \\ \Pi^{(2)'}X_t + e_t, e_t \sim N(0, \Omega^{(2)}); & t \geq \tau \end{array} \right\} \tag{10}$$

The breakdate τ is assumed to be unknown, and is thus treated as a parameter to be estimated.

We estimate the parameters of the model in (10) using Bayesian methods. We assume a normal prior for the intercept/slope parameters, $\text{vec}(\Pi^{(i)}) \sim N(\underline{\Pi}, \underline{\Sigma})$, $i = 1, 2$, an inverted Wishart distribution for the variance–covariance parameters, $\Omega^{(i)^{-1}} \sim W(\underline{\nu}, \underline{S})$, and a uniform prior for the breakdate, $\tau \sim U(\kappa T, (1-\kappa)T)$, where κ determines the fraction of the sample period over which a structural break is allowed to occur. While the joint posterior density of $(\Pi^{(1)}, \Pi^{(2)}, \Omega^{(1)}, \Omega^{(2)}, \tau)$ is not available analytically, it can be simulated via the Gibbs sampler (Gelfand and Smith, 1990). In particular, we obtain 10,000 draws from the Gibbs sampler for each parameter, after discarding an initial 5000 draws to ensure convergence. The details of the Gibbs sampling procedure are provided in the Appendix.

Given draws from the joint posterior density for the reduced-form parameters, we can solve for the structural propagation parameters and the structural variance parameters using the long-run identification procedure. Then, we can form the actual and counterfactual variances $\Gamma_0^{(i,j)}$ by mixing the variance parameters from structural regime i with the propagation parameters in structural regime j, where $i, j = 1, 2$. Doing this for each draw from the Gibbs sampler provides us with posterior densities for the counterfactual variances that are neither conditional on model parameters nor conditional on the timing of the unknown breakdate, but take uncertainty about these parameters and the breakdate into account.

5. EMPIRICAL RESULTS

In this section we present the empirical results for our investigation of the sources of the volatility reduction in US real GDP growth based on the long-run structural VAR model presented in the previous section. First, we describe the details of the data, model specification, and priors used for estimation. Second, we report the estimation results for the structural break and model parameters. Third, we report the findings for the counterfactual experiments based on the Bayesian analysis.

5.1. Data, Model Specification, and Priors

The data series are 100 times the log first differences of US real GDP and the level of the US civilian unemployment rate. We sample seasonally adjusted measures at a quarterly frequency for the period 1960 : Q1 to 2005 : Q4. The raw data were obtained from the St Louis Fed website (http://www.stls.frb.org/fred/).

Beyond which variables to include in the VAR model, the main specification issue is how many lags to include in estimation. We report results for a model with four lags of quarterly data, although we note that the results are qualitatively very similar both for models with fewer lags and for models with more lags.

We consider noninformative priors for simplicity, although the results are robust to more informative priors based on training sample information. The priors are 'noninformative' in the

DOI: 10.1002/jae

sense that the priors for the VAR parameters would have resulted in posterior means and standard deviations that are the same as OLS estimates and standard errors if the breakdate were known, although because we also assume a noninformative (i.e., uniform) prior for the breakdate, the results are not identical to OLS. For the intercept/slope parameters, the hyperparameters for the normal prior are $\underline{\Pi} = 0$ and $\underline{\Sigma} = (1.0 \times 10^6) \times I$, with rejection sampling based on the largest modulus of the eigenvalues for the companion form representation of the VAR to ensure stationarity. For the variance–covariance parameters, the hyperparameters for the Wishart prior are $\underline{\nu} = 0$ and $\underline{S} = (1.0 \times 10^6) \times I$, which is a flat, but improper prior (i.e., it does not integrate to one) for the precision matrix $\Omega^{(i)^{-1}}$. For the unknown breakdate, the hyperparameter for the uniform prior is $\kappa = 0.15$, which corresponds to equal weights for a breakdate sometime between 1966 : Q3 and 1998 : Q4.

5.2. Bayesian Estimation Results

We first look at the timing of the structural break. Figure 1 displays the posterior density and cumulative distribution for the breakdate parameter τ. It is clear from the sharpness of the posterior density and the steepness of the corresponding cumulative distribution that the data are highly informative about the presence of a structural break in the VAR parameters sometime between 1982 and 1988. In particular, if there were no structural break, the posterior would be relatively flat, like the uniform prior. Also, it is clear from the precision of the posterior density that it would be robust to a wide range of truncations of possible breakdates implied by the hyperparameter κ for the uniform prior.[6] The posterior includes 1984, which is often cited as the most likely date of a volatility reduction in output growth. Meanwhile, it should be noted that the dates for the possible timing of the structural break reflect that the model allows for a change in mean and variance parameters for both output growth and the unemployment rate, rather than simply a change in output growth volatility.

Given the results for the structural break, we next look at the estimates for structural shocks and propagation before and after the break. First, in terms of structural shocks, Table I presents quartiles for the posterior distributions of the aggregate supply (AS) and aggregate demand (AD) shock standard deviations, σ_{AS} and σ_{AD}. The results suggest that the structural break corresponds to a reduction in the size of structural shocks, at least in terms of their impact on output growth. The quartile values for the posterior distributions of σ_{AS} and σ_{AD} both fall by half, although the reduction in aggregate demand shocks is smaller in absolute magnitude. Second, in terms of propagation, Figure 2 displays the median and quartiles of the posterior distribution of the impulse response functions, $\frac{\partial y_{t+q}}{\partial \varepsilon_t^{\mathrm{AS}}}, \frac{\partial y_{t+q}}{\partial \varepsilon_t^{\mathrm{AD}}}$, for $q = 1$ to 40.[7] The impulse response functions have the same general shape as reported in Blanchard and Quah (1989). Aggregate supply shocks generate persistent long-run effects, while aggregate demand shocks generate a hump-shaped response that decays to zero over a business cycle horizon. Notably, the impulse response functions look reasonably similar before and after the structural break. However, the relatively wide quartile bands imply a fair degree of uncertainty about whether there was a change in propagation.

