Research Division
Federal Reserve Bank of St. Louis
Working Paper Series

A Flexible Finite-Horizon Alternative to Long-run Restrictions with an Application to Technology Shocks

Neville Francis
Michael T. Owyang
and
Jennifer E. Roush

Working Paper 2005-024E
http://research.stlouisfed.org/wp/2005/2005-024.pdf

July 2008

FEDERAL RESERVE BANK OF ST. LOUIS
Research Division
P.O. Box 442
St. Louis, MO 63166

The views expressed are those of the individual authors and do not necessarily reflect official positions of the Federal Reserve Bank of St. Louis, the Federal Reserve System, or the Board of Governors.

Federal Reserve Bank of St. Louis Working Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Federal Reserve Bank of St. Louis Working Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors.

A Flexible Finite-Horizon Alternative to Long-run Restrictions with an Application to Technology Shocks*

Neville Francis†
Michael T. Owyang‡
and
Jennifer E. Roush§

keywords: productivity, structural VAR, long-run restrictions

July 16, 2008

Abstract

Recent studies using long-run restrictions question the validity of the technology-driven real business cycle hypothesis. We propose an alternative identification that maximizes the contribution of technology shocks to the forecast-error variance of labor productivity at a long, but finite, horizon. In small-sample Monte Carlo experiments, our identification outperforms standard long-run restrictions by significantly reducing the bias in the short-run impulse responses and raising their estimation precision. When applied to the data, the hours response is shown to be sensitive to the contribution of non-technology shocks to the variance of productivity at long horizons. We conclude that long-run restrictions aimed at isolating the effects of technology shocks on productivity beyond business cycle frequencies do not provide information sufficient to robustly predict short-run movements in labor hours. [JEL: C32, C50, E32]

---

*Kristie M. Engemann provided research assistance. This paper benefited from conversations with Dave Altig, David Bowman, Jon Faust, John Fernald, Jonas Fisher, Oscar Jorda, Dale Henderson, Jim Nason, Marco Del Negro, Valerie Ramey, Barbara Rossi, Chris Sims, and Jonathan Wright. The authors thank seminar participants at the Federal Reserve Banks of Atlanta and New York and Duke University for comments. Views expressed here are the authors' alone and do not reflect the opinions of the Federal Reserve Bank of St. Louis, the Federal Reserve Board of Governors, or the Federal Reserve System. All remaining errors are our own.

†Department of Economics. University of North Carolina. Gardener Hall CB# 3305. Chapel Hill, NC 27599. nrfranci@email.unc.edu.

‡Research Division. Federal Reserve Bank of St. Louis. P.O. Box 442, St. Louis, MO 63166-0442. owyang@stls.frb.org.

§Corresponding author. Board of Governors of the Federal Reserve System, Mailstop 74, 20th Street and C St., NW, Washington, DC 20551. jennifer.e.roush@frb.gov.

1 Introduction

By their nature, long-run restricted structural vector autoregressions (LRs) are subject to the criticism that restrictions on infinite-order lag polynomials are ill-suited to samples of realistic proportions [see, for example, Sims (1972); Faust (1996); and Faust and Leeper (1997)]. In finite samples, measures of the VAR moving-average parameters at very long horizons are imprecise; when relied on for identification, this parameter uncertainty translates into potentially spurious inference. Using Monte Carlo methods, Erceg, Guerrieri, and Gust (EGG, 2005) and Chari, Kehoe, and McGrattan (CKM, 2007) assessed the extent of these small-sample estimation problems. These papers simulate repeated small samples from variations of the standard Real Business Cycle (RBC) model and apply the LR identification to obtain hypothetical small-sample distributions of the impulse responses to technology shocks. Both studies conclude that LR impulse responses can be substantially biased, either in sign or in magnitude.

Recently, long-run restrictions have attracted renewed attention, this time as a means for identifying technology shocks in VARs [see, in particular, Galí (1999)]. In these papers, identification is based on the assumption that the unit root in labor productivity arises exclusively from technology shocks. Results from some of these studies have led some to question the notion that technological innovation is the preeminent force behind business cycle fluctuations. Positive technology shocks identified in this manner in U.S. data yield a decline in hours, apparently contradicting the theoretical predictions of a broad class of RBC models. This result has initiated some controversy, with a number of studies offering conflicting evidence based on alternative specifications of the non-productivity component of Galí's empirical model.[1] This paper focuses instead on the identifying assumption regarding the estimated long-run productivity process.

[1] The decline in hours is obtained when hours enter the VAR in first differences (Galí, 1999) or are demographically adjusted (Francis and Ramey, 2008). Christiano, Eichenbaum, and Vigfusson (CEV, 2004) argue that *per capita* labor is bounded and cannot have a unit root, and, if assumed stationary, responds to a technology shock positively on impact. Basu, Fernald, and Kimball (2006) and Shea (1999) use different techniques to identify technology and conclude that the hours response is negative.

We offer an alternative approach to identification with the intent of addressing some of the aforementioned shortcomings associated with LR in small-sample estimation. When applied to technology shocks, our methodology preserves the association between technology and productivity at frequencies beyond typical business cycles. Specifically, we identify the technology shock as that associated with the maximum forecast-error variance share (Max Share) in labor productivity at a long, finite horizon.[2]

The Max Share approach has several potential advantages over the conventional LR approach. First, in Monte Carlo experiments, our approach explains more of the spectrum of productivity at all frequencies than the LR approach. Second, by focusing on a finite horizon, we hope to gain estimation precision over LR, which relies on much longer horizon parameter estimates. Third, our approach allows us to include labor productivity in levels in the VAR, implying the estimated parameters are superconsistent and, therefore, likely to be less biased in small samples. Fourth, in place of the restriction that the unit root in productivity is driven *exclusively* by technology, our approach imposes a weaker restriction that the forecast-error variance in productivity at long horizons is *dominated* by the technology shock. Thus, we essentially allow other shocks to influence labor productivity at all (finite) horizons over which we employ the Max Share algorithm.

The mechanics of this methodology are similar to those introduced by Faust (1998); however, the current application to technology shocks is substantively different in a number of important ways. To our knowledge, we are the first to recognize the suitability of the Max Share approach as a finite-horizon alternative for identification of structural VARs with long-run restrictions. In its original context, Faust identified monetary policy shocks using only the robust predictions of structural VARs identified with short-run restrictions. The two approaches also differ conceptually. Whereas Faust used his objective function as a robustness check of the claim that monetary policy shocks explain only a small portion of

[2] While our application is to technology shocks, the identification can be applied to any case in which a dominant driving process exists. For example, in work following an earlier draft of this paper, Barsky and Sims (2006) adopt our approach to identify information shocks that are orthogonal to contemporaneous output but have permanent effects on future output.

output variability, we use our objective function as a necessary condition for identification.[3] Finally, we take advantage of the methodology as a way to obtain better small-sample estimation properties than in long-run structural vector autoregressions (SVARs), something not considered in either Faust or in related work by Uhlig (2004).

The Max Share approach is also similar in spirit to the medium-term driving forces recently proposed by Uhlig (2004) and Comin and Gertler (2006). However, a fundamental difference between these and the Max Share approach is that the latter allows the data to determine the relative importance of technology at a predetermined horizon instead of specifying its relative importance at the outset. For instance, Uhlig estimates a model in which technology shocks are identified by a process that explains all of the h-step-ahead forecast revision of labor productivity for some fixed $0 < h < \infty$. Our approach, on the other hand, utilizes a maximization routine for horizons up to and including h. We find this more palatable because, in the RBC world, technology explains all of the forecast-error variance, at best, only at $h = \infty$; under the Uhlig assumption the spectrum may be radically shifted in ways that potentially violate the underlying RBC assumption.

Using data simulated from an off-the-shelf RBC model, we find that the Max Share approach exhibits less bias (measured by the deviation between the median response and the theoretical response) and less uncertainty (measured by the width of the 68 percent error bands) than the LR approach. Moreover, these advantages are found to be robust to alternative specifications of the theoretical technology and non-technology shocks. Results using the Max Share approach are consistent with Galí's original finding that hours decline after a technology shock. However, relaxing the Galí assumption by allowing non-technology shocks to have nontrivial effects on labor productivity at sufficiently long horizons can qualitatively alter this short-run hours response.

In the next section, we present the Max Share identification approach. We then compare

[3] Faust makes no presumption – theoretical or otherwise – that the identifying restriction imposed by the optimization criterion necessarily holds in the data-generating process. In contrast, the objective function in our approach serves a fundamental role as a substitute for the restriction that the long-run variance of labor productivity is primarily driven by technology shocks.

the small-sample performances of the Max Share and LR approaches using Monte Carlo methods and a standard RBC model. In the remainder of the paper, we apply the Max Share approach to postwar U.S. data and examine the robustness of the LR findings to our relaxation of the original identifying assumption. Finally, we incorporate the additional restriction that hours respond positively to a technology shock to examine whether this causes a significant shift in the associated share of the maximum forecast-error variance.

2 Identification

The assumption identifying technology shocks using the LR approach arises from a class of RBC models in which *log* labor productivity, x_t, can be decomposed into two orthogonal components, technology z_t and non-technology η_t, in the following manner:

$$x_t = z_t + \eta_t. \tag{1}$$

Since all processes except technology are assumed stationary, the unit root in productivity must arise from z_t.[4] Note that equation (1) resembles the Beveridge-Nelson decomposition, where η_t can be thought of as a composite non-technology shock including fiscal, monetary, and tax shocks. This condition provides the foundation for both the standard LR identification and our finite-horizon Max Share identification.

2.1 The LR Identification

Assume that the data-generating process can be approximated by the following linear model:

$$A(L)y_t = \varepsilon_t,$$

[4] This arises from the steady-state condition $X = W = \alpha Z \left(\frac{k}{N}\right)^{1-\alpha}$ relating labor productivity X to wages W, where k is the ratio of capital to technology, N is labor, α is the marginal productivity of labor, and Z is the level of technology.

where $A(L) = \sum_{i=0}^{p} A_i L^i$ is a matrix polynomial in the lag operator, L; ε_t is a structural innovation; $E(\varepsilon_t \varepsilon_t') = I$; and y_t is an $n \times 1$ vector of period-t macroeconomic variables with labor productivity ordered first and entered in *differences*.

To estimate this model, we begin with the reduced-form VAR:

$$B(L) y_t = \mu_t, \tag{2}$$

where $B(L) = \sum_{i=0}^{p} B_i L^i$, $B_0 = I$, and $E(\mu_t \mu_t') = V$. The goal is to find a rotation of the moving-average representation of the VAR:

$$y_t = C(L) A_0^{-1} A_0 \mu_t,$$

which identifies the i.i.d. structural shocks of the model:

$$\varepsilon_t = A_0 \mu_t,$$

where $C(L) = B(L)^{-1}$ and A_0 is the contemporaneous structural parameter matrix. Identification in the conventional long-run SVAR framework is accomplished by imposing restrictions on the effect of the jth shock on the ith variable at an infinite horizon, given by $\left[C(1)A_0^{-1}\right]_{i,j}$, where neutrality implies the restriction $\left[C(1)A_0^{-1}\right]_{i,j} = 0$ for some j. The preceding argument reflects the key identifying assumption in Galí (1999) that the technology shocks are the only influence on long-run labor productivity and is formalized below.

