What Happens After A Technology Shock?*

Lawrence J. Christiano[†], Martin Eichenbaum[‡] and Robert Vigfusson[§]

March 16, 2003

Abstract

This paper examines the economic impact of permanent shocks to technology. We argue that a positive technology shock drives hours worked up. This contrasts sharply with results in the literature. The evidence suggests that those results are an artifact of overdifferencing the hours worked data.

More generally, we find that positive technology shocks have qualitative effects that students of real business cycle models would anticipate: in addition to driving hours worked up, they also lead to increases in productivity, output, consumption and investment, while generating a decline in inflation. Nevertheless, we find that technology shocks play a relatively small role in business fluctuations. Their importance is greater in the lower frequency components of the data.

*Christiano and Eichenbaum thank the National Science Foundation for financial assistance. The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any person associated with the Federal Reserve System. We are grateful for discussions with Susanto Basu and Valerie Ramey.

[†]Northwestern University and NBER.

[‡]Northwestern University and NBER.

[§]Board of Governors of the Federal Reserve System. (Email robert.j.vigfusson@frb.gov)

1 Introduction

Standard real business cycle models imply that hours worked rise after a permanent shock to technology. Despite the a priori appeal of this prediction, there is a large and growing literature which argues it is inconsistent with the data. This literature uses reduced form time series methods in conjunction with minimal identifying assumptions that hold across large classes of models to estimate the actual effects of a technology shock. The results reported in this literature are important because they call into question basic properties of many structural business cycle models.

Consider for example the widely cited paper by Gali (1999). His basic identifying assumption is that innovations to technology are the only shocks that have an effect on the long run level of labor productivity. Gali (1999) reports that hours worked *fall* after a positive technology shock. The fall is so long and protracted that, according to his estimates, technology shocks are a source of negative correlation between output and hours worked. Because hours worked are in fact strongly procyclical, Gali concludes that some other shock or shocks must play the predominant role in business cycles with technology shocks best playing only a minor role. Moreover, he argues that standard real business cycle models shed little light on whatever small role technology shocks do play because they imply that hours worked rise after a positive technology shock. In effect, real business cycle models are doubly dammed: they address things that are unimportant, and they do it badly at that. Other recent papers reach conclusions that complement Gali's in various ways (see, e.g., Shea (1998), Basu, Kimball and Fernald (1999), and Francis and Ramey (2001).) In view of the important role attributed to technology shocks in business cycle analyses of the past two decades, Francis and Ramey perhaps do not overstate too much when they say (p.2) that Gali's argument is a '...potential paradigm shifter'.

Not surprisingly, the 'damning' result that hours worked fall after a positive technology shock has attracted a great deal of attention. Indeed, there is a growing literature aimed at constructing general equilibrium business cycle models that can account for this result. Gali (1999) and others have argued that the most natural explanation is based on sticky prices. Others, like Francis and Ramey (2001) and Vigfusson (2002), argue that this finding is consistent with real business cycle models, modified to allow for richer sets of preferences and technology, such as habit formation and investment adjustment costs.[1]

We do not build a model that can account for the 'damning' result. Instead we challenge the result itself. Using the same identifying assumption as Gali (1999), Gali, Lopez-Salido, and Valles (2002) and Francis and Ramey (2001) we re-examine the cyclical role played by technology shocks. Our main findings can be summarized as follows. First, the weight of the evidence is consistent with the view that a positive technology shock drives hours worked *up,* not down. In addition it leads to a rise in output, average productivity, investment, and consumption. Second, permanent technology shocks play a very small role in business cycle fluctuations. According to our estimates, they account for less than 10% of the cyclical variance of output, hours worked, investment, the federal funds rate, and consumption. Somewhat surprisingly, they play a larger role in the cyclical variation of inflation (32%).

[1] Other models which can account for the Gali (1999) finding are contained in Christiano and Todd (1996) and Boldrin, Christiano and Fisher (2001).

Third, we find that technology shocks play a much more important role in the low frequency component of aggregate fluctuations.

In sum, we find that a permanent shock to technology has qualitative consequences that a student of real business cycles would anticipate.[2] But, those shocks do not appear to be a major source of business cycle fluctuations. Instead, they are quantitatively important at frequencies of the data that a student of traditional growth models might anticipate.