[6] However, as with classical estimation of unknown breakdates, it is important to have a sufficient number of observations on either side of a breakdate to avoid irregularities in the likelihood function when a small number of observations can be overfit by a heavily parameterized model.

[7] To conserve space and because our primary focus is on output growth volatility, we do not report the impulse response functions for the unemployment rate.

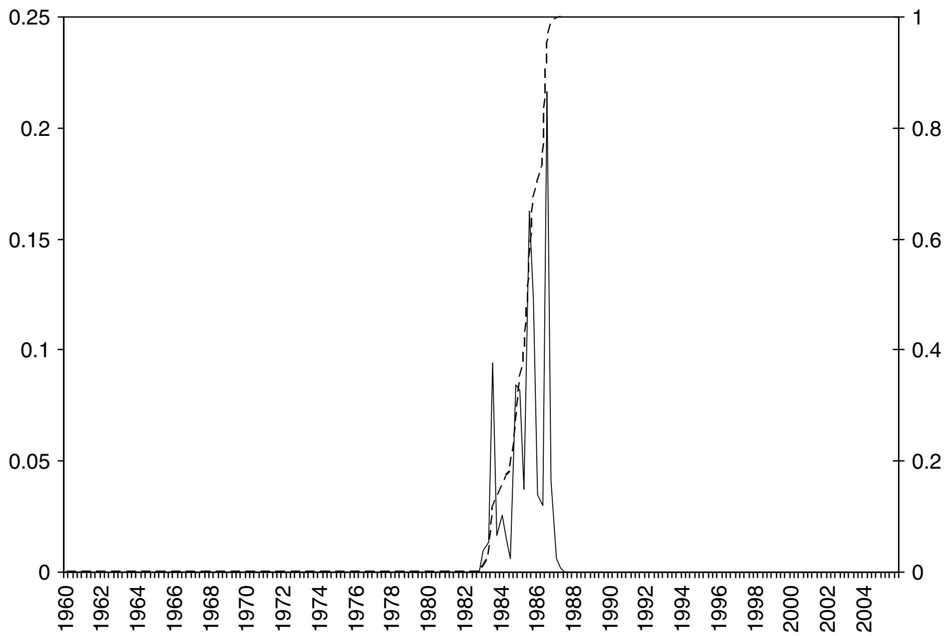

Figure 1. Posterior density for the timing of the structural break. *Note*: The x-axis units are quarters in the sample period. The y-axis units refer to frequency as a fraction of the total number of simulations. The solid line is the simulated posterior density for the unknown breakdate. The dashed line is the corresponding cumulative distribution (right axis)

Table I. Quartiles of posterior distributions for standard deviations of aggregate supply and aggregate demand shocks

Period	25th percentile	Median	75th percentile
Aggregate supply			
1	2.79	3.38	3.77
2	1.47	1.68	1.84
Aggregate demand			
1	1.08	1.92	2.71
2	0.37	0.71	1.08

Note: The structural shocks are normalized in terms of their initial impact on output growth. The standard deviations are expressed in terms of quarterly percentage points. Period 1 refers to the pre-break period and Period 2 to the post-break period.

5.3. Results for Counterfactual Experiments

Given the estimation results for the structural VAR, it is useful to consider counterfactual experiments to determine whether the change in the size of structural shocks could have generated the entire reduction in output growth volatility or, alternatively, whether even small changes in propagation could have generated large changes in output growth volatility due to the nonlinear impact of propagation on variance. Meanwhile, given uncertainty about the extent of a change