Assumption A.1 The unit root in productivity is solely attributable to the technology shock z. That is,

$$\left[C(1)A_0^{-1}\right]_{i=1, j \neq i} = 0, \tag{3}$$

where $i = 1$ represents labor productivity growth ordered first and $j \neq i$ indicates all non-technology shocks. Equation (3) is isomorphic to the assumption that the zero-frequency spectrum of labor productivity growth is attributable entirely to technology.

2.2 Finite-Horizon Max Share Identification

As in Galí (1999), our objective is to isolate technology shocks by characterizing their effect on productivity at long horizons. However, instead of imposing the unit root in labor productivity by estimating a VAR with the growth rate of productivity, we identify the technology shock by maximizing the forecast-error variance share of productivity at long, finite horizons. This approach follows a suggestion in Uhlig (2004) and is adapted from methods introduced in Faust (1998). We begin by introducing the methodology and then discuss its suitability and practicality.

2.2.1 Methodology

In the Max Share identification, all variables including labor productivity enter the VAR in *log-levels*.[5] The method is operationalized by first expressing the h-step-ahead forecast error for y as a function of realized reduced-form errors:

$$y_{t+h} - \widehat{y}_{t+h} = \sum_{\tau=0}^{h-1} C_\tau \mu_{t+h-\tau}, \tag{4}$$

where $\widehat{y}_{t+h}$ is the h-step-ahead forecast of y conditional on time-t information. Next, we define an arbitrary orthonormal matrix D, which obtains an alternative linear representation of the reduced-form model:

$$y_{t+h} - \widehat{y}_{t+h} = \sum_{\tau=0}^{h-1} C_\tau D D' \mu_{t+h-\tau}.$$

[5] Labor productivity must be entered in differences in LR. In principle, the Max Share identification can be used on systems in which labor productivity enters either in levels or in differences. Sims, Stock, and Watson (1990) advocate estimating systems in levels when the true order of integration is unknown. We estimate the VAR in levels under the assumption that technology may *not* have a unit root. As an alternative, strong beliefs can be incorporated into a prior.

Then the h-step-ahead forecast-error variance share for a particular variable i attributable to a particular shock j in this new representation is

$$\omega_{ij}\left(\alpha(h)\right) = \frac{e_i'\left[\sum_{\tau=0}^{h-1} C_\tau D e_j e_j' D' C_\tau'\right] e_i}{e_i'\left[\sum_{\tau=0}^{h-1} C_\tau \Omega_\mu C_\tau'\right] e_i} = \frac{e_i'\left[\sum_{\tau=0}^{h-1} C_\tau \alpha\alpha' C_\tau'\right] e_i}{e_i'\left[\sum_{\tau=0}^{h-1} C_\tau \Omega_\mu C_\tau'\right] e_i}, \tag{5}$$

where e_i is an $n \times 1$ indicator vector that picks out the impulse vector α, the ith column vector of D.

The technology shock is identified by solving the following maximization problem over all possible α, for a given value of h:

$$\max_{\alpha} \omega_{1j}\left(\alpha(h)\right), \tag{6}$$

with the additional normalization $\alpha'\alpha = 1$. In application, we ensure that α belongs to an orthonormal matrix by performing the optimization problem on an othogonalized impulse-response-generating matrix, $\widetilde{C}_\tau = C_\tau H$, where H is obtained by a Cholesky decomposition of Ω_μ. Thus, the identified technology shock, $\varepsilon^{tech} = \alpha' H^{-1}\mu_\tau$, is orthogonal to other shocks in the system. The restriction that α has unit length ensures that the technology shocks have unit variance. The horizon h at which the forecast-error variance for labor productivity is maximized is chosen exogenously. While h is initially fixed at 10 years, we later consider the effect of varying h from 1 to 20 years. As in Faust (1998), sign and shape restrictions on the impulse response of variable i at horizon(s) h can also be incorporated by solving the optimization problem (6) with an additional constraint of the form $e_i' C_t \alpha \geq 0$. In later sections, we will consider such sign restrictions.

2.2.2 Suitability

This subsection considers the appropriateness of the Max Share approach as an empirical identification of technology shocks. In particular, we ask whether a VAR specified in levels of productivity can be used to identify the stochastic technology process. As justification, we appeal to the Beveridge-Nelson decomposition and standard asymptotic theory of non-stationary processes.

When productivity x_t is $I(1)$, it can be decomposed into the sum of a trend component x_t^T and a stationary component x_t^S:

$$x_t = x_t^T + x_t^S,$$

where $x_t^T = E(x_\infty | x_t, x_{t-1}, ...)$. By definition, permanent shocks affect only the trend component, which can be rewritten as

$$x_\infty = x_t + \sum_{i=1}^{\infty} \Delta x_{t+i}.$$

Since the first difference Δx_t is stationary, it can be expressed in its fundamental form as $\psi(L)\varepsilon_t$, or, alternately,

$$E(\Delta x_{t+i} | H_t(\Delta x)) = \sum_{k=0}^{\infty} \psi_{k+i} \varepsilon_{t-k},$$

where $H_t(\Delta x)$ is the information set containing past Δx. Substitution yields

$$x_t^T = x_t + \left(\sum_{i=1}^{\infty} \psi_i\right) \varepsilon_t + \left(\sum_{i=2}^{\infty} \psi_i\right) \varepsilon_{t-1} + \left(\sum_{i=3}^{\infty} \psi_i\right) \varepsilon_{t-2} + ... \tag{7}$$

Equation (7) is essentially the Beveridge-Nelson decomposition, which redefines x_t as the sum of a random walk (x_t^T) and a stationary component [see proposition 17.2, page 504 of Hamilton (1994)]. As t increases, only the random walk component matters for the distribution of x_t. Thus, under assumption A.1 above, technology dominates the forecast-error variance of the log-level of productivity at suitably long horizons, h. This is the foundation for the identifying assumption imposed by the Max Share approach.[6]

[6] Recall that the variance of a unit root process, $var(x_t^T) = t\sigma_\varepsilon$, increases with t. Hence, the variance of x_t is dominated by its trend component for large t. Similarly, when x_t is included in a VAR, the forecast-error variance of x_t grows unbounded as the forecast horizon h increases. At sufficiently long horizons, the forecast-error variance is dominated by the non-stationary component [Lütkepohl (1993), p. 377].

2.2.3 Practical Considerations

The LR identification imposes infinite-horizon restrictions on parameters that, in small samples, may be imprecisely estimated. Thus, LR often performs poorly on simulated data for sample sizes comparable to the U.S. postwar sample (see EGG and CKM). The Max Share approach, on the other hand, may demonstrate less small-sample bias, in part because it relies on restrictions at a finite horizon.[7] The intuition behind this conjecture is adapted from arguments made by Sims (1972) and, more recently, Faust and Leeper (1997) and CEV (2007) regarding the differences between short- and long-run restrictions. Similar arguments have been made in favor of medium-run restrictions [e.g., Uhlig (2004); Khan and Tsoukalas (2005)].[8]

The root of all of the impulse response bias – regardless of the identification methodology – lies in the misspecification of the estimated VAR relative to the underlying data-generating process. In simulated data exercises, the VAR representation of the RBC model calls for a VAR(∞); however, in small samples, the econometrician must truncate the VAR to finite order q. In data, the econometrician uses some criteria to determine the lag order, which may or may not be the appropriate specification. In these cases, the model may be misspecified, leading to a misidentification of the shocks and bias in the implied responses. The degree of bias depends on the manner of identification.

For example, LR places a restriction on $C(1)$, the infinite sum of the transformed VAR coefficients. The econometrician, however, does not use the true $C(1)$; instead, the estimated $\widehat{C}(1)$, which is known to be misspecified, is used and produces a bias in the identification of the shock.[9] CEV (2007) argue that using only short-run restrictions minimizes the effect of

[7] Because the Max Share identification is the product of a maximization problem, the nature of the identification changes with the subsample, which prevents a formal proof of bias reduction in small samples. Unlike in LR, the rotation is not preserved as the sample size increases (Mittnik and Zadrozny, 1993).

[8] Dupor and Kiefer (2007) take a different approach with a similar flavor. They place restrictions on the finite-horizon effects estimated via local projections (Jorda, 2005).

[9] An informal proof of this conjecture can be found in CEV (2007). Moreover, CEV and others argue that identification using very-long-horizon restrictions can be problematic because the spectral mass near the zero frequency can be small.

potential misspecification because short-run restrictions do not require $\widehat{C}(1)$ for identification. However, short-run restrictions are not necessarily implied by the theory as there is no link between the RBC notion of technology and the short-run identifying restrictions. The goal, then, is to resolve the tension between the theoretical consistency of LR while reducing the effect of potential misspecification bias in the identification caused by truncating the VAR. Max Share, by comparison, imposes restrictions on $\sum_{\tau=0}^{h-1} C_\tau \alpha \alpha' C_\tau'$. Shortening the horizon reduces the potential misspecification bias but preserves the theoretical interpretation of the identification.

The preceding discussion suggests that errors in estimating $C_\tau A_0^{-1}$ may influence the identification, making it possible to attribute too much of the forecast-error variance in productivity to technology. While both LR and Max Share may, to some extent, suffer from this problem, the risks in the LR approach are symmetric around zero. In the Max Share approach, however, the maximization algorithm shrinks the risk of underestimating relative to that of overestimating. Given this trade-off, we must determine whether the Max Share identification yields a net advantage in small samples. In the next section, we measure the net effect of employing the Max Share identification by comparing the small-sample performance of the LR and Max Share identifications in Monte Carlo experiments against a known data-generating process.

3 Monte Carlo Experiments

This section outlines the Monte Carlo methods used to gauge the ability of the Max Share and LR approaches to identify technology shocks in small samples. We calibrate a neoclassical growth model in which equation (1) holds to obtain theoretical impulse responses to a technology shock. The model details and a proof of the existence of its VAR representation appear in the appendix. Theoretical impulse responses and simulated data are generated for a variety of parameterizations of the model.[10] Table 1 presents the sets of parameter values used to simulate the model. Our parameterizations are similar to those used by EGG (2005) in their benchmark model without capital utilization to match moments in U.S. data.

We use this model to simulate 174 observations of each variable, equal to the size of our postwar quarterly dataset.[11] We then estimate a reduced-form VAR and obtain the structural impulse responses by solving the maximization problem,

$$\max_{\alpha} \alpha' V \alpha = \max_{\alpha} \alpha' \frac{e_i' \left[\sum_{\tau=0}^{h-1} \widetilde{C}_\tau \widetilde{C}_\tau' \right] e_i}{e_i' \left[\sum_{\tau=0}^{h-1} \widetilde{C}_\tau \Omega_\mu \widetilde{C}_\tau' \right] e_i} \alpha, \tag{8}$$

adapted from equation (6).[12] The posterior distributions for the impulse responses are simulated from 1,000 draws utilizing an uninformative prior [see Sims and Zha (1999)]. To account for potential asymmetries in the impulse responses, we retain the median, 16th, and 84th percentiles for each response. This process is repeated to obtain 1,000 median estimates and error bands, each corresponding to what an econometrician would estimate given a single set of data. The average of these 1,000 median statistics and error bands, which can be interpreted as the expected value of the econometrician's estimates, is reported and compared with their theoretical counterparts. We also analyze the correlation between the theoretical shocks and those estimated by both LR and Max Share.

