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
Jennifer E. Roush
and
Riccardo DiCecio

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

April 2005
Revised December 2010

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

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

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

keywords: productivity, structural VAR, long-run restrictions

December 14, 2010

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. Unlike its long-run restriction counterpart, when our Max Share identification technique is applied to U.S. data it delivers the robust result that hours worked responds negatively to positive technology shocks.

[JEL: C32, C50, E32]

*Julieta Caunedo and 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, Elena Pesavento, 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.

†Corresponding author. 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.

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

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

1 Introduction

By their nature, long-run restricted structural vector autoregressions 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) assesses 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 long-run (LR) identification to obtain hypothetical small-sample distributions of the impulse responses to technology shocks. Both studies conclude that impulse responses identified by LR restrictions can be substantially biased, either in sign or in magnitude.

Recently, LR restrictions have attracted renewed attention 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 using U.S. data yield a decline in hours, apparently contradicting the theoretical predictions of a broad class of RBC models.[1] 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.[2] This paper focuses instead on the identifying assumption regarding the estimated long-run productivity process.

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

[1] Basu, Fernald, and Kimball (2006) and Shea (1999) use different techniques to identify technology and conclude that the hours response is negative.

[2] A number of papers [including Galí and Rabanal (2005) and Christiano, Eichenbaum, and Vigfusson (CEV, 2004)] have focused on the stationarity of hours as the key determinant of the sign of the hours response to a technology shock. A decline in hours is obtained when hours are first-differenced [Galí (1999)], detrended [Fernald (2007)], or demographically-adjusted [Francis and Ramey (2009)]. CEV (2004) argue that *per capita* labor is bounded and cannot have a unit root. If assumed stationary, responds to a technology shock positively on impact.

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.[3]

The Max Share approach has several potential advantages over the conventional LR approach. First, by focusing on a finite horizon, we hope to gain estimation precision over LR, which relies on much longer horizon parameter estimates. Second, 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. Third, 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 output variability, we use our objective function as a necessary condition for identification.[4] 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

[3]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.

[4]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.

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 and a standard medium-scale DSGE model with sticky prices, 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. These advantages are found to be robust to alternative specifications of the theoretical technology and non-technology shocks. 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. Results using the Max Share approach applied to U.S. data are consistent with Galí's original finding that hours decline after a technology shock.

In the next section, we present the Max Share identification approach. We then compare the small-sample performances of the two identification approaches using data simulated from the RBC and sticky price models. 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

In this section, we provide an overview of the mechanics behind the standard LR and our Max Share identifications. Both approaches isolate the (primary) driver of long-run productivity trend. While the identifying approaches differ in their implementation, the key assumption behind both the Max Share and LR identification is consistent with that imposed by Galí:

Assumption 1: Technology shock is the sole contributor of long-run labor productivity shifts. All other structural innovations having transitory effects on labor productivity.

The assumption identifying technology shocks arises from a broad class of models in which *log* labor productivity, x_t, can be decomposed into two orthogonal components: an unobserved random walk trend component, x_t^T, which we will call technology, and an unobserved cyclical process, x_t^S.[5] The trend-cycle decomposition of productivity is

$$x_t = x_t^T + x_t^S, \tag{1}$$

where

$$x_t^T = x_{t-1}^T + \eta_t, \tag{2}$$

$\eta_t \sim iid\ N(0, \sigma_\eta^2)$, and x_t^S is stationary and ergodic. Since all processes except technology are assumed stationary, the unit root in productivity must arise from x_t^T.[6] Thus, under Assumption 1, technology will necessarily dominate the forecast-error variance of the log-level of productivity at suitably long forecast horizons.[7] Given that ν_t can be thought of as a amalgamated non-technology shock (composed of fiscal, monetary, and tax shocks), Assumption 1 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:

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

[5] Throughout, we will assume that labor productivity is $I(1)$ and that the system is estimated with a single stochastic trend.

[6] This arises from the steady-state condition $X = W = \alpha x_t^T \left(\frac{k}{N}\right)^{1-\alpha}$ 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 x_t^T is the level of technology.

[7] Recall that the variance of a unit root process, $var(x_t^T) = t\sigma_\eta$, increases with t. 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].

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:

$$B(L)y_t = \mu_t, \tag{3}$$

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. Conventional long-run identification [e.g., Blanchard and Quah (1989); Shapiro and Watson (1988)] imposes restrictions on the effect of the jth shock on the ith variable at an infinite horizon. This is implemented through restrictions on $\left[C(1)A_0^{-1}\right]_{i,j}$, where neutrality implies the restriction $\left[C(1)A_0^{-1}\right]_{i,j} = 0$ for some j. Formally, we have

$$\left[C(1)A_0^{-1}\right]_{i=1,j\neq i} = 0, \tag{4}$$

where $i = 1$ represents labor productivity growth ordered first and $j \neq i$ indicates all non-technology shocks.[8]

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 long-run restrictions, we identify the

[8] Equation (4) is isomorphic to the assumption that the zero-frequency spectrum of labor productivity growth is attributable entirely to technology [see also DiCecio and Owyang (2010)].

technology shock by maximizing the forecast-error variance share of productivity at long, finite horizons. This approach is consistent with suggestions in Uhlig (2004) and CEV and is adapted from methods introduced in Faust (1998). We begin by introducing the methodology and then discuss its practicality.

2.2.1 Methodology

In the Max Share identification, all variables including labor productivity enter the VAR in *log-levels*.[9] The method is operationalized by first expressing the h-step-ahead forecast error for y as a function of realized reduced-form errors:

$$y_{t+h} - \widehat{y}_{t+h} = \sum_{\tau=0}^{h-1} C_\tau \mu_{t+h-\tau}, \tag{5}$$

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

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

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

$$\omega_{ij}\left(\alpha(h)\right) = \frac{e_i' \left[\sum_{\tau=0}^{h-1} C_\tau \alpha \alpha' C_\tau'\right] e_i}{e_i' \left[\sum_{\tau=0}^{h-1} C_\tau \Omega_\mu C_\tau'\right] e_i}, \tag{6}$$

where e_i is an $n \times 1$ indicator vector that picks out the impulse vector $\alpha = De_j$, 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:

$$\max_\alpha \omega_{1j}\left(\alpha(h)\right) = \max_\alpha \alpha' \frac{e_i' \left[\sum_{\tau=0}^{h-1} \widetilde{C}_\tau \widetilde{C}_\tau'\right] e_i}{e_i' \left[\sum_{\tau=0}^{h-1} \widetilde{C}_\tau \Omega_\mu \widetilde{C}_\tau'\right] e_i} \alpha, \tag{7}$$

[9] To execute LR, labor productivity must be entered in differences. 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. As an alternative, strong beliefs can be incorporated into a prior.

with the additional normalization $\alpha'\alpha = 1$.[10] 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.