DOI: 10.1002/jae

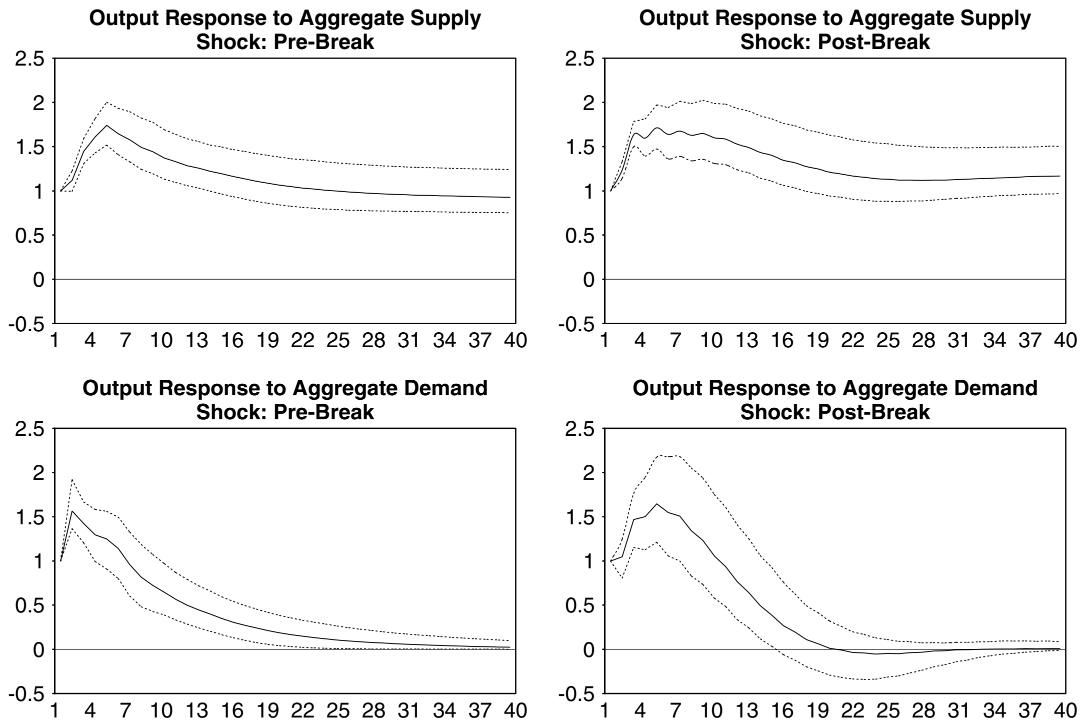

Figure 2. Impulse response functions. *Note*: The x-axis units are quarters after a shock. The y-axis units refer to impact of the shock on output. Results are based on simulated posterior distributions. Solid lines represent median responses. Dashed lines represent upper and lower quartile responses

in propagation, it is particularly important to consider the Bayesian approach to counterfactual analysis because it captures the extent to which an imprecisely estimated change in propagation could have generated a large volatility reduction, even if posterior means suggest little reduction.

As discussed above, the counterfactual experiments involve mixing shock variance and propagation parameters from before and after the structural break and solving for the resulting unconditional variances. For ease of interpretation of units, we report our inferences in terms of the standard deviations rather than variances. Also, while the structural VAR model includes the unemployment rate, we focus on the results for output growth. Specifically, we consider $\sigma_{\Delta y}^{(i,j)}$, which is the square root of the (1,1) element of $\Gamma_0^{(i,j)}$ in (8), where $i, j = 1, 2$. Table II presents quartiles for the posterior distributions of the actual and counterfactual standard deviations of output growth.

The first and second rows of Table II report results for the actual standard deviations, $\sigma_{\Delta y}^{(1,1)}$ and $\sigma_{\Delta y}^{(2,2)}$, in the pre-break and post-break periods, respectively. The results confirm the large reduction in output growth volatility after the structural break. Specifically, the median estimate fell by roughly 50%, while the upper quartile for $\sigma_{\Delta y}^{(2,2)}$ is well below the lower quartile for $\sigma_{\Delta y}^{(1,1)}$. These results are evident in Figure 3, which displays the posterior densities for the actual and counterfactual standard deviations of output growth. Almost all of the posterior density for the post-break standard deviation lies below the density for the pre-break standard deviation.

Table II. Quartiles of posterior distributions for standard deviations of real GDP growth

Propagation	Shocks	25th percentile	Median	75th percentile
Actual				
1	1	4.37	4.66	5.00
2	2	2.08	2.25	2.46
Counterfactuals				
2	1	4.57	5.02	5.86
1	2	2.11	2.32	2.68
Individual shock counterfactuals				
1	AS 1 AD 2	3.46	4.02	4.60
1	AD 1 AS 2	2.69	3.20	3.87

Note: Standard deviations are expressed in terms of quarterly percentage points. Period 1 refers to the pre-break period and Period 2 to the post-break period. AS refers to aggregate supply. AD refers to aggregate demand.

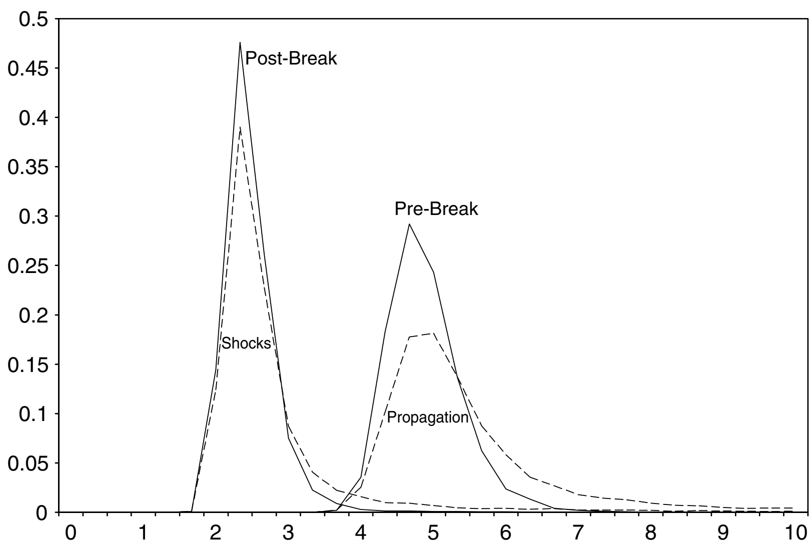

Figure 3. Posterior densities for actual and counterfactual standard deviations of output growth. *Note*: The x-axis units are values for the standard deviations expressed in terms of quarterly percentage points. The y-axis units refer to frequency as a fraction of the total number of simulations. Solid lines denote the simulated posterior densities for the pre-break and post-break periods. Dashed lines denote the simulated posterior densities for the counterfactuals, where 'Shocks' refers to post-break shocks and pre-break propagation, and 'Propagation' refers to post-break propagation and pre-break shocks