3.1 Benchmark Results

The benchmark empirical model is a VAR(4) with the logs of labor productivity, hours, the consumption-output ratio, and the investment-output ratio. Each reduced-form VAR is

[10] The model produces data that are a deviation around the steady-state growth path. To facilitate comparison with existing empirical work, this transformation must be reversed. This avoids over-differencing productivity using the transformed data.

[11] An important issue is whether the estimated impulse responses for Max Share exhibit a reduction in bias from the theoretical responses as the sample size increases. *In results not discussed here, we found that increasing the sample size reduced the bias between the estimated and theoretical responses.* These results are available upon request from the corresponding author.

[12] As described in the appendix in Faust (1998), the maximization problem is solved by α^*, the eigenvector associated with the maximum eigenvalue of V, where $\widetilde{C}_\tau = C_\tau H$, and H is the Cholesky decomposition of Ω_μ.

estimated using conditional maximum likelihood. Log productivity enters the VAR in first differences for the LR but enters in levels for the Max Share approach. In the benchmark specification, the AR(1) coefficient of technology (ρ_z) is set to one and all other stochastic processes have AR(1) coefficients of 0.6. We assume a maximization horizon h for the Max Share approach of 10 years (i.e., the technology shock is chosen to be that which maximizes the forecast-error variance share at a horizon of 10 years) before considering alternative horizons below.

Figure 1 presents the impulse responses to a 100-basis-point shock to technology. The thick solid line depicts the theoretical impulse responses. The Max Share median responses are shown by the thick dashed lines, with the shaded areas representing the accompanying 68 percent error bands. The LR median responses and their error bands are shown by dotted lines. In the theoretical model, a positive technology shock leads to an immediate increase in labor productivity, hours, and the investment-output ratio and a decline in the consumption-output ratio.

Our results corroborate EGG's findings that LR biases the median responses but preserves their qualitative nature. The Max Share impulse responses match the theoretical impulse responses qualitatively and display less bias than the LR responses. For each variable, the Max Share approach yields an impulse response biased toward zero.[13] Although the theoretical impulse responses are near the upper tail of the 68 percent probability intervals for both methods, the Max Share responses are considerably closer to the theoretical responses for the first two years following the shock. The probability intervals from the Max Share model are also narrower than their LR counterparts over this horizon.

In addition to the responses to the shocks, we can compare the time series of identified shocks to the time series from the generated data. Table 2 presents the correlations between the model-generated and the estimated technology shocks from both identifications. Our

[13] This bias toward zero may, in part, be caused by our choice to estimate the VAR with productivity in levels. We explore this possibility below. The bias may also result from Max Share not *imposing* a unit root response of productivity to a technology shock.

benchmark results are in the first row of the table, where a higher correlation suggests a more accurate identification of the time series of shocks. The median correlation for the Max Share shocks is greater (about 0.9) than that for LR shocks (about 0.6). Additionally, the median correlation for the LR model lies in the far left tail of the distribution for the Max Share correlations.

Figure 2 revisits the cause for the small-sample bias reduction that Max Share enjoys by examining the spectral decomposition of productivity attributed to technology. The figure plots the share of productivity's variance explained by the Max Share approach divided by that for LR at various frequencies. At low frequencies, the share of productivity's forecast-error variance explained by the Max Share technology shock is greater than that for the LR shock. At these frequencies, one expects the explanatory power of the technology shock to be substantial. Taken together with the preceding results on bias and correlation of the theoretical and estimated technology shocks, the higher explanatory power of the Max Share shock suggests the technology shock is better identified by the finite-horizon alternative.

3.2 Increasing the Importance of Non-Technology Shocks

To this point, we have assumed that technology is the only process influencing the long-horizon behavior of labor productivity. Some recent studies [e.g., Francis and Ramey (2005) and Uhlig (2004)] argue that other shocks – capital tax shocks, for example – may contribute to the variance of long-run labor productivity. In this section, we consider the case where the non-technology shocks are allowed to play a greater role in determining labor productivity at long horizons. Specifically, the non-technology stochastic processes — namely, government spending, capital, and/or labor taxes — are assumed to be highly persistent, with their innovation variances set equal to the variance of technology. In this case, technology remains the source of the unit root in productivity, preserving equation (3). However, increasing the persistence and variances of the non-technology processes implies that other factors may have greater influence on labor productivity at horizons beyond the business cycle. This can be a source of possible contamination, making it more difficult for the Max Share approach to isolate the technology process. On the other hand, the LR restriction is still valid at the infinite horizon, potentially giving LR an advantage over Max Share, all else equal.

Figure 3 shows the response for this parameterization when technology has a unit root, all non-technology processes have AR(1) coefficients of 0.98, and all stochastic processes have equal variances. Again, the Max Share impulse responses demonstrate less bias than the LR and have narrower error bands. An interesting finding is that the LR error bands for labor hours now include zero, making the sign on the hours response statistically indeterminate. Moreover, the second row of Table 2 shows that the Max Share-identified shocks are, on average, more closely correlated with the model-generated shocks. Therefore, even in the presence of more influential, potentially contaminating non-technology components, the Max Share identification still outperforms the conventional identification approach.

3.3 Stationary Technology

Although the RBC theory assumes that technology shocks exclusively drive the unit root in labor productivity, there are reasons to doubt this assumption. Unit root tests are understood to have low power in finite samples [Blough (1992); Cochrane (1991); Christiano and Eichenbaum (1990); Stock (1990); and Sims (1989)]. In these cases, pretesting might falsely accept the presence of a unit root, leading the econometrician to misspecify labor productivity in differences. A stationary but highly persistent technology process violates equation (3), rendering the LR identification invalid. Perhaps for this reason, several researchers – including Comin and Gertler (2006) and Uhlig (2004) – identify business cycle drivers at medium-term frequencies, abstracting from any unit root assumption.

The Max Share identification, on the other hand, may still be valid even when labor productivity is stationary. In this case, we can no longer appeal to the Beveridge-Nelson decomposition; but, if technology shocks play a dominant role at the horizon of interest, Max Share remains applicable. Uhlig (2004) uses a theoretical model and finds that, under a number of parameterizations, technology does dominate the forecast-error variance of labor productivity at horizons of three to ten years.

To this end, we rerun the Monte Carlo experiments from above, reparameterizing the persistence of the shocks. Figure 4 depicts the case in which the AR(1) coefficients on the technology and non-technology components are 0.98 and 0.6, respectively (Case 1).[14] Here, the bias from the LR identification may be due to both over-differencing productivity and small-sample error. Figure 4 shows that the Max Share approach still outperforms the LR. In particular, the median Max Share responses for all variables are less biased than the LR responses during the first year. Additionally, the 68 percent probability intervals for Max Share are much narrower than similar bands for LR.

Figure 5 examines the case when both approaches are potentially misspecified: The AR(1) coefficients on all stochastic processes are 0.98 (Case 2). Max Share may also be contaminated by persistent non-technology shocks. Even in this case, however, the Max Share identification outperforms the conventional LR identification. Thus, even when the Max Share model is potentially misspecified, it exhibits less bias than the LR approach.

The final two rows of Table 2 present the correlations with the theoretical shocks under the stationary technology cases. Under both parameterizations, the Max Share shocks are more highly correlated – up to 1.5 times – with the model shocks than those identified by LR. These correlations are similar in magnitude to those found in the unit root cases, suggesting that the misspecification is not too costly in terms of the size of the bias and/or the correlation of the shocks.

3.4 Alternative Specification and Bias Analysis

In the previous subsections, we argued that Max Share outperforms the LR approach, in part, because it avoids misspecification error associated with the low power of unit root tests when technology is highly persistent but stationary. However, one might wonder how Max

[14] For this and all other parameterizations that follow, the variances of the innovations correspond to their values in the benchmark model.

Share performs when the VAR is specified with productivity in differences and technology indeed has a unit root. When technology is known to yield a unit root in productivity, the differencing resolves the coefficient bias that may be introduced when the VAR is estimated in levels. Figure 6 plots the impulse responses for both approaches when productivity enters the VAR in differences. Again, the impulse responses from the Max Share approach exhibit less bias and less uncertainty than the LR approach.[15]

All of the previous results have been obtained with an exogenously chosen Max Share horizon, h, of 10 years. Here, we evaluate the robustness of the Max Share algorithm to variation in the forecast horizon. Figure 7 summarizes the bias properties of the Max Share approach when h varies from 5 to 80 quarters for the benchmark parameterization. The three lines show the average absolute deviation of the median response from the theoretical model response over the first four quarters, expressed as a percentage of the true model response. The dotted and dash-dotted lines depict the bias of the Max Share approach when productivity enters the VAR in differences and levels, respectively. The solid line shows a similar measure for LR, which assumes a long-run horizon of 80 quarters throughout.

The bias shown by Max Share is clearly smaller than that of the LR for all variables and all horizons. As expected, the Max Share bias increases with h but at a surprisingly slow rate. The productivity bias for Max Share(80) is approximately 16 percent, well below that for LR (31 percent) but only slightly higher than for Max Share(5) (13 percent). In this model environment, the bias improvement of the Max Share approach is robust to the choice of the maximization horizon.

Figure 7 also confirms that, in practice, the advantage of utilizing the differences specification in the Max Share approach even when technology is a unit root may be small. For this parameterization of the theoretical model, the bias in the impulse responses from the difference specification is indeed smaller than that for levels, as the former reflects an

[15] This exercise demonstrates the flexibility of the Max Share approach. In cases in which the order of integration is unknown, the model can be estimated in levels. When the series of interest (in this case, productivity) is known to have a unit root, estimation can be performed in differences. In both cases, the identification produces less bias than the LR approach.

ex ante knowledge of the true order of integration. However, this exercise shows that the improvement associated with *imposing* the unit root by differencing is not large (even when true) and may not outweigh the potential increase in bias in cases in which the order of integration is unknown.

4 Max Share Identification in the Data

Having evaluated the small-sample performance of our identification scheme through Monte Carlo experiments, we turn to the data and estimate a four-variable, four-lag VAR. The data are quarterly series from 1959:1 to 2002:4 for private business productivity, private business hours, real consumption as a share of output, and real investment as a share of output.[16] All variables enter in log levels. Raw data are taken from the Bureau of Economic Analysis and the Bureau of Labor Statistics. As in the Monte Carlo section above, the error bands are computed for 68 percent coverage using methods detailed in Sims and Zha (1999).

4.1 Baseline Results

The dashed lines in the left column of Figure 8 present the median impulse response to a one-standard-deviation technology shock, with the shaded areas representing the 68 percent probability intervals.[17] The thick solid lines in these figures represent sign-restricted responses discussed below. In response to a positive technology shock, both consumption and investment increase. Labor hours fall for the first few quarters and eventually rise above zero. Table 3 compares the forecast-error variance shares for output and hours attributable to technology for both Max Share and LR. While the share of output variance is large at

[16] As some variation exists in the data, we estimated versions of the VAR in which consumption is composed of nondurables, services, and government spending; investment is composed of private investment plus durables; and hours and productivity measures are adjusted for demographic components (Francis and Ramey, 2008). The impulse responses were qualitatively similar across these models and results are available upon request.