Since we make the same fundamental identification assumption as Gali (1999), Gali, Lopez-Salido, and Valles (2002) and Francis and Ramey (2001), the key questions is: What accounts for the difference in our findings? By construction, the difference must be due to different maintained assumptions. As it turns out a key culprit is how we treat hours worked. If we assume, as do Gali and Francis and Ramey, that per capita hours worked is a difference stationary process, and work with the growth rate of hours (*the growth rate specification)*, then we too find that hours worked *falls* after a positive technology shock. But if we assume that per capita hours worked is a stationary process and work with the level of hours worked (*the level specification*), then we find the opposite: hours worked *rise* after a positive technology shock.

Classical hypothesis tests do not yield much information about which specification is correct. Standard tests cannot reject the null hypothesis that per capita hours worked are difference stationary. But standard tests also cannot reject the null hypothesis that hours worked are stationary. This is not surprising in light of the large literature that documents the difficulty of distinguishing between a difference stationary stochastic process and a persistent stationary process. Univariate classical tests are simply not very informative under these circumstances.

So we are faced with two alternatives: either hours worked rise after a positive technology shock, or they fall. Each answer is based on a different statistical model, each of which appears to be defensible on classical grounds. To judge between the competing specifications, we assess their relative plausibility. To this end, we ask 'which specification has an easier time explaining the observation that hours worked falls under one specification, and rises under the other?' Using this criterion, we find that the level specification is preferred.

We now discuss the results that lead to this conclusion. First, the level specification encompasses the growth rate specification. We show this by calculating what an analyst who adopts the growth rate specification would find if our estimated level specification were true. For reasons discussed below, by differencing hours worked this analyst commits a specification error. We find that such an analyst would, on average, infer that hours worked *fall* after a positive technology shock even though they *rise* in the true data generating process. Indeed the extent of this fall is very close to the actual decline in hours worked implied by the estimated growth rate specification. In addition, the level specification easily encompasses the impulse responses of the other relevant variables.

Second, the growth rate specification also encompasses the level specification results, but for very different reasons. Proceeding as above, we calculate what an analyst who adopts the level specification would find if our estimated growth rate specification was true. Here

[2] That the consequences of a technology shock resemble those in a real business cycle model may well reflect that the actual economy has various nominal frictions, and monetary policy has successfully mitigated those frictions. See Altig, Christiano, Eichenbaum and Linde (2002) for empirical evidence in favor of this interpretation.

the analyst is not committing a specification error. It is simply that a true restriction is not being imposed. At least in large samples, he will uncover the correct parameter values, including the unit root. So, if we abstract from small sample considerations, the estimated growth rate specification implies that an econometrician working with the level specification would find that hours worked fall with a positive technology shock. This suggests that the growth rate specification is unlikely to encompass our empirical level specification results.

Of course, in principle small sample considerations do matter because of finite-sample bias or sampling uncertainty. We find that the finite-sample bias is small: the analyst would on average conclude that hours worked fall for a substantial period of time after a positive technology shock, just as they do in the data generating process. At the same time, we find that sampling uncertainty implied by the growth rate specification is large. Indeed it is sufficiently large to enable the growth rate specification to encompass our level results.

To quantify the relative plausibility of the level and growth rate specifications, we compute the type of posterior odds ratio considered in Christiano and Ljungqvist (1988). The basic idea is that the more plausible of the two specifications is the one that has the easiest time simultaneously explaining the level and growth rate specification results. The specific 'results' that we focus on are the average percentage change in hours in the first six periods after a technology shock. This response is positive in one specification but negative in the other. Supposing that we have diffuse priors over the two specifications, the posterior odds ratio is roughly two to one in favor of the level specification. On this basis we conclude that the weight of the evidence favors the view that hours worked rise after a permanent technology shock.

In our view, the evidence in favor of the level model is even greater than what the previous paragraph suggests. Per capita hours worked is bounded above. As is well-known, this is difficult to reconcile with a difference stationary specification. In light of this, our prior in favor of the level specification is considerably higher than 50 percent. So, our posterior odds in favor of that specification are considerably higher than 2.

We assess the robustness of the level representation against alternative models of the low frequency component of per capita hours worked. In particular, we consider the possibility that there is a quadratic trend in hours worked. We show that there is a trend specification which has the implication that hours worked drops after a positive shock to technology. Using the methodology described above, we argue that the preponderance of the evidence favors the level specification relative to this alternative trend specification.

