

Working Paper Series



This paper can be downloaded without charge from:
<http://www.richmondfed.org/publications/>



THE FEDERAL RESERVE BANK OF RICHMOND
RICHMOND ■ BALTIMORE ■ CHARLOTTE

Sales, Inventories, and Real Interest Rates: A Century of Stylized Facts*

Luca Benati
University of Bern[†]

Thomas A. Lubik
Federal Reserve Bank of Richmond[‡]

April 2012
Working Paper No. 12-02

Abstract

We use Bayesian time-varying parameters structural VARs with stochastic volatility to investigate changes in both the reduced-form and the structural correlations between business inventories and either sales growth or the real interest rate in the United States during both the interwar and the post-WWII periods. We identify four structural shocks by combining a single long-run restriction to identify a permanent output shock as in Blanchard and Quah (1989), with three sign restrictions to identify demand- and supply-side transitory shocks. We produce several new stylized facts which should inform the development of new models of inventories. In particular, we show that *(i)* during both the interwar and the post-WWII periods, the structural correlation between inventories and the real interest rate conditional on identified interest rate shocks is systematically positive; *(ii)* the reduced-form correlation between the two series is positive during the post-WWII period, but in line with the predictions of theory it is robustly negative during the interwar era; and *(iii)* during the interwar era, the correlations between inventories and either of the two other series exhibits a remarkably strong co-movement with output at the business-cycle frequencies.

Keywords: Bayesian VARs; stochastic volatility; time-varying parameters; structural VARs; long-run restrictions; sign restrictions; inventories; monetary policy; monetary regimes.

JEL codes: C11; C32; E32.

*The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Department of Economics, University of Bern, Schanzeneckstrasse 1, CH-3001 Bern, Switzerland. Email: luca.benati@vwi.unibe.ch

[‡]Research Department, Federal Reserve Bank of Richmond, 701 E Byrd Street, Richmond, VA 23218, USA. E-mail: Thomas.Lubik@rich.frb.org

1 Introduction

Inventories have been at the center of modern economics ever since its inception as a quantitative discipline. Starting with accelerator models designed to capture the inventory cycle, inventory behavior has attracted an enormous amount of attention in macroeconomics. The chief reason is that while inventory investment makes up only a small fraction of GDP, roughly one-half of a percent, it can contribute up to 90 percent of its cyclical variation.¹ At the same time, the inventory literature is rife with puzzles in the sense that theoretical inventory models cannot reproduce salient empirical facts in the data. Maccini, Moore, and Shaller (2010) classify two sets of puzzles. The traditional puzzles tend to describe the unconditional properties of inventories, such as the relative volatilities of inventories, production and sales, their persistence, and their comovement relationships with potential determinants. The second set of puzzles concerns the relationship between interest rates and inventories. While theory derives clear predictions - in particular, increases in interest rates lower inventory investment as the cost of holding inventories rises - this relationship is much more difficult to establish empirically.

In this paper we discuss one explanation of why these puzzles exist, namely the surprising difficulty of establishing invariant stylized facts over almost a century of U.S. data. We focus on the relationship between final sales, inventories and interest rates and utilize data from 1919 on. Using Bayesian time-varying parameter VARs with stochastic volatility we study changes in both the reduced-form and the structural correlations between these variables in the United States during both the interwar and the post-WWII periods. Our estimates are based on the identification of four structural shocks, whereby we combine a single long-run restriction to identify a permanent output shock with three sign restrictions to identify transitory demand- and supply-side shocks.

Our main findings are the following. We first show that the absence of a negative correlation between inventories and interest rates only pertains to the post-WWII period. During the interwar era, the correlation is, in fact, strongly negative. Second, we find that the behavior of this relationship is asymmetric over the business cycle and over subperiods. During the interwar period, the interest rate-inventories correlation exhibits positive co-movement with real output at business-cycle frequencies. Although the correlation is overall systematically negative, it is relatively less negative in business-cycle upswings, and even more negative during downturns. During the post-WWII period, on the other hand, the correlation does not exhibit any clear-cut systematic pattern of co-movement with the cyclical component of real output. Finally, estimates from the structural VAR specification show that the negative reduced-form correlation of the interwar era is induced by demand- and supply-side transitory shocks, and, especially at longer horizons, by the permanent output shock. The contribution of interest rate shocks, on the other hand, is uniformly positive. As

¹See Blinder and Maccini (1991a) and Blinder and Maccini (1991b).

for the post-WWII era, the positive reduced-form correlation between inventories and the interest rate can be traced back to the positive structural correlations induced by interest rate shocks, and, to a lesser extent, by demand-side transitory shocks and the permanent output shock.

We also look at the relationship between inventories and sales. During the interwar era, the correlation is mostly positive. At business-cycle frequencies it exhibits strong negative co-movement with the cyclical component of real output. During the post-WWII period, and in line with the evidence reported in Wen (2005), the correlation is mostly negative at very short horizons, and uniformly positive at longer horizons. Evidence from the structural VAR suggests that the positive correlation of the interwar era is due to the demand shock, and in later periods of the sample, due to the permanent output shock, whereas the correlation conditional on the other two shocks is negative. As for the post-WWII period, the evolution of the reduced-form correlation closely mirrors the evolution of the structural correlation conditional on the permanent output shock, being strongly positive both in the earlier and in the most recent parts of the sample at all horizons, and, at shorter horizons, being instead comparatively smaller, and sometimes negative, during the Great Inflation years.

Our results suggest that identified interest rate shocks have systematically and robustly induced a positive correlation between business inventories and the real interest rates during both the interwar and the post-WWII periods. The puzzle of the absence of a negative correlation between the two series over the post-WWII period identified in the previous literature is therefore deeper, and more intriguing, than previously thought. Even conditional on a structural shock for which both economic theory, and simple intuition, suggest that the correlation should be negative, our results highlight that, in fact, it has been systematically positive. It follows that the reduced-form relationship that masks this conditional behavior is therefore due to the changing nature of the comovement with output and the importance of the underlying shocks. In that sense, the interwar period was, in fact, markedly different from the post-WWII era.

We see our paper as making the following contributions to the literature. First, we establish that the reduced form relationships between inventories and salient variables has varied quite substantially since the Great Depression. While it is possible to identify periods of stable reduced-form relationships, the changes are often sufficiently large that we find it difficult to treat the stylized facts as time-invariant. We argue that this presents a distinct challenge to theoretical models, which is not easy to resolve.

We conjecture that a solution of this issue may be found by recognizing that a positive structural innovation to the interest rate has two effects. First, by raising the real interest rate, it causes, *ceteris paribus*, optimal inventories holdings to be reduced, which turns inventory investment negative. Second, by engineering a decline in output growth, it causes an unanticipated fall in sales, which translates into unplanned inventories accumulation. Ultimately, whether an interest rate shock

will cause an increase or a decrease in inventories will crucially hinge upon which of these two effects dominates. If the latter effect turns out to be stronger, an interest rate shock will be associated with an increase in the real interest rate, and with an accumulation of inventories. Crucially, different sample and historical periods may have been associated with the dominance of either the first or the second effect, and this might explain why different periods have been associated with correlations with a different sign.

Our second contribution is to the broader empirical inventory literature, in that we introduce Bayesian time-varying VAR techniques, and the use of long time series. To the very best of our knowledge, this is the first study of inventory relationships which allows for time-variation in both the coefficients and the covariance matrix of innovations, which is based on structural VAR methods, and which utilizes long time series, in particular including the interwar period. Our paper is thus part of a recent literature which, following Cogley and Sargent (2002)'s seminal contribution, estimates Bayesian VARs with time-varying parameters.²

It has been noted before that the behavior of inventories is different across sample periods. Perhaps most notably, changes in inventory management are often associated with changes in the behavior of other aggregate time series, as, for instance, in the Great Moderation.³ Much of the discussion thus focuses on selected episodes, but does not take into account whether the identified stylized facts are stable over other periods. One approach to the issue of sample selection is to identify time periods of interest, such as the shift from the Great Inflation to the Great Moderation. This approach rests on the ability of the researcher to identify such episodes in the data. This may be plausible in the case of the Great Moderation, but is less convincing otherwise. An alternative is to instead compute statistics of interest using a rolling window approach. In this paper, we take the latter to its logical conclusion and estimate a Bayesian VAR with time-variation in both the lag-coefficients and the covariance matrix of reduced-form innovations, for sales, inventories, inflation and nominal interest rates.

The remainder of the paper is organized as follows. The next section describes the Bayesian methodology we use to estimate the time-varying parameters VARs with stochastic volatility. We detail our identification strategy in the case of the structural VAR, and we discuss the methodology to compute the VAR's impact matrix. Section 3 discusses estimation results and evidence from the reduced-form specification, while Section 4 examines the structural evidence. We particularly focus on the effects of the identified monetary policy shocks and use the structural evidence to understand time variation in the reduced-form correlations. Section 5 concludes.

²See Cogley and Sargent (2005), Primiceri (2005), and Benati and Goodhart (2011).

³See Kahn, McConnell, and Perez-Quiros (2002).

2 Methodology

2.1 A Bayesian Time-Varying Parameter VAR with Stochastic Volatility

We specify the following time-varying parameter VAR(p) model:

$$Y_t = B_{0,t} + B_{1,t}Y_{t-1} + \dots + B_{p,t}Y_{t-p} + \epsilon_t \equiv X_t'\theta_t + \epsilon_t. \quad (1)$$

The notation is standard. The vector $Y_t \equiv [\Delta i_t, \Delta s_t, \pi_t, r_t]'$ collects the data series of interest. Δs_t is the log-difference of real sales, while π_t is inflation, computed as the log-difference of the relevant price index. We use the GNP deflator for the interwar period, and the GDP deflator for the post-WWII years. r_t is the short-term rate, specifically, the three-month commercial paper rate and the three-month Treasury bill rate, respectively, for the two sample periods.⁴ Finally, Δi_t is the change in real inventories normalized by potential output⁵. For the interwar period we use Balke and Gordon's (1986) estimate of potential GNP, whereas for the post-WWII years we use the *Congressional Budget Office's* estimate of potential GDP. For a complete description of the data and of their sources, see Appendix A.

The overall sample periods are 1919Q1-1941Q4 for the interwar period and 1949Q1-2011Q1 for the post-WWII era. For both sample periods we use the first 10 years of data to compute the Bayesian priors. Consequently, the effective sample periods are 1929Q1-1941Q4 and 1959Q1-2011Q1, respectively. As it is customary in the literature on Bayesian time-varying parameter VARs⁶, we set the lag order to $p = 2$. We collect the VAR's time-varying coefficients—that is, the elements of the matrices $B_{0,t}$, $B_{1,t}$, ..., $B_{p,t}$ —in the vector θ_t , and we postulate that they evolve according to:

$$p(\theta_t \mid \theta_{t-1}, Q) = I(\theta_t) f(\theta_t \mid \theta_{t-1}, Q), \quad (2)$$

with $I(\theta_t)$ being an indicator function that rejects unstable draws, thus enforcing a stationarity constraint on the VAR. $f(\theta_t \mid \theta_{t-1}, Q)$ is given by:

$$\theta_t = \theta_{t-1} + \eta_t, \quad (3)$$

with $\eta_t \sim N(0, Q)$.

The VAR's reduced-form innovations in (1) are assumed to be zero-mean, normally distributed. We factor the time-varying covariance matrix Ω_t as:

$$\text{Var}(\epsilon_t) \equiv \Omega_t = A_t^{-1} H_t (A_t^{-1})', \quad (4)$$

⁴We do not annualize the short rate in order to make its scale comparable to that of inflation. Specifically, if R_t is the relevant short-term rate, reported as, say, 10%, r_t is computed as $r_t = (1 + R_t/100)^{1/4} - 1$.

⁵We discuss in detail the motivation behind this scaling assumption in the data Appendix A.

⁶See e.g. Cogley and Sargent (2002), Cogley and Sargent (2005), Primiceri (2005), Benati (2008), and Benati and Goodhart (2011).

where the matrices H_t and A_t are defined as:

$$H_t \equiv \begin{bmatrix} h_{1,t} & 0 & 0 & 0 \\ 0 & h_{2,t} & 0 & 0 \\ 0 & 0 & h_{3,t} & 0 \\ 0 & 0 & 0 & h_{4,t} \end{bmatrix} \quad A_t \equiv \begin{bmatrix} 1 & 0 & 0 & 0 \\ \alpha_{21,t} & 1 & 0 & 0 \\ \alpha_{31,t} & \alpha_{32,t} & 1 & 0 \\ \alpha_{41,t} & \alpha_{42,t} & \alpha_{43,t} & 1 \end{bmatrix}. \quad (5)$$

The $h_{i,t}$'s are assumed to evolve as geometric random walks:

$$\ln h_{i,t} = \ln h_{i,t-1} + \nu_{i,t}. \quad (6)$$

For future reference, we define $h_t \equiv [h_{1,t}, h_{2,t}, h_{3,t}, h_{4,t}]'$. As in Primiceri (2005), we postulate that the non-zero and non-unity elements of the matrix A_t , which we collect in the vector $\alpha_t \equiv [\alpha_{21,t}, \alpha_{31,t}, \dots, \alpha_{43,t}]'$, evolve as driftless random walks:

$$\alpha_t = \alpha_{t-1} + \tau_t. \quad (7)$$

Finally, we assume the vector $[u_t', \eta_t', \tau_t', \nu_t']'$ to be distributed as:

$$\begin{bmatrix} u_t \\ \eta_t \\ \tau_t \\ \nu_t \end{bmatrix} \sim N(0, V), \text{ with } V = \begin{bmatrix} I_4 & 0 & 0 & 0 \\ 0 & Q & 0 & 0 \\ 0 & 0 & S & 0 \\ 0 & 0 & 0 & Z \end{bmatrix} \text{ and } Z = \begin{bmatrix} \sigma_1^2 & 0 & 0 & 0 \\ 0 & \sigma_2^2 & 0 & 0 \\ 0 & 0 & \sigma_3^2 & 0 \\ 0 & 0 & 0 & \sigma_4^2 \end{bmatrix}, \quad (8)$$

where u_t is such that $\epsilon_t \equiv A_t^{-1} H_t^{\frac{1}{2}} u_t$.

We impose a block-diagonal structure on V for parsimony, since the model is already quite heavily parameterized. Moreover, allowing for a completely generic correlation structure among different sources of uncertainty would preclude any structural interpretation of the innovations, as discussed in Primiceri (2005). Finally, following Primiceri (2005) we adopt the additional simplifying assumption of a block-diagonal structure for S :

$$S \equiv \text{Var}(\tau_t) = \text{Var}(\tau_t) = \begin{bmatrix} S_1 & 0_{1 \times 2} & 0_{1 \times 3} \\ 0_{2 \times 1} & S_2 & 0_{2 \times 3} \\ 0_{3 \times 1} & 0_{3 \times 2} & S_3 \end{bmatrix}, \quad (9)$$

with $S_1 \equiv \text{Var}(\tau_{21,t})$, $S_2 \equiv \text{Var}([\tau_{31,t}, \tau_{32,t}]')$, and $S_3 \equiv \text{Var}([\tau_{41,t}, \tau_{42,t}, \tau_{43,t}]')$. This implies that the non-zero and non-unity elements of A_t which belong to different rows evolve independently. As discussed in Primiceri (2005, Appendix A.2), this assumption drastically simplifies inference, since it allows us to perform Gibbs sampling on the non-zero and non-unity elements of A_t equation by equation.

2.2 Estimation and Identification

We estimate (1)-(9) via standard Bayesian methods. Appendix B discusses our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the

posterior distribution of the hyperparameters and the states conditional on the data. We identify four structural shocks by combining a single long-run restriction, in order to identify a permanent output shock as in Blanchard and Quah (1989), with three sign restrictions to identify demand- and supply-side transitory shocks.

As a preliminary assessment of the plausibility of our long-run assumption, we perform standard unit root tests on the logarithms of real sales and real output.⁷ We find that we cannot reject the null of a unit root for either variable at conventional significance levels. We also assess whether the two series are cointegrated. We can reject the null of a unit root in the difference of these variables. The rejection is very strong for the post-WWII period, whereas it is weaker for the interwar period. We therefore proceed under the assumption that, for either period, real sales and real output share a common stochastic trend. We thus identify permanent output shocks based on the restriction that they are the only shocks exerting a permanent impact on log sales.

The remaining three shocks are identified based on a standard set of sign restrictions reported in Table 1. We identify the interest rate shock based on the restriction that it has a non-negative impact on the interest rate, and a non-positive impact on both inflation and real sales growth. A demand shock is postulated to have a non-negative impact on either the interest rate, inflation, or real sales growth. Finally, a transitory supply shock is differentiated from the other two shocks because it is assumed to induce negative co-movement between inflation and real sales, whereas its impact on the interest rate is left unconstrained. These restrictions are the same used by Benati (2008) and Benati and Goodhart (2011), the only difference being that both papers featured output growth instead of real sales growth. These identifying assumptions are compatible with a wide range of macroeconomic models.

An important aspect of our identification strategy is that, although we constrain the sign of the response of sales growth to the three transitory shocks, we instead leave the response of the change in inventories unconstrained. As for the permanent output shock, the responses of all variables to it are left unconstrained by assumption. This implies that the structural correlations between inventories and sales, and between inventories and the real interest rate—that is, the correlations *conditional* on either of the four structural shocks—are left entirely unconstrained. These correlations will be one of the key objects of our investigation.

2.3 Computing the Structural Impact Matrix

For each quarter, and for each draw from the ergodic distribution, we compute the time-varying structural impact matrix, $A_{0,t}$, by combining the methodology proposed by Rubio-Ramirez, Waggoner, and Zha (2005) for imposing sign restrictions, and the procedure proposed by Gali and Gambetti (2009) for imposing long-run restrictions within a time-varying parameters VAR context. Specifically, let $\Omega_t = P_t D_t P_t'$ be the

⁷All of these results are available from the authors upon request.

eigenvalue-eigenvector decomposition of the VAR’s time-varying covariance matrix Ω_t , and let $\tilde{A}_{0,t} \equiv P_t D_t^{\frac{1}{2}}$. We draw an $N \times N$ matrix, K , from a standard-normal distribution and compute the QR decomposition of K , that is, we compute matrices Q and R such that $K = Q \cdot R$. The “intermediate estimate” of the time-varying structural impact matrix can then be computed as $\bar{A}_{0,t} = \tilde{A}_{0,t} \cdot Q'$.