[17] All of the Max Share results shown are based on $h = 10$; similar results were obtained for horizons of 5, 15, and 25 years. Although the width of the error bands for the Max Share identification increases with h, they are always narrower than those obtained from LR.

most horizons under both identifications, technology typically explains only a minority share of the variance in hours. This result is consistent with CEV (2004) and suggests that technology is not an important driver of the positive correlation in output and hours at business cycle frequencies.

The right column of Figure 8 displays the impulse responses to a technology shock identified by LR with the associated 68 percent error bands. The median predictions for all of the variables except hours are similar for both the LR and Max Share identifications. The *median* response of hours on impact is positive under LR but not statistically distinguishable from zero. The hours response to the Max Share shock, on the other hand, is significantly negative on impact. In addition, the error bands associated with the Max Share are everywhere narrower.

As we have previously noted, the identifying restriction imposed by Max Share depends on the horizon for which the forecast-error variance share is maximized. To this end, we assess the Max Share impulse responses' sensitivity to the forecast horizon by varying h between 10 and 25 years and find them to be qualitatively similar across horizons. As might be expected, the width of the error bands for each response grows as the optimization horizon increases. This result points to the difficulty that longer-horizon restrictions yield more uncertainty in their corresponding short-run predictions.

4.2 Incorporating Sign Restrictions

A salient advantage of the Max Share approach is that it allows for exhaustive robustness analysis across a broad class of models. While our finding that hours respond negatively to a positive technology shock seems to corroborate Galí's original result, many papers have questioned the robustness of this prediction.[18] In particular, it is not hard to imagine that a small modification to the identifying assumption – perhaps corresponding to a more

[18] A number of papers [including Galí and Rabanal (2005) and CEV (2004)] have focused on the stationarity of hours as the key determinant of the sign of the hours response to a technology shock. We address this issue in the next subsection.

accommodative monetary policy or a greater influence of non-technology factors – could yield a different qualitative prediction for hours. Taking all such possibilities into account, a more complete robustness test asks: *(1) Is it possible to identify a technology shock that yields a positive response in hours? and (2) If so, what are the features of such a shock?*

To this end, we reestimate the Max Share model with the additional restriction that hours respond positively on impact to a positive technology shock.[19] We then examine whether this has a discernible effect on the forecast-error variance share attributable to technology. The left column in Figure 8 shows the median impulse responses to a technology shock when the additional sign restriction is imposed.[20] The addition of the short-run sign restriction preserves the shape but shifts the point estimates of the responses for each variable. In particular, the median hours response shifts upward above the 68 percent probability interval of the unrestricted hours model for the first five or so quarters. The positive restriction on hours also raises the short-term consumption and investment responses. The net result of restricting technology to raise hours is apparently to amplify technology's effect on short-run output. But how effective is this "restricted" shock at explaining cyclical fluctuations?

Table 4 displays the maximum forecast-error variance share values estimated at horizons of 10 and 25 years. The last column presents the share from the model estimated with the sign restriction on hours. The share of labor productivity fluctuations explained by technology declines when hours are restricted to respond positively to a technology shock. Moreover, the maximum attainable forecast-error variance share in the restricted model is always less than the 16th quantile of the unrestricted model. Thus, a positive hours response is attainable, but only when the importance of technology shocks significantly diminishes at these horizons.[21] This suggests that a positive hours response may also result when non-

[19]Dedola and Neri (2007) also use an agnostic approach to identify technology shocks using only sign restrictions on the impulse responses. Their results, however, are not directly comparable to ours as they make no restrictions on technology's contribution to productivity fluctuations.

[20]Section 2.2.1 above explains how sign restrictions are implemented in Max Share. The additional restriction on hours is not an overidentifying restriction suitable to a likelihood ratio test. The Max Share identifying assumption (8) is sufficient to identify technology, but the system as a whole is underidentified with or without the additional sign restriction.

[21]This does not necessarily imply that non-technology shocks influence long-run productivity. The odds

technology shocks are influential to labor productivity at an infinite horizon. In other words, the exclusivity assumption in the LR model may be playing an important role in obtaining the negative hours prediction. These results indicate that models identified by restrictions made at long horizons may contain only limited information about short-run movements in hours.[22]

4.3 The Stationarity of Hours

Figure 8 corroborates others' finding that hours respond positively to an LR-identified technology shock when entered into the VAR in levels [e.g., CEV (2004)]. In this subsection, we consider whether this result is sensitive to the inclusion of an informative prior that accounts for potential non-stationarity and cointegration in hours and other series. Other papers focus on varying the hours specification between levels and differences; however, differencing alone does not allow for cointegration between hours and other variables in the system. By estimating the VAR with the Sims-Zha (1998) prior, we may address the robustness question without imposing either stationarity or non-stationarity ex ante. The Sims-Zha prior is a shrinkage prior that places a probability on the presence of unit roots and cointegration in the VAR. This feature may be important for this application (and especially to the identification of the shock with LR) if technology and hours are non-stationary or share a common trend with other variables in the system.

Figure 9 shows the responses to technology shocks identified by both the LR and Max Share approaches using the Sims-Zha prior. In the LR model with the Sims-Zha prior, hours respond negatively on impact even though they enter in *levels*, suggesting the short-term response of hours may be contaminated either by its low-frequency movements or by its co-movement with productivity.[23] The conclusion we draw from this result is that

ratios for the restricted and unrestricted models were indistinguishable, preventing any conclusions based on how well the model fits the data.

[22] Similar differences are attainable from alternative calibrations of the theoretical model presented in the Monte Carlo section of this paper.

[23] In results not shown, a negative hours response was also obtained when using the demographically

specification of non-productivity variables in the VAR is important for the LR identification. Max Share, on the other hand, appears robust to these specification changes, producing a negative hours response regardless of how hours are specified or whether an informative prior is included.

5 Conclusion

We propose an alternative method for identifying shocks in VARs in which long-run restrictions have been ordinarily used. This methodology has the advantage of being robust to relaxing key assumptions about the data-generating process while maintaining the spirit of long-run restrictions. When applied to technology, the shock is identified as that which yields the maximum forecast-error variance share of productivity at some predetermined, yet finite, horizon.

Applied to artificial small samples generated from an off-the-shelf RBC model, the Max Share identification outperforms the standard LR identification. In particular, our identification reduces the bias of estimated impulse responses relative to theoretical responses. In addition, the Max Share impulse responses are more precisely estimated than those identified using the identifying restrictions proposed by Galí (1999). We also find that the Max Share technology shocks are more highly correlated with the theoretical shocks than those identified by LR. These results reveal a clear improvement over the LR estimates in small samples.

For U.S. postwar data, the Max Share model predicts a negative short-run response in hours, confirming the original LR finding of Galí (1999) and others. However, a positive hours response is attainable if a greater role for non-technology shocks is allowed. Unfortunately, neither model can be rejected based on posterior odds. Nevertheless, our results suggest that the rejection of the RBC framework on the basis of the qualitative response

adjusted hours series, with or without the Sims-Zha prior. These results are available from the corresponding author upon request.

in hours depends critically on the assumption that technology has exclusive influence on long-run productivity.

When we view our model and the infinite-horizon models as a class, our results can be interpreted as demonstrating the limitations of long- (or infinite-) horizon restrictions in predicting short-run movements in hours. A modest, empirically reasonable adjustment to the assumption regarding the long-run importance of non-technology factors yields different predictions for the direction of the hours response. In light of these findings, we advocate a more flexible identification environment such as the one proposed here.

Table 1		
Parameter Values Used in Model Simulation		
Parameter	Description	Value
α	capital share	0.36
δ	quarterly depreciation rate	0.02
β	discount factor	1/1.03
Φ	preference parameter	1
ρ_z	autocorrelation of technology shock	1(0.98)
ρ_k	autocorrelation of capital tax shock	0.6(0.98)
ρ_n	autocorrelation of labor tax shock	0.6(0.98)
ρ_g	autocorrelation of government spending shock	0.6(0.98)
$\overline{g}/\overline{y}$	steady-state ratio of government to output	0.03
$\overline{n}$	steady-state labor	1/3
$\overline{\varsigma}_k$	steady-state capital tax rate	0.38
$\overline{\varsigma}_n$	steady-state labor tax rate	0.22
σ_z	technology shock standard deviation	0.0148
σ_{τ_k}	capital tax shock standard deviation	0.008(0.0148)
σ_{τ_n}	labor tax shock standard deviation	0.052(0.0148)
σ_g	government spending shock standard deviation	0.016(0.0148)

Notes: The parameters are for the RBC model augmented with preference shocks, capital, and labor income taxes outlined in Appendix 1. The setup abstracts from international markets (imports and exports) and variable capital utilization. Numbers outside parentheses are for the benchmark model, and numbers inside parentheses are for alternative model simulations in sections 3.3 and 3.4.

Table 2				
Correlation Between Estimated and Model Technology Shocks				
Parameterization		$16th$ percentile	median	$84th$ percentile
$\rho_z = 1; \rho_\eta = 0.6$	LR	0.29	0.61	0.82
	Max Share	0.80	0.89	0.93
$\rho_z = 1; \rho_\eta = 0.98$	LR	0.24	0.54	0.76
	Max Share	0.64	0.82	0.93
$\rho_z = 0.98; \rho_\eta = 0.6$	LR	0.22	0.57	0.79
	Max Share	0.79	0.88	0.94
$\rho_z = 0.98; \rho_\eta = 0.98$	LR	0.23	0.54	0.76
	Max Share	0.64	0.82	0.92

Notes: We generate artificial data from the RBC model using the AR(1) values for technology and non-technology shocks in column 1. For each parameterization, we employ our empirical identification techniques (LR and Max Share) and back out implied technology shocks for each of 1,000 draws of the simulated data. The correlations between the estimated shocks and the artificial shocks are then calculated for each of the 1,000 draws. The median, 16th, and 84th percentiles from the posterior distributions are used for the correlations for each artificial sample as described in section 3.

Table 3				
Contribution of Technology to Forecast-Error Variance				
	Output		Hours	
Horizon	Max Share	LR	Max Share	LR
1 year	0.90	0.73	0.03	0.34
5 years	0.95	0.86	0.21	0.51
10 years	0.96	0.93	0.29	0.51
20 years	0.95	0.52	0.42	0.52

Notes: Table 3 shows the forecast-error variance share of output and hours due to each method's identified technology shock at several horizons. The numbers for Max Share are calculated with maximization at an horizon of 10 years.

Table 4		
Comparing Shares of Productivity with and without Sign Restrictions		
	Hours Unrestricted	$\text{Hours}_0 \geq 0$
$h = 10$	.951	.876**
$h = 25$	.938	.914*

Notes: Table 4 shows the forecast-error variance share attributed to technology shocks identified by Max Share, estimated with an uninformative prior using a maximization horizon of $h = 10$ and 25 years. The two columns reflect the effect of identification when the hours response is unrestricted and when it is restricted to be non-negative on impact. **Outside 95 percent interval of unrestricted model; *Outside 85 percent interval of the unrestricted model.