The third and fourth rows of Table II report results for the counterfactual standard deviations, $\sigma_{\Delta y}^{(1,2)}$ and $\sigma_{\Delta y}^{(2,1)}$, corresponding to a change in propagation alone and a change in shocks alone, respectively. A change in propagation alone is entirely unsuccessful at generating the actual reduction in variance. Indeed, the median of $\sigma_{\Delta y}^{(1,2)}$ is actually above that of the actual standard deviation, $\sigma_{\Delta y}^{(1,1)}$, in the pre-break period. A change in shocks alone is able to generate

DOI: 10.1002/jae

a volatility reduction, with the median value close to that of the actual standard deviation, $\sigma_{\Delta y}^{(2,2)}$, in the post-break period. Again, these results are evident in Figure 3, which shows that the counterfactual standard deviation given a change in shocks alone is almost as precisely estimated as the actual post-break standard deviation, but the counterfactual standard deviation given a change in propagation alone is less precisely estimated than the pre-break standard deviation.

While the results in Table II and Figure 3 clearly suggest that it is a change in shocks, not propagation, that could have generated the entire reduction in output growth volatility, the Bayesian analysis also allows for more precise inferences along these lines. In particular, as discussed in Section 3, we can calculate ratios of functions of the actual and counterfactual variances and evaluate the probabilities that these ratios are less than a specified fixed value. Based on this idea, Figure 4 displays probabilities that counterfactual reductions in the standard deviation of output growth are larger than fixed percentages of the actual reduction in the standard deviation of output growth. At one extreme, we consider 0% of the actual change (i.e., no change). The probability that a change in shocks alone would have reduced volatility is high at about 0.9, while the probability that a change in propagation alone would have reduced volatility is low at about 0.3. At the other extreme, we consider 100% of the actual change. The probability that a change in shocks alone would have reduced volatility by more than the actual change is just below 0.5, which would correspond to 'even odds' that one reduction is larger than the other. The probability that a change in propagation alone would have reduced volatility by more than the actual change is actually 0 (i.e., the Bayesian analysis suggests that there is no way a change in propagation alone would have reduced volatility more than the actual change). In between the extremes, these probabilities are useful for illustrating the economic significance of a change in volatility. In particular, there is a reasonably high probability (above 0.8) that the counterfactual change in volatility given a change in shocks alone would have lowered volatility by more than 75% of the actual reduction. Conversely, there is only a small probability (less than 0.05) that the counterfactual change in volatility given a change in propagation alone would have lowered volatility by even 25% of the actual reduction.

Given the primary role played by shocks in the overall volatility reduction, we examine the relative importance of AS and AD shocks separately. The fifth and sixth rows of Table II report results for the counterfactual standard deviations, $\sigma_{\Delta y}^{(1,\mathrm{AS1/AD2})}$ and $\sigma_{\Delta y}^{(1,\mathrm{AD1/AS2})}$, corresponding to a change in AS shocks alone and a change in AD shocks alone, respectively. The results suggest that AS shocks are more successful than AD shocks at generating a large reduction in volatility. The median for $\sigma_{\Delta y}^{(1,\mathrm{AS1/AD2})}$, which corresponds to a change in AD shocks alone, is only a bit lower than that of the actual standard deviation, $\sigma_{\Delta y}^{(1,1)}$, in the pre-break period. By contrast, the median for $\sigma_{\Delta y}^{(1,\mathrm{AD1/AS2})}$, which corresponds to a change in AS shocks alone, is closer to that of the actual standard deviation, $\sigma_{\Delta y}^{(2,2)}$, in the post-break period, although the lower quartile for $\sigma_{\Delta y}^{(1,\mathrm{AD1/AS2})}$ is still above the upper quartiles for $\sigma_{\Delta y}^{(2,2)}$ and $\sigma_{\Delta y}^{(2,1)}$, suggesting that AS shocks could not have generated the entire volatility reduction on their own. These results are evident in Figure 5, which displays the posterior densities for the actual and individual shock counterfactual standard deviations of output growth. Much of the posterior density for the counterfactual standard deviation given a change in AS shocks alone lies below the density for the counterfactual standard deviation given a change in AD shocks alone,

DOI: 10.1002/jae

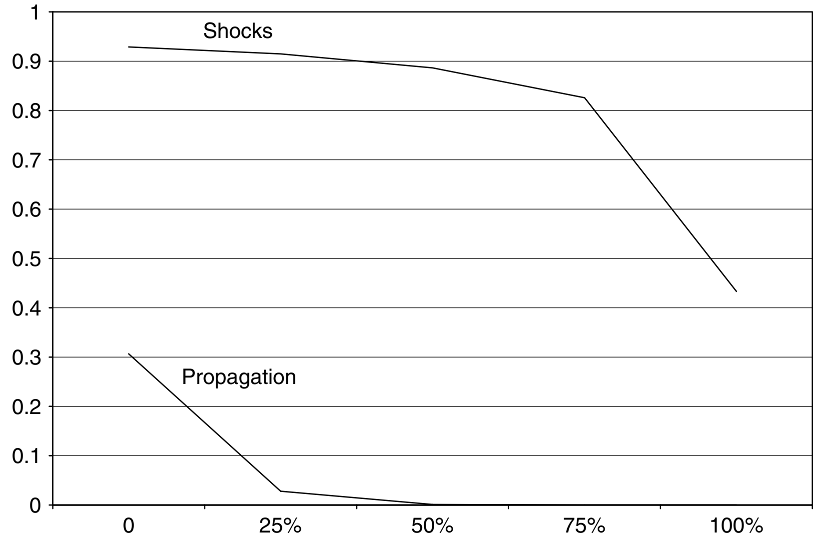

Figure 4. Probabilities of counterfactual changes in volatility. *Note*: The x-axis units are percentages of the actual change in the standard deviation of output growth. The y-axis units are probabilities. Solid lines denote the probabilities that the counterfactual reductions in volatility are larger than fixed percentages of the actual reduction in volatility, where 'Shocks' refers to post-break shocks and pre-break propagation, and 'Propagation' refers to post-break propagation and pre-break shocks

although these counterfactuals are clearly less precisely estimated than the actual standard deviations.