2.2.2 Practical Considerations

LR has previously been found perform poorly in small samples of length comparable to the U.S. postwar sample [see EGG and CKM]. The Max Share approach may demonstrate *less* small-sample bias, in part because it relies on restrictions at a finite horizon.[11] Sims (1972) and, more recently, Faust and Leeper (1997) and CEV (2007) make similar arguments 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)].[12]

The econometrician typically uses some criteria to determine the lag order, which may or may not be the appropriate specification. For example, the VAR representation of many theoretical models is a VAR(∞), which necessitates a truncation of the lag order for estimation. If the model is misspecified, the implied responses will be biased, the degree of bias depending on the manner of identification. There may be key differences in the bias that arise from using either LR or Max Share in small samples. LR, for example, places a restriction on $C(1)$, the infinite sum of the transformed/rotated VAR coefficients. The econometrician, however, does not use the true $C(1)$; instead, she uses the estimated $\widehat{C}(1)$ which may bias the identification of the shock.[13] Our conjecture is that imposing restrictions on $\sum_{\tau=0}^{h-1} C_\tau \alpha\alpha' C'_\tau$ rather than $C(1)$ can

[10] 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.

[11] Unlike LR, the Max Share rotation is not preserved as the sample size increases [Mittnik and Zadrozny (1993)] making the formal proof of bias reduction problematic.

[12] 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)].

[13] An informal proof of this conjecture can be found in CEV (2007). CEV (2007) contend that using short-run restrictions minimizes the effect of potential misspecification because short-run restrictions do not require $\widehat{C}(1)$

reduce – but not eliminate – misspecification bias while preserving the theoretical interpretation of the identification.[14]

The preceding discussion suggests that errors in estimating $C_\tau A_0^{-1}$ may confound the identification, making it possible to attribute too much of the FEV 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 the FEV share 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 known data-generating processes.

3 Monte Carlo Experiments

This section outlines the Monte Carlo methods used to assess the ability of the two identification approaches for small samples. We use simulated data from a neoclassical growth model (RBC) and a New Keynesian sticky price (NK) model. For both models, equation (1) holds, allowing us to obtain theoretical impulse responses to a technology shock, but the hours response varies across models. For the RBC model, hours rise in response to a positive technology shock; for the NK model, hours fall. Theoretical impulse responses and simulated data are generated for parameterizations of the two models which vary the persistence of the nontechnology component. Model details and parameterizations appear in the appendices.

Each Monte Carlo iteration consists of 202 simulated observations of each of the variables.[15] We estimate a reduced-form VAR and obtain structural impulse responses both by using LR restrictions and by Max Share. The posterior distributions for the impulse responses are simulated from 1,000 draws utilizing a Normal–Inverse-Wishart prior centered on the OLS estimates [see Sims and Zha (1999)]. To account for potential asymmetries in the impulse responses, we retain the median,

for identification. 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.

[14] In addition to bias caused by misspecifying the VAR, the Max Share identification also may be subject to bias caused by choosing h too small. It is not clear *ex ante* which effect will dominate in small samples.

[15] In results not reported here, we found that increasing the sample size reduced the bias between the estimated and theoretical impulse responses for both identification mothods. These results are available upon request.

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. Analysis of the posterior coverage follows. We also analyze the correlation between the theoretical shocks and those estimated by both LR and Max Share.

3.1 Benchmark Results

The benchmark RBC model is a four-variable VAR(4) with the logs of labor productivity, hours, the consumption-output ratio, and the investment-output ratio. The benchmark NK model is a seven-variable VAR(7) with the logs of labor productivity, hours, the consumption-output ratio, the investment-output ratio, inflation, the real wage, and a short-term interest rate. Each reduced-form VAR is estimated using maximum likelihood. Log productivity enters the VAR in first differences for the LR but enters in levels for the Max Share approach.[16] The benchmark RBC model is parameterized to be consistent with EGG; the AR(1) coefficient of technology (ρ_z) is set to 1 and non-technology shocks have AR(1) coefficients of 0.6. The NK model is calibrated to the estimated values in Smets and Wouters (2007a) but adjusted for unit root technology. For Max Share, we assume a maximization horizon h 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); we consider alternative horizons below.

Figures 1 and 2 present the impulse responses to a 100-basis-point technology shock for both the RBC model and the NK model, respectively. For each figure, the thick solid lines depict the theoretical impulse responses. The average of the Max Share median responses across Monte Carlo iterations are shown by the thick dashed lines, with the shaded areas representing the accompanying 68-percent coverage bands. The average of the LR median responses and their error bands are shown by dotted lines. In the RBC model, a positive technology shock leads to an immediate increase in both labor productivity and hours. In the NK model, a positive technology shock leads to an increase in labor productivity but a short-run decline in hours.

[16] In results not reported here, we considered the effect of entering productivity in differences for the Max Share identification. The results were qualitatively similar and available upon request.

For the RBC model, 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, especially at short horizons. For each variable, the Max Share approach yields an impulse response biased toward zero. 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.

For the NK model, the Max Share responses are typically less biased. In addition, Max Share correctly identifies a decline in hours worked on impact; LR's point response for hours is negative but shows no statistically important change in hours at any horizon. As it is the series often used to differentiate the RBC and NK models, LR's inability to identify a decline in hours on impact may be problematic.

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 1 presents the correlations between the model-generated and the estimated technology shocks from both identifications. Our benchmark results are in the first and third rows 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 0.90 and 0.74 for the RBC and NK models, respectively) than that for LR shocks (about 0.55 and 0.48 for the RBC and NK models, respectively). Additionally, the median correlation for the LR model lies in the far left tail of the distribution for the Max Share correlations.

We also investigate the share of the forecast error variance explained by Max Share and LR in the frequency domain. Figures 3 and 4 plot the spectral decomposition of the share of productivity's variance explained by technology under both identification schemes along with the true values from the RBC and NK models, respectively. For the RBC model, both LR and Max Share attribute less of the FEV to technology than the truth, but Max Share is closer to the truth than LR for all variables. For the NK model, the advantage for Max Share is less clear. Taken together with the preceding results, however, this suggests the technology shock is better identified by the finite-horizon alternative Max Share.

3.2 Increasing the Importance of Non-Technology Shocks

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 allow the non-technology shocks to play a greater role in determining labor productivity at long horizons. Specifically, the non-technology stochastic processes — e.g., government spending, capital, and/or labor taxes — are assumed to be highly persistent, with their innovation variances set equal to the variance of technology. Technology, however, remains the source of the unit root in productivity, consistent with (4). Increasing the persistence and variances of the non-technology processes allows them to 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 either identification approach to isolate the technology process. Because (4) still holds, LR is still valid at the infinite horizon, potentially giving LR an advantage over Max Share, all else equal.

Figure 5 shows the responses for this parameterization of the RBC model when technology has a unit root, all non-technology processes have AR(1) coefficients of 0.98, and all stochastic processes have equal variances. Figure 6 shows the responses for the NK model with more persistent nontechnology shocks (see Table A.2 in the appendix for details of the parameterization). Perhaps surprisingly, these results are similar to the previous parameterizations – the Max Share impulse responses demonstrate less bias than the LR responses. An interesting finding is that, for the RBC model, the LR error bands for hours now include zero, making the sign on the hours response statistically indeterminate. For the NK model, the sign for LR is indeterminate while Max Share is weakly negative (the boundary of the 68 percent probability interval is right at zero). Obviously, increasing the importance of nontechnology shocks diminishes the power of both identification methods, but Max Share typically remains less biased overall. The second and fourth rows of Table 1 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 Bias and Coverage Analysis

The performance of Max Share depends on choosing h, the identification horizon, large enough to pick out the correct shock but small enough to minimize the misspecification bias. All of the previous results have been obtained with an exogenously chosen Max Share horizon, h, of 10 years, Here, we ask how much does changing the identification horizon affect the results? To answer this, we compute the bias of the responses identified by Max Share while varying the forecast horizon. Figures 7 and 8 summarize the bias properties of the Max Share approach when h varies from 5 to 80 quarters for the benchmark parameterization. The bias is computed over the first four quarters, expressed as a percentage of the true model response. The dash-dotted lines depict the bias of the Max Share approach; the solid line shows a similar measure for LR which is constant by definition. For both models, the bias shown by Max Share is clearly smaller than that of the LR for all variables and all identification horizons. The Max Share bias is relatively constant over a wide range of h. For example, for the RBC model, 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 these model environments, the bias improvement of the Max Share approach is robust to choosing relatively short identification horizons.