Following Gali and Gambetti (2009, Section II), we then compute a local approximation to the matrix of the cumulative impulse response functions (IRFs) to the VAR’s structural shocks as:⁸

$$\bar{C}_{t,\infty} = \underbrace{[I_N - B_{1,t} - \dots - B_{p,t}]^{-1}}_{C_{0,t}} \bar{A}_{0,t}, \quad (10)$$

where I_N is the $N \times N$ identity matrix. We then rotate the matrix of the cumulative IRFs via an appropriate Householder matrix⁹ H in order to introduce zeros in all of the second row of $\bar{C}_{t,\infty}$, that is, in the row corresponding to log real sales. The second row of the resulting local approximation to the matrix of the cumulative IRFs,

$$C_{t,\infty} = \bar{C}_{t,\infty} H = C_{0,t} \bar{A}_{0,t} H = C_{0,t} A_{0,t}, \quad (11)$$

is given by $C_{t,\infty}^2 = [x \ 0'_{(N-1) \times 1}]$, with $0_{(N-1) \times 1}$ being a vector of $(N - 1)$ zeros, and x being a non-zero entry.

This implies that the first shock is the only one exerting a long-run impact on the level of log sales, and therefore on the level of log output. If the resulting structural impact matrix $A_{0,t} = \bar{A}_{0,t} H$ satisfies the sign restrictions we store it. Otherwise we discard it, and we repeat the procedure until we obtain an impact matrix which satisfies both the sign restrictions and the long-run restriction at the same time.

3 Reduced-Form Correlations

In discussing our estimation results we proceed in several steps. We first present estimates for the time-varying reduced-form correlations of the forecast errors. This should give us an idea of how the variables have co-moved over the sample period, and in particular of how such co-movements may have changed over time. This establishes our set of stylized facts since the late 1920s. In the subsequent Section we discuss the same type of evidence, but this time conditional on the identified structural shocks. Finally, we use the structural estimation evidence in order to shed light on the sources of time-variation in the reduced-form correlations. We find it convenient throughout to discuss the two sample periods separately.

⁸The only difference with respect to Gali and Gambetti (2009) is that they compute the local approximation to the matrix of the cumulative IRFs based on the companion form of the VAR, whereas we compute it directly based on the VAR itself.

⁹We compute the Householder matrix *via* Algorithm 5.5.1 of Golub and VanLoan (1996).

3.1 The Inter-War Period

Conceptually in line with Cogley, Primiceri, and Sargent (2010), Figures 1 and 2 show, for the interwar period, evidence of changes in the correlation between the reduced-form forecast errors for the change in real inventories and either the ex post real interest rate, or real sales growth, at various horizons. Specifically, the top rows in the two figures show the median and the one- and two-standard deviation percentiles of the posterior distribution of the correlations between the forecast errors for the relevant series generated by the time-varying VAR at each point in time; the middle rows show the fractions of the draws from the posterior distribution for which the correlation is positive; and the bottom rows show the business-cycle components of the series shown in the middle rows, together with the business-cycle component of the logarithm of real GNP.¹⁰ For each quarter, the correlations between the forecast errors for the relevant series at the various horizons generated by the time-varying VAR were computed based on the estimated covariance matrix of the forecast errors, following Cogley, Primiceri, and Sargent (2010).¹¹

Starting from the relationship between the change in real inventories and the ex post real interest rate, the evidence reported in the top row of Figure 1 clearly suggests that the correlation is consistently negative over the entire sample period and for the different forecasting horizons. This impression is confirmed by the results reported in the middle row, although the fraction of draws for which the correlation is positive exhibits significant fluctuations over the sample period, both on impact and at all horizons. In particular, the fraction of draws reaches a minimum in April 1933, when the U.S. went off the Gold Standard, and it achieves local maxima in 1932 and 1936.

To gain further insight into the driving forces behind these patterns, we decompose the fraction of draws into cyclical and trend components. When we look at the business cycle component, that is, at the component associated with frequencies between 6 quarters and 8 years, we see that, at all horizons, the business-cycle component of the fraction of draws for which the correlation is positive exhibits a striking positive correlation with the business-cycle component of the logarithm of real GNP. On the other hand, casual inspection does not suggest any systematic lead or lag patterns between the series. Although the correlation is systematically negative during the en-

¹⁰Business-cycle components have been extracted via the band-pass filter proposed by Christiano and Fitzgerald (2003). Following established conventions in business-cycle analysis—see e.g. Baxter and King (1999) and Stock and Watson (1999)—business-cycle frequencies have been defined as those pertaining to fluctuations with frequencies of oscillation between 6 quarters and 8 years.

¹¹See expressions (9) to (11) of Cogley, Primiceri, and Sargent (2010). An important point to stress here is that, given the two-sided nature of the estimates of the time-varying VAR produced by the Gibbs sampler, the VAR-generated forecast errors we are working with (and therefore their estimated covariance matrices) should only be regarded as approximations to the authentic objects we would obtain if we had estimated the VAR based on recursive samples. Such an approximation is routinely used in the literature—beyond Cogley, Primiceri, and Sargent (2010), see e.g. Benati and Surico (2008)—because of the significant computational burden associated with recursive estimation of time-varying parameters VARs.

tire period, it also tends, equally systematically, to be less negative during economic upswings, and even more negative during downswings. This suggests a fundamental asymmetry in the way economic upswings and downturns evolve, and thus a challenge for theoretical business cycle models. Furthermore, as we now discuss with respect to the relationship between the change in real inventories and real sales growth, this systematic pattern of variation at the business-cycle frequencies does not appear to be an isolated result.

The evidence reported in the top row of Figure 2 on the relationship between the change in real inventories and real sales growth suggests that the correlation is mostly positive over the sample period, and for the different forecasting horizons. The exceptions are the very early quarters of the sample and the very last quarter. We note, however, that in most cases a zero correlation is contained within the one-standard deviation bands. In the middle panel, the fraction of draws shows significant variation over the sample period, although, consistent with the evidence we just discussed, most of the draws fall on the positive side. When we decompose the fraction of draws into a trend and a business cycle component, we see that, at the business-cycle frequencies, the fraction of draws for which the correlation is positive exhibits a remarkably strong negative correlation with the logarithm of real GNP.

We can summarize our findings as follows. In the inter-war period, the real rate and inventories strongly comove negatively, as theory would suggest. Increases in real rates raise the cost of maintaining inventory holdings. What underlies this reduced-form correlation is a stable low-frequency relationship, whereas time-variation is generated by strongly pro-cyclical business cycle components. In economic expansions, the business cycle component reduces the effect of a negative long-run relationship. This finding is related to Wen's (2005) argument that the inventory facts are conditional on the sampling frequency. The correlation between sales and inventory growth rates is positive, although not as strongly as for the other variables. This pattern is largely driven by the trend component. The business cycle component is strongly countercyclical, to the effect that trend and business-cycle movements can both reinforce and dampen each other in different periods. This points towards a theory that offers predictions for both short-run and medium-run business cycle movements.¹² Finally, the negative correlation between the business cycle component of GDP and the sales-inventories correlation suggests a more dominant role for demand factors.¹³ We will delve deeper into this based on our structural identification scheme.

3.2 The Post-WWII Period

We now turn to the post-WWII period, which has been the main focus of previous work in the literature. We report the same kind of evidence as for the first sample period in Figures 3 and 4. We note that the correlation between the reduced-form

¹²See Comin and Gertler (2006).

¹³See e.g. West (1995).

forecast errors for the real ex post interest rate and the change in real inventories at various horizons is almost uniformly positive over the entire sample period and for the various horizons. The relationship is stronger at longer horizons, and much less so for the contemporaneous correlation (that is, on impact). The correlation is weakest (and negative) in the very first quarters of the sample, for which a minority of the draws from the posterior distribution is positive. Similarly, the correlation on impact is mildly negative during the middle of the 1980s, when a slight majority of the draws is associated with a negative correlation.

Our finding of a positive correlation between interest rates and inventories is consistent with the previous empirical literature on this topic—see Ramey and West (1999)—and constitutes one of the puzzles in the inventory literature.¹⁴ What the overall positive correlation over this time period masks, however, is the substantial time-variation in this pattern, which has not been documented before. These movements are especially apparent on impact, when the correlation turns from negative to strongly positive during the period up to the Volcker disinflation. It then changes sign during the first half of the 1980s, after which it exhibits a strong hump-shaped pattern, with a peak of one around the turn of the century. The contemporaneous correlation shows a dramatic decrease during subsequent years, with the fraction of draws for which the correlation is positive falling below 50 percent in the last few quarters of the sample. Evidence for other horizons is weaker, but it replicates the very broad features of the evidence on impact.

As in Figures 1 and 2, in the bottom panel of Figure 3 we plot the business-cycle component of the fraction of draws for which the correlation has been positive, together with the business-cycle component of log real GDP. In sharp contrast with the interwar period, at the business-cycle frequencies the fractions of draws for which the correlation has been positive does not exhibit any clear-cut, stable relationship with the logarithm of real GDP, comoving with it sometimes positively, and sometimes negatively, but without any systematic discernible pattern. This again presents a challenge for theoretical models, in terms of explaining both a positive long-run relationship and time variation in the short-run together with an apparent disconnect with movements in GDP.

We now turn to the relationship between inventories and sales growth (see Figure 4). We find that, first, the correlation between the two series' forecast errors is predominantly negative on impact, but it becomes more and more positive at longer horizons. This is conceptually in line with the evidence reported in Wen (2005). Second, there is a significant extent of time-variation in the correlations between the forecast errors at the various horizons. This is especially apparent from the middle row of Figure 4. Specifically, both on impact and at all subsequent horizons, the fraction of draws for which the correlation is positive has exhibits a broadly V-shaped pattern over the sample period, although with significant short-run fluctuations around such a secular movement.

¹⁴See e.g. Maccini, Moore, and Shaller (2010).

Some of these fluctuations coincide with well-known events in U.S. post-WWII monetary history. The fraction of draws for which the correlation is positive exhibits significant decreases during the collapse of Bretton Woods and the Volcker disinflation of the early 1980s, which is especially apparent for horizons up to one year ahead. This provides evidence that a non-negligible portion of the time-variation in the correlation between the forecast errors documented in Figure 4 originates from some of the same fundamental macroeconomic forces which have shaped U.S. post-WWII macroeconomic dynamics, specifically the evolution of the U.S. monetary regime. Contrary to the previous results, the fraction of draws on impact shows a much different pattern than those at longer horizons, which suggests underlying structural differences.

Finally, as shown in the bottom row, the frequency-domain-based decomposition of the evolution of the fraction of draws points towards a positive co-variation with the logarithm of real GDP at business-cycle frequencies since the early 1980s, whereas evidence for the former period is not clear-cut. At best, there is some mild negative co-variation. This suggests a possible element of continuity between the earliest part of the post-WWII period and the interwar era. As we discussed in the previous subsection, the business-cycle component of the fraction of draws exhibits a remarkably strong negative correlation with the cyclical component of log real GDP. This indicates that the most recent period represents a discontinuity with the pattern that prevailed during previous decades.

In summary, we find that the real interest rate-inventory growth correlation is positive for the post-WWII period - which is in line with the previous literature -, but markedly different from the inter-war period. The sales-inventories correlation is mildly positive, but there are substantial differences over the various time horizons, and substantial time-variation in the business cycle component. In particular, the impact period exhibits different patterns along these dimensions than the other time horizons. We also note that the Volcker disinflation marks a turning point for many of these statistics. We return to this issue, and how it relates to the causes of the Great Moderation below.

4 Structural Evidence

The correlation patterns in the data that we discussed in the previous section represent a set of stylized facts, some of which are well known in the literature. However, our approach of using a time-varying parameters model reveals a substantial extent of variation in these facts since the end of the 1920s. We now impose our structural identification scheme on the reduced-form VAR, and we interpret the reduced-form findings in light of our structural identification. To recap, we identify a permanent output shock based on a long-run restriction, and three temporary shocks - an interest-rate shock, a demand shock, and a supply shock - based on sign restrictions. We first discuss the dynamic behavior of sales, inventories and the real interest rate

by analyzing the IRFs to these shocks. We then compute structural forecast error correlation and use these to dig deeper into the previously discussed stylized facts.

4.1 Impulse Response Functions

4.1.1 The Inter-War Period

Figures 5 to 8 show the normalized IRFs of the change in inventories, real sales growth, and the ex post real interest rate to each of the four identified structural shocks over the inter-war period for the first and the last quarter of the sub-sample (1929Q4 and 1941Q4, respectively), and for 1933Q3, in the middle of the Great Depression.¹⁵ A permanent output shock leads to a temporary fall in the ex post real rate, and a temporary increase in sales growth. The former responses are reasonably tightly estimated, while the error bands of the sales growth responses include zero both at the beginning and at the end of the sample. However, the median estimates show an increase on impact, while there is some variation in inventory growth behavior in subsequent periods.

In response to a positive interest rate shock, the ex post real rate rises (recall that its median impact response has been normalized to unity), but swiftly returns to zero within a year. The estimated median response of sales growth is negative (which we have imposed in identification), while that of inventories is positive, although the extent of uncertainty with respect to the latter is large enough that it is not possible to make statements with any confidence. While the response of sales is consistent with a priori theorizing - a monetary contraction reduces sales -, the response of inventories is less obvious. One possibility is that the economic contraction leads to a build-up in unsold final goods and thus inventory accumulation. We note, however, that the reduced-form correlation between inventories and the ex post real rate is negative during the inter-war period (see Figure 1). This suggests that another underlying driving process explains the co-variation of these two variables.

A positive demand shock (see Figure 7) causes a temporary increase in the change in inventories, and a temporary decrease with a hump-shaped adjustment pattern in the ex post real rate. Similarly, a transitory supply shock leads to temporary increases in both sales growth and the ex post real rate, whereas for inventories uncertainty is once again so large that we cannot rule out a zero impact at any horizon. Median estimates, however, show a temporary decrease in the change in inventories.

¹⁵The IRFs to a permanent output shock are normalized so that the median of the distribution of the cumulative IRFs of sales growth is equal to one (which implies that the median long-run impact of the shock on the logarithm of real GNP is, likewise, equal to one). The IRFs to an interest rate shock are normalized so that the median impact at zero on the ex post real rate is equal to one. Finally, the IRFs to the transitory supply and demand shocks are normalized so that the median impact at zero on sales growth at zero is equal to one. We compute the IRFs using the Monte Carlo integration procedure described in Appendix C, which allows us to tackle the uncertainty originating from future time-variation in the VAR's structure.

We thus find that an unconditional negative comovement between the interest rate and inventories can be traced back to either temporary demand or supply shocks. These shocks are separately identifiable through their impact on sales relative to inventories. We found that the reduced-form correlation between these two series in the inter-war period is largely positive, which would point to a more dominant role for demand shocks (see Figure 11). Finally, none of the IRFs for either of the three series exhibits any discernible time-variation over the sample period.

4.1.2 The Post-WWII Period

We report the normalized IRFs of the change in inventories, real sales growth, and the ex post real interest rate to each of the four identified structural shocks in Figures 9-12 for selected quarters. On impact, a permanent output shock raises both sales growth and the ex post real rate, which then reverts back to zero. Whereas the responses of sales growth are quite tightly estimated and do not exhibit a significant extent of variation across time periods, this is often not the case for the ex post real rate. In particular, as the bottom row of Figure 9 shows, around the time of the Great Inflation the extent of uncertainty is significantly larger than either before or after. The response of the change in inventories, on the other hand, exhibits some time-variation. It is consistently negative on impact over the entire post-WWII period, albeit with varying degrees of uncertainty. In the latter half of the sample period, a pronounced hump-shaped response emerges, whereas in the earlier part the adjustment is more monotonic. What is striking is the difference with respect to the inter-war period, when the response on impact is positive. Similarly, the behavior of the ex post real rate is different between the two sub-samples, whereas the response of sales does not differ much.

Turning to the interest rate shock in Figure 10, none of the variables' impulse responses exhibit significant time-variation. The ex post real rate rises on impact, and then reverts monotonically to zero. The speed of mean-reversion appears to be smaller in the first decades of the sample, but evidence on this is weak, given the large extent of uncertainty. Mean-reversion is significantly faster for sales growth, which jumps down on impact, then oscillates around zero for a few quarters. Finally, the response for inventories is positive on impact, which then reverts to zero monotonically. Overall, the post-war impulse responses exhibit the same pattern as those for the pre-war period, thereby reflecting the positive reduced-form correlation between the two series.

A positive demand shock (see Figure 11) is estimated to lead to a small and statistically insignificant decrease in the ex post real rate, the magnitude of which remains essentially unchanged over the sample period. The evolution of the posterior distributions of the impulse responses suggests that the dynamics of the real rate are not significantly different from zero in the latter part of the sample, and comparatively much stronger in the earlier part. As for sales growth, responses on impact tend to be

positive, but insignificant, with a more drawn-out behavior at the end of the sample. The responses of inventories show a sizeable extent of time-variation, with the impact being negative in the early part of the sample, and then turning positive around the time of the Great Inflation. However, the extent of uncertainty is substantial. Once again, the corresponding patterns for the inter-war period are broadly similar, although the behavior of inventories and sales resemble more closely the latter part of the post-war sample.

Finally, the responses to the transitory supply shock reported in Figure 12 point towards no time-variation for any of the three variables. Sales rise, inventories fall, and the real rate increases significantly and persistently. Sales revert back to zero after one quarter, while the response of inventories is more drawn-out, with a pronounced hump-shaped pattern in the latter part of the sample. This mirrors the behavior of inventories in response to a permanent supply shock. Comparison to the pre-war responses do not reveal substantial differences.