Figure 6 displays the probabilities that the counterfactual reductions in volatility implied by changes in one type of shock alone are larger than fixed percentages of the actual reduction in the standard deviation of output growth. For both AS and AD shocks, there is a high probability that the corresponding counterfactual change would have reduced volatility, with the probability above 0.9 for AS shocks and above 0.8 for AD shocks. At the same time, there is a low probability that either AS shocks alone or AD shocks alone would have generated the entire reduction, with the probability about 0.1 for AS shocks and about 0.05 for AD shocks. There is a higher probability of a large volatility reduction for AS shocks alone than for AD shocks alone, with 'even odds' that a change given AS shocks alone would have generated about 60% of the actual reduction and a change given AD shocks alone would have generated about 25% of the actual reduction.

6. ROBUSTNESS

There are some issues with Blanchard and Quah's (1989) structural VAR based on a long-run restriction that potentially argue for consideration of a short-run structural VAR instead. First, the long-run identification presented above is predicated on the assumption that there is only one type of AS shock and one type of AD shock, where 'type' is defined in terms of the impact of the shock on the dynamic relationship between output and unemployment. If there are multiple types of AS and AD shocks, then it would be important to include additional variables and allow for additional shocks in the structural VAR. However, it is, arguably, much more practical to

DOI: 10.1002/jae

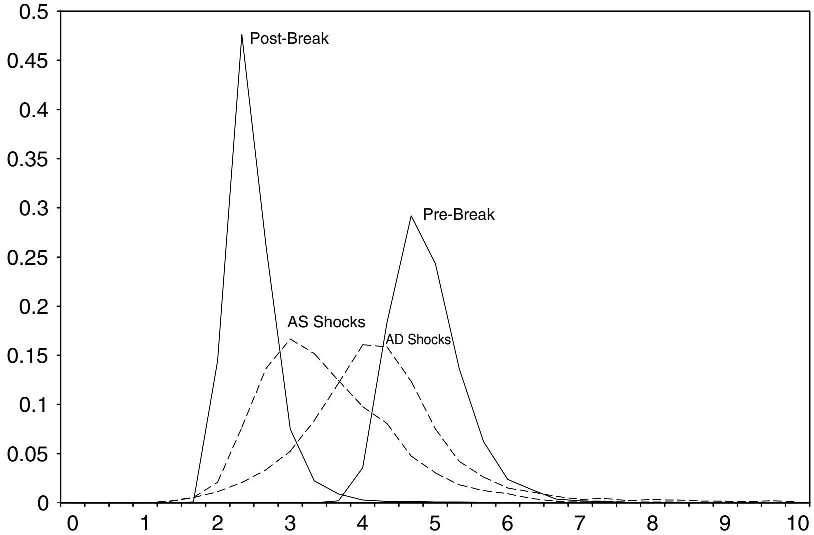

Figure 5. Posterior densities for actual and individual shock counterfactual standard deviations of output growth. *Note*: The x-axis units are values for the standard deviations expressed in terms of quarterly percentage points. The y-axis units refer to frequency as a fraction of the total number of simulations. Solid lines denote the simulated posterior densities for the pre-break and post-break periods. Dashed lines denote the simulated posterior densities for the counterfactuals, where 'AS Shocks' refers to post-break aggregate supply shocks and pre-break aggregate demand shocks and propagation, and 'AD Shocks' refers to post-break aggregate demand shocks and pre-break aggregate supply shocks and propagation

include a larger number of variables in a short-run structural VAR and still identify key structural shocks. Second, from an econometric perspective, identification based on the spectral density at frequency zero, as is done with long-run restrictions, is much weaker than identification based on the short-run variance–covariance matrix of forecast errors (see Faust and Leeper, 1997, on this point). In particular, estimates of the spectral density at frequency zero can be highly sensitive to the number of lags included in estimation, while estimates of the short-run variance–covariance matrix are somewhat more stable.

With these concerns in mind, we conduct a robustness check on our main model by also considering a short-run monetary VAR along the lines of what was considered in Boivin and Giannoni (2003) and Ahmed *et al*. (2004). The short-run monetary VAR model includes output growth, consumer-price inflation, commodity-price inflation, and the monetary policy instrument, and identifies monetary shocks on the basis of the short-run restriction of delayed responses for output and prices to monetary shocks. Specifically, in terms of the notation in Section 4, the impact matrix β_0 for the structural model is assumed to be lower triangular, with the policy instrument variable being the last element of Y_t. Identification is based on the Choleski factorization of the variance–covariance matrix Ω, with normalization of the size of shocks being in terms of their initial impact on their associated variable (e.g., the monetary shock in terms of the monetary policy variable).