We might also be interested in how the bias changes over the response horizon. Figures 9 and 10 show the mean bias of the identified responses over the first 20 quarters for the baseline parameterizations of the RBC and NK models, respectively. For the RBC model, the productivity response identified by Max Share exhibit less bias on average than those identified by LR. For the other variables, Max Share exhibits less bias at short horizons – horizons typically used to distinguish the models. For the NK model, the bias advantage of Max Share is smaller but still apparent.[17]

These difference in bias also affect the true coverage. In the analysis above, we reported the 68 percent nominal coverage – that is, the means of the interior 68 percent of the posterior distributions of the impulse responses. These nominal coverages may not necessarily contain the true (model) response with 68 percent probability. In fact, the bias in the responses for both Max Share and LR can affect the true coverage leading us to ask how well the 68-percent nominal coverage represents

[17] Because the NK responses are often close to zero, the reported percentage biases as functions of the true values are much larger than those for the RBC model.

the true coverage. Figures 11 and 12 display the percentage of Monte Carlo iterations in which the 68 percent nominal coverage band contains the true response for the baseline parameterizations of the RBC and NK models, respectively. Because of the bias in both identified responses, the true coverage – especially at short horizons – can be much less than the 68 percent nominal coverage. For most horizons and for short horizons in particular, Max Share's 68 percent nominal coverage bands contain the truth more often than LR's.

4 Max Share Identification in U.S. Data

Having evaluated the small-sample performance of our identification scheme through Monte Carlo experiments, we turn to U.S. data and estimate a four-variable VAR(4). The data are quarterly series from 1959:I to 2009:III for private business productivity, private business hours, real consumption as a share of output, and real investment as a share of output.[18] 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 for U.S. Data

The dashed lines in the left column of Figure 13 present the median impulse response to a one-standard-deviation technology shock, with the shaded areas representing the 68 percent probability intervals.[19] 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.

The right column of Figure 13 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

[18] 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, 2009). The impulse responses were qualitatively similar across these models.

[19] All of the Max Share results shown are based on $h = 40$ quarters; similar results were obtained for horizons of 20, 60, and 100 quarters. Although the width of the error bands for the Max Share identification increases with h, they are always narrower than those obtained from LR.

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.

Table 2 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 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.

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 for U.S. Data

The Max Share approach also allows for exhaustive robustness analysis across a broad class of models. While our finding that hours respond negatively to a positive technology shock appears to corroborate Galí's original result, many papers have questioned the robustness of this prediction. In particular, one can imagine that a small modification to the identifying assumption – perhaps corresponding to a more 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.[20] We then examine whether this has a discernible effect on the FEV share attributable to technology. The left column in Figure 13

[20] Dedola and Neri (2007) propose 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's FEV.

also shows the median impulse responses to a technology shock when the additional sign restriction is imposed.[21] 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 3 displays the maximum forecast-error variance share values estimated at horizons of 40 and 100 quarters. 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.[22] This suggests that a positive hours response may also result when non-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.

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.

[21] 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 (7) with an additional constraint of the form $e_i' C_t \alpha \geq 0$. The additional restriction on hours is not an overidentifying restriction suitable to a likelihood ratio test. The Max Share identifying assumption (7) is sufficient to identify technology, but the system as a whole is underidentified with or without the additional sign restriction.

[22] This does not necessarily imply that non-technology shocks influence long-run productivity. The odds ratios for the restricted and unrestricted models were indistinguishable, preventing any conclusions based on how well the model fits the data.

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 off-the-shelf RBC and NK models, 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 appear to, on average, less biased 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 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.

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] Blanchard, Olivier J. and Quah, Danny. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review,* September 1989, 79(4), pp. 655-673.

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

[6] 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.

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

[8] DiCecio, Riccardo and Owyang, Michael. "Identifying Technology Shocks in the Frequency Domain." Federal Reserve Bank of St. Louis Working Paper 2010-025A

[9] 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.

[10] Dupor, Bill and Kiefer, Leonard. "Executing Long Run Restrictions." Manuscript, October 2007.

[11] 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.

[12] 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.

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

[14] 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.

[15] Fernald, John. "Trend Breaks, Long-Run Restrictions, and Contractionary Technology Improvements." *Journal of Monetary Economics* 54(8), November 2007, 2467-2485

[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." *Journal of Money Credit and Banking*, September 2009, Volume 41, Issue 6, pp 1047-1263.

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

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

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

[23] 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.

[24] Shapiro, Matthew D. and Watson, Mark W. "Sources of Business Cycle Fluctuations." NBER Macroeconomics Annual, 1988, Vol. 3, pp. 111-148.

[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.; 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.

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

[29] Smets, Frank and Wouters, Rafael. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach."" The *American Economic Review*, June 2007, 97(3), pp. 586-606.

[30] Smets, Frank and Wouters, Rafael. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." ECB Working Paper #722, February 2007.

[31] 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.

[32] 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.

[33] 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.

Table 1				
Correlation Between Estimated and Model Technology Shocks				
RBC Model				
Parameterization	Identification	$16th$ percentile	median	$84th$ percentile
Baseline	LR	0.10	0.55	0.81
	Max Share	0.85	0.90	0.93
Persistent Nontechnology	LR	0.09	0.50	0.74
	Max Share	0.66	0.80	0.88
NK Model				
Parameterization	Identification	$16th$ percentile	median	$84th$ percentile
Baseline	LR	0.15	0.48	0.68
	Max Share	0.63	0.74	0.82
Persistent Nontechnology	LR	0.11	0.42	0.46
	Max Share	0.42	0.60	0.72

Table 1: We generate1000 draws of artificial data from each model. All technology processes have unit roots. Model parameters are given in the tables in the appendices. For each parameterization, the implied technology shocks are identified using LR and Max Share. 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 2				
Contribution of Technology to FEV for U.S. Data				
	Output		Hours	
Horizon	Max Share	LR	Max Share	LR
4 quarters	0.82	0.58	0.02	0.50
20 quarters	0.90	0.77	0.16	0.61
40 quarters	0.94	0.94	0.20	0.58
80 quarters	0.92	0.97	0.31	0.57

Table 2: 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 40 quarters.