The remainder of this paper is organized as follows. Section 2 discusses our strategy for identifying the effects of a permanent shock to technology and presents the results from a bivariate analysis using data on hours worked and the growth rate of labor productivity. Later we show that on some dimensions inference is sensitive to only including two variables in the analysis. But the bivariate systems are useful because they allow us to highlight the basic issues in a simple setting and they allow us to compare our results to a subset of the results in the literature. Section 3 reports our encompassing results and the posterior odds ratio for the bivariate systems. In Section 4 we expand the analysis to include more variables. Here, we establish the benchmark system that we use later to assess the cyclical effects of technology shocks. Section 5 explores robustness of our analysis to the possible presence of deterministic trends. In addition, we examine the subsample stability of our time series model. In Section 6 we report our findings regarding the overall importance of technology shocks in cyclical fluctuations. Finally, Section 7 contains concluding remarks.

2 Identifying the Effects of a Permanent Technology Shock

In this section, we discuss our strategy for identifying the effects of permanent shocks to technology. We follow Gali (1999), Gali, Lopez-Salido, and Valles (2002) and Francis and Ramey (2001) and adopt the identifying assumption that the only type of shock which affects the long-run level of average productivity is a permanent shock to technology. This assumption is satisfied by a large class of standard business cycle models. See for example the real business cycle models in Christiano (1988), King, Plosser, Stock and Watson (1991) and Christiano and Eichenbaum (1992) which assume that technology shocks are a difference stationary process. The models have the property that technology shocks are the only disturbance that affects labor productivity in the long run.[3]

As discussed below, we use reduced form time series methods in conjunction with our identifying assumption to estimate the effects of a permanent shock to technology. An advantage of this approach is that we do not need to make all the usual assumptions required to construct Solow-residual based measures of technology shocks. Examples include corrections for labor hoarding, capital utilization, and time-varying markups.[4] Of course there exist models which do not satisfy our identifying assumption. For example, the assumption is not true in an endogenous growth model where *all* shocks affect productivity in the long run. Nor is it true in an otherwise standard model when there are permanent shocks to the tax rate on capital income. These caveats notwithstanding, we proceed as in the literature.

We estimate the dynamic effects of a technology shock using the method proposed in Shapiro and Watson (1988). The starting point of the approach is the relationship:

$$\Delta f_t = \mu + \beta(L)\Delta f_{t-1} + \tilde{\alpha}(L)X_t + \varepsilon_t^z. \tag{1}$$

Here f_t denotes the log of average labor productivity and $\tilde{\alpha}(L)$, $\beta(L)$ are polynomials of order q and $q-1$ in the lag operator, L, respectively. Also, Δ is the first difference operator and we assume that Δf_t is covariance stationary. The white noise random variable, ε_t^z, is the innovation to technology. Suppose that the response of X_t to an innovation in some non-technology shock, ε_t, is characterized by $X_t = \gamma(L)\varepsilon_t$, where $\gamma(L)$ is a polynomial in non-negative powers of L. We assume that each element of $\gamma(1)$ is non-zero. The assumption that non-technology shocks have no impact on f_t in the long run implies the following restriction on $\tilde{\alpha}(L)$:

$$\tilde{\alpha}(L) = \alpha(L)(1-L), \tag{2}$$

where $\alpha(L)$ is a polynomial of order $q-1$ in the lag operator. To see this, note first that the only way non-technology shocks can impact on f_t is by their effect on X_t, while the long-run

[3] If these models were modified to incorporate permanent shocks to agents' preferences for leisure or to government spending, these shocks would have no long run impact on labor productivity, because this is determined by the discount rate and underlying rate of growth of technology.

[4] See Basu, Fernald and Kimbal (1999) for an interesting application of this alternative approach.

impact of a shock to ε_t on f_t is given by:

$$\frac{\tilde{\alpha}(1)\gamma(1)}{1-\beta(1)}.$$

The assumption that Δf_t is covariance stationary guarantees $|1-\beta(1)| < \infty$. This, together with our assumption on $\gamma(L)$, implies that for the long-run impact of ε_t on f_t to be zero it must be that $\tilde{\alpha}(1) = 0$. This in turn is equivalent to (2).