The main finding in this section is the comovement pattern between inventories and the real interest rate. Their responses to an identified interest rate shock suggest that the puzzle of a positive correlation between the two series over the post-WWII period does not uniquely pertain to their reduced-form relationship. Arguably, this is perhaps more surprising than the positive reduced-form correlation. In principle, a positive reduced-form correlation is compatible with a negative correlation conditional on interest rate shocks, which is, in fact, what we would expect to find based on theoretical considerations. A positive reduced-form correlation can result from the fact that the negative correlation generated by interest rate shocks is dominated by the positive correlation induced by at least one of the other shocks and that the fractions of variance of inventories and the real interest rate explained by the shocks generating positive conditional correlations are sufficiently large. Our results suggest that even interest rate shocks generate a positive co-movement between the two series, which deepens the inventory-interest rate puzzle.

4.2 Structural Correlations

We now turn to an analysis of the structural correlations between the series, that is, to the correlations conditional on the identified structural shocks. The evidence on the evolution of the reduced-form correlations between inventories and either sales growth or the real rate naturally raises two questions. First, which structural shocks generate such correlations in the data? Second, what are the causes of time variation in the correlations? More specifically, is the time-variation due to changes in the relative importance of the different shocks hitting the economy, or due to changes in the way their impact is propagated through the system? In this section we address the first question, whereas we answer the latter in the next .

4.2.1 The Inter-War Period

Figures 13 and 14 show the structural correlations between inventories and either the ex post real interest rate or sales growth, for both the long run and one quarter ahead. We also report the fraction of draws for which the structural correlations are positive. We only consider structural correlations at horizons greater than zero because, on impact, the correlation between two series conditional on a single shock is by construction always equal to one. At horizons greater than zero, on the other hand, the dynamics of the system do matter, so that even correlations conditional on a single shock are no longer necessarily equal to one.

As documented above, the reduced-form correlation between the change in inventories and the ex post real rate is systematically negative during the entire inter-war period. Figure 13 shows that the negative correlation is due to transitory demand and supply shocks, which induce a uniformly and substantially negative correlation over the entire period. This is re-enforced by the permanent output shock, for which the conditional correlation turns substantially negative after the early 1930s, and especially at longer horizons. In line with our discussion of the IRFs, the correlation conditional on interest rate shocks is systematically positive over the entire sample period. This highlights once again the puzzling nature of the relationship between interest rates and business inventories. It is also apparent from the figures that the years leading up to the abandonment of the Gold Standard exhibit a significantly different relationship between real rates and inventories under permanent supply shocks, whereas the effect on the other identified shocks is much more muted.

Figure 14 shows a similar set of graphs for the correlation between inventories and sales. The mostly positive correlation of the inter-war era is largely due to the permanent output shock, especially after April 1933. As for the contribution of the demand shock, it is close to borderline. In particular, it is slightly positive at very short horizons, but essentially neutral at longer horizons. Finally, the interest rate shock and the transitory supply shock systematically induce a negative correlation, which is not strong enough to overturn the contributions of the other two shocks.

4.2.2 The Post-WWII Period

We now look at the structural correlations for the variables of interest in post-war data, which are reported in Figures 15 and 16. Although the reduced-form correlation between the real rate and inventories is almost uniformly positive during this period, Figure 15 depicts a complex picture for its structural sources. A positive correlation is induced over the entire sample period, and over all horizons, by the interest rate shock, which particularly at shorter horizons generates a conditional correlation of almost one. The permanent output shock plays a supporting role at short horizons, but less so at longer horizons. The correlation it induces also exhibits a significant extent of time-variation. The demand shock contributes to an overall positive correlation only during the period of the Great Moderation and, to a lesser extent, before the collapse

of Bretton Woods. During the Great Inflation the contribution of this shock to the overall correlation is in fact negative. This is consistent with the idea, espoused in Maccini, Moore, and Schaller (2004) that the monetary regime may play a crucial role in determining the comovement between inventories and real rates.

We previously found that the reduced-form correlation between inventories and sales is strongly positive at most horizons, with the exception of the very short ones. In Figure 16, we note that both the interest rate and the transitory supply shock generate a uniformly negative correlation at all horizons over the sample period. The demand shock contributes positively to the correlation before the beginning of the Volcker stabilization, and especially during the Great Inflation years. At the same time, the evolution of the reduced-form correlation closely mirrors the evolution of the correlation conditional on the permanent output shock. It is very strongly positive at the beginning of the sample, then it oscillates around zero around the time of the Great Inflation, before turning positive during the most recent part of the sample. In the long run, this large degree of time variation in the structural correlation disappears in favor of a robustly positive correlation.

A comparison between these patterns across our two main sample periods shows that the role of the identified interest rate shock in generating a negative correlation between sales and inventories and a positive correlation between real rates and inventories is the same. Similarly, transitory supply shocks imply the same reduced-form correlation in the pre-war and post-war samples. The differences in the reduced-form behavior across periods are thus largely explained by the permanent supply shocks and, to a lesser extent, by the demand shocks. The permanent shock induces a positive correlation between sales and inventories in both sample periods, but switches sign in the conditional correlation between interest rate and inventories between the two periods, as do demand shocks. Moreover, we find that during times of extreme economic turbulence, such as the abandonment of both the Gold Standard and the Bretton Woods system of fixed exchange rates, the structural correlations shift, in the extreme case changing their signs. The most prominent example is the one conditional on the permanent productivity shock. Inventory models that attempt to capture this feature of the data might therefore have to incorporate regime shifts.

4.3 Understanding Time-Variation in the Correlations

Our final exercise attempts to delve deeper into the underlying causes of the changes in the reduced-form and structural correlations over the two sample periods. Specifically, we attempt to disentangle the contribution of the time-varying VAR-coefficients and the time-varying components of the variance-covariance matrix. These results are reported in Figure 17. For both sample periods we show the medians of the posterior distributions of the VAR's time-varying coefficients (the $B_{j,t}$'s), of the logarithms of the diagonal elements of the matrix H_t (the $h_{i,t}$'s), and of the off-diagonal elements of the matrix A_t (the $\alpha_{hk,t}$'s).

As for the inter-war period, the first column points towards essentially no time-variation in the VAR coefficients. Moreover, there is no time-variation for three volatilities out of four. The main source of variation pertains instead to the off-diagonal elements of the covariance matrix. This implies that the VAR coefficients cannot have played a significant role in fostering changes in the two correlations. On the other hand, it is not possible to further disentangle the separate roles played by changes in the structural shocks' volatilities, and changes in the way the shocks have impacted the economy.¹⁶ What we can conclude is that the VAR coefficients played essentially no role in causing the variation of the reduced-form correlations. Turning to the post-WWII era, the bottom row of Figure 17 shows time-variation in all of the VAR elements, although the visual evidence clearly suggests that the extent of variation for the VAR's coefficients is comparatively smaller than for elements of the covariance matrix.

We also conducted analogous experiments for the structural correlations¹⁷. In the pre-war period, the impact of changes in the VAR's coefficients is negligible. In the post-WWII period, on the other hand, for the correlation between inventories and the real rate the impact is virtually zero conditional on the interest rate shock. It is minor, and only at longer horizons, for the one conditional on the transitory supply shock, and significant at longer horizons for the demand non-interest rate shock. On the other hand, it is quite substantial at all horizons for the permanent output shock. For the correlation between inventories and sales, the impact is negligible for all shocks, and at all horizons.

To summarize, evidence for the pre-war period points towards no role of changes in the VAR coefficients in driving changes in either of the two correlations. Instead, these originate from changes in both the volatilities of structural shocks, and especially the way in which such shocks have impacted the economy. For the post-war period, changes in both the shocks' volatilities and the way in which they impact the variables play a dominant role, but it is possible to detect some impact of changes in the VAR's coefficients for the correlation between inventories and the real rate at longer horizons.

¹⁶Note that from equation (4), and from the definition of the structural impact matrix $A_{0,t}$, we have that under the assumption of *unit-variance* structural shocks $\Omega_t = A_t^{-1}H_t(A_t^{-1})' = A_{0,t}A_{0,t}'$. However, this is equivalent to

$$\Omega_t = A_t^{-1}H_t(A_t^{-1})' = A_{0,t}A_{0,t}' = \hat{A}_{0,t}V_t'\hat{A}_{0,t}' = \hat{A}_{0,t} \begin{bmatrix} v_{1,t} & 0 & 0 & 0 \\ 0 & v_{2,t} & 0 & 0 \\ 0 & 0 & v_{3,t} & 0 \\ 0 & 0 & 0 & v_{4,t} \end{bmatrix} \hat{A}_{0,t}', \quad (12)$$

where the $v_{j,t}$'s are the non-unitary, and possibly time-varying volatilities of the structural shocks.

$\hat{A}_{0,t}$ is the structural impact matrix associated with these shocks, for *any* values of the $v_{j,t}$'s. Time-variation in the elements of the matrices A_t and H_t during the pre-war era cannot be automatically mapped into time-variation in the $v_{j,t}$'s and the elements of $\hat{A}_{0,t}$, due to an identification problem.

¹⁷These results are available from the authors upon request.

5 Conclusions

We use Bayesian time-varying VAR methods to study the statistical relationship between inventories, sales and the real interest rate. Reduced-form estimates detail a substantial degree of time-variation both between the pre-war and post-war sample periods, but also within them. Our main finding is that the positive correlation between inventories and interest rates is robust over the entire sample. While this stylized fact is well-known for the post-WWII period, we further establish that the reduced-form correlation between these variables is negative in the inter-war period. However, this pattern is driven demand and supply shocks. Once we condition this correlation on the identified interest-rate shock, the positive correlation emerges. This highlights that one of the salient inventory puzzles is more pertinent than the previous literature let on. Moreover, we also demonstrate the need to establish conditional stylized facts since the changing patterns of demand and supply shocks render the notion of fixed relationships meaningless.

As a conjecture for a theoretical mechanism at play we suggest the following. An unanticipated structural innovation to the interest rate has two effects. First, it increases both nominal and real interest rates. The latter effect causes, *ceteris paribus*, a decrease in firms' desired level of inventories, and therefore a negative change in inventories. Second, the interest rate shock will cause an unexpected fall in sales, and therefore an undesired accumulation of inventories. The notion that an interest rate shock causes a positive or a negative change in inventories crucially hinges upon which of these two effects dominates: if the extent of the undesired accumulation of inventories due to the recessionary effect of the interest rate shock is larger than the extent of their planned decrease due to the higher real interest rate, an interest rate shock will be associated with an increase in inventories, rather than a decrease.

A The Data

Before describing the data in more detail, we first want to address the issues of trending behavior in the raw series. To the very best of our knowledge, no previous study has performed any normalisation of the change in real inventories. Within the present context, however, expressing the change in real inventories in absolute terms rather than in percentage terms, as is the case for sales, would make the interpretation of our results especially difficult. Economic growth causes the variance of the change in real inventories to increase without bounds, thus automatically introducing a systematic element of distortion in any comparison over time and across quarters. Consider for example the response of the economy to an identified interest rate shock. Since the change in sales is expressed in percentage terms, once the IRFs have been appropriately normalised on the interest rate, comparing the response of sales across quarters is appropriate and meaningful. This is however not the case for the change in real inventories: in this case, what the IRFs are telling us is by how much, in real dollars, inventories would have changed in each quarter in response to a normalized interest rate shock. Since the economy in, say, 2009Q1 was significantly larger than it had been in 1959Q1, comparing the response to a normalized interest rate shock in the two quarters of the change in inventories expressed in constant dollars does not provide any meaningful information. This automatically implies that, for our results to be interpretable, the change in real inventories has to be normalized in such a way as to eliminate the impact of economic growth. We have chosen to normalize it by potential, rather than by actual output, in order to avoid distorting the cyclical properties of the resulting series.

A.1 Interwar period

Quarterly seasonally adjusted series for real GNP (the acronym is RGNP72), real potential GNP (TRGNP), the GNP deflator (GNPDEF), nominal GNP (GNP), the change in nominal business inventories (DBUSINOM), and the commercial paper rate (CPRATE) are all from Balke and Gordon (1986). The sample period is 1919Q1-1941Q4. Our sample period is dictated by data availability, as the inventories series from Balke and Gordon (1986) starts in 1919Q1 and is not available for the period 1942Q1-1946Q4. Real sales have been computed as the difference between nominal GNP and the change in nominal inventories, deflated by the GNP deflator, whereas the change in real inventories has been computed as the ratio between the change in nominal inventories and the GNP deflator (the series thus obtained is near-identical to Balke and Gordon's series for the change in real inventories, DBUSI72).

A.2 Post-WWII period

A quarterly seasonally adjusted series for the GDP deflator is from Table 1.1.9 of the *Bureau of Economic Analysis' National Income and Product Accounts* (hence-

forth, *NIPA*). Quarterly seasonally adjusted series for nominal GDP and the change in nominal inventories are from Table 1.1.5. of the *NIPA*. The series for potential GDP (‘GDPPOT: Real Potential Gross Domestic Product, U.S. Congress: Congressional Budget Office, Budget and Economic Outlook, Quarterly, Billions of Chained 2005 Dollars’) is from FRED II. The series for real GDP has been computed as the ratio between nominal GDP and the GDP deflator. The series thus obtained is numerically near-identical to the chain-weighted series for real GDP, GDPC96, which can be found, e.g., at the St. Louis Fed’s internet data portal, FRED II (‘GDPC96: Real Gross Domestic Product, 3 Decimal, U.S. Department of Commerce: Bureau of Economic Analysis, Gross Domestic Product, Seasonally Adjusted Annual Rate, Quarterly, Billions of Chained 2005 Dollars’). This justifies our use of the Congressional Budget Office’s chain-weighted estimate of potential real GDP in order to normalize the change in real inventories. The series for the change in real inventories has been computed as the ratio between the change in nominal inventories and the GDP deflator. Real sales have been computed as the difference between nominal GDP and the change in nominal inventories, deflated by the GDP deflator. A monthly seasonally unadjusted series for the 3-Month Treasury Bill: Secondary Market Rate (acronym is TB3MS) is from the Board of Governors of the Federal Reserve System, and it has been converted to the quarterly frequency by taking averages within the quarter.

B Details of the Markov-Chain Monte Carlo Procedure

We estimate (1)-(8) via Bayesian methods. The next two subsections describe our choices for the priors, and the Markov-Chain Monte Carlo algorithm we use to simulate the posterior distribution of the hyperparameters and the states conditional on the data, while the third section discusses how we check for convergence of the Markov chain to the ergodic distribution.

B.1 Priors

For the sake of simplicity, the prior distributions for the initial values of the states— θ_0 and h_0 —which we postulate all to be normal, are assumed to be independent both from each other, and from the distribution of the hyperparameters. In order to calibrate the prior distributions for θ_0 and h_0 we estimate a time-invariant version of (1) based on the first 10 years of data, and we set

$$\theta_0 \sim N \left[\hat{\theta}_{OLS}, 4 \cdot \hat{V}(\hat{\theta}_{OLS}) \right] \quad (\text{B1})$$

where $\hat{V}(\hat{\theta}_{OLS})$ is the estimated asymptotic variance of $\hat{\theta}_{OLS}$. As for h_0 , we proceed as follows. Let $\hat{\Sigma}_{OLS}$ be the estimated covariance matrix of ϵ_t from the time-invariant

VAR, and let C be its lower-triangular Cholesky factor, i.e., $CC' = \hat{\Sigma}_{OLS}$. We set

$$\ln h_0 \sim N(\ln \mu_0, 10 \times I_N) \quad (\text{B2})$$

where μ_0 is a vector collecting the logarithms of the squared elements on the diagonal of C . As stressed by Cogley and Sargent (2005), “a variance of 10 is huge on a natural-log scale, making this weakly informative” for h_0 .

Turning to the hyperparameters, we postulate independence between the parameters corresponding to the two matrices Q and A , which is an assumption we adopt uniquely for reasons of convenience. Further, we make the following, standard assumptions. The matrix Q is postulated to follow an inverted Wishart distribution,

$$Q \sim IW(\bar{Q}^{-1}, T_0) \quad (\text{B3})$$

with prior degrees of freedom T_0 and scale matrix $T_0\bar{Q}$. In order to minimize the impact of the prior, thus maximizing the influence of sample information, we set T_0 equal to the minimum value allowed, the length of θ_t plus one. As for \bar{Q} , we calibrate it as $\bar{Q} = \gamma \times \hat{\Sigma}_{OLS}$, setting $\gamma = 1.0 \times 10^{-4}$, the same value used in Cogley and Sargent (2005), and a slightly more “conservative” prior (in the sense of allowing for less random-walk drift) than the 3.5×10^{-4} used by Cogley and Sargent (2005). As for α , we postulate it to be normally distributed with a “large” variance,

$$f(\alpha) = N(0, 10000 \cdot I_{N(N-1)/2}). \quad (\text{B4})$$

Finally, as for the variances of the stochastic volatility innovations, we follow Cogley and Sargent (2002, 2005) and we postulate an inverse-Gamma distribution for $\sigma_i^2 \equiv \text{Var}(\nu_{i,t})$:

$$\sigma_i^2 \sim IG\left(\frac{10^{-4}}{2}, \frac{1}{2}\right) \quad (\text{B5})$$

B.2 Simulating the posterior distribution

We simulate the posterior distribution of the hyperparameters and the states conditional on the data via the following MCMC algorithm, as found in Cogley and Sargent (2005). In what follows, x^t denotes the entire history of the vector x up to time t , i.e., $x^t \equiv [x'_1, x'_2, \dots, x'_t]'$ —while T is the sample length.