Table 3		
Comparing Shares of Productivity for U.S. Data		
	w/o Sign Restrictions	w/ Sign Restrictions
	Hours Unrestricted	$\text{Hours}_0 \geq 0$
$h = 40$	.928	.872**
$h = 100$	.906	.900*

Table 3: Forecast-error variance share for U.S. data attributed to technology shocks identified by Max Share, estimated with an uninformative prior using a maximization horizon of h=40 and 100 quarters. 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 90 percent interval of unrestricted model using a one tailed-test. *Outside 85 percent interval of the unrestricted model using a one-tailed test.

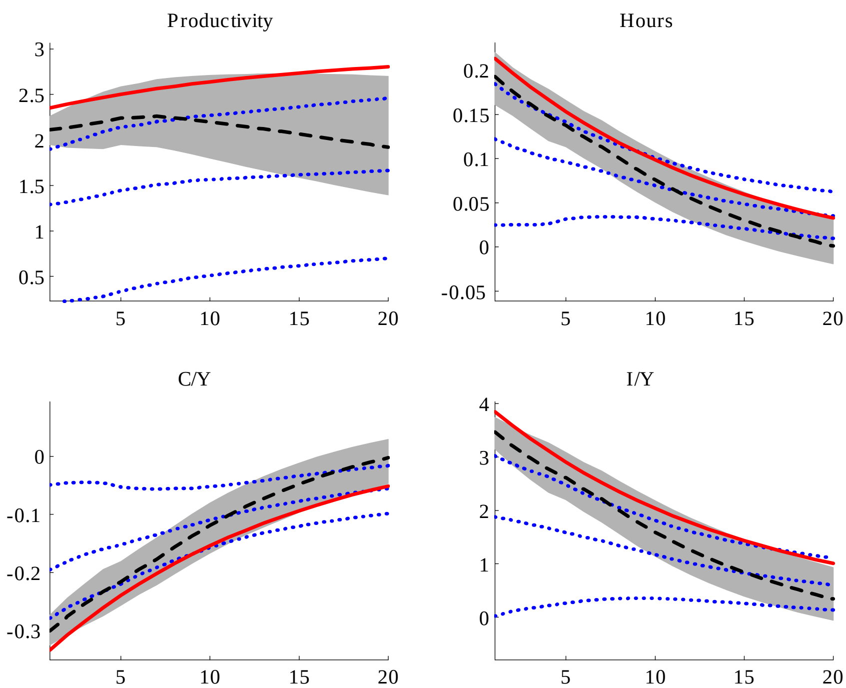

Figure 1: Impulse responses to a technology shock in the benchmark RBC model.

Theoretical responses [with AR(1) technology coefficient $\rho_z = 1.0$ and nontechnology 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 202 observations and 1,000 draws from the posterior distributions for the impulse responses.

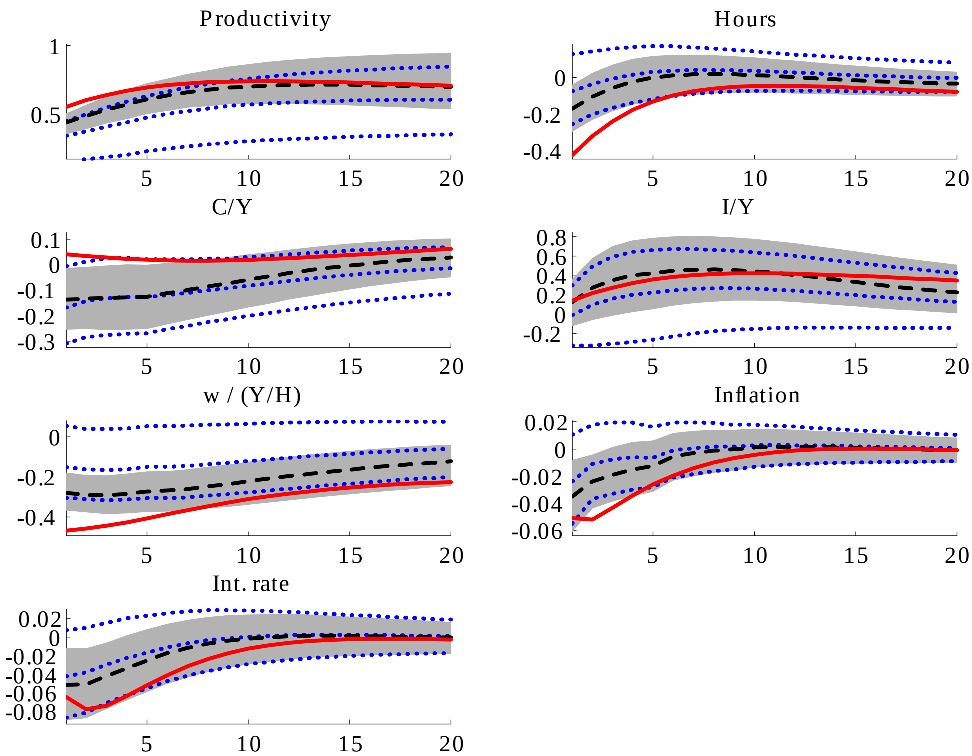

Figure 2: Impulse responses to a technology shock in the benchmark NK model. Theoretical responses are shown by thick solid lines. Median and 68 percent probability intervals for Max Share from Monte Carlo experiments are shown with dashed lines and shaded areas. LR median responses and 68 percent probability 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 202 observations and 1,000 draws from the posterior distributions for the impulse responses.

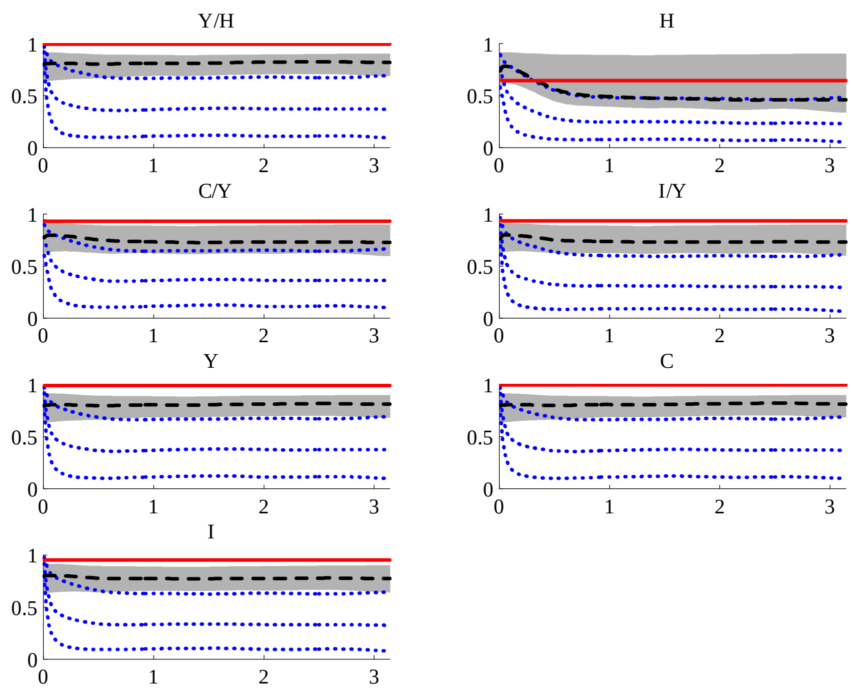

Figure 3: FEV share of productivity for the RBC model at various frequencies.