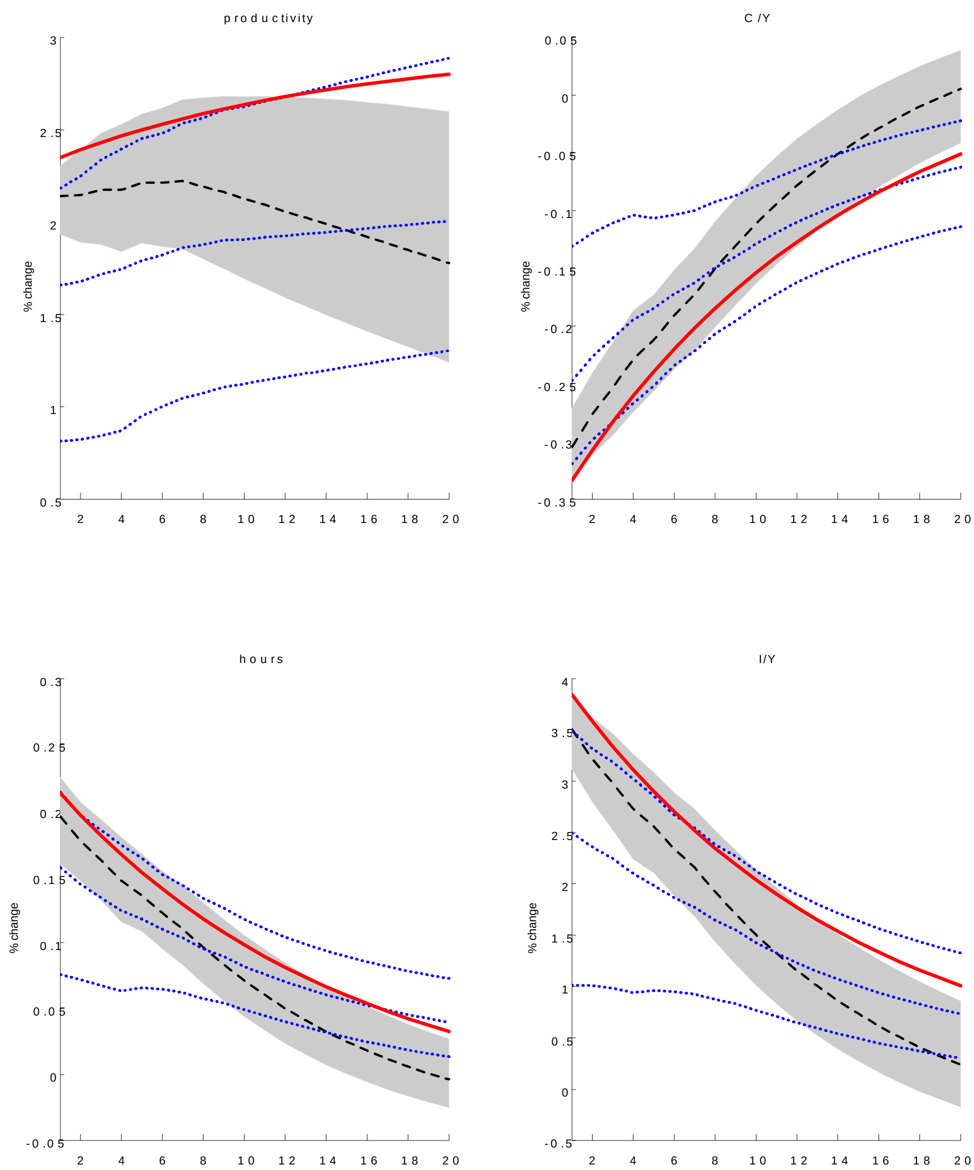

Figure 1: Impulse responses to a technology shock in the benchmark model. Theoretical responses [with AR(1) technology coefficient $\rho_z = 1.0$ and non-technology AR(1) coefficients $\rho = 0.6$] are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability intervals are shown by dotted lines. Median estimates and error bands are averages across 1,000 estimates, each representing what an econometrician would estimate based on a sample with 174 observations and 1,000 draws from the posterior distributions for the impulse responses.

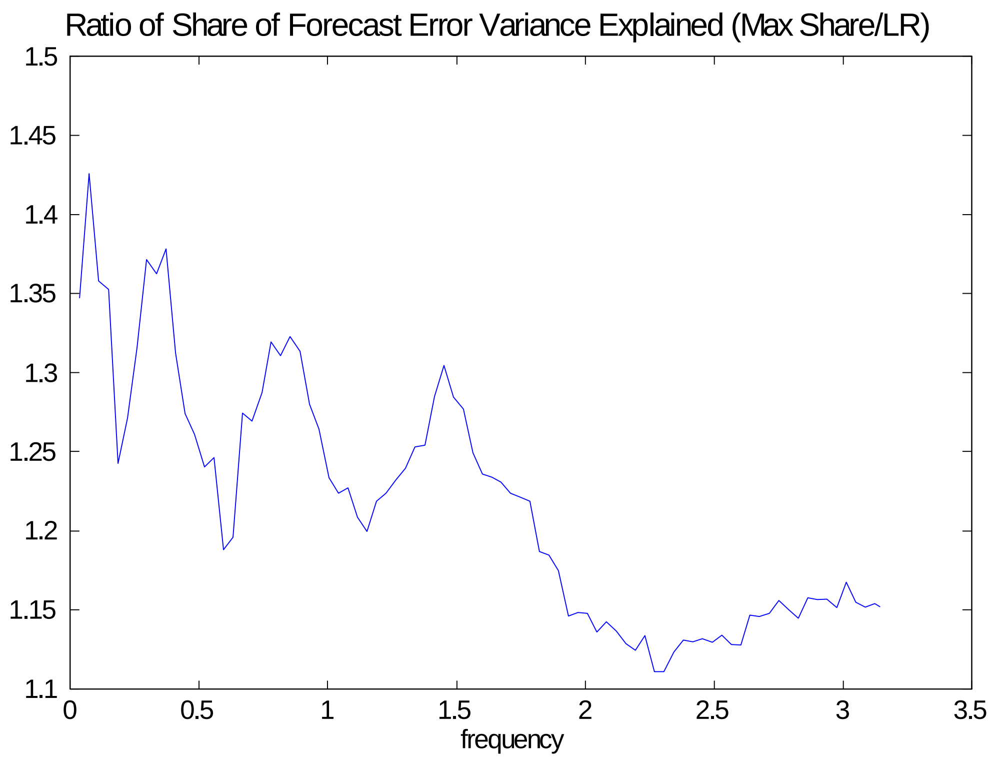

Figure 2: Share of the variance of productivity explained by the Max Share versus LR methods at various frequencies.

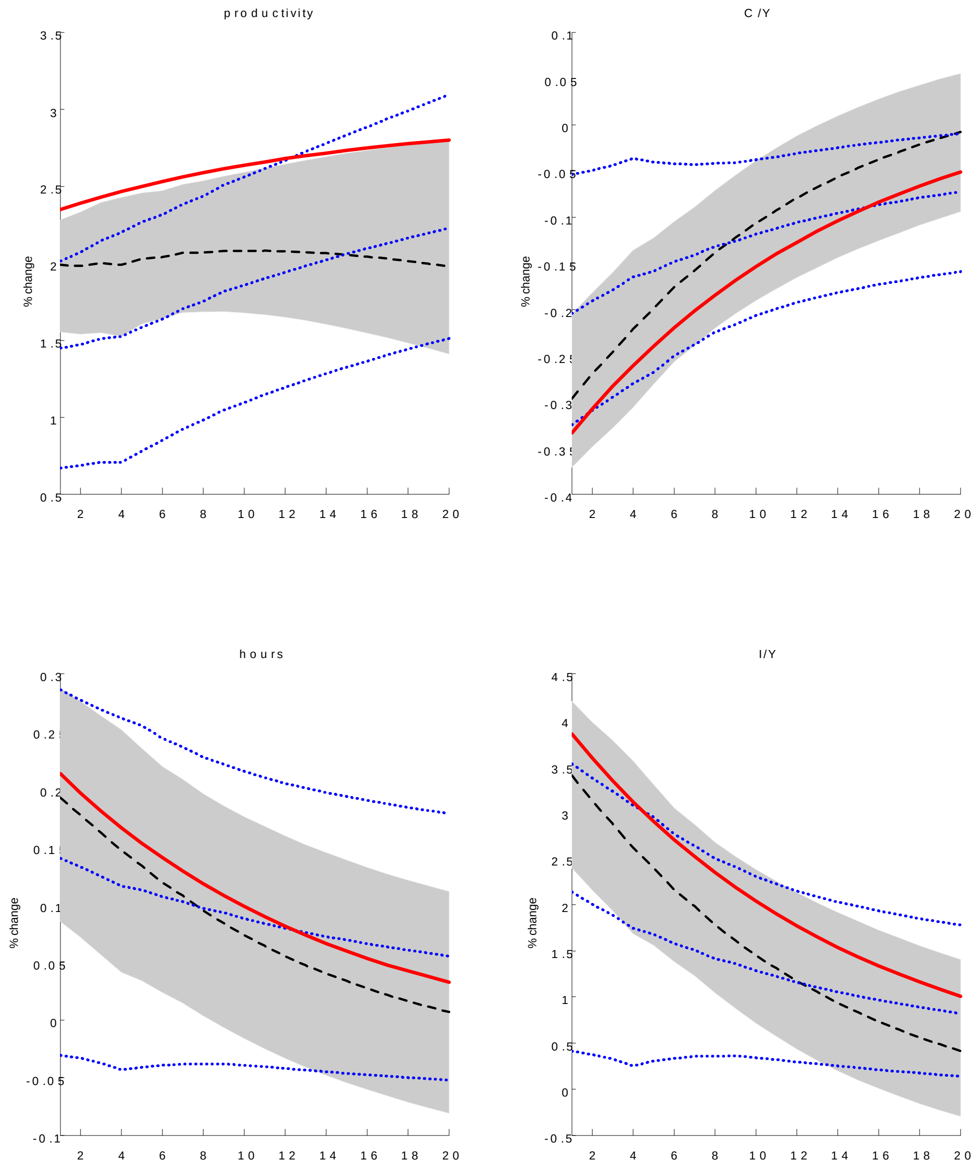

Figure 3: Impulse responses to a technology shock in the presence of important non-technology shocks [AR(1) technology coefficient $\rho_z = 1.0$, non-technology AR(1) coefficients $\rho = 0.98$, and the four shock variances equal 0.0148]. Theoretical responses are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability intervals are shown by dotted lines. Median estimates and error bands are averages across 1,000 estimates, each representing what an econometrician would estimate based on a sample with 174 observations and 1,000 draws from the posterior distributions for the impulse responses.