(a) *Drawing the elements of θ_t* Conditional on Y^T , α , and H^T , the observation equation (1) is linear, with Gaussian innovations and a known covariance matrix. Following Carter and Kohn (2004), the density $p(\theta^T | Y^T, \alpha, H^T)$ can be factored as

$$p(\theta^T | Y^T, \alpha, H^T) = p(\theta_T | Y^T, \alpha, H^T) \prod_{t=1}^{T-1} p(\theta_t | \theta_{t+1}, Y^T, \alpha, H^T). \quad (\text{B6})$$

Conditional on α and H^T , the standard Kalman filter recursions nail down the first element on the right hand side of (A6), $p(\theta_T | Y^T, \alpha, H^T) = N(\theta_T, P_T)$, with P_T being

the precision matrix of θ_T produced by the Kalman filter. The remaining elements in the factorization can then be computed via the backward recursion algorithm found, e.g., in Kim and Nelson (1999), or Cogley and Sargent (2005, appendix B.2.1). Given the conditional normality of θ_t , we have

$$\theta_{t|t+1} = \theta_{t|t} + P_{t|t}P_{t+1|t}^{-1}(\theta_{t+1} - \theta_t). \quad (\text{B7})$$

$$P_{t|t+1} = P_{t|t} - P_{t|t}P_{t+1|t}^{-1}P_{t|t}, \quad (\text{B8})$$

which provides, for each t from $T-1$ to 1, the remaining elements in (1), $p(\theta_t|\theta_{t+1}, Y^T, \alpha, H^T) = N(\theta_{t|t+1}, P_{t|t+1})$. Specifically, the backward recursion starts with a draw from $N(\theta_T, P_T)$, call it $\tilde{\theta}_T$. Conditional on $\tilde{\theta}_T$, (A7)-(A8) give us $\theta_{T-1|T}$ and $P_{T-1|T}$, thus allowing us to draw $\tilde{\theta}_{T-1}$ from $N(\theta_{T-1|T}, P_{T-1|T})$, and so on until $t = 1$.

(b) *Drawing the elements of H_t* Conditional on Y^T , θ^T , and α , the orthogonalised innovations $u_t \equiv A(Y_t - X_t'\theta_t)$, with $\text{Var}(u_t) = H_t$, are observable. Following Cogley and Sargent (2002), we then sample the $h_{i,t}$'s by applying the univariate algorithm of Jacquier, Polson, and Rossi (1994) element by element.¹⁸

(c) *Drawing the hyperparameters* Conditional on Y^T , θ^T , H^T , and α , the innovations to θ_t and to the $h_{i,t}$'s are observable, which allows us to draw the hyperparameters, namely the elements of Q and the σ_i^2 , from their respective distributions.

(d) *Drawing the elements of α* Finally, conditional on Y^T and θ^T the ϵ_t 's are observable, satisfying

$$A\epsilon_t = u_t, \quad (\text{B9})$$

with the u_t being a vector of orthogonalized residuals with known time-varying variance H_t . Following Cogley and Sargent (2005), we interpret (B9) as a system of unrelated regressions. The first equation in the system is given by $\epsilon_{1,t} \equiv u_{1,t}$, while the following equations can be expressed as transformed regressions as

$$\begin{aligned} \left(h_{2,t}^{-\frac{1}{2}}\epsilon_{2,t}\right) &= -\alpha_{2,1}\left(h_{2,t}^{-\frac{1}{2}}\epsilon_{1,t}\right) + \left(h_{2,t}^{-\frac{1}{2}}u_{2,t}\right) \\ \left(h_{3,t}^{-\frac{1}{2}}\epsilon_{3,t}\right) &= -\alpha_{3,1}\left(h_{3,t}^{-\frac{1}{2}}\epsilon_{1,t}\right) - \alpha_{3,2}\left(h_{3,t}^{-\frac{1}{2}}\epsilon_{2,t}\right) + \left(h_{3,t}^{-\frac{1}{2}}u_{3,t}\right) \\ &\quad \dots \\ \left(h_{N(N-1)/2,t}^{-\frac{1}{2}}\epsilon_{N(N-1)/2,t}\right) &= -\alpha_{N(N-1)/2,1}\left(h_{N(N-1)/2,t}^{-\frac{1}{2}}\epsilon_{1,t}\right) - \dots \\ &\quad \dots - \alpha_{N(N-1)/2,N(N-1)/2}\left(h_{N(N-1)/2,t}^{-\frac{1}{2}}\epsilon_{N(N-1)/2,t}\right) + \left(h_{N(N-1)/2,t}^{-\frac{1}{2}}u_{N(N-1)/2,t}\right) \end{aligned} \quad (\text{B10})$$

where the residuals are independent standard normal. Assuming normal priors for each equation's regression coefficients, the posterior is also normal, and can be computed via equations (77) of (78) in Cogley and Sargent (2005, section B.2.4).

¹⁸For details, see Cogley and Sargent (2005, Appendix B.2.5).

Summing up, the MCMC algorithm simulates the posterior distribution of the states and the hyperparameters, conditional on the data, by iterating on (a)-(d). In what follows, we use a burn-in period of 50,000 iterations to converge to the ergodic distribution, and after that we run 10,000 more iterations sampling every 10th draw in order to reduce the autocorrelation across draws.¹⁹

C Computing Generalised Impulse Response Functions

Here we describe the Monte Carlo integration procedure we use in Section 4.1 in order to compute the generalised IRFs to the structural shocks.

Randomly draw the current state of the economy at time t from the Gibbs sampler’s output. Given the current state of the economy, repeat the following procedure 100 times.

- Draw four independent $N(0, 1)$ variates (the four structural shocks), and based on the relationship $\epsilon_t = A_{0,t}e_t$, with $e_t \equiv [e_t^{SP}, e_t^R, e_t^D, e_t^{ST}]'$ —where e_t^{SP} , e_t^R , e_t^D , and e_t^{ST} are the permanent output shock, and the interest rate, demand non-interest rate, and transitory supply structural shocks, respectively—compute the reduced-form shocks ϵ_t at time t .
- Simulate both the VAR’s time-varying parameters and the covariance matrix of its reduced-form innovations, Ω_t , 40 quarters into the future. Based on the simulated Ω_t , randomly draw reduced-form shocks from $t+1$ to $t+40$. Based on the simulated θ_t , and on the sequence of reduced-form shocks from t to $t+40$, compute simulated paths for the four endogenous variables. Call these simulated paths as $\hat{X}_{t,t+40}^j$, with $j = 1, \dots, 100$.
- Repeat the same procedure based on *exactly* the same simulated paths for the VAR’s time-varying parameters, the θ_t ; the same reduced-form shocks at times $t+1$ to $t+40$; and the same structural shocks e_t^{SP} , e_t^R , e_t^D , and e_t^{ST} at time t , with the only difference that, in order to compute the GIRF to shock e_t^x , with $x = SP, R, D, ST$, you set $e_t^x = 1$. Call these simulated paths as $\tilde{X}_{t,t+40}^j$, with $j = 1, \dots, 100$.

For each of the 100 iterations define $irf_{t,t+40}^j \equiv \hat{X}_{t,t+40}^j - \tilde{X}_{t,t+40}^j$. Finally, compute each of the 1,000 generalised IRFs as the mean of the distribution of the $irf_{t,t+40}^j$ ’s.

¹⁹In this we follow Cogley and Sargent (2005). As stressed by these authors, however, this has the drawback of “increasing the variance of ensemble averages from the simulation”.

References

- BALKE, N., AND R. GORDON (1986): “Appendix B: Historical Data,” in *Gordon, R.J., ed. (1986), The American Business Cycle: Continuity and Changes, The University of Chicago Press.*
- BAXTER, M., AND R. KING (1999): “Approximate Band-Pass Filters for Economic Time Series: Theory and Applications,” *Review of Economics and Statistics*, 81(4), 575–593.
- BENATI, L. (2008): “The Great Moderation in the United Kingdom,” *Journal of Money, Credit, and Banking*, 39(1), 121–147.
- BENATI, L., AND C. GOODHART (2011): “Monetary Policy Regimes and Economic Performance: The Historical Record, 1979–2008,” in *B. Friedman, B., and Woodford, M. (eds.), Handbook of Monetary Economics, Volume 3, North Holland.*
- BENATI, L., AND P. SURICO (2008): “Evolving U.S. Monetary Policy and the Decline of Inflation Predictability,” *Journal of the European Economic Association, Proceedings, April-May 2008*, 6(2-3), 634–646.
- BLANCHARD, O. J., AND D. QUAH (1989): “The Dynamic Effects of Aggregate Demand and Supply Disturbances,” *American Economic Review*, 79(4), 655–673.
- BLINDER, A., AND L. J. MACCINI (1991a): “The Resurgence of Inventory Research: What Have We Learned?,” *Journal of Economic Surveys*, 5, 291–328.
- (1991b): “Taking Stock: A Critical Assessment of Recent Research on Inventories,” *Journal of Economic Perspectives*, 5, 73–96.
- CARTER, C. K., AND R. P. KOHN (2004): “On Gibbs Sampling for State Space Models,” *Biometrika*, 81, 541–553.
- CHRISTIANO, L., AND T. FITZGERALD (2003): “The Band-Pass Filter,” *International Economic Review*, 44(2), 435–465.
- COGLEY, T., AND T. J. SARGENT (2002): “Evolving Post-WWII U.S. Inflation Dynamics,” in *B. Bernanke and K. Rogoff, eds. (2002), NBER Macroeconomics Annuals 2001.*
- (2005): “Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII U.S.,” *Review of Economic Dynamics*, 8(April), 262–302.
- COGLEY, T. W., G. E. PRIMICERI, AND T. J. SARGENT (2010): “Inflation-Gap Persistence in the U.S.,” *American Economic Journal: Macroeconomics*, January.

- COMIN, D., AND M. GERTLER (2006): “Medium-Term Business Cycles,” *American Economic Review*, 96(3), 523–551.
- GALI, J., AND L. GAMBETTI (2009): “On the Sources of the Great Moderation,” *American Economic Journal: Macroeconomics*, 1(1), 26–57.
- GOLUB, G., AND C. VANLOAN (1996): *Matrix Computations, III Edition*. Johns Hopkins University Press, Baltimore.
- JACQUIER, E., N. G. POLSON, AND P. ROSSI (1994): “Bayesian Analysis of Stochastic Volatility Models,,” *Journal of Business and Economic Statistics*, 12, 371–418.
- KAHN, J., M. MCCONNELL, AND G. PEREZ-QUIROS (2002): “On the Causes of the Increased Stability of the U.S. Economy,” *Federal Reserve Bank of New York, Economic Policy Review*, 8(1), 183–202.
- KIM, C. J., AND C. NELSON (1999): *State-Space Models with Regime Switching*. Cambridge, Mass., The MIT Press.
- MACCINI, L. J., B. MOORE, AND H. SCHALLER (2004): “The Interest Rate, Learning, and Inventory Investment,” *American Economic Review*, 94(5), 1303–1327.
- MACCINI, L. J., B. MOORE, AND H. SHALLER (2010): “A Unified Solution to Inventory Puzzles,” , mimeo.
- PRIMICERI, G. E. (2005): “Time Varying Structural Vector Autoregressions and Monetary Policy,” *The Review of Economic Studies*, 72, 821–852.
- RAMEY, V. A., AND K. D. WEST (1999): “Inventories,” in *John B. Taylor and Michael Woodford eds., Handbook of macroeconomics, Vol. 1B. Amsterdam: North-Holland*, pp. 863–923.
- RUBIO-RAMIREZ, J., D. WAGGONER, AND T. ZHA (2005): “Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference,” *Review of Economic Studies*, 77(2), 665–696.
- STOCK, J., AND M. WATSON (1999): “Business Cycle Fluctuations in U.S. Macroeconomic Time Series,” in *Taylor, J. B., and Woodford, M., eds., Handbook of Macroeconomics*, Amsterdam, North Holland.
- WEN, Y. (2005): “Understanding the inventory cycle,” *Journal of Monetary Economics*, 52, 1533–1555.
- WEST, K. D. (1995): “Inventory Models,” in *M. Hashem Peseran and Michael R. Wickens, eds., Handbook of applied econometrics, Vol. I, Oxford: Blackwell*, pp. 188–220.

Table 1: Sign Restrictions			
Variable:	Shock:		
	ϵ_t^R	ϵ_t^D	ϵ_t^S
<i>Interest rate</i>	≥ 0	≥ 0	?
<i>Inflation</i>	≤ 0	≥ 0	≤ 0
<i>Real sales growth</i>	≤ 0	≥ 0	≥ 0
? = left unconstrained			

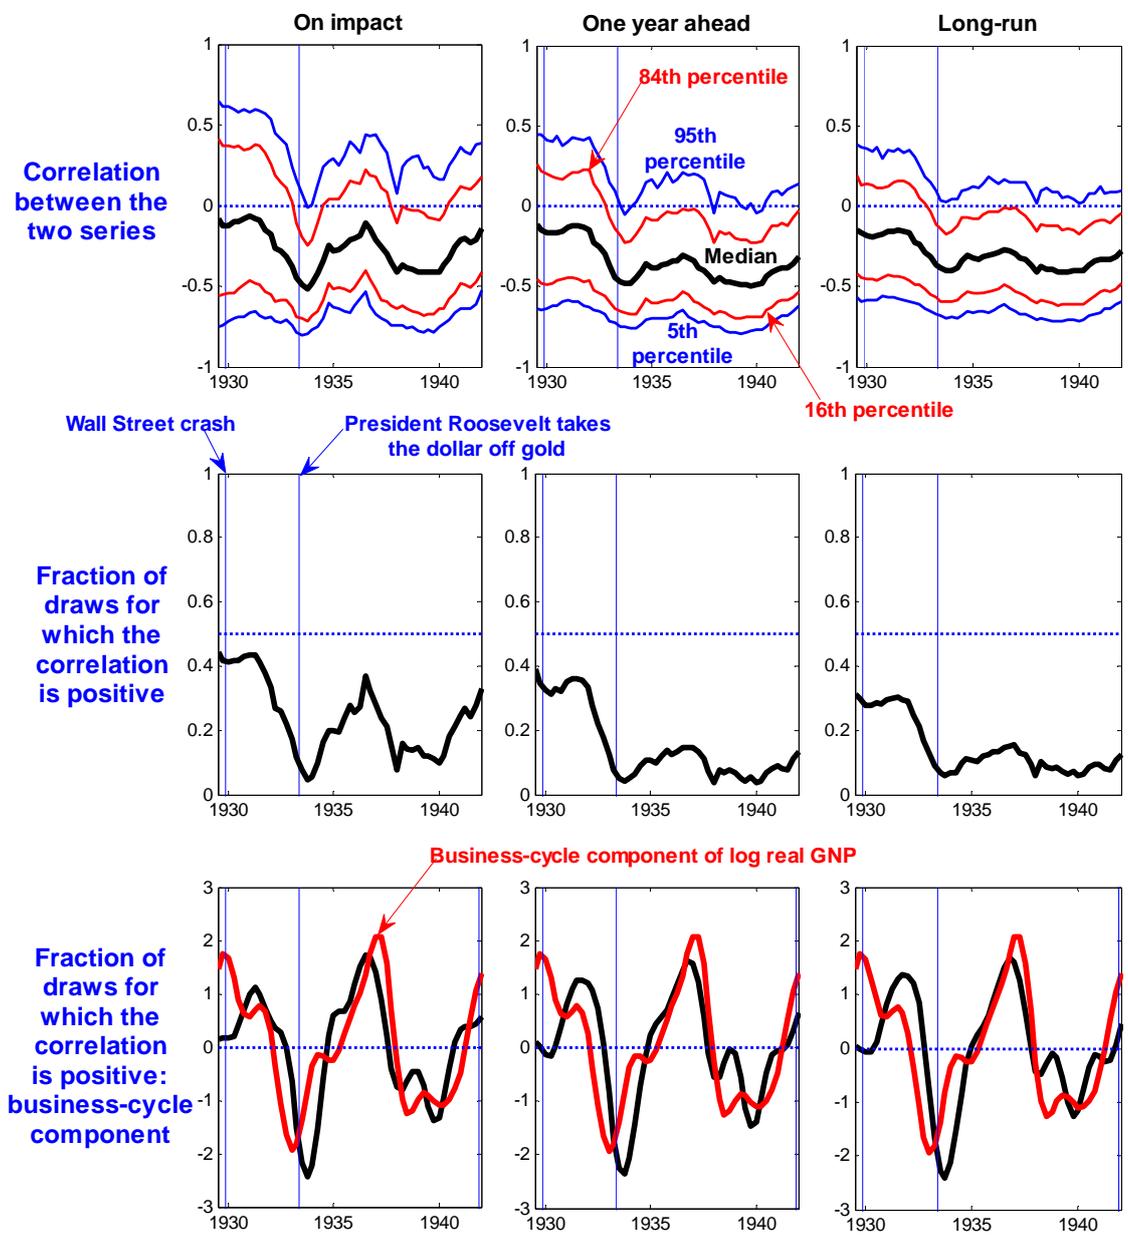


Figure 1 United States, interwar period: correlation between the reduced-form forecast errors for the real *ex post* interest rate and the change in real inventories at various horizons, and fractions of draws for which the correlation is positive