True FEV (solid), Max Share (dashed), and LR (dotted) methods using simulated data using the benchmark parameterizations). Shaded region and dotted lines represent 68 percent coverages for Max Share and LR, respectively.

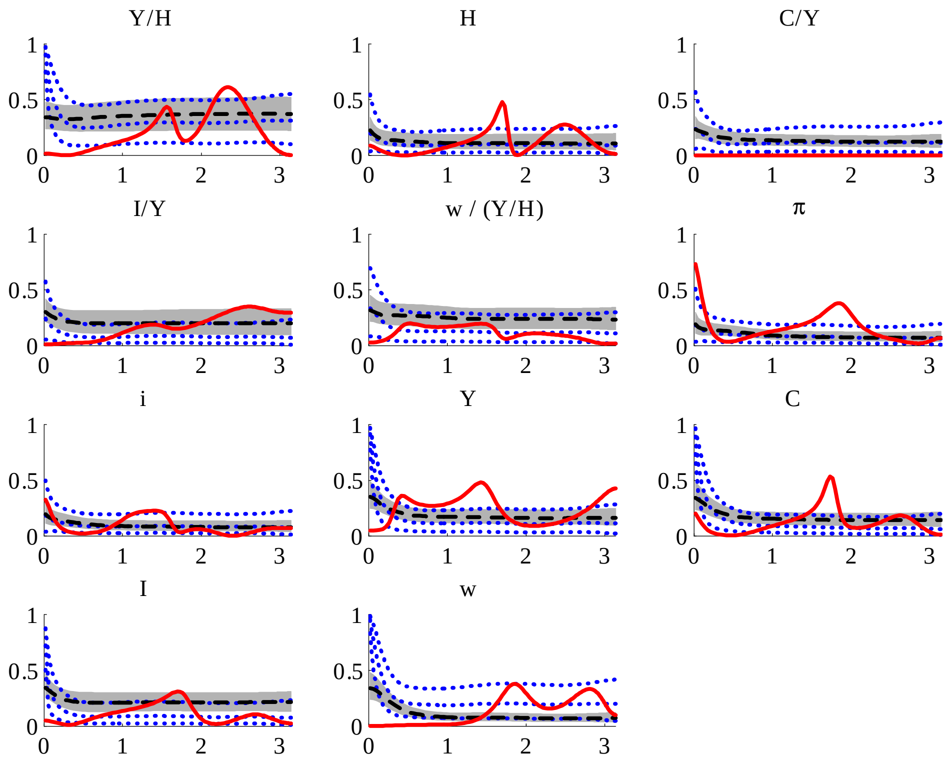

Figure 4: FEV share of productivity for the NK model at various frequencies.

True FEV (solid), Max Share (dashed), and LR (dotted) methods using simulated data using the benchmark parameterizations). Shaded region and dotted lines represent 68 percent coverages for Max Share and LR, respectively.

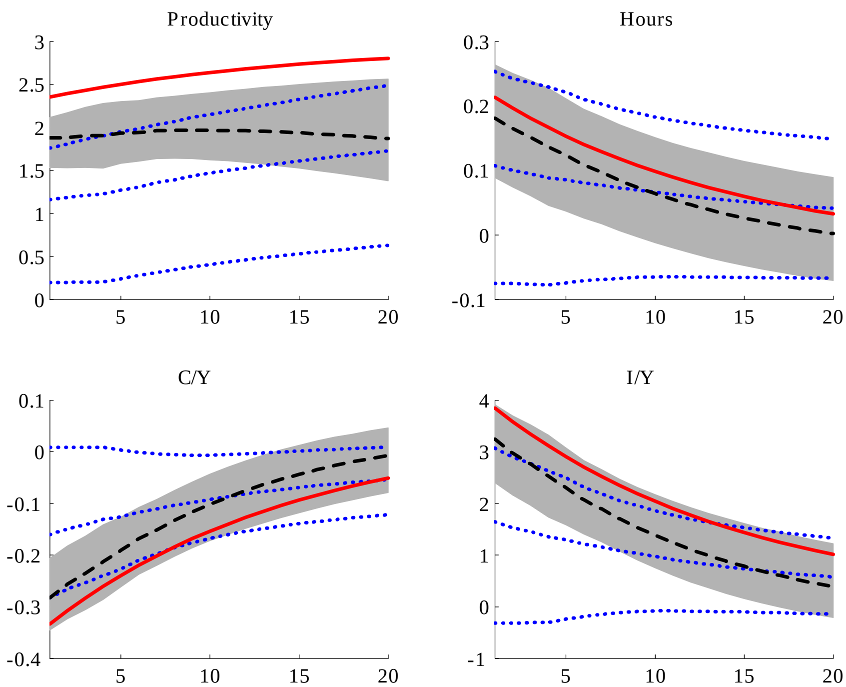

Figure 5: Impulse responses for the RBC model with persistent non-technology shocks.

Data is simulated with AR(1) technology coefficient $\rho_z = 1.0$, nontechnology 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 202 observations and 1,000 draws from the posterior distributions for the impulse responses.

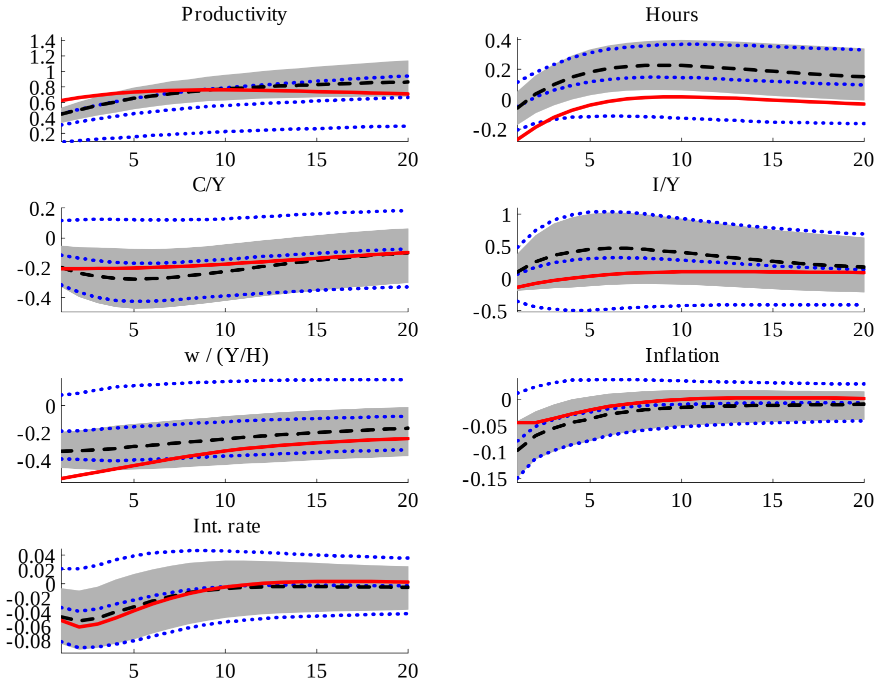

Figure 6: Impulse responses in the NK model with persistent nontechnology shocks. Theoretical responses with AR(1) technology coefficient, $\rho_z = 1.0$, and other nontechnology AR(1) coefficients calibrated to Smets and Wouters (2007a) are shown by thick solid lines. Median and 68 percent probability 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 202 observations and 1,000 draws from the posterior distributions for the impulse responses.