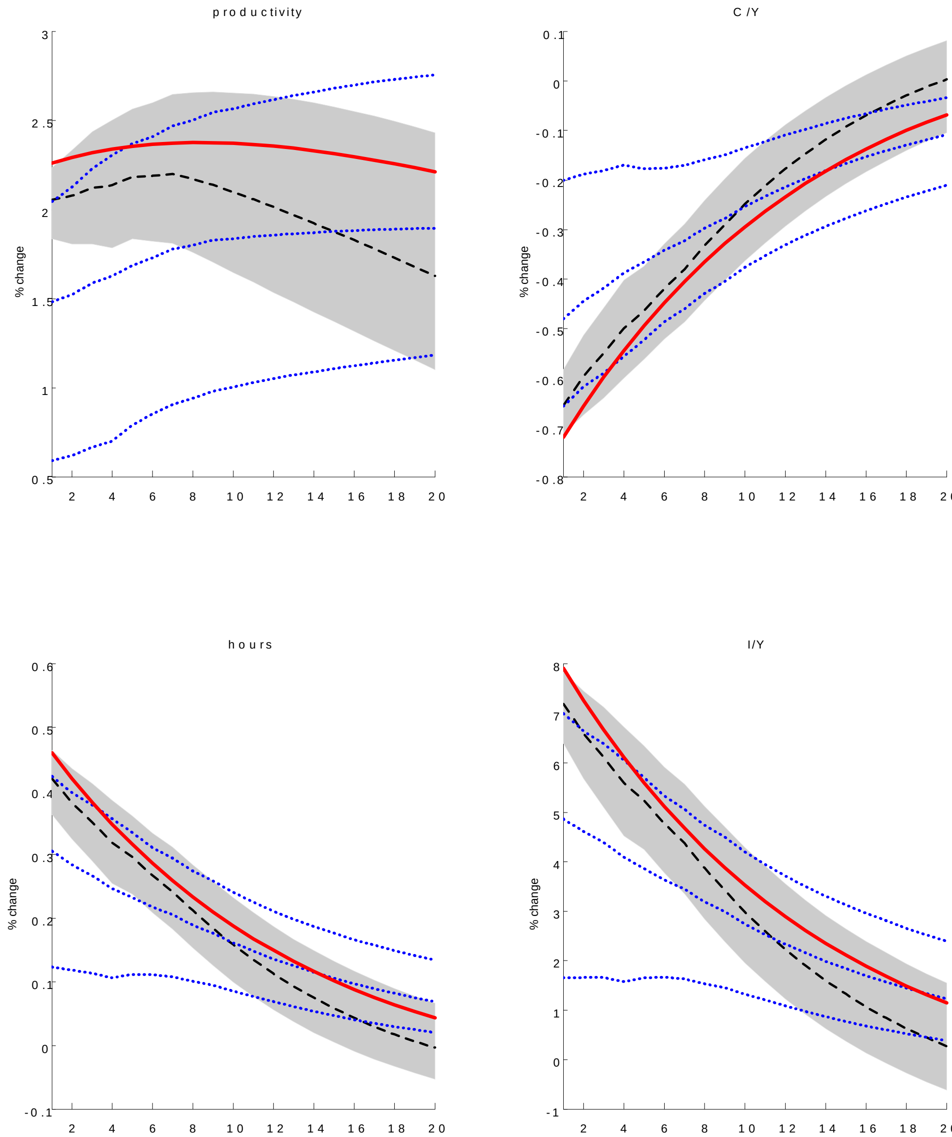

Figure 4: Impulse responses to a stationary technology shock (Case 1). Theoretical responses [with AR(1) technology coefficient $\rho_z = 0.98$ and non-technology AR(1) coefficients $\rho = 0.6$] are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability intervals are shown by dotted lines. Median estimates and error bands are averages across 1,000 estimates, each representing what an econometrician would estimate based on a sample with 174 observations and 1,000 draws from the posterior distributions for the impulse responses.

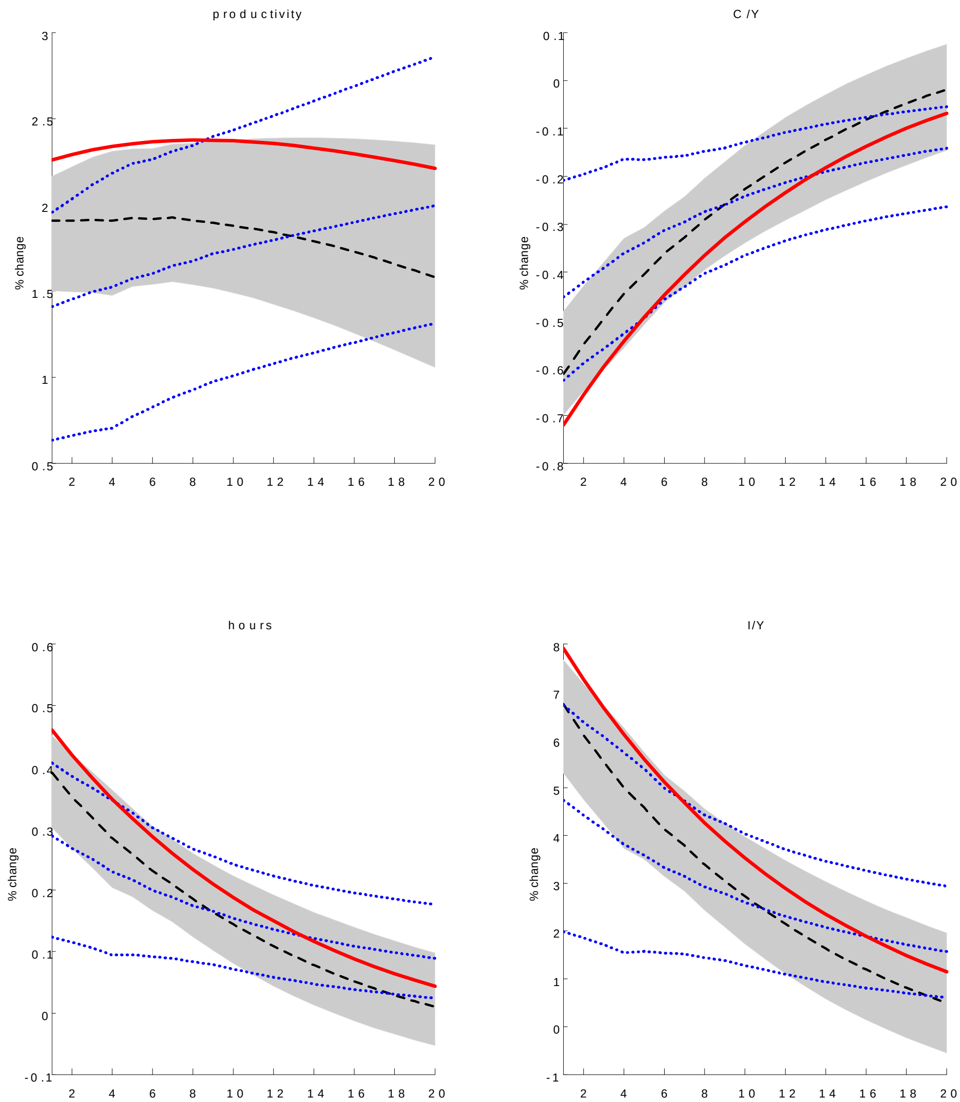

Figure 5: Impulse responses to a stationary technology shock (Case 2). Theoretical responses [with AR(1) technology coefficient $\rho_z = 0.98$ and non-technology AR(1) coefficients $\rho = 0.98$] are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability intervals are shown by dotted lines. Monte Carlo estimates are based on 1,000 draws of samples with 174 observations.

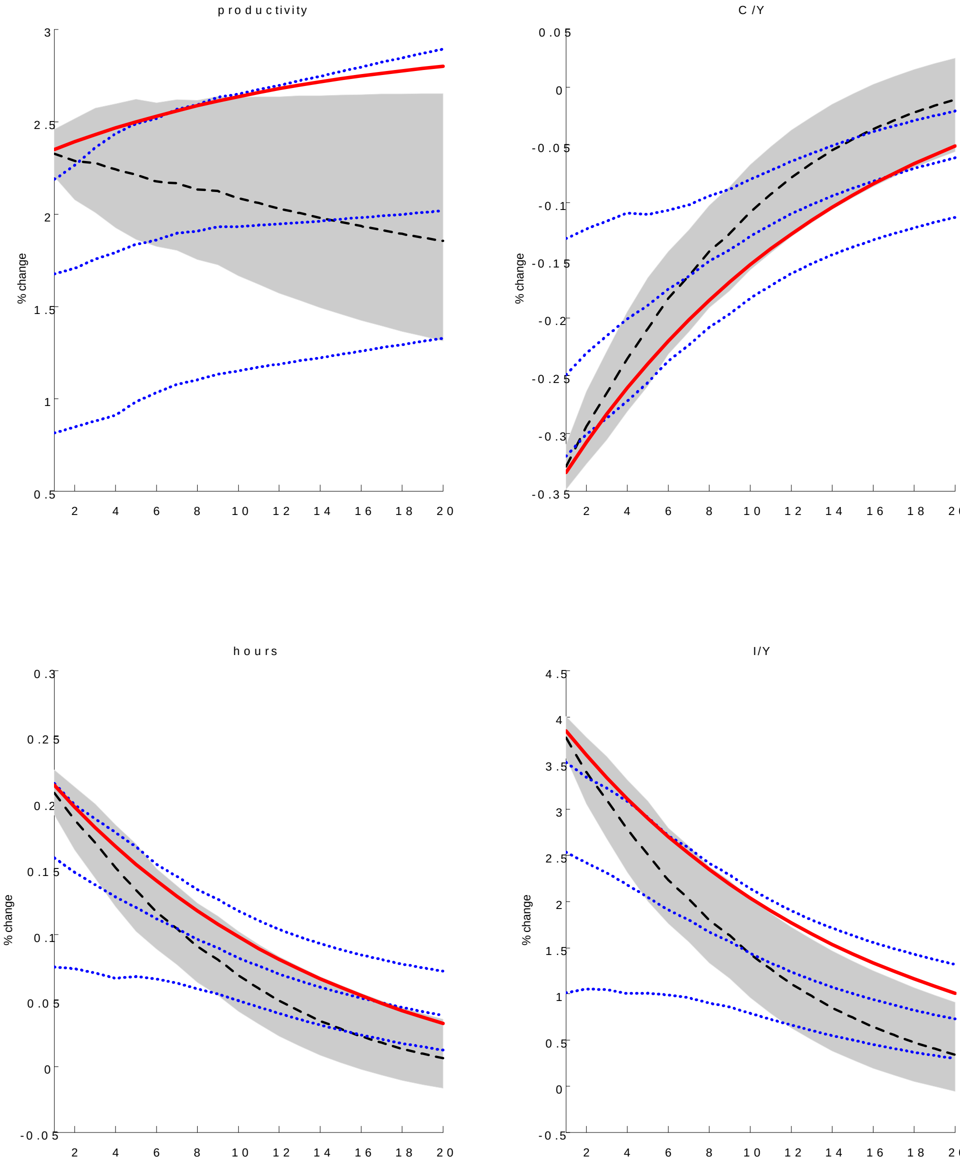

Figure 6: Impulse responses to a technology shock in the benchmark model when productivity is in differences in the VAR. Theoretical responses [with AR(1) technology coefficient $\rho_z = 1.0$ and non-technology AR(1) coefficients $\rho = 0.6$] are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability intervals are shown by dotted lines. Median estimates and error bands are averages across 1,000 estimates, each representing what an econometrician would estimate based on a sample with 174 observations and 1,000 draws from the posterior distributions for the impulse responses.