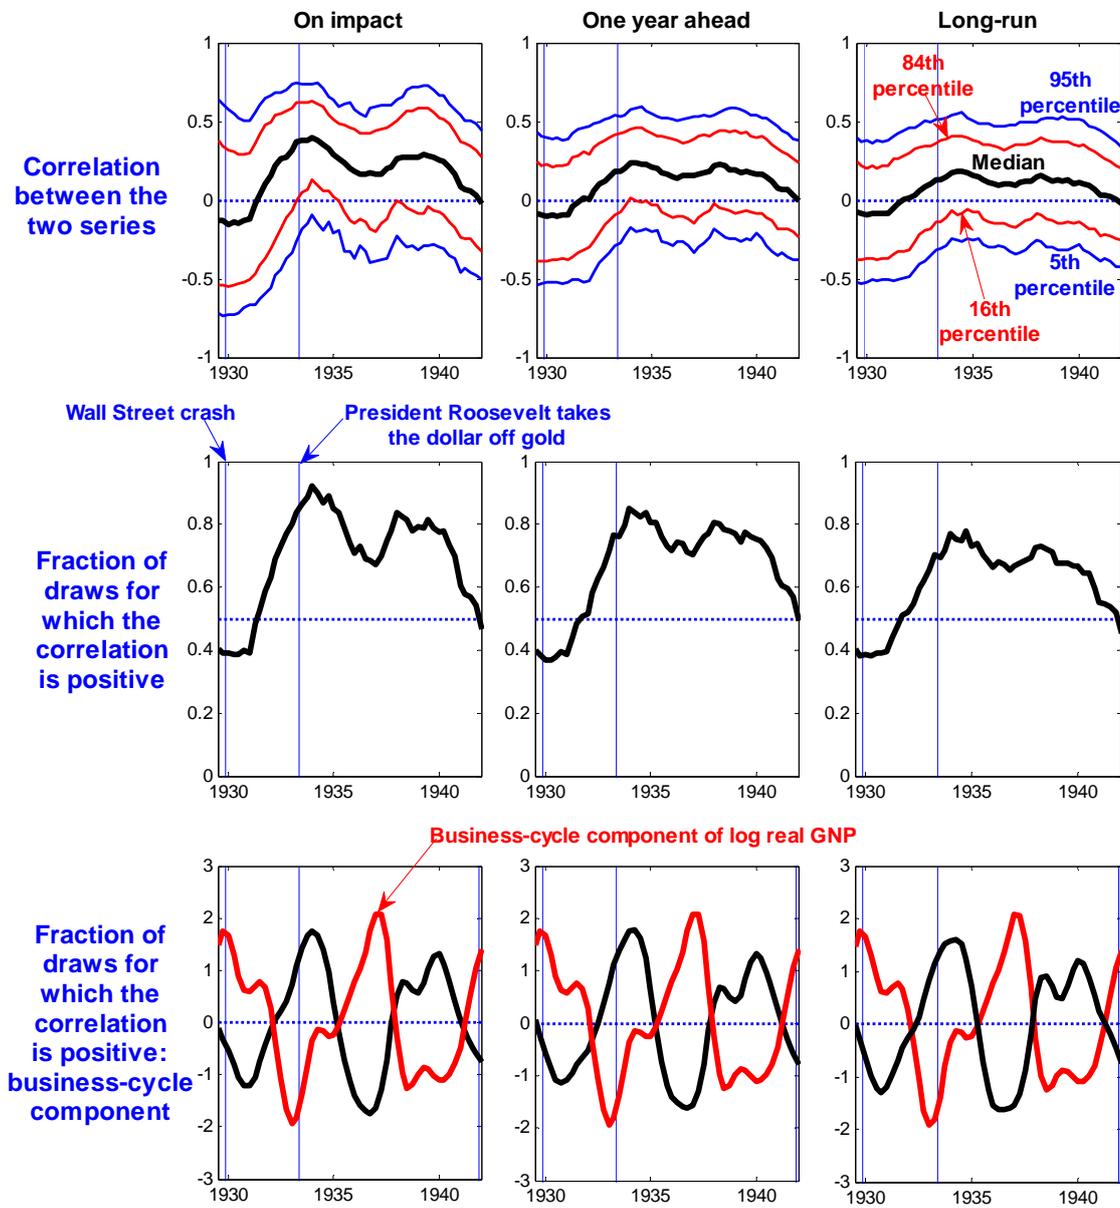


Figure 2 United States, interwar period: correlation between the reduced-form forecast errors for real sales growth and the change in real inventories at various horizons, and fractions of draws for which the correlation is positive

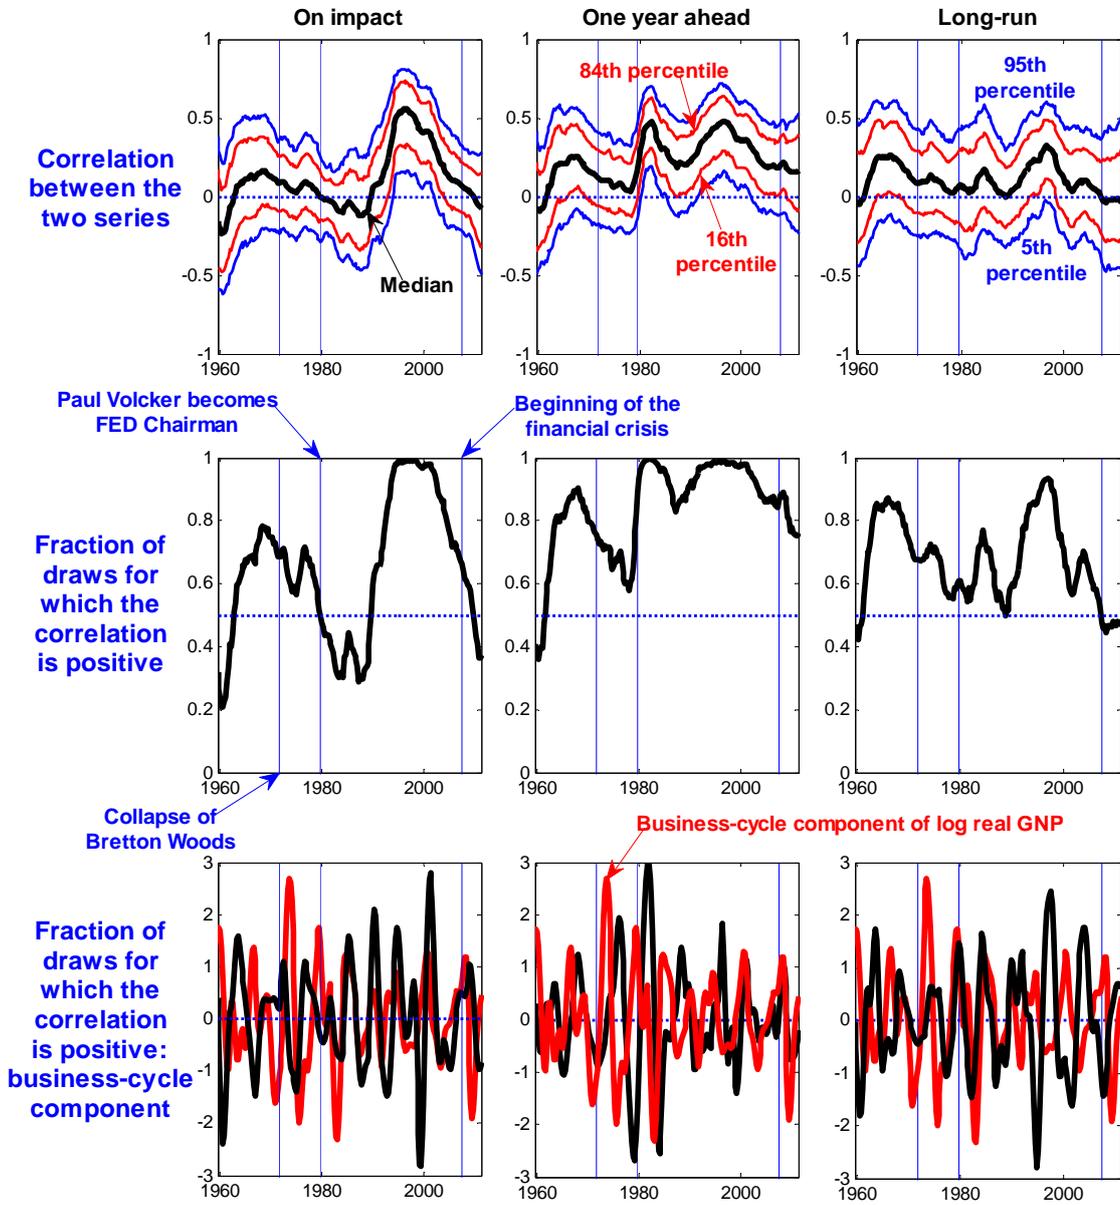


Figure 3 United States, post-WWII period: correlation between the reduced-form forecast errors for the real *ex post* interest rate and the change in real inventories at various horizons, and fractions of draws for which the correlation is positive

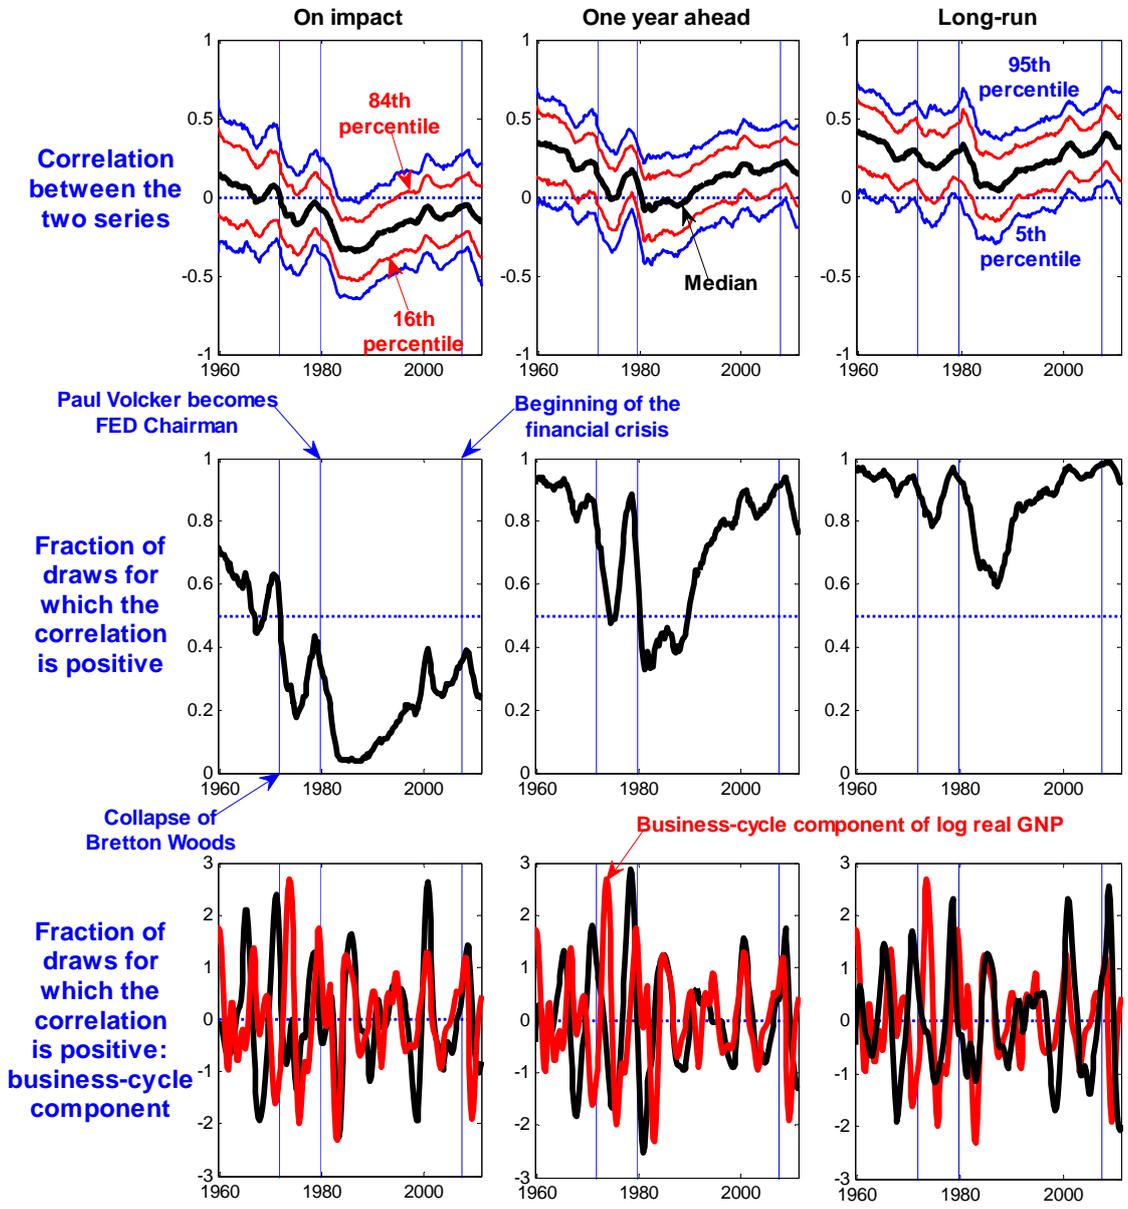


Figure 4 United States, post-WWII period: correlation between the reduced-form forecast errors for real sales growth and the change in real inventories at various horizons, and fractions of draws for which the correlation is positive

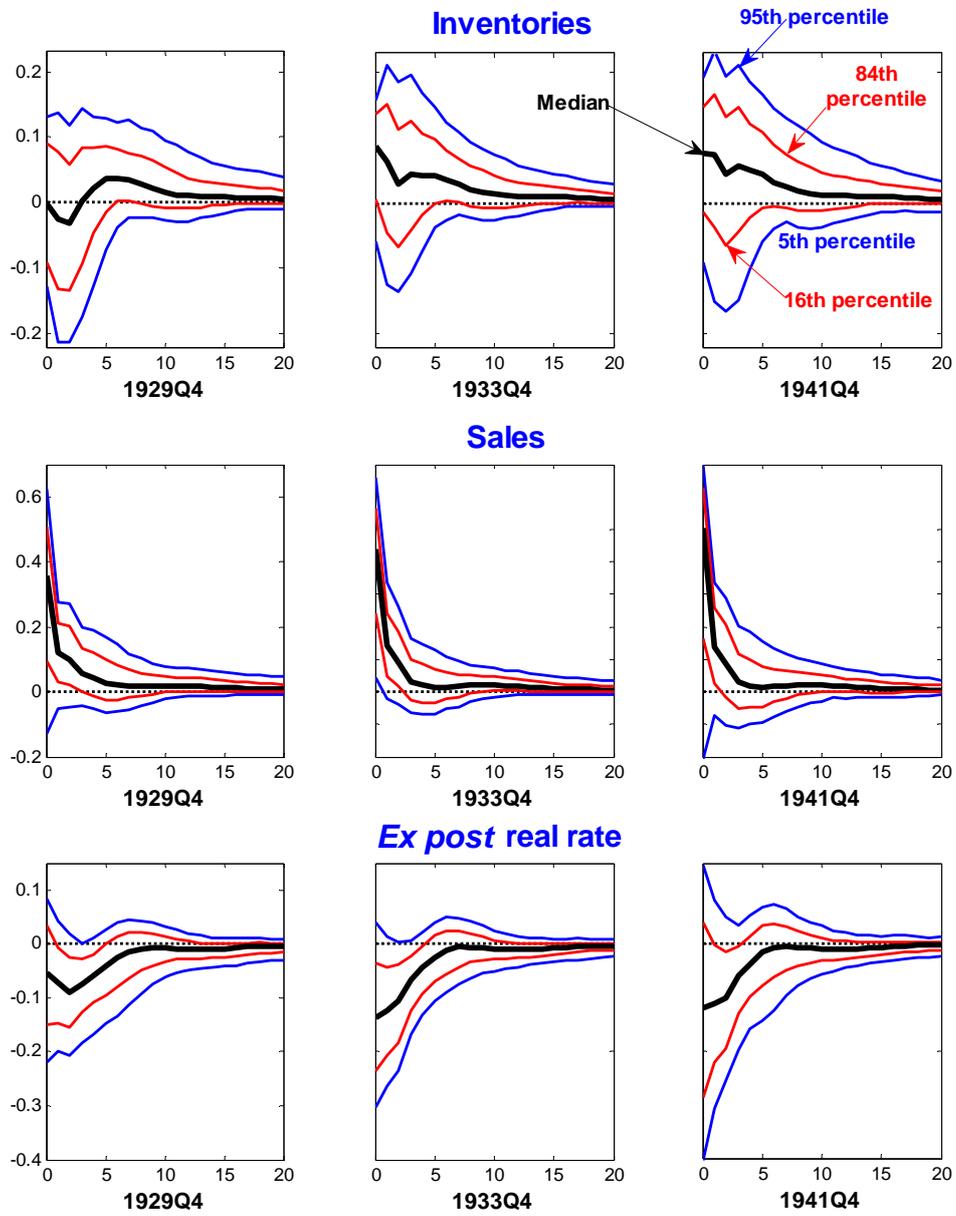


Figure 5 United States, interwar period: impulse-response functions to a permanent output shock in selected quarters

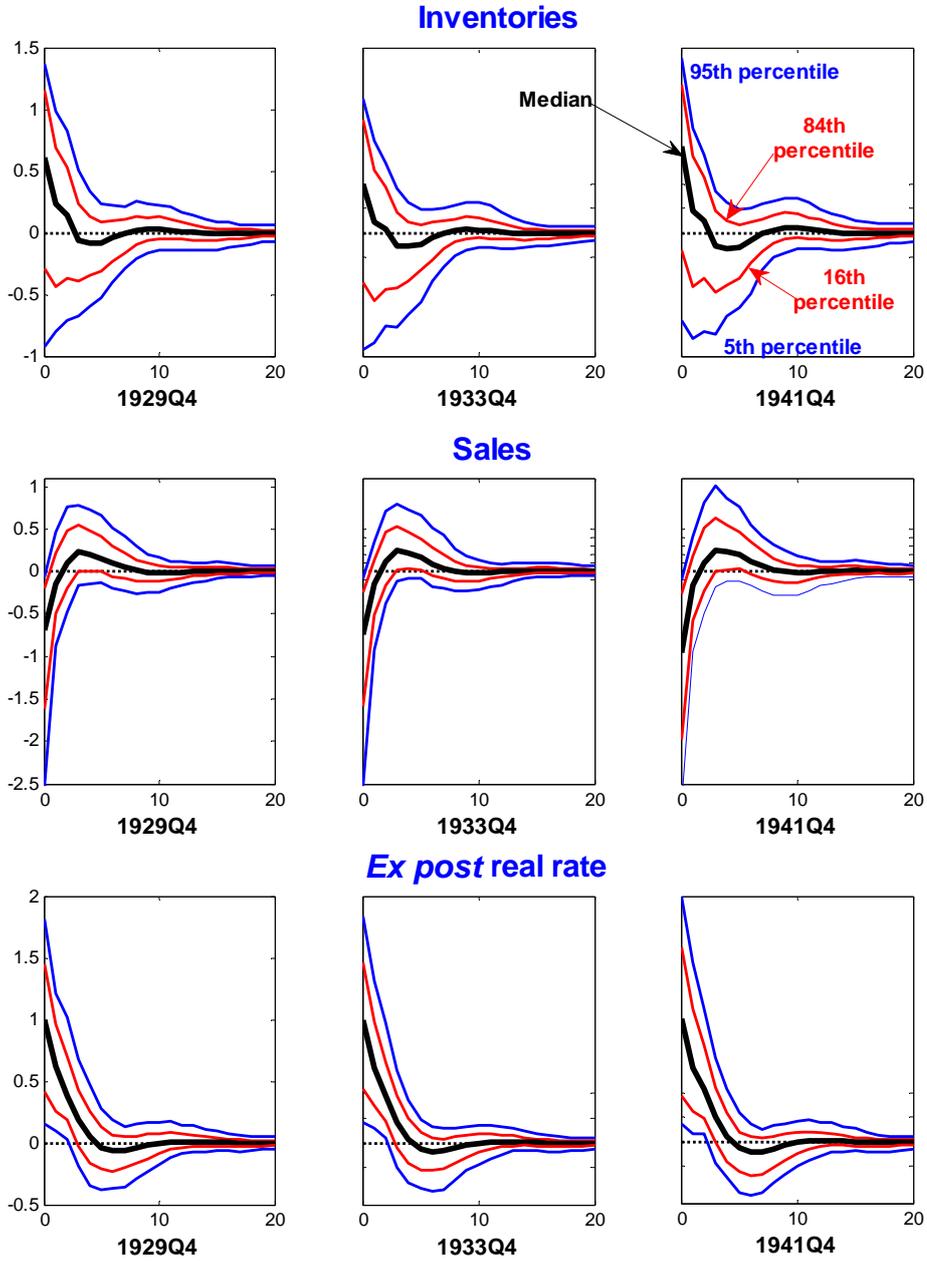


Figure 6 United States, interwar period: impulse-response functions to an interest rate shock in selected quarters