Substituting (2) into (1) yields the relationship:

$$\Delta f_t = \mu + \beta(L)\Delta f_{t-1} + \alpha(L)\Delta X_t + \varepsilon_t^z. \tag{3}$$

We obtain an estimate of ε_t^z by using (3) in conjunction with estimates of μ, $\beta(L)$ and $\alpha(L)$. If one of the shocks driving X_t is ε_t^z,then X_t and ε_t^z will be correlated. So, we cannot estimate the parameters in $\beta(L)$ and $\alpha(L)$ by ordinary least squares. Instead, we apply the standard instrumental variables strategy used in the literature. In particular, we use as instruments Δf_{t-s} and X_{t-s}, $s = 1, 2, ...,q$.

Given an estimate of the shocks in (3), we obtain an estimate of the dynamic response of f_t and X_t to ε_t^z as follows. We begin by estimating the following q^{th} order vector autoregression:

$$Y_t = \alpha + B(L)Y_{t-1} + u_t,\ Eu_tu_t' = V, \tag{4}$$

where

$$Y_t = \begin{pmatrix} \Delta f_t \\ X_t \end{pmatrix},$$

and u_t is the one-step-ahead forecast error in Y_t. Also, V is a positive definite matrix. The parameters in this VAR, including V, can be estimated by ordinary least squares applied to each equation. In practice, we set $q = 4$. The fundamental economic shocks, e_t, are related to u_t by the following relation:

$$u_t = Ce_t,\ Ee_te_t' = I.$$

Without loss of generality, we suppose that ε_t^z is the first element of e_t. To compute the dynamic response of the variables in Y_t to ε_t^z, we require the first column of C. We obtain this by regressing u_t on ε_t^z by ordinary least squares. Finally, we simulate the dynamic response of Y_t to ε_t^z. For each lag in this response function, we computed the centered 95 percent Bayesian confidence interval using the approach for just-identified systems discussed in Doan (1992).[5]

3 Bivariate Results

This section reports results based on a simple, bivariate VAR in which f_t is the log of business labor productivity. The second element in Y_t is the log of hours worked in the business sector

[5]This approach requires drawing $B(L)$ and V repeatedly from their posterior distributions. Our results are based on 2500 draws.

divided by a measure of the population.[6] Our data on labor productivity growth and per capita hours worked are displayed in the first row of Figure 1. The properties of the hours worked data are discussed below.

Consider our results for the sample period, 1948Q1-2001Q4, the longest period for which data is available on the variables in our VAR. We refer to this as the long sample. The start of this sample period coincides with the one in Francis and Ramey (2001) and Gali (1999). The former work, as we do, with per capita hours worked, while Gali (1999) works with total hours worked. We prefer to work with per capita hours worked since this is the object that appears in most general equilibrium business cycle models. Since much of the business cycle literature works with post-1959 data, we also consider a second sample period given by 1959Q1-2001Q4. We refer to this as the short sample.

Panel A of Figure 2 displays the response of log output and log hours to a positive technology shock, based on the longer sample. A number of interesting results emerge here. First, the impact effect of the shock on output and hours is positive (1.17% percent and 0.34%, respectively) after which both rise in a hump shaped pattern. The responses of both output and hours are statistically significantly different from zero over the 20 quarters displayed. Second, in the long run, output rises by 1.33%. By construction the long run effect on hours worked is zero. Third, since output rises by more than hours, labor productivity also rises in response to a positive technology shock.

Panel B of Figure 2 displays the analog results for the shorter sample period. As before, the impact effect of the shock on output and hours is positive (0.94% and 0.14%, respectively), after which both rise in a hump-shaped pattern. The long run impact of the shock is to raise output by 0.96%. Again average productivity rises in response to the shock and there is no long run effect on hours worked. The rise in output is statistically different from zero at all horizons displayed. The rise in hours is statistically significantly different from zero between one and three years after the shock. So regardless of which sample period we use, the same picture emerges: a permanent shock to technology drives hours, output and average productivity up.

The previous results stand in sharp contrast to the literature according to which hours worked fall after a positive technology shock. The difference cannot be attributed to our identifying assumptions or the data that we use. To see this, note that we reproduce the bivariate-based results in the literature if we assume that X_t in (1) and (3) corresponds to the growth rate of hours worked, rather than the level of hours worked. The two panels in Figure 3 display the analog results to those in Figure 2, with this change in the definition of X_t.