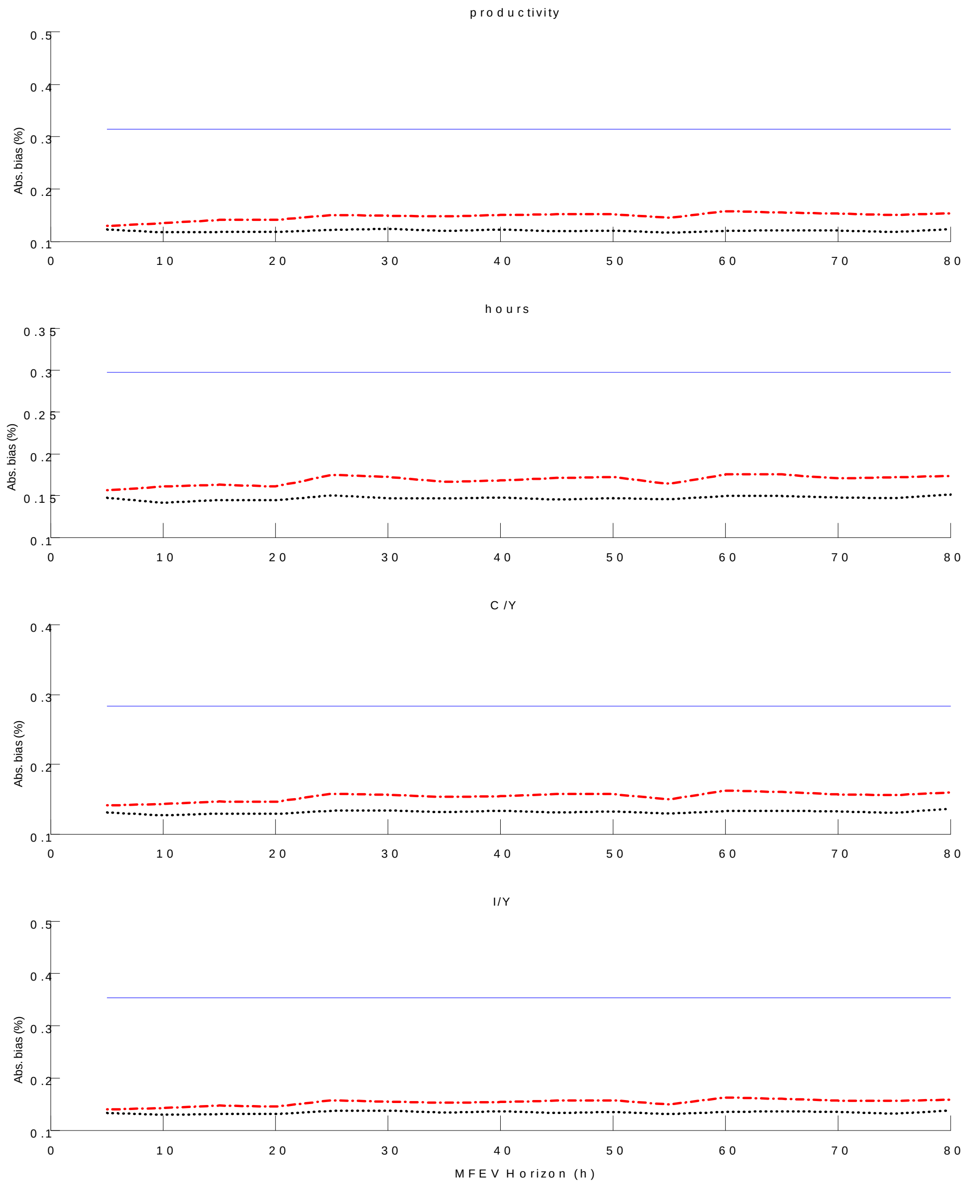

Figure 7: Average absolute bias in LR (solid) versus Max Share across alternative maximization horizons when productivity is specified in levels (dash-dotted) and in differences (dotted) in the Max Share VAR. Bias is measured as the absolute difference between the median Max Share (or LR) and theoretical responses, averaged over the first four quarters. The theoretical model is the benchmark model. The underlying Max Share and LR responses are averages across 1,000 median estimates, each representing what an econometrician would estimate based on a sample with 174 observations.

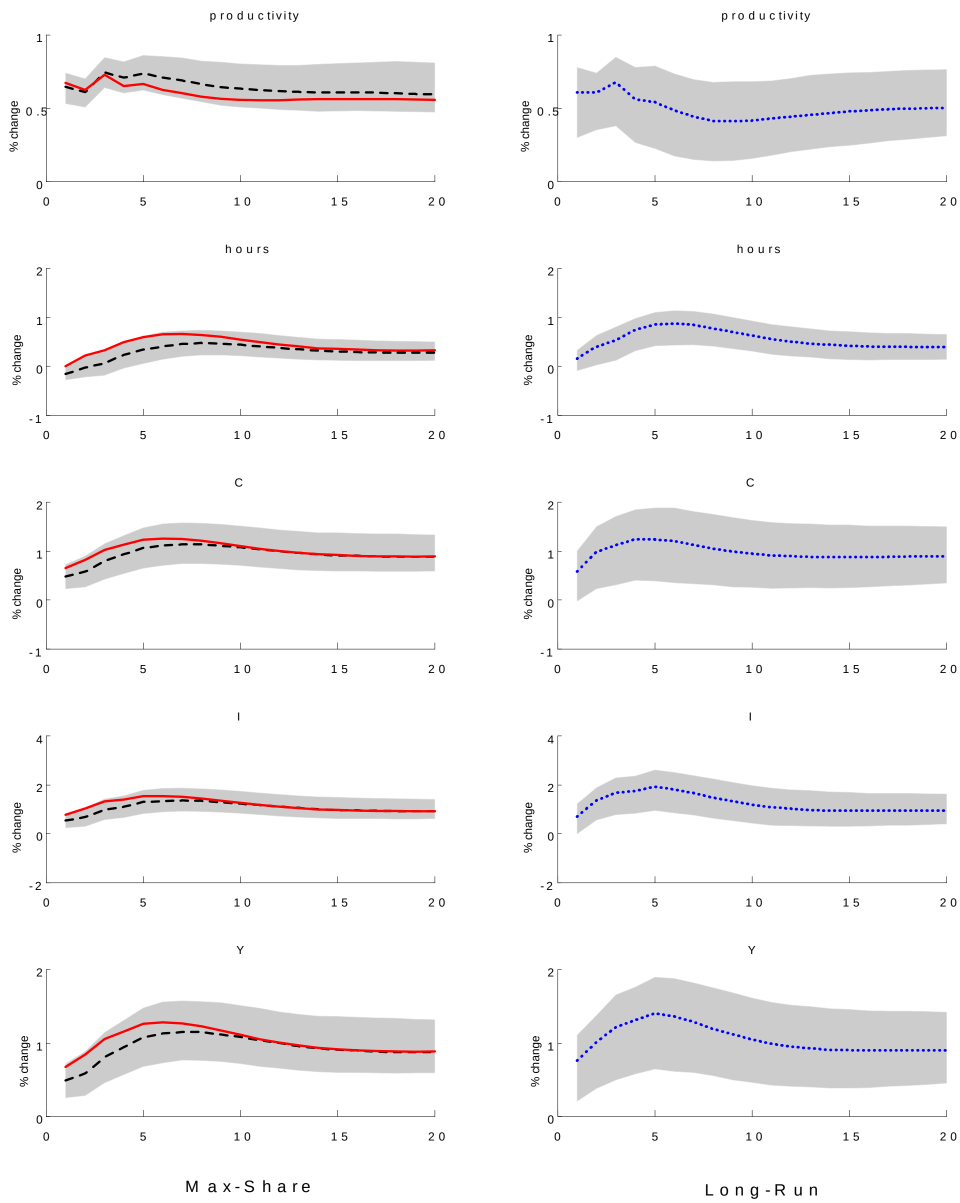

Figure 8: Impulse responses to a technology shock in the Max Share (on left) and the LR models (on right) estimated without a prior using quarterly data1959:1 - 2002:4 (shaded areas are 68 percent probability interval), as well as predicted responses when hours are restricted to be positive on impact in the Max Share model (thick solid line on left).

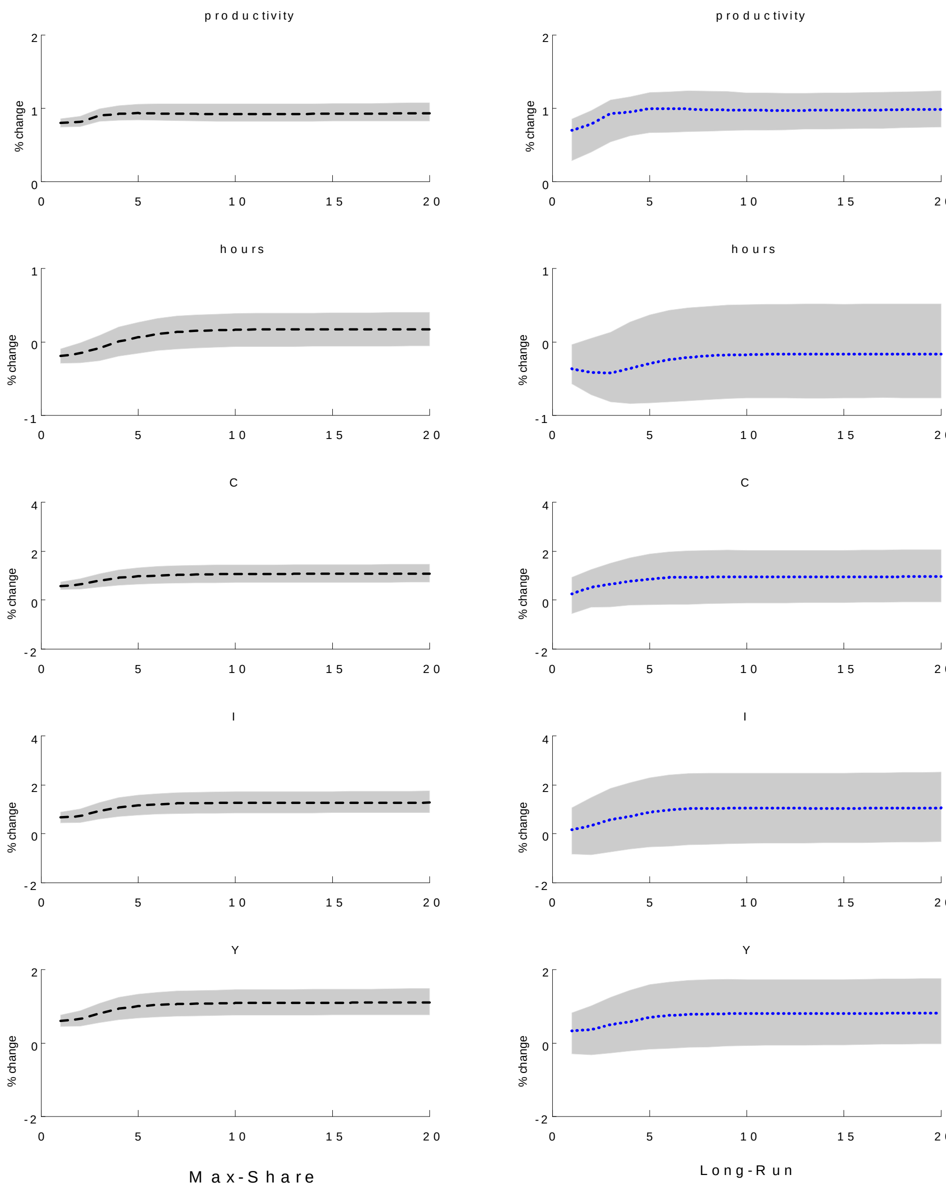

Figure 9: Impulse responses to a technology shock using quarterly data estimated with Sims-Zha prior in the Max Share (on left) and the LR (on right) models. Error bands are 68 percent intervals from posterior probability distributions (shaded areas).