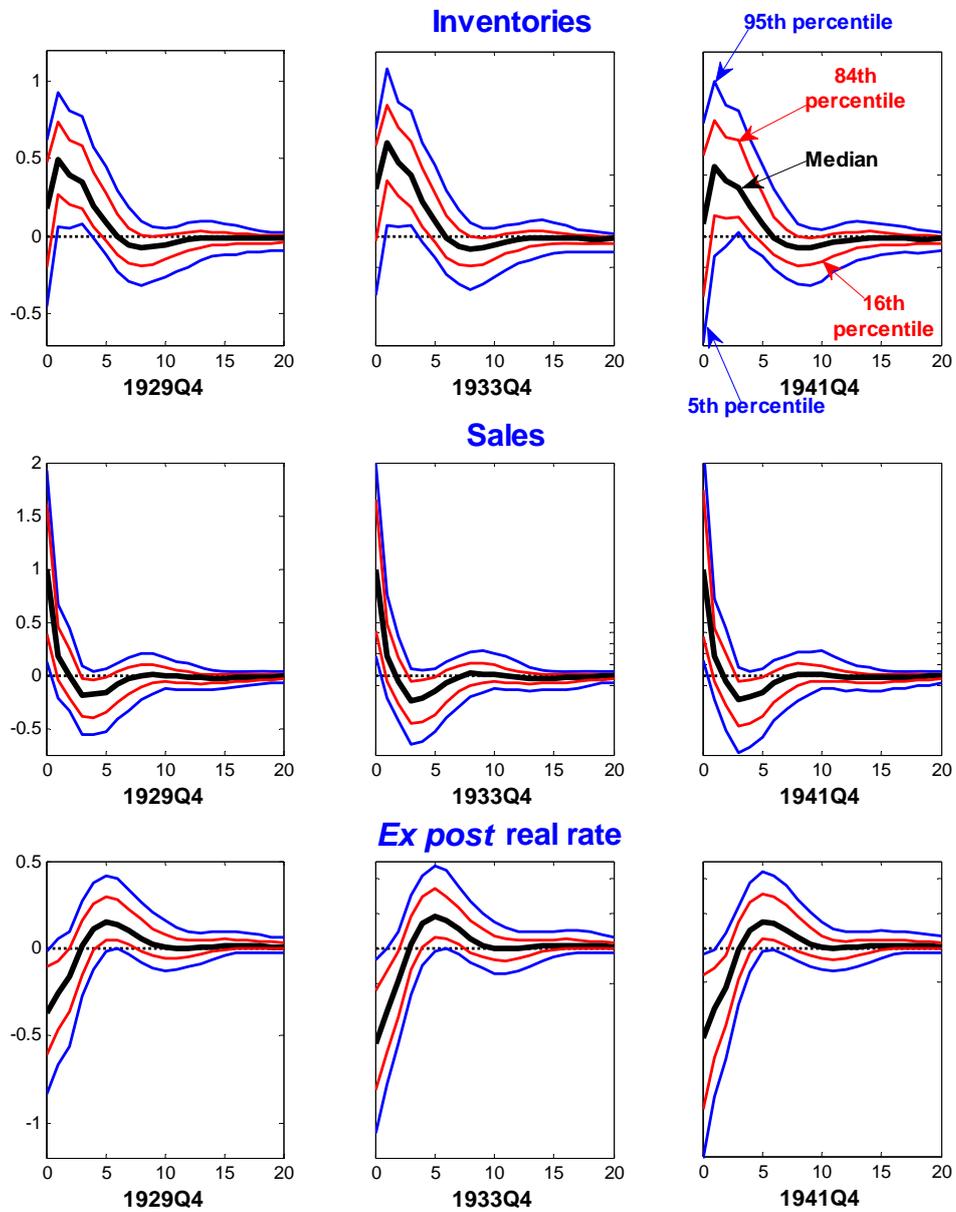


Figure 7 United States, interwar period: impulse-response functions to a demand non-interest rate shock in selected quarters

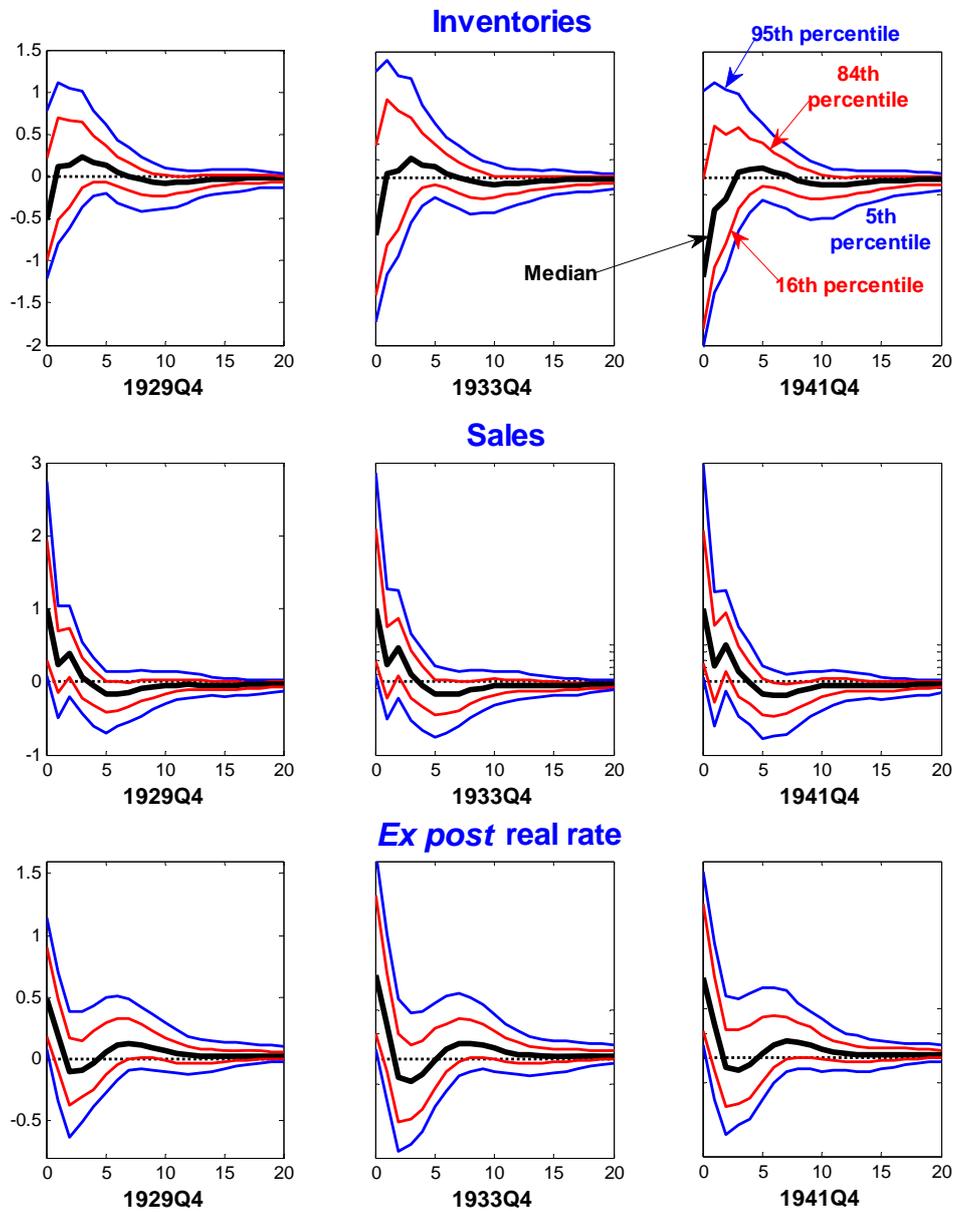


Figure 8 United States, interwar period: impulse-response functions to a transitory supply shock in selected quarters

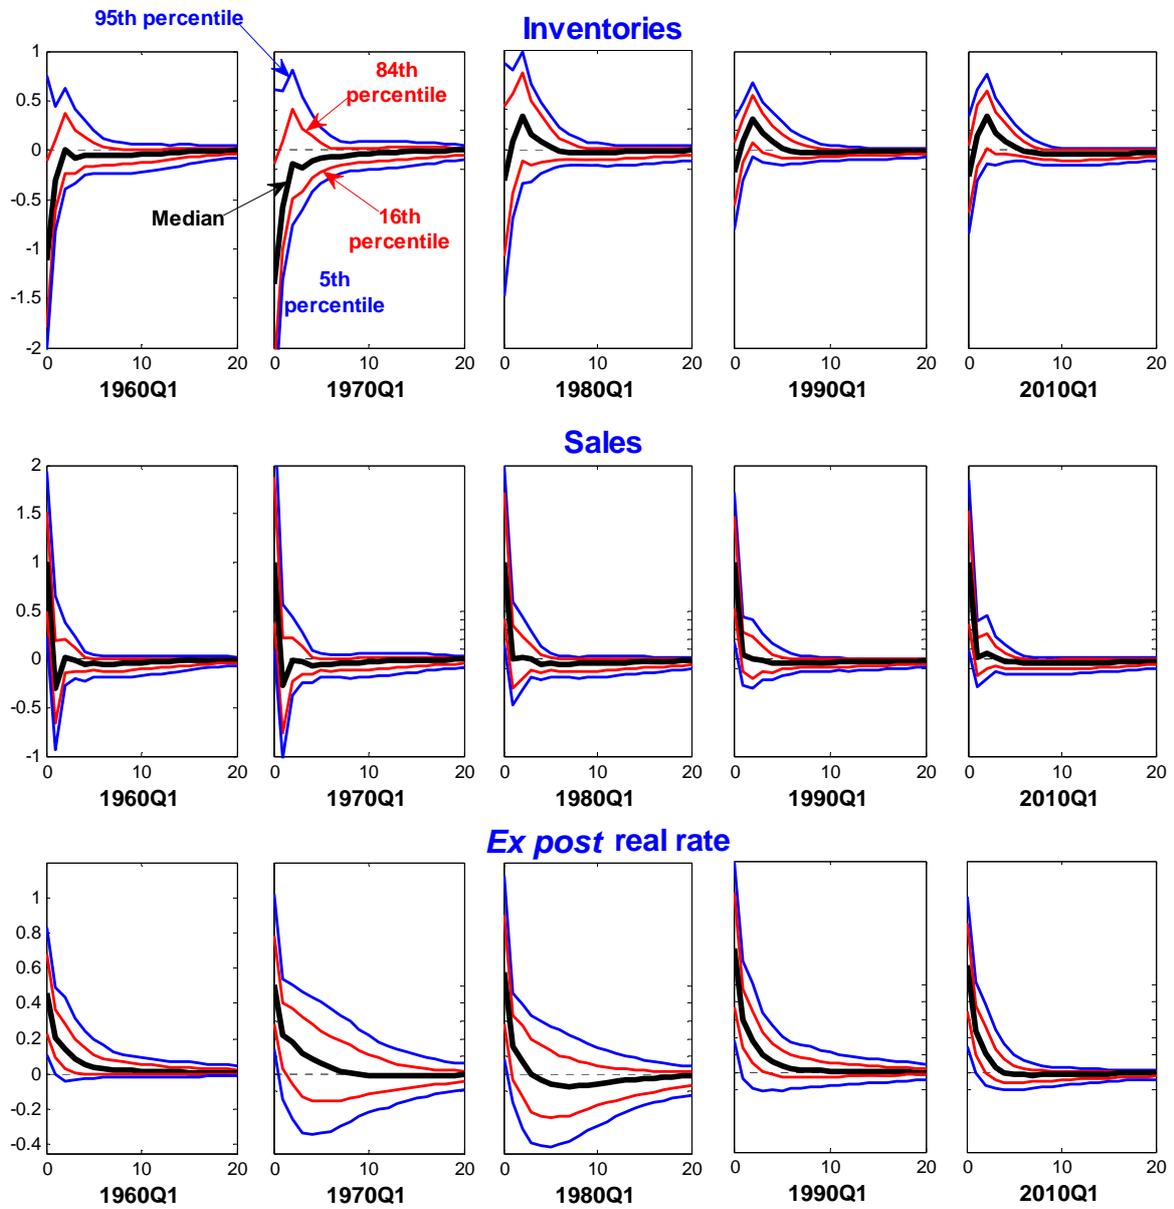


Figure 9 United States, post-WWII period: impulse-response functions to a permanent output shock in selected quarters

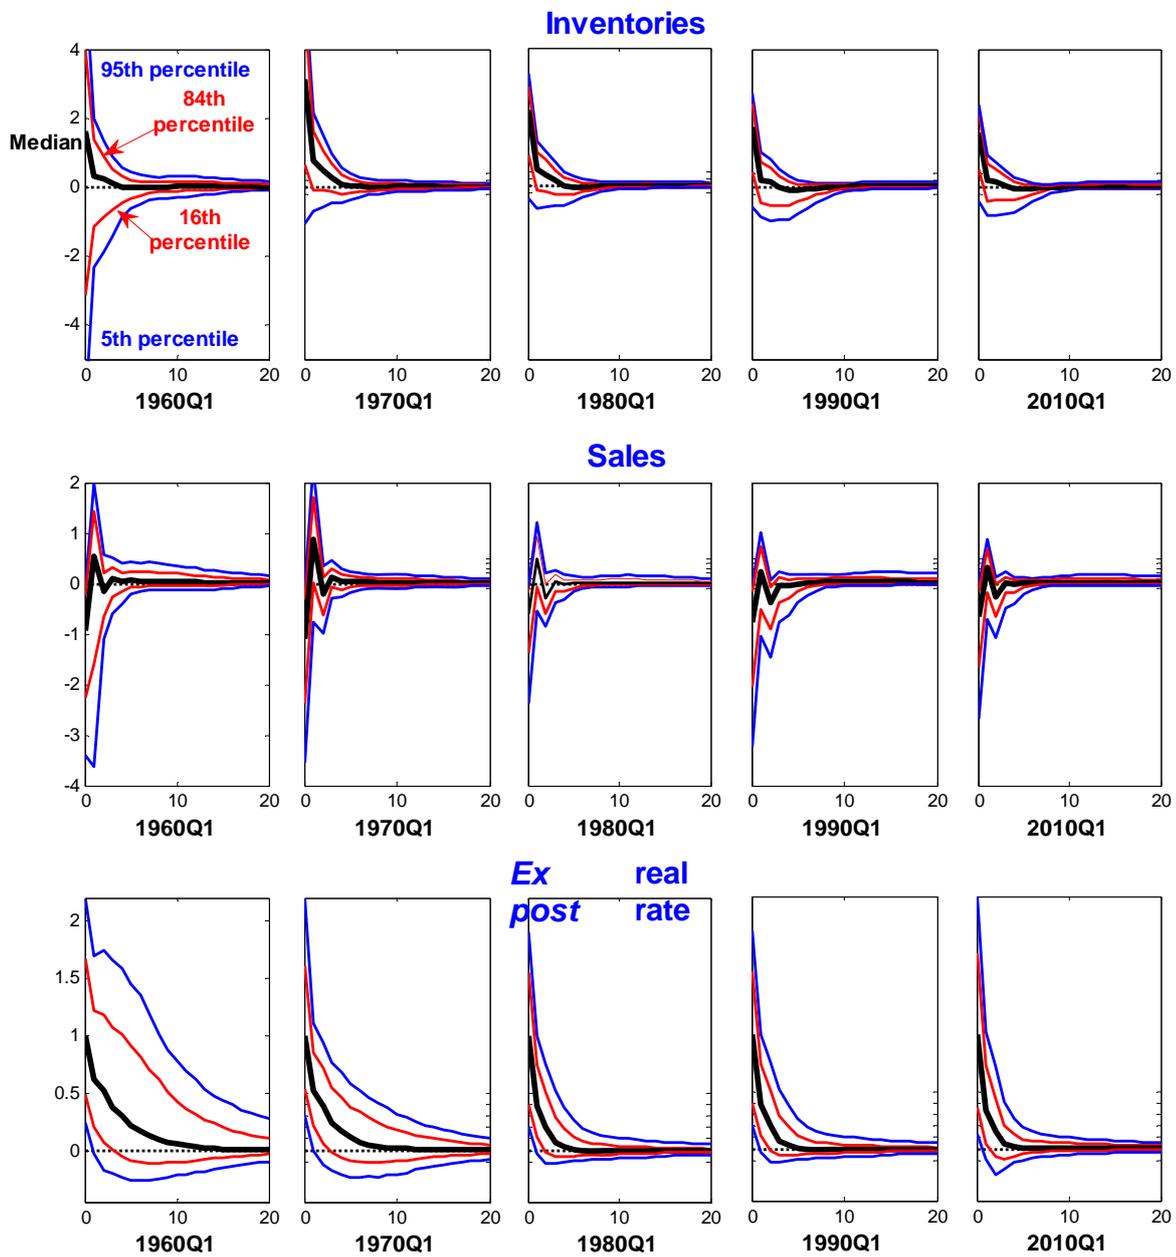


Figure 10 United States, post-WWII period: impulse-response functions to an interest rate shock in selected quarters