The data series for output growth is the same as before. The series for the inflation variables are 100 times the log first differences of the US CPI and the Commodity Research Bureau Spot Price Index for all Commodities, respectively. The series for the monetary policy instrument is the level

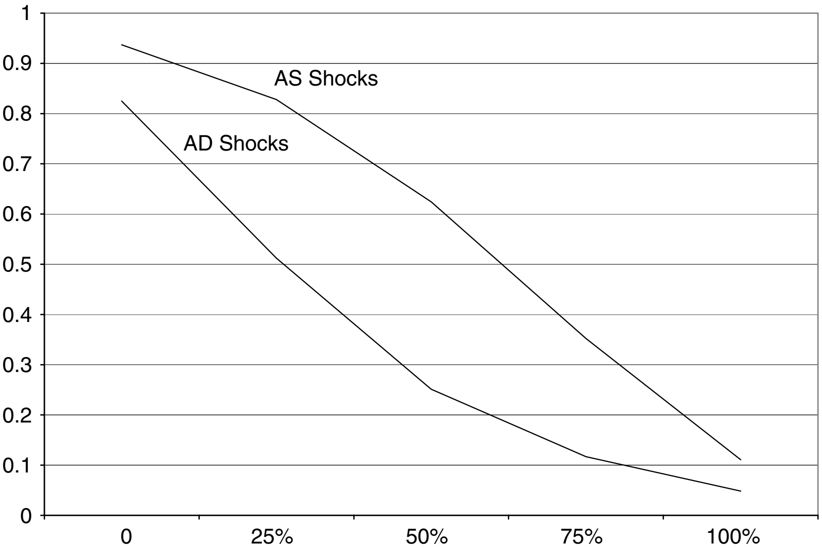

Figure 6. Probabilities of individual shock counterfactual changes in volatility. *Note*: The x-axis units are percentages of the actual change in the standard deviation of output growth. The y-axis units are probabilities. Solid lines denote the probabilities that the counterfactual reductions in volatility are larger than fixed percentages of the actual reduction in volatility, where 'AS Shocks' refers to post-break aggregate supply shocks and pre-break aggregate demand shocks and propagation, and 'AD Shocks' refers to post-break aggregate demand shocks and pre-break aggregate supply shocks and propagation

of the federal funds rate. As before, we sample seasonally adjusted (when applicable) measures at a quarterly frequency for the period of 1960 : Q1 to 2005 : Q4 and the raw data were obtained from the St Louis Fed website (http://www.stls.frb.org/fred/). The number of lags and the priors are the same as before, with only a change in the number of variables in the VAR model.

Table III presents quartiles for the posterior distributions of the actual and counterfactual standard deviations of output growth given estimation based on the short-run monetary VAR. The results are qualitatively very similar to those in Table II. A change in propagation alone generates a standard deviation that is close to the standard deviation in the pre-break period, while a change in shocks alone generates a standard deviation that is close to the standard deviation in the post-break period. Also, consistent with the findings for AD shocks, monetary shocks appear to generate only a small reduction in volatility. These results are qualitatively very similar to the findings in Stock and Watson (2002), Boivin and Giannoni (2003), and Ahmed *et al*. (2004), but our results provide statistical credibility for the idea that shocks, not propagation, could have generated the volatility reduction in US output growth.

7. CONCLUSIONS

We have used counterfactual experiments to study the sources of the 'Great Moderation' in US real GDP growth volatility in the 1980s. The contribution of our paper is to use Bayesian analysis to make inferences about the counterfactual variance of real GDP growth. Contrary to an existing literature, which constructs this counterfactual variance using classical estimation and

DOI: 10.1002/jae

Table III. Quartiles of posterior distributions for standard deviations of real GDP growth based on a short-run monetary VAR

Propagation	Shocks	25th percentile	Median	75th percentile
Actual				
1	1	4.72	5.06	5.52
2	2	2.26	2.46	2.71
Counterfactuals				
2	1	4.60	5.00	5.52
1	2	2.46	2.67	2.94
Individual shock counterfactuals				
1	Monetary 1 Other shocks 2	3.05	3.33	3.73
1	Other shocks 1 Monetary 2	4.34	4.65	5.05

Note: Standard deviations are expressed in terms of quarterly percentage points. Period 1 refers to the pre-break period and Period 2 to the post-break period. The results for this table are based on a short-run monetary VAR model of output growth, consumer price inflation, commodity price inflation, and the federal funds rate, with four lags and monetary shocks identified from the restriction that they impact output growth, consumer price inflation, and commodity price inflation only with a lag.

point estimates, the Bayesian analysis provides posterior densities that give a sense of statistical precision about counterfactual quantities and allow us to evaluate the probability of economically significant changes in those quantities. The results support the notion that the volatility reduction in output growth was driven by smaller shocks hitting the economy rather than changes in the propagation of those shocks. Indeed, we have the extremely strong finding that there is zero probability that a change in propagation alone would have generated a volatility reduction as large as what actually occurred. Furthermore, in terms of smaller shocks, we find that a change in the size of aggregate demand shocks and monetary shocks would not have generated much of a volatility reduction on their own. Thus, our results provide further statistical support, beyond what was suggested by classical estimation and point estimates alone, for the 'good luck' hypothesis in explaining the Great Moderation.