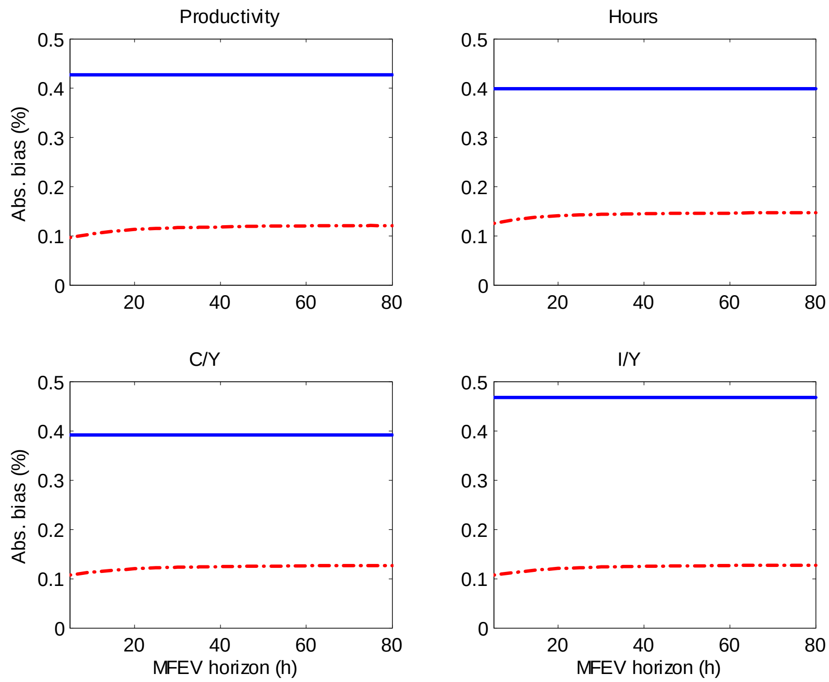

Figure 7: Average absolute bias in LR (solid) versus Max Share (dash-dotted) across alternative maximization horizons for data simulated from the RBC model.

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 202 observations.

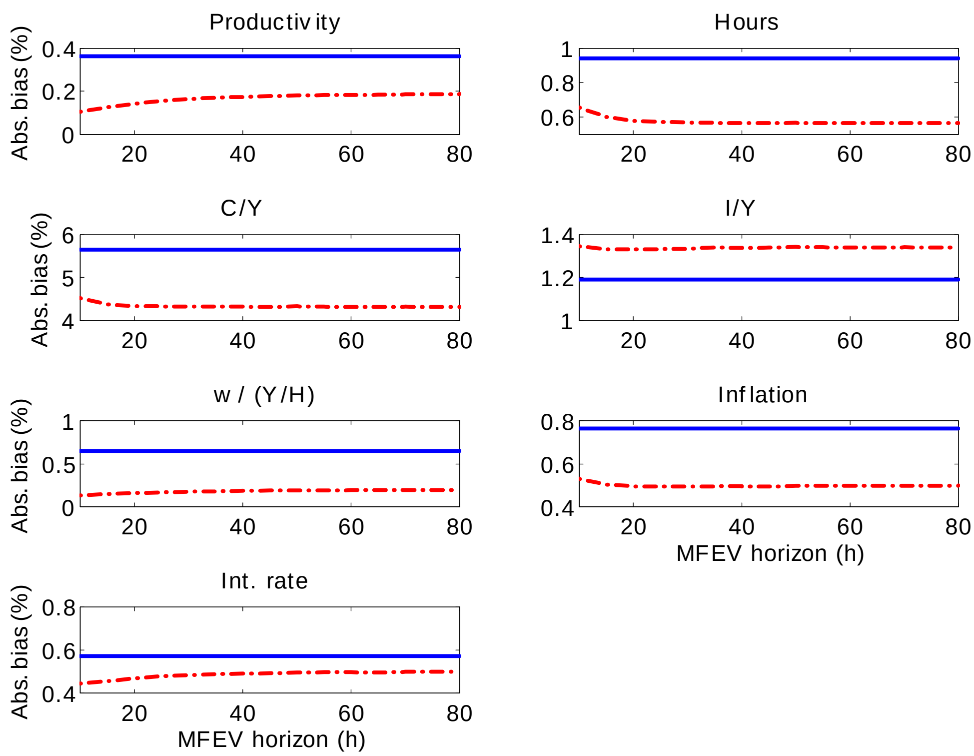

Figure 8: Average absolute bias in LR (solid) versus Max Share (dash-dotted) across alternative maximization horizons for data simulated from the NK model.

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 202 observations.

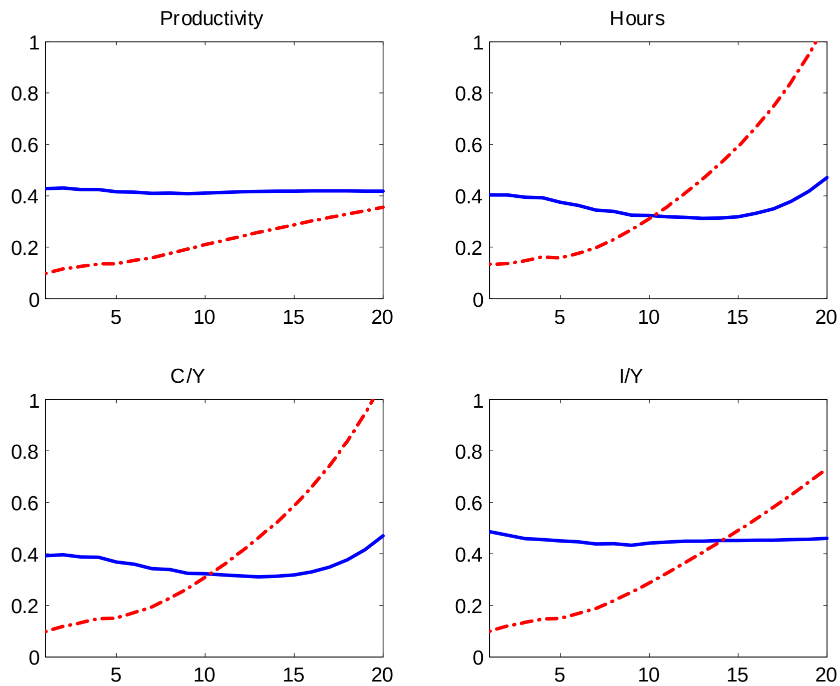

Figure 9: Average absolute bias for data simulated from the RBC model.

The bias is the average over 1000 MC iterations of the baseline parameterization and is computed as a percentage of the true value. The solid line is the bias for the LR identification. The dash-dot line is the bias for Max Share.