6 Appendix: Theoretical Model

In this appendix, we outline the RBC model used to generate the data for the Monte Carlo experiments. Households choose consumption, C_t, labor, N_t, and investment, I_t, to maximize the expected present-discounted value of utility:

$$U(C_t, N_t) = \sum_{t=1}^{\infty} \beta^{t-1} \left[\ln(C_t) + \Phi \ln(1 - N_t) \right],$$

subject to a standard budget constraint:

$$C_t + I_t = (1 - \varsigma_{nt}) W_t N_t + (1 - \varsigma_{kt}) r_t K_t + \delta \varsigma_{kt} K_t - \Psi_t;$$

the equation characterizing the evolution of capital, K_t:

$$K_{t+1} = (1 - \delta) K_t + I_t;$$

an economy-wide resource constraint:

$$C_t + I_t + G_t \leq Y_t;$$

and a government spending constraint:

$$G_t = \varsigma_{nt} W_t N_t + \varsigma_{kt} (r_t - \delta) K_t + \Psi_t,$$

where r_t is the pre-tax return on capital, W_t is the real wage rate, δ is the depreciation rate, β is the discount factor, Ψ_t is a lump-sum tax, ς_{nt} is the tax on labor, and ς_{kt} is the tax on capital income. Consumers own the capital and rent it to firms. The government balances its budget each period and finances its spending through a combination of lump-sum taxes and distortionary labor and capital income taxes. Tax rates on capital and labor income are stochastically determined by $\tau_{it} = \rho_i \tau_{it-1} + \sigma_{\tau_i} \varepsilon_{\tau_i}$, for $i = k, n$ where $\tau_{it} = \ln(\varsigma_{it}) - \ln(\overline{\varsigma_i})$, and $\overline{\varsigma_i}$ are the steady-state values.[24] The steady-state deviation of government purchases, g_t, has a similar first-order autoregressive process. Finally, output is determined by a Cobb-Douglas production technology:

$$Y_t = (Z_t N_t)^{\alpha} K_t^{1-\alpha},$$

where Z_t is an exogenous process for labor-augmenting technological innovation, $z_t = \rho_z z_{t-1} + \sigma_z \varepsilon_{z_t}$ is the log of technology, and $\varepsilon_z \backsim i.i.d.N\left(0, \sigma_z^2\right)$.

The model is solved by first eliminating non-stationarities arising from technology by dividing Y_t, K_{t-1}, I_t, C_t, G_t, W_t, and Ψ_t by Z_t. Next, the necessary first-order and steady-state conditions are computed based on selected parameter values. The model is log-linearized around the steady-state growth paths, and the recursive equilibrium law of motion is solved using the method of undetermined coefficients. A more detailed explanation of

[24] The steady-state value for the ratio of government to output deserves special mention. The difference between private output and the sum of private consumption and investment is treated as exogenous government consumption [see Figure 8 and the accompanying text in Uhlig (2003)]. Under this simplifying assumption, the international sector and government investment are not explicitly modeled, although they may, in fact, be relevant in the transmission of technology.

this procedure for solving dynamic stochastic models can be found in Uhlig (1999).

To ensure the VAR representation exists under each parameterization, the model is written in its VARMA form. We can then verify that the MA portion is invertible. Here, we present the derivation of the VARMA representation. Given the recursive solution

$$\psi_t = p\psi_{t-1} + Q\xi_t, \tag{9}$$

$$y_t = W\psi_t + S\xi_t, \tag{10}$$

where ψ_t is a vector of endogenous state variables (in our case, capital, k, is the lone endogenous state variable), ξ_t is a vector of exogenous state variables (technology, A, government shock, g, and capital and labor tax shocks, τ_k and τ_n), and y_t is a vector of other endogenous variables (output, hours, consumption, and investment). The endogenous variables used to estimate the VAR are labor productivity (output-hours ratio), hours, consumption-output ratio, and investment-output ratio. Since these variables are basic transforms of the underlying variables, invertibility remains; of course we verified that this is indeed the case. The scalar p and the vectors Q, W, and S are determined by simulating the model, conditional on the parameter values from Table 1.

Substituting (9) into (10) yields

$$y_t = pW\psi_{t-2} + WQ\xi_t + S\xi_t. \tag{11}$$

Realize that

$$W\psi_{t-1} = w_{t-1} - S\xi_{t-1}, \tag{12}$$

and substitute (12) into (11). Collecting terms yields:

$$y_t - py_{t-1} = S\xi_t + (WQ - pS)\xi_{t-1}. \tag{13}$$

We can rewrite this as:

$$S^{-1}y_t - (S^{-1}p)y_{t-1} = \xi_t + S^{-1}(WQ - pS)\xi_{t-1}, \tag{14}$$

$$D(L)y_t = C(L)\xi_t.$$

Finally, given the parameterizations from Table 1, we ensured that the roots of $C(L) = I + [S^{-1}(WQ - pS)]L$ lie outside the unit circle as required for invertibility.

References

[1] Barsky, Robert B. and Sims, Eric R. "Information Shocks, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." Manuscript, University of Michigan, October 2006.

[2] Basu, Susanto; Fernald, John G.; and Kimball, Miles S. "Are Technology Improvements Contractionary?" *American Economic Review*, December 2006, *96*(5), pp. 1418-48.

[3] Blough, Stephen R. "The Relationship Between Power and Level for Generic Unit Root Tests in Finite Samples." *Journal of Applied Econometrics*, July-September 1992, *7*(3), pp. 295-308.

[4] Chari, V.V.; Kehoe, Patrick J.; and McGrattan, Ellen R. "Are Structural VARs with Long-Run Restrictions Useful in Developing Business Cycle Theory?" Staff Report 364, Federal Reserve Bank of Minneapolis, May 2007.

[5] Christiano, Lawrence J. and Eichenbaum, Martin. "Unit Roots in Real GNP: Do We Know, and Do We Care?" *Carnegie-Rochester Conference Series on Public Policy*, Spring 1990, *32*(0), pp. 7-61.

[6] Christiano, Lawrence J.; Eichenbaum, Martin; and Vigfusson, Robert. "What Happens After a Technology Shock?" Manuscript, August 2004.

[7] Christiano, Lawrence J.; Eichenbaum, Martin; and Vigfusson, Robert. "Assessing Structural VARs." *NBER Macroeconomics Annual 2006*, Volume 21, Eds. Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, 2007, pp. 1-72.

[8] Cochrane, John H. "A Critique of the Application of Unit Root Tests." *Journal of Economic Dynamics and Control*, April 1991, *15*(2), pp. 275-84.

[9] Comin, Diego and Gertler, Mark. "Medium-Term Business Cycles." *American Economic Review*, June 2006, *96*(3), pp. 523-51.

[10] Dedola, Luca and Neri, Stafano. "What Does a Technology Shock Do? A VAR Analysis with Model-Based Sign Restrictions." *Journal of Monetary Economics*, March 2007, *54*(2), pp. 512-49.

[11] Dupor, Bill and Kiefer, Leonard. "Executing Long Run Restrictions." Working Paper, October 2007.

[12] Erceg, Christopher J.; Guerrieri, Luca; and Gust, Christopher. "Can Long-Run Restrictions Identify Technology Shocks?" *Journal of the European Economic Association*, December 2005, *3*(6), pp. 1237-78.

[13] Faust, Jon. "Theoretical Confidence Level Problems with Confidence Intervals for the Spectrum of a Time Series." International Finance Discussion Paper No. 575, Board of Governors of the Federal Reserve System, December 1996.

[14] Faust, Jon. "The Robustness of Identified VAR Conclusions about Money." *Carnegie-Rochester Conference Series on Public Policy*, December 1998, *49*(0), pp. 207-44.

[15] Faust, Jon and Leeper, Eric M. "When Do Long-Run Identifying Restrictions Give Reliable Results?" *Journal of Business and Economic Statistics*, July 1997, *15*(3), pp. 345-53.

[16] Francis, Neville and Ramey, Valerie A. "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." *Journal of Monetary Economics*, November 2005, *52*(8), pp. 1379-99.

[17] Francis, Neville and Ramey, Valerie A. "Measures of Per Capita Hours and their Implications for the Technology-Hours Debate." Manuscript, June 2008.

[18] Galí, Jordi. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, March 1999, *89*(1), pp. 249-71.

[19] Galí, Jordi and Rabanal, Pau. "Technology Shocks and Aggregate Fluctuations: How Well Does the Real Business Cycle Model Fit Postwar U.S. Data?" *NBER Macroeconomics Annual 2004*, Volume 19, Eds. Mark Gertler and Kenneth Rogoff, 2005, pp. 225-88.

[20] Hamilton, James D. *Time Series Analysis.* Princeton, NJ: Princeton University Press, 1994.

[21] Jorda, Oscar. "Estimation and Inference of Impulse Responses by Local Projections." *American Economic Review*, March 2005, *95*(1), pp. 161-82.

[22] Khan, Hashmat and Tsoukalas, John. "Technology Shocks and UK Business Cycles." Working Paper, December 2005.

[23] Lütkepohl, Helmut. *Introduction to Multiple Time Series Analysis*, Second Edition. New York: Springer-Verlag, 1993.

[24] Mittnik, Stefan and Zadrozny, Peter A. "Asymptotic Distributions of Impulse Responses, Step Responses, and Various Decompositions of Estimated Linear Dynamic Models." *Econometrica*, July 1993, *61*(4), pp. 857–70.

[25] Shea, John. "What Do Technology Shocks Do?" *NBER Macroeconomics Annual 1998*, Volume 13, Eds. Ben S. Bernanke and Julio J. Rotemberg, 1999, pp. 275-310.

[26] Sims, Christopher A. "The Role of Approximate Prior Restrictions in Distributed Lag Estimation." *Journal of the American Statistical Association*, March 1972, *67*(337), pp. 169-75.

[27] Sims, Christopher A. "Modeling Trends." Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, Discussion Paper 22, December 1989.

[28] Sims, Christopher A.; Stock, James H.; and Watson, Mark W. "Inference in Linear Time Series Models with Some Unit Roots." *Econometrica*, January 1990, *58*(1), pp. 113-44.

[29] Sims, Christopher A. and Zha, Tao. "Bayesian Methods for Dynamic Multivariate Models." *International Economic Review*, November 1998, *39*(4), pp. 949-68.

[30] Sims, Christopher A. and Zha, Tao. "Error Bands for Impulse Responses." *Econometrica*, September 1999, *67*(5), pp. 1113-55

[31] Stock, James H. "Unit Roots in Real GNP: Do We Know, and Do We Care?: A Comment." *Carnegie-Rochester Conference Series on Public Policy*, Spring 1990, *32*(0), pp. 63-82.

[32] Uhlig, Harald. "A Toolkit for Analysing Nonlinear Dynamic Stochastic Models Easily." *Computational Methods for the Study of Dynamic Economies,* Eds. Ramon Marimon and Andrew Scott, New York: Oxford University Press, 1999, pp. 30-61.

[33] Uhlig, Harald. "How Well Do We Understand Business Cycles and Growth? Examining the Data with a Real Business Cycle Model." 2002 Ottobeuren Conference, Volume 32, Eds. Wolfgang Franz, Hans Jürgen Ramser, and Manfred Stadler, 2003, pp. 295-319.

[34] Uhlig, Harald. "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association*, April/May 2004, *2*(2/3), pp. 361-71.