APPENDIX

A.1. Identification of Structural VAR via Long-Run Restrictions

The structural VAR in (6) is identified using a long-run restriction, as in Blanchard and Quah (1989). Long-run identification proceeds by assuming that there is no long-run effect of the aggregate demand shock on the level of output (y_t). Technically, the spectral density of the aggregate demand component of output growth is equal to zero at frequency zero. Given (5) and (6), the spectral density for X_t at frequency zero is proportional to the long-run variance–covariance matrix, denoted Λ:

$$\Lambda = \Psi(1)\Omega\Psi(1)' = \beta(1)D\beta(1)' \tag{A.1}$$

In order for the aggregate demand component of output growth to have a spectral density of zero at frequency zero, $\beta(1)$ must be lower triangular. That is, the aggregate demand shock

DOI: 10.1002/jae

does not contribute to the permanent movements in output. In practice, identification proceeds by constructing Λ using $\Psi(1)$ and Ω from the reduced-form model. Then, the Choleski factorization of Λ provides a unique lower triangular matrix that is equivalent to $\beta(1)D^{1/2}$.

Given $\beta(1)D^{1/2}$, full identification of the structural model follows from the relationship between the reduced-form forecast errors and the structural shocks implied by the assumption that forecast errors and structural shocks are both uncorrelated across time:

$$e_t = \beta_0 \varepsilon_t \tag{A.2}$$

Then, substituting (A.2) into (5) and comparing to (6) implies

$$\Psi(L)\beta_0 = \beta(L) \tag{A.3}$$

The condition (A.3) is general and holds for the following case:

$$\Psi(1)\beta_0 = \beta(1) \tag{A.4}$$

Rearranging and multiplying both sides by $D^{1/2}$, we can determine the un-normalized impact matrix $\theta_0 = \beta_0 D^{1/2}$:

$$\theta_0 = \Psi(1)^{-1}\beta(1)D^{1/2} \tag{A.5}$$

Then, given (A.5) and the normalization of shocks in terms of their initial impact on output growth, we can solve for β_0 and D as follows:

$$\beta_0 = \begin{bmatrix} 1 & 1 \\ \theta_{0,21}/\theta_{0,11} & \theta_{0,22}/\theta_{0,12} \end{bmatrix}, D = \begin{bmatrix} \theta_{0,11}^2 & 0 \\ 0 & \theta_{0,12}^2 \end{bmatrix} \tag{A.6}$$

A.2. Bayesian Estimation of the Reduced-Form VAR via the Gibbs Sampler

We begin by writing the reduced-form VAR model in (9) in matrix notation. Define the vectors $\Delta\tilde{y} = (\Delta y_1, \ldots, \Delta y_T)'$, $\tilde{u} = (u_1, \ldots, u_T)'$, and $\tilde{e}_i = (e_{i1}, \ldots, e_{iT})'$ and form the matrices $\mathrm{Y} = [\,\Delta\tilde{y} \quad \tilde{u}\,]$, $\mathrm{X} = \lfloor\, \tilde{1} \quad Y_{-1} \quad . \quad . \quad Y_{-p} \,\rfloor$, and $e = [\,\tilde{e}_1 \quad \tilde{e}_2\,]$, where Y_{-k} holds the kth lag of $\Delta\tilde{y}$ and $\tilde{u}$. The model in (9) is then equivalently written as

$$\mathrm{Y} = \mathrm{X}\Pi + e \tag{A.7}$$

Again, we are interested in the case where the parameters of the reduced-from VAR undergo a one-time structural break at the breakdate τ. Thus the model in (A.7) can be partitioned into the two subsamples:

$$\begin{aligned} \mathrm{Y}^{(1)} &= \mathrm{X}^{(1)}\Pi^{(1)} + e^{(1)} \\ \mathrm{Y}^{(2)} &= \mathrm{X}^{(2)}\Pi^{(2)} + e^{(2)} \end{aligned} \tag{A.8}$$

where $\mathrm{Y}^{(1)}$, $\mathrm{X}^{(1)}$, $\Pi^{(1)}$ and $e^{(1)}$ hold the first $\tau - 1$ rows of Y, X, Π and e respectively, while $\mathrm{Y}^{(2)}$, $\mathrm{X}^{(2)}$, $\Pi^{(2)}$ and $e^{(2)}$ hold last the $T - (\tau - 1)$ rows of Y, X, Π and e respectively.

DOI: 10.1002/jae

The model in (A.8) can be estimated in a Bayesian framework using the Gibbs sampler. In particular, starting with arbitrary initial values for $\Omega^{(1)}$, $\Omega^{(2)}$ and τ, the Gibbs sampler proceeds by iterating the following three steps:

1. Generate $\text{vec}(\Pi^{(1)})$ and $\text{vec}(\Pi^{(2)})$ from their conditional posterior density given previous values for $\Omega^{(1)}$, $\Omega^{(2)}$ and τ. This density is given by

$$\text{vec}(\Pi^{(i)})|\Omega^{(i)}, \tau, Y^{(i)}, X^{(i)} \sim N(\overline{\Pi}^{(i)}, \overline{\Sigma}^{(i)})$$

where $\overline{\Sigma}^{(i)} = \left(\underline{\Sigma}^{-1} + \Omega^{(i)-1} \otimes X^{(i)\prime}X^{(i)}\right)^{-1}$, $\overline{\Pi}^{(i)} = \overline{\Sigma}^{(i)} \times (\underline{\Sigma}^{-1}\underline{\Pi} + (\Omega^{(i)-1} \otimes X^{(i)\prime})\text{vec}(Y^{(i)}))$. A generated value for $\Pi^{(i)}$ is then formed from the generated value of $\text{vec}(\Pi^{(i)})$.
2. Generate $\Omega^{(1)}$ and $\Omega^{(2)}$ from their conditional posterior distribution given previous values for $\Pi^{(1)}$, $\Pi^{(2)}$ and τ. This density is given by

$$\Omega^{(i)^{-1}}|\Pi^{(i)}, \tau, Y^{(i)}, X^{(i)} \sim W(\overline{\nu}^{(i)}, \overline{S}^{(i)})$$