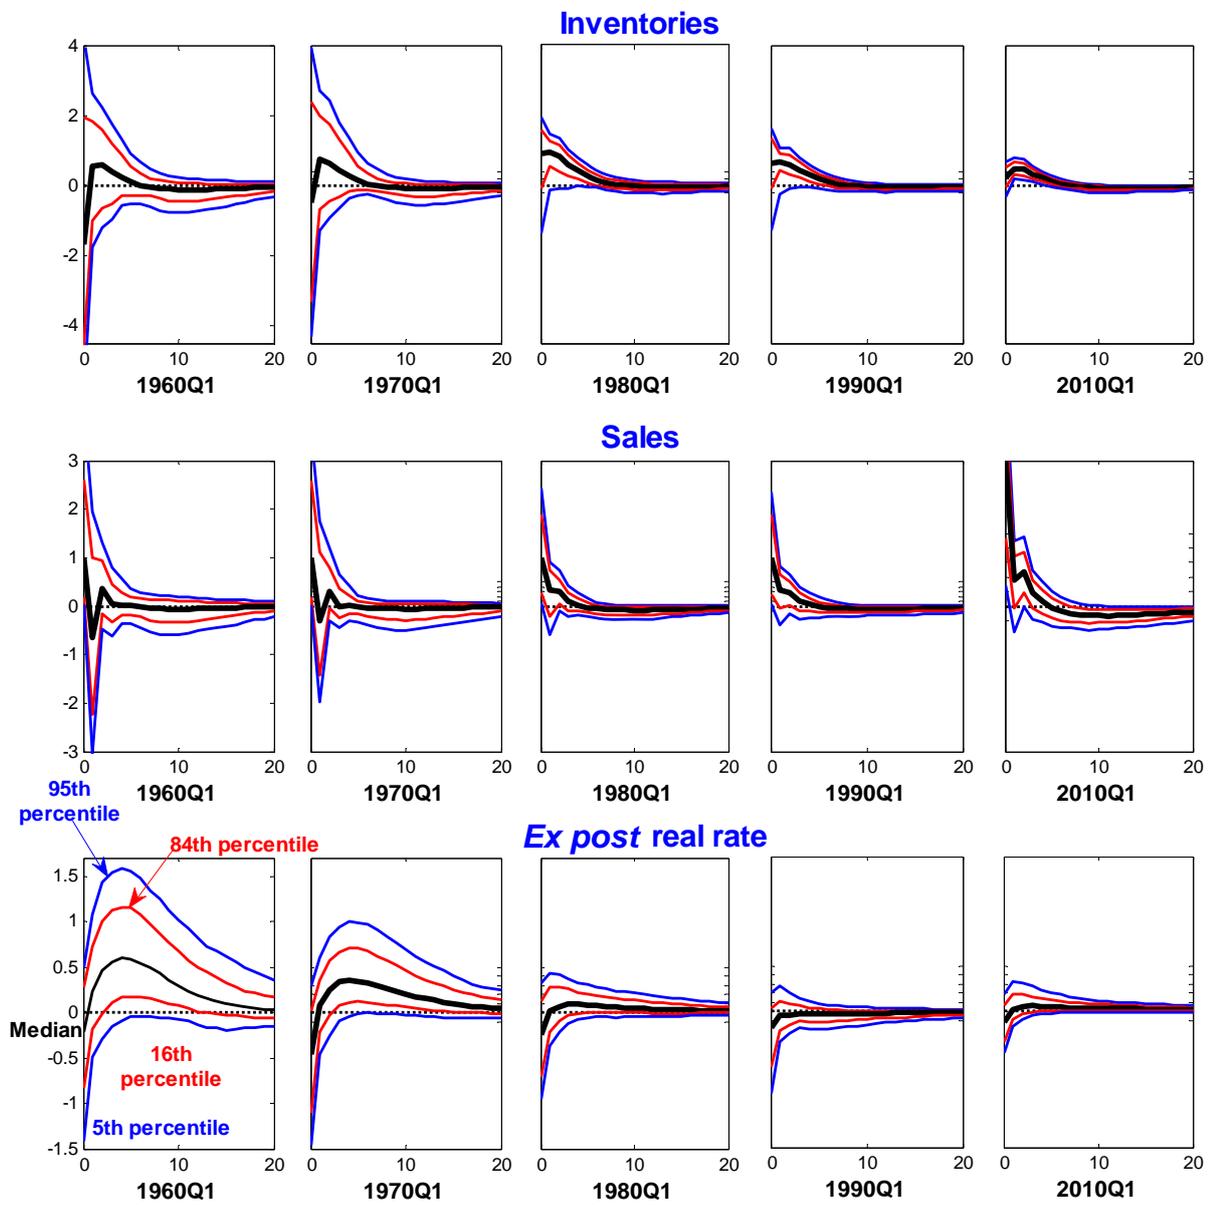


Figure 11 United States, post-WWII period: impulse-response functions to a demand non-interest rate shock in selected quarters

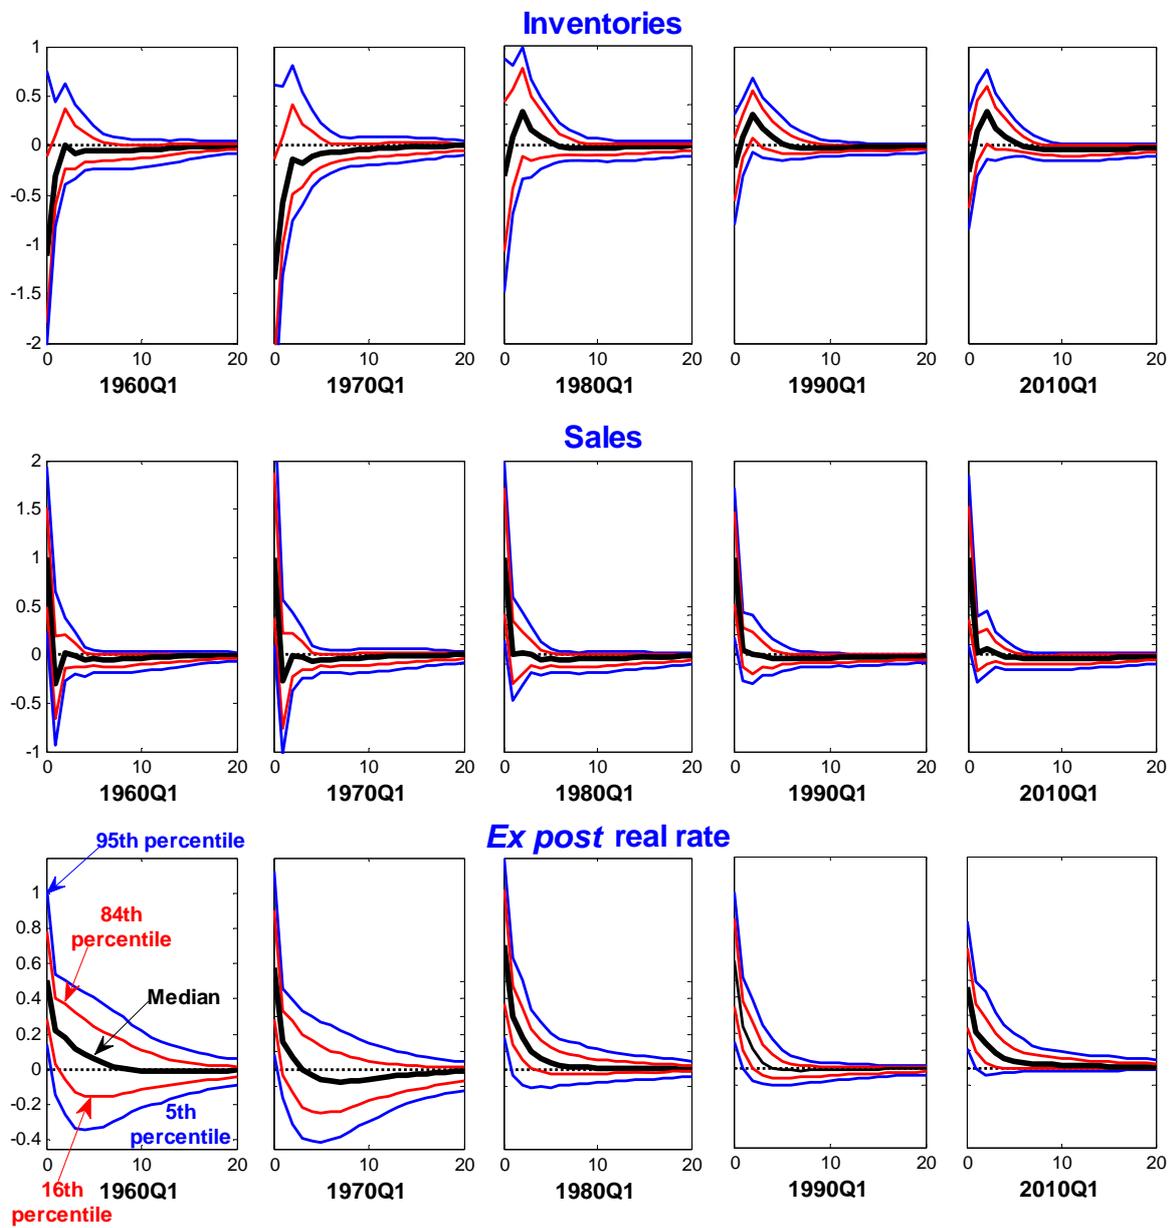


Figure 12 United States, post-WWII period: impulse-response functions to a transitory supply shock in selected quarters

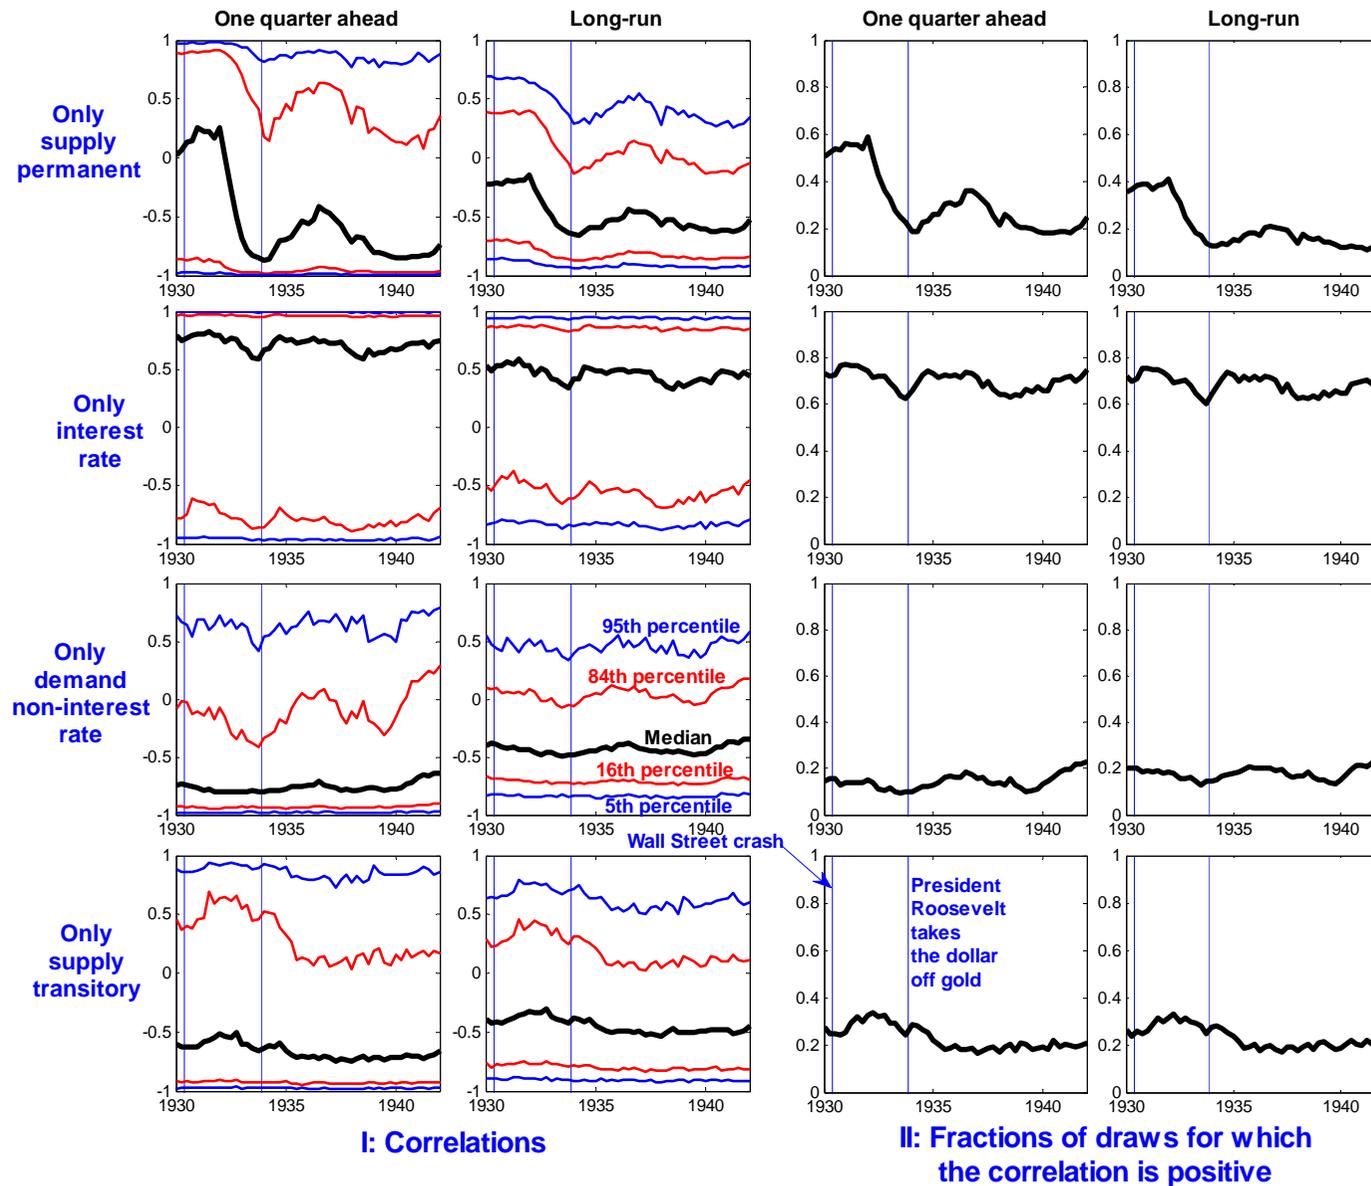


Figure 13 United States, interwar period: correlation between real *ex post* interest rate and the change in real inventories conditional on individual shocks

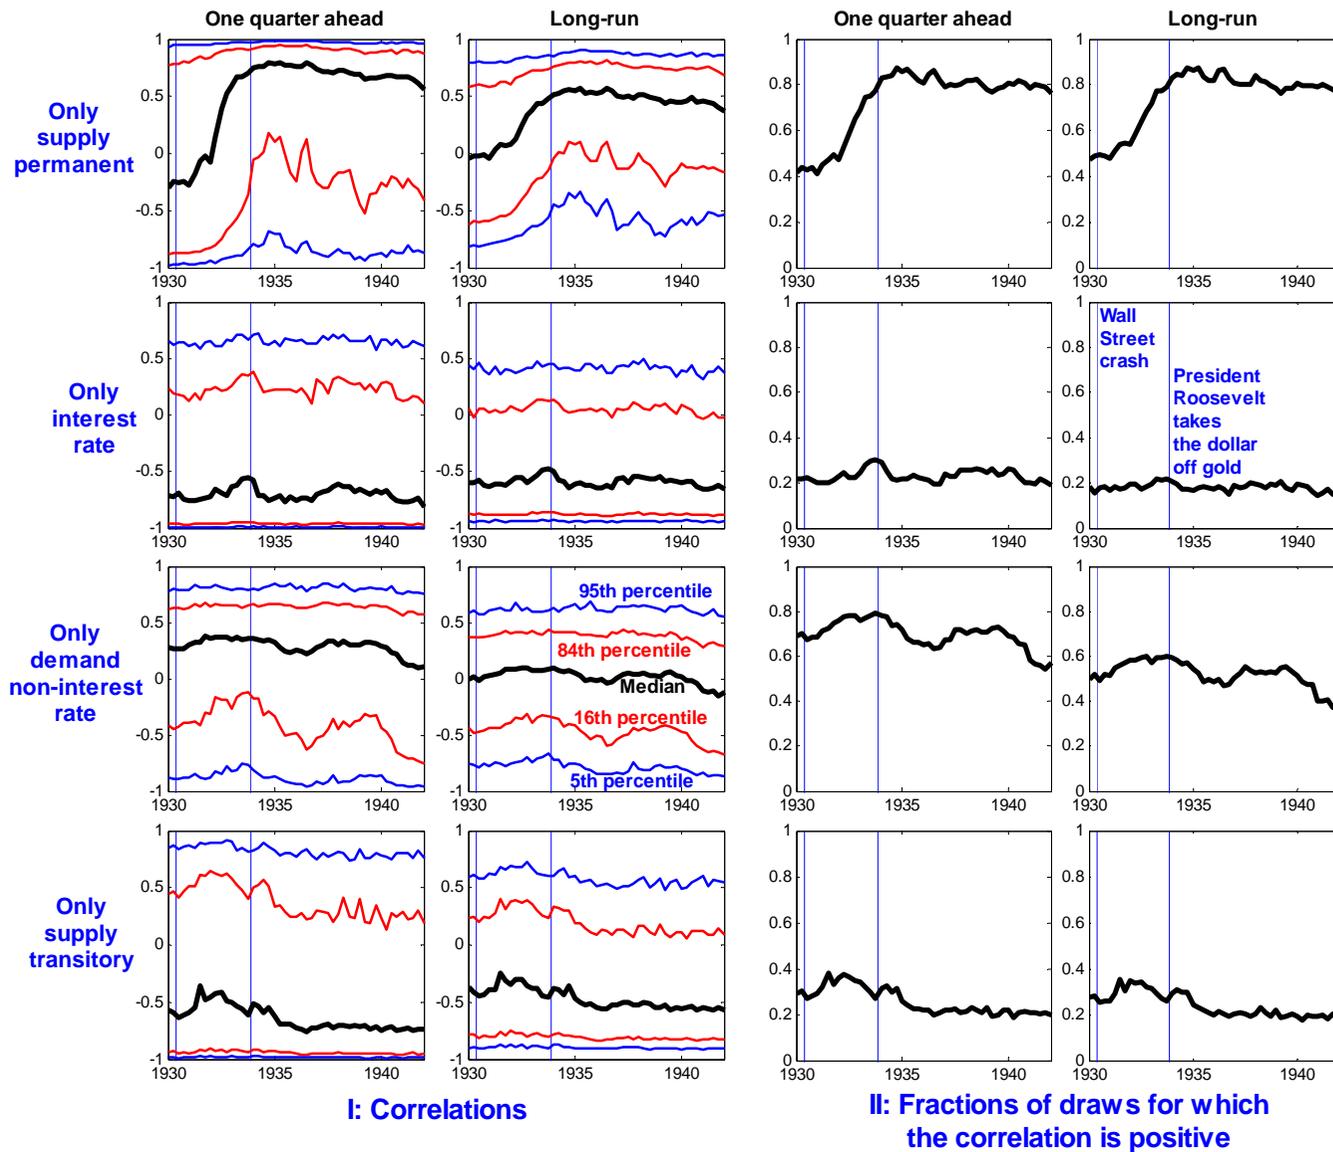


Figure 14 United States, interwar period: correlation between real sales growth and the change in real inventories conditional on individual shocks