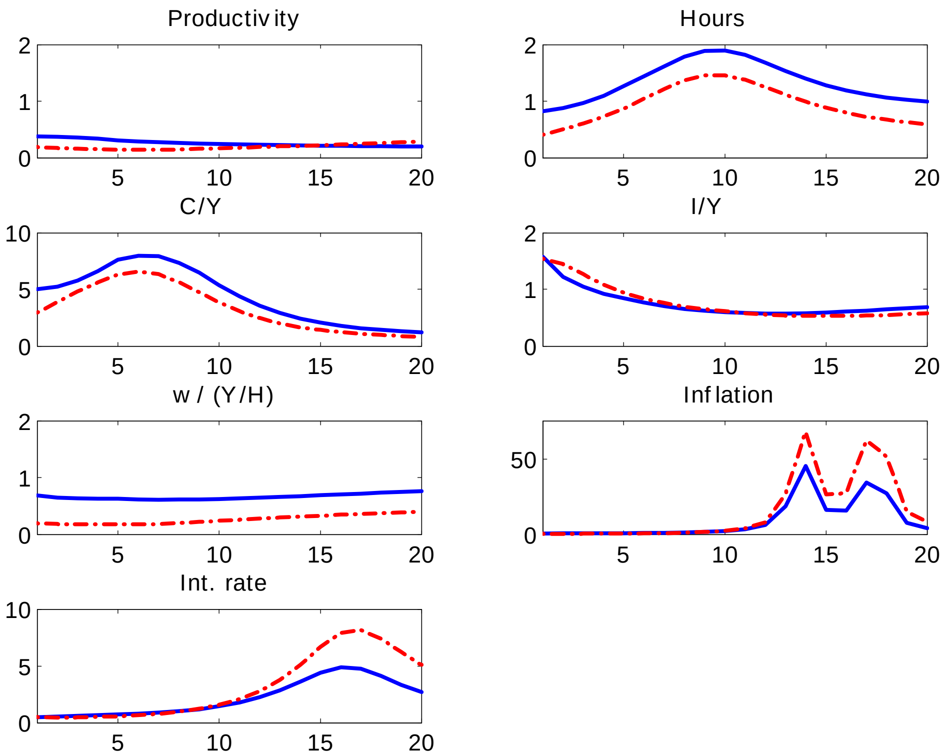

Figure 10: Average absolute bias for data simulated from the NK model.

The bias is the average over 1000 MC iterations of the baseline parameterization and is computed as a percentage of the true value. The solid line is the bias for the LR identification. The dash-dot line is the bias for Max Share.

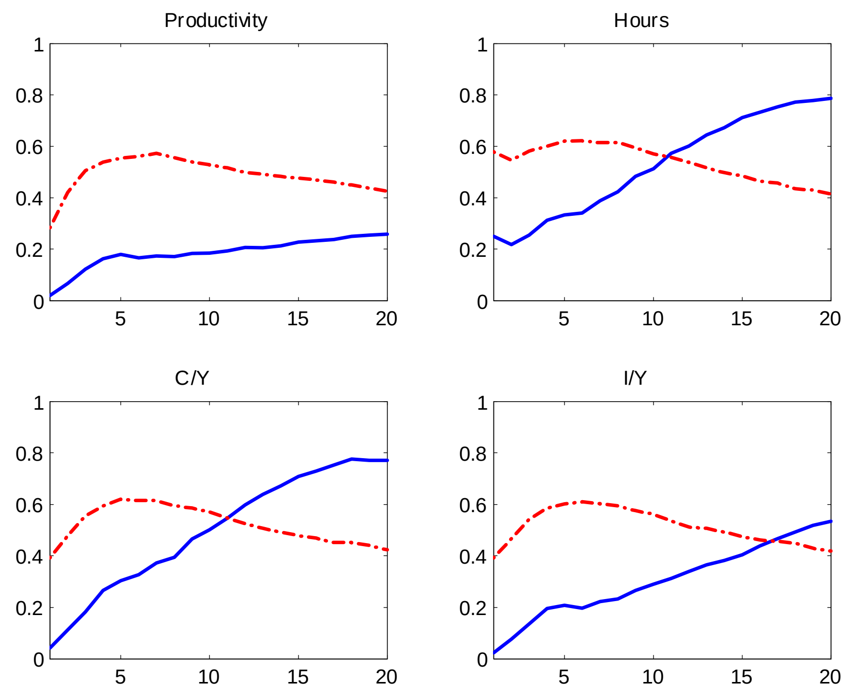

Figure 11: Small sample 68 percent coverage for the RBC model responses.

Percent of MC iterations for which the true response lies within the 68 percent nominal converage interval. The data is 1000 MC iterations generated from the baseline parameterization of the NK model. The solid line is for the LR identification; the dash-dot line is for Max Share.

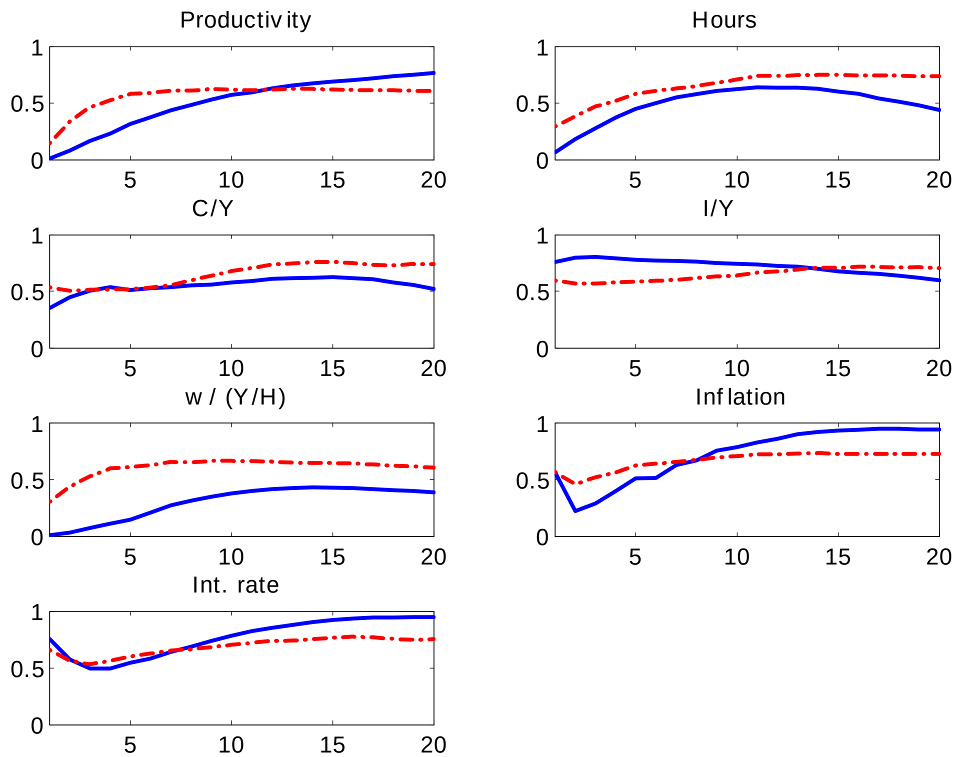

Figure 12: Small sample 68 percent coverage for the NK model responses.

Percent of MC iterations for which the true response lies within the 68 percent nominal converage interval. The data is 1000 MC iterations generated from the baseline parameterization of the NK model. The solid line is for the LR identification; the dash-dot line is for Max Share.