where $\overline{\nu}^{(1)} = \underline{\nu} + \tau - 1$, $\overline{\nu}^{(2)} = \underline{\nu} + T - (\tau - 1)$ and $\overline{S} = (\underline{S}^{-1} + e'e)$.
3. Generate τ from its conditional posterior distribution given previous values of $\Pi^{(1)}$, $\Pi^{(2)}$, $\Omega^{(1)}$, and $\Omega^{(2)}$. This density is given by

$$p(\tau|\Pi^{(1)}, \Pi^{(2)}, \Omega^{(1)}, \Omega^{(2)}, Y, X) \propto L(\tau|\Pi^{(1)}, \Pi^{(2)}, \Omega^{(1)}, \Omega^{(2)}, Y, X)$$

where the likelihood is evaluated for all possible breakdates between κT and $(1-\kappa)T$, with $L(\tau|\Pi^{(1)}, \Pi^{(2)}, \Omega^{(1)}, \Omega^{(2)}, Y, X)$ equal to the joint normal density of the data Y given $\Pi^{(1)}$, $\Pi^{(2)}$, $\Omega^{(1)}$, $\Omega^{(2)}$ and τ.

After a sufficient number of initial draws, the subsequent draws from the Gibbs sampler will no longer reflect the arbitrary starting values, but will behave like draws from the joint posterior density, $p(\Pi^{(1)}, \Pi^{(2)}, \Omega^{(1)}, \Omega^{(2)}, \tau|Y, X)$.

ACKNOWLEDGEMENTS

We would like to thank Arabinda Basistha, Ryan Compton, Robert Rich, Eric Zivot, the editor, and two anonymous referees for helpful comments and suggestions. Responsibility for any errors is our own. Morley acknowledges support from the Weidenbaum Center on the Economy, Government, and Public Policy. Much of this paper was written while Piger was a Senior Economist at the Federal Reserve Bank of St Louis. The views expressed in this paper should not be interpreted as those of the Weidenbaum Center, the Federal Reserve Bank of St Louis, or the Federal Reserve System.

REFERENCES

Ahmed S, Levin A, Wilson BA. 2004. Recent US macroeconomic stability: good policy, good practices, or good luck? *Review of Economics and Statistics* **86**: 824–832.

Bernanke B. 2004. The Great Moderation. Meetings of the Eastern Economic Association, Washington, DC.

Blanchard O, Quah D. 1989. The dynamic effects of aggregate demand and supply disturbances. *American Economic Review* **79**: 655–673.

Blanchard O, Simon J. 2001. The long and large decline in US output volatility. *Brookings Papers on Economic Activity* 135–174.

Boivin J, Giannoni M. 2003. Has monetary policy become more effective? NBER Working Paper no. 9459.

Chauvet M, Potter S. 2001. Recent changes in the US business cycle. *The Manchester School* **69**: 481–508.

Choi K, Jung C. Forthcoming, The sources of the decline in US output volatility. *Contemporary Economic Policy*.

Dueker M, Nelson C. 2003. Business cycle detrending of macroeconomic data via a latent business cycle index. Working Paper 2002-025B, Federal Reserve Bank of St Louis.

Faust J, Leeper EM. 1997. When do long-run identifying restrictions give reliable results? *Journal of Business and Economic Statistics* **15**: 345–353.

Gelfand AE, Smith AFM. 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* **85**: 398–409.

Gordon R. 2006. What caused the decline in US business cycle volatility?' NBER Working Paper no. 11777.

Herrera AM, Pesavento E. 2005. The decline in US output volatility: structural changes and inventory investment. *Journal of Business and Economic Statistics* **23**: 462–472.

Kahn J, McConnell MM, Perez-Quiros G. 2002. On the causes of the increased of the US economy, *Economic Policy Review* **8**: 183–202.

Kim C-J, Nelson CR. 1999. 'Has the US economy become more stable? A Bayesian approach based on a Markov-switching model of the business cycle. *Review of Economics and Statistics* **81**: 608–616.

Kim C-J, Nelson CR, Piger J. 2004. The less-volatile US economy: a Bayesian investigation of timing, breadth, and potential explanations. *Journal of Business and Economic Statistics* **22**: 80–93.

Leeper EM, Zha T. 2003. Modest policy interventions. *Journal of Monetary Economics* **50**: 1673–1700.

Lucas R. 1976. Econometric policy evaluation: a critique. *Journal of Monetary Economics*, Supplementary Series **1**: 19–46.

McConnell MM, Perez-Quiros G. 2000. Output fluctuations in the United States: what has changed since the early 1980s? *American Economic Review* **90**: 1464–1476.

Niemira MP, Klein PA. 1994. *Forecasting Financial and Economic Cycles*. Wiley: New York.

Rudebusch GD. 2005. Assessing the Lucas Critique in monetary policy models. *Journal of Money, Credit, and Banking* **37**: 245–272.

Sensier M, van Dijk D. 2004. Testing for volatility changes in US macroeconomic time series. *Review of Economics and Statistics* **86**: 833–839.

Simon J. 2000. Essays in empirical macroeconomics. PhD thesis, MIT, Cambridge, MA.

Sims C, Zha T. 2006. Were there regime switches in US monetary policy? *American Economic Review* **96**: 54–81.

Stock JH, Watson MW. 2002. Has the business cycle changed and why? *NBER Macroeconomics Annual* **17**: 159–218.

DOI: 10.1002/jae