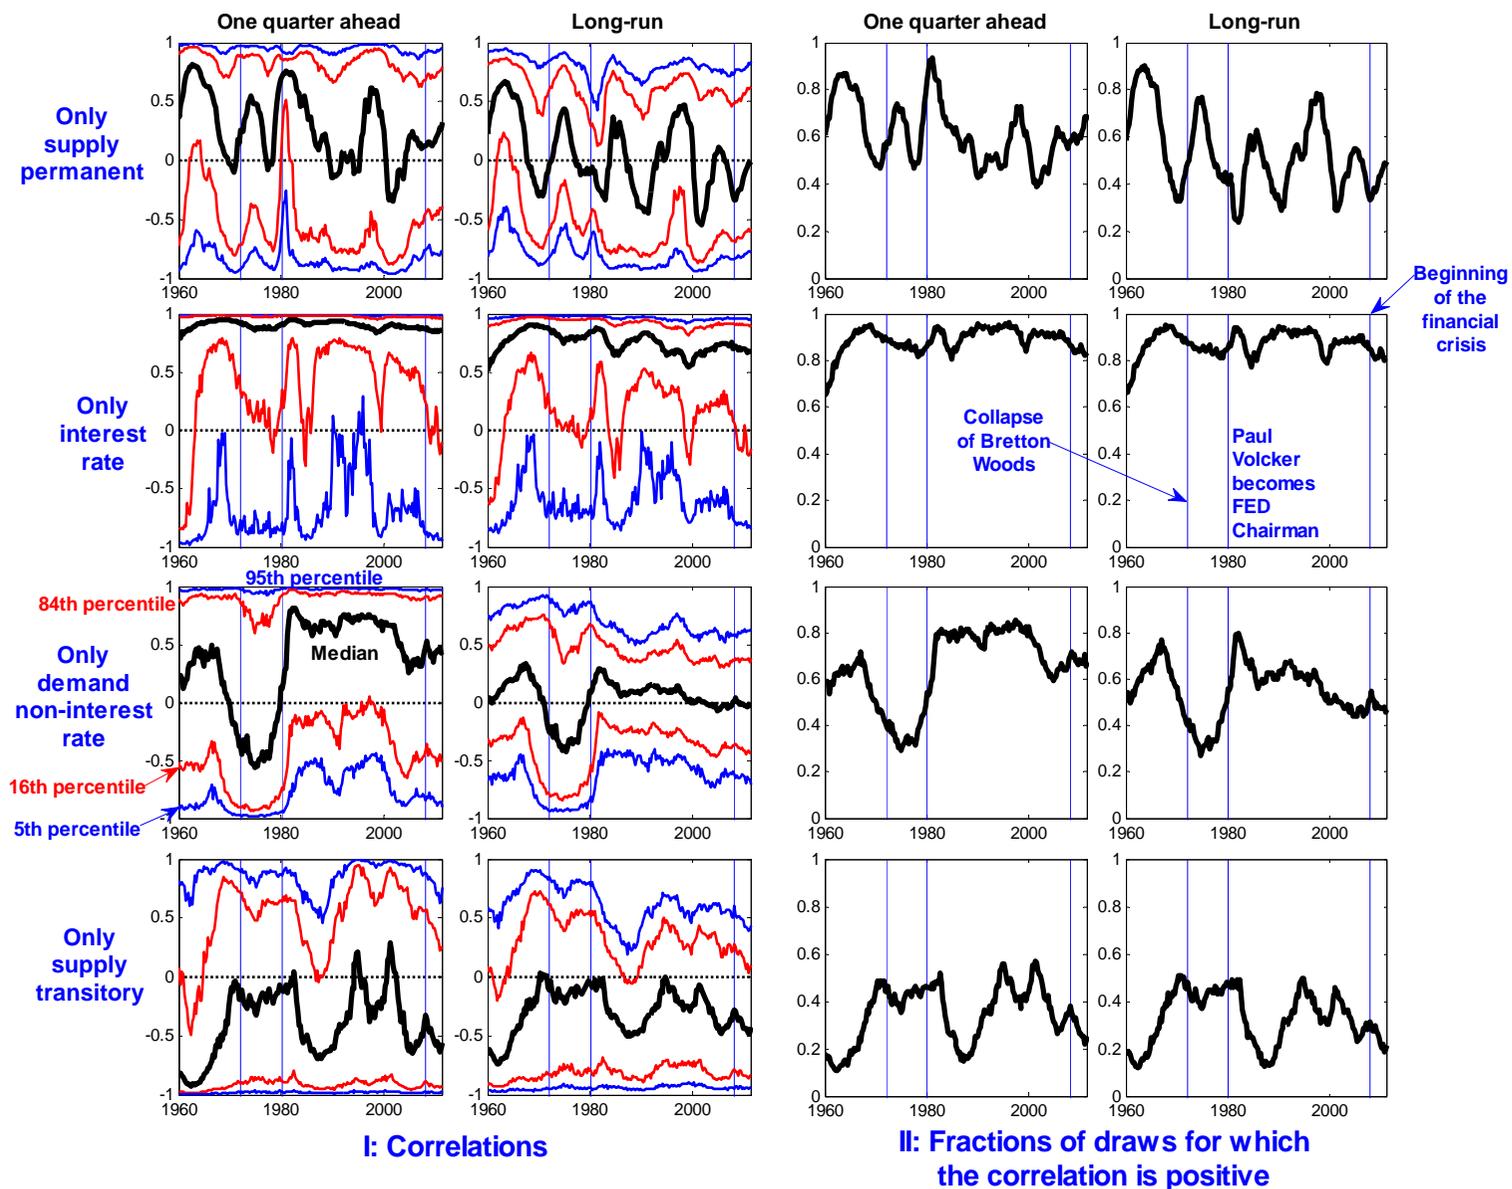


Figure 15 United States, post-WWII period: correlation between real *ex post* interest rate and the change in real inventories conditional on individual shocks

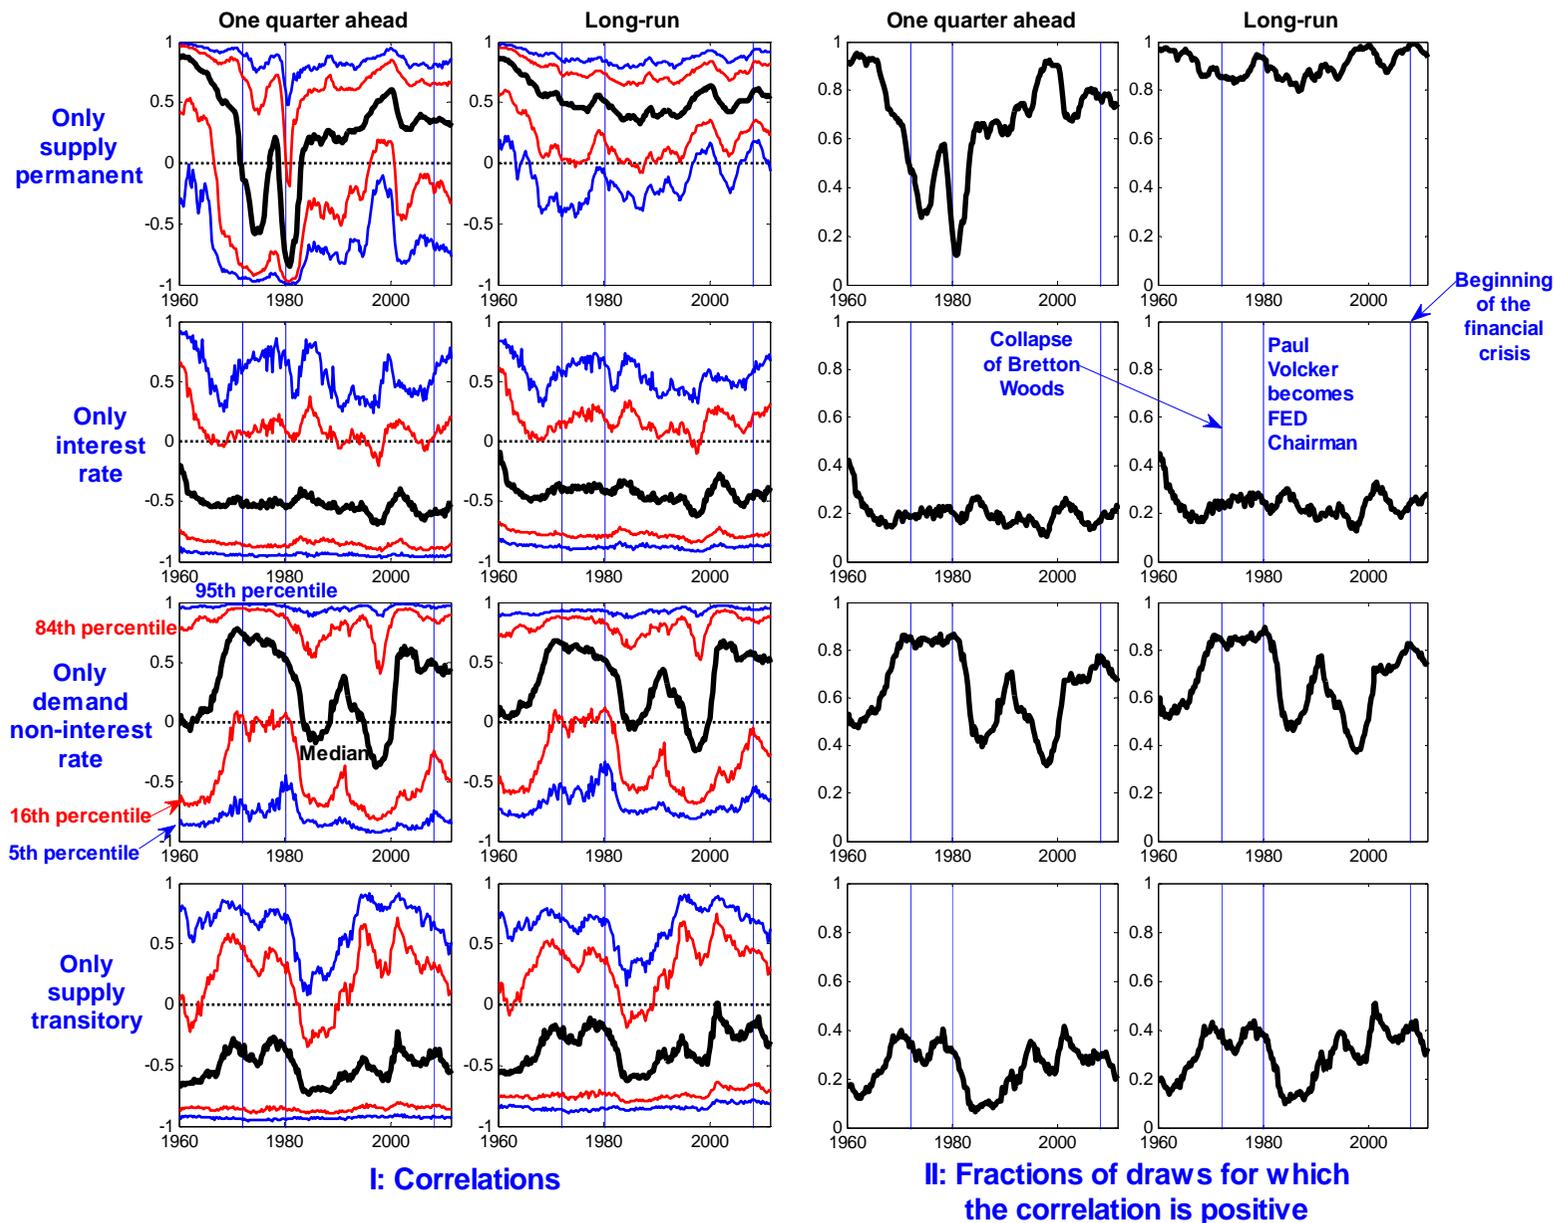


Figure 16 United States, post-WWII period: correlation between real sales growth and the change in real inventories conditional on individual shocks

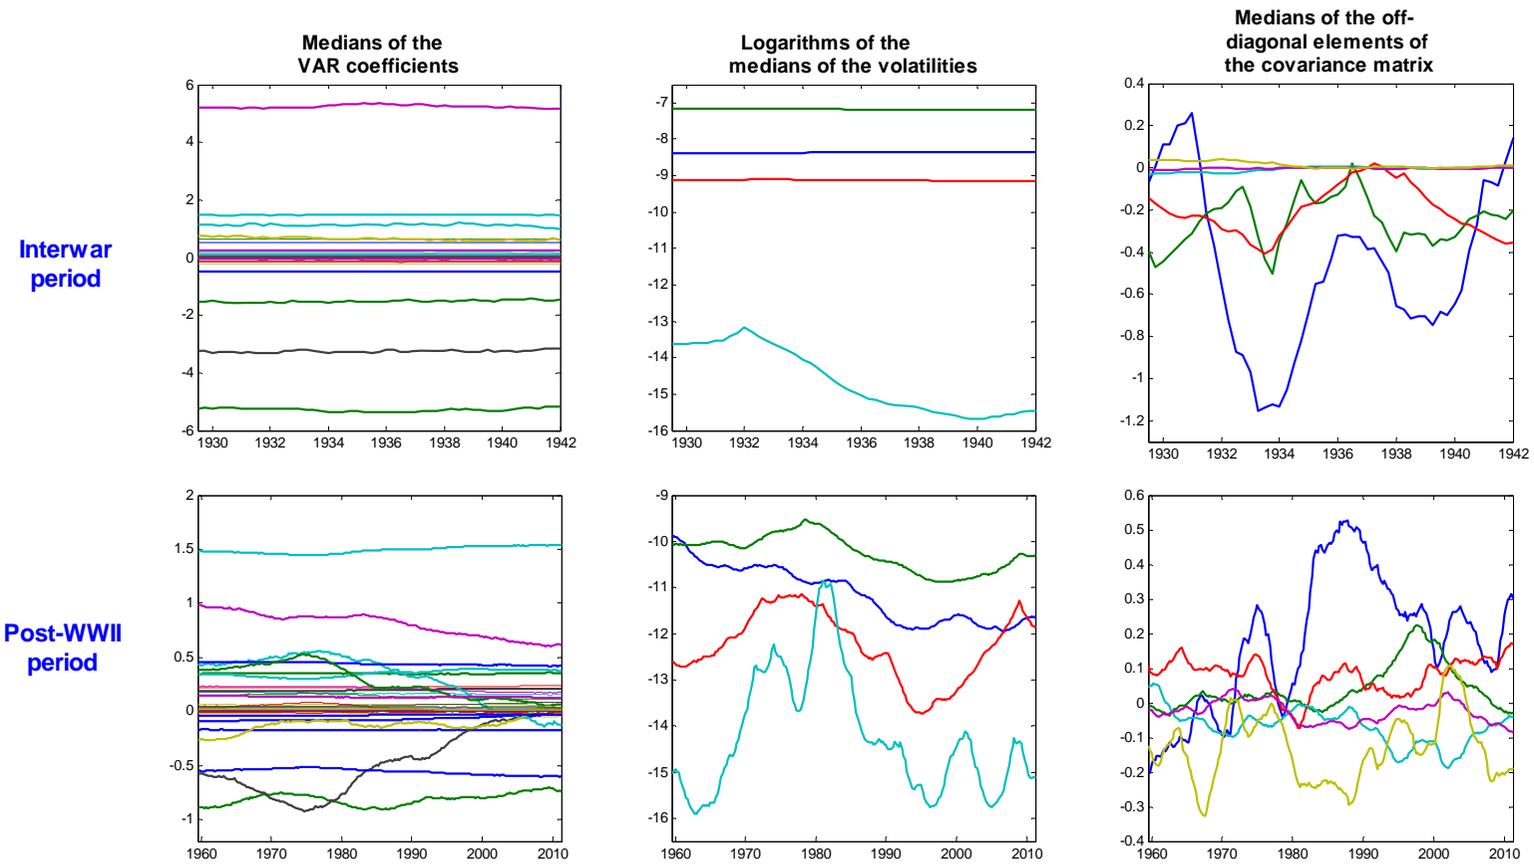


Figure 17 United States: the evolution of the elements of the VAR