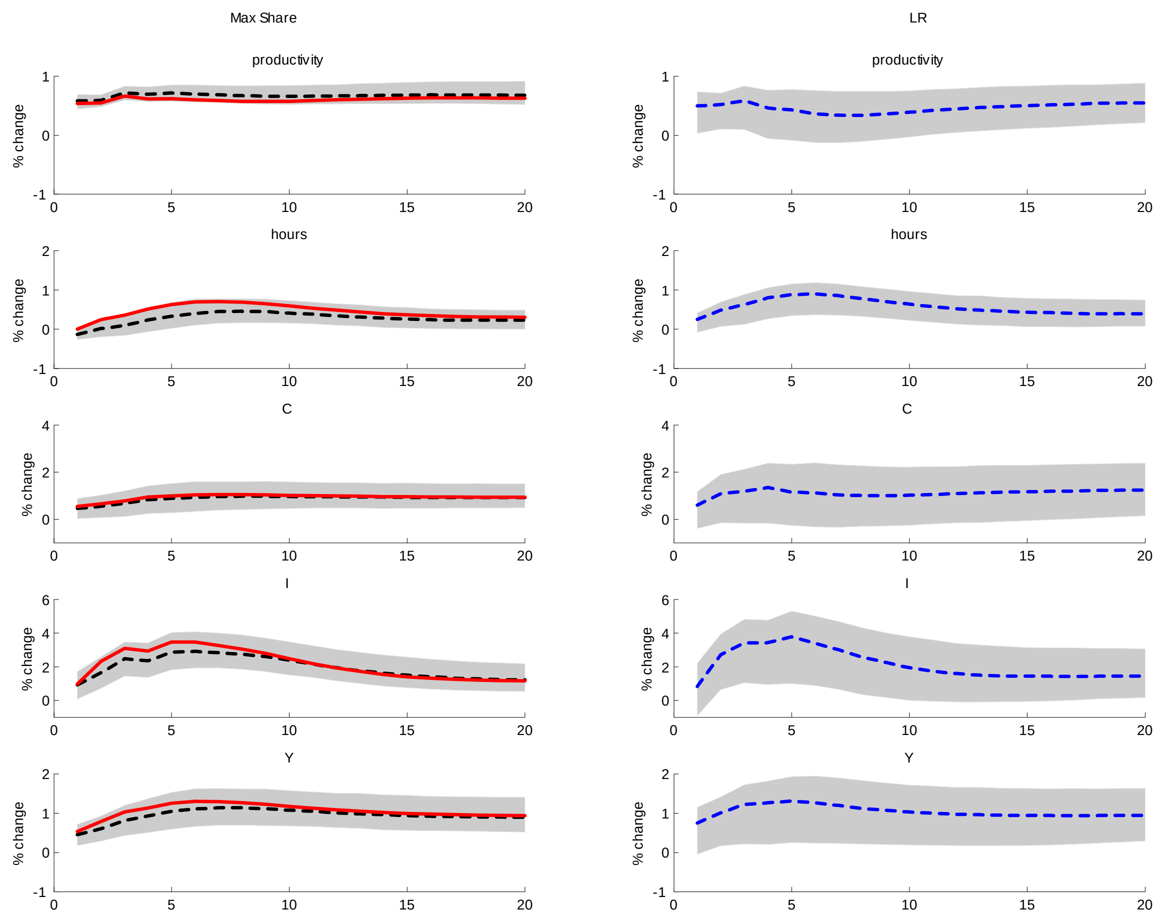

Figure 13: Impulse responses to a technology shock in the Max Share (on left) and the LR models (on right) using quarterly data1959:I - 2009:III.

Shaded areas are 68 percent Max Share probability intervals, as well as predicted responses when hours are restricted to be positive on impact in the Max Share model (thick solid line on left).

A Appendix

A.1 RBC 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.[23]

[23] 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

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 $\varepsilon_z \backsim i.i.d.N\left(0, \sigma_z^2\right)$. Table A.1 presents the sets of parameter values used to simulate the neoclassical growth model and the stick price model, respectively. Parameterizations for the neoclassical growth model are similar to those used by EGG (2005) in their benchmark model without capital utilization to match moments in U.S. data.

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.[24] A more detailed explanation of 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

$$\psi_t = p\psi_{t-1} + Q\xi_t, \tag{8}$$

$$y_t = W\psi_t + S\xi_t, \tag{9}$$

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,

[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.

[24] The models produce 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.

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 A.1.

Substituting (8) into (9) yields

$$y_t = pW\psi_{t-2} + WQ\xi_t + S\xi_t. \tag{10}$$

Realize that

$$W\psi_{t-1} = w_{t-1} - S\xi_{t-1}, \tag{11}$$

and substitute (11) into (10). Collecting terms yields:

$$y_t - py_{t-1} = S\xi_t + (WQ - pS)\xi_{t-1}. \tag{12}$$

We can rewrite this as:

$$S^{-1}y_t - (S^{-1}p)y_{t-1} = \xi_t + S^{-1}(WQ - pS)\xi_{t-1}, \tag{13}$$

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

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

A.2 NK Model

The New Keynesian (NK) model used in the Monte Carlo section is taken from Smets and Wouters (2007a).[25] The model is a standard-medium scale DSGE model which comprises intermediate and final goods sectors, labor unions and packers, utility maximizing households, and fiscal and monetary authorities. The model includes a number of rigidities. For example, prices are sticky, set prices via a partial indexation Calvo scheme; household utility is subject to external habit

[25] See the appendix of Smets and Wouters (2007b) for a more detailed description of the model.

Table A.1		
Parameter Values Used in Simulation (RBC Model)		
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
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
$\bar{\varsigma}_k$	steady-state capital tax rate	0.38
$\bar{\varsigma}_n$	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)

Table A.1: 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 A.2		
Parameter Values Used in Simulation (NK Model)		
ρ_a	persistence of technology shock	1.0
ρ_b	persistence of risk premium shock	0.18
ρ_g	persistence of government spending shock	0.6 (0.97)
ρ_i	persistence of investment-specific shock	0.6 (0.71)
ρ_r	persistence of monetary policy shock	0.12
ρ_p	persistence of price mark-up shock	0.6 (0.90)
ρ_w	persistence of wage mark-up shock	0.6 (0.97)
ρ_{ga}	feedback from technology to govt spending shock	0.01 (0.52)

Table A.2: Notes: The parameters are for the sticky price model outlined in Appendix 2. 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.

formation; and firms face an investment adjustment cost. Households supply labor to a union which differentiates the labor services, set wages *á la* Calvo (with partial indexation), and offers labor services to labor packers. The central bank follows a Taylor rule.

The model has seven shocks: neutral technology, investment-specific technology, risk premium, government spending, monetary policy, wage mark-up, price, mark-up, and feedback between the productivity shock to the fiscal policy shock. In the original Smets and Wouters' model, the price and wage mark-ups are ARMA, and all other shocks are AR(1). We augment the model for unit root neutral technology to conform with Assumption 1.

For our benchmark simulation, we calibrate the model to either the calibrated values or the modes of the posterior distributions from Smets and Wouters (2007a). Smets and Wouters' estimated values of the AR parameters for the shock processes may be highly persistent. We use two sets of parameters to calibrate the shock processes. These parameter values for the shock processes are shown in Table A.2. Values in parentheses are for the alternative parameterization reflecting more persistent nontechnology shocks.