According to the point estimates displayed in Panels A and B of Figure 3, a positive shock to technology induces a rise in output, but a persistent decline in hours worked.[7] Confidence intervals are clearly very large. Still, the initial decline in hours worked is statistically

[6] Our data were taken from the DRI Economics database. The mnemonic for business labor productivity is LBOUT. The mnemonic for business hours worked is LBMN. The business hours worked data were converted to per capita terms using a measure of the civilian population over the age of 16 (mnemonic, P16).

[7] For the long sample, the contemporaneous effect of the shock is to drive output up by 0.56% and hours down by 0.31%.The long run effect of the shock is to raise output by 0.84% and hours worked by 0.06%. For the short sample, the contemporaneous effect of the shock is to raise output 0.43% and reduce hours worked by 0.30%. The long run effect of the shock is to raise output by 0.74% and hours worked by 0.05%.

significant. This is consistent with the bivariate analysis in Gali (1999) and Francis and Ramey (2001).

The question is: Which results are more plausible, those based on the level specification or the growth rate specification? We turn to this question in the next section.

4 Analyzing the Bivariate Results

The previous section presented conflicting answers to the question: how do hours worked respond to a positive technology shock? Each answer is based on a different statistical model, corresponding to whether we assume that hours worked are a difference stationary or stationary in levels. To determine which answer is more plausible, we need to select between the underlying models (i.e., between the level and growth rate specifications). The first subsection below addresses the issue using classical diagnostic tests and shows that they do not convincingly discriminate between the competing models. The second and third sections approach the issue using encompassing methods.

4.1 Classical Diagnostic Tests

We begin by testing the null hypothesis of a unit root in hours worked using the Augmented Dickey Fuller (ADF) test. For both sample periods, this hypothesis cannot be rejected at the 10% significance level.[8] Evidently we cannot rule out the growth rate specification, at least based on this test. Of course it is well known that standard unit root tests have very poor power properties relative to the alternative that the time series in question is a persistent stationary stochastic process. So while it always true that failure to reject a null hypothesis does not mean we can reject the alternative, this caveat is particularly relevant in the present context.

To test the null hypothesis that per capita hours is a stationary stochastic process (with no time trend) we use the KPSS test (see Kwiatkowski et al (1992)).[9] For the short sample period, we cannot reject the null hypothesis at a 5% level of significance using standard asymptotic critical values.[10] For the long sample period, we can reject the null hypothesis at the 5% significance level standard asymptotic critical values. However, it is well known that the KPSS test (and close variants like the Leybourne and McCabe (1994) test) rejects the null hypothesis of stationarity too often if the data generating process is a persistent but stationary time series.[11] It is common practice to use size-corrected critical values that are constructed using data simulated from a particular data generating process.[12] We did so

[8]For the long and short sample, the ADF test statistic is equal to -2.46 and -2.49,respectively. The critical value corresponding to a 10% significance level is -2.57.

[9]In implementing this test we set the number of lags in our Newey-West estimator of the relevant covariance matrix to eight.

[10]The value of the KPSS test statistic is 0.4. The asymptotic critical values corresponding to 10% and 5% significance levels are 0.347 and 0.46,respectively.

[11]See Table 3 in Kwiatkowski et al. (1992) and also Caner and Kilian (1999) who provide a careful assessment of the size properties of the KPSS and Leybourne and McCabe tests.

[12]Caner and Kilian (1999) provide critical values relevant for the case in which the data generating process is a stationary AR(1) with an autocorrelation coefficient of 0.95. Using this value we fail to reject, at the

using the level specification VAR estimated over the long sample. Specifically, using this as the data generating process, we generated 1000 synthetic data sets, each of length equal to the number of observations in the long sample period, 1948-2001.[13] For each synthetic data set we constructed the KPSS test statistic. In 90% and 95% of the data sets, the KPSS test statistic was smaller than 1.89 and 2.06, respectively. The value of this statistic computed using the actual data over the period 1948-2001 is equal to 1.24. Thus we cannot reject the null hypothesis of stationarity at conventional significance levels.

4.2 The Impact of Different Specification Errors: A Priori Considerations

The preceding subsection showed that conventional classical methods are not useful for selecting between the level and growth specifications of our VAR. An alternative way to select between the competing specifications is to use an encompassing criterion. Under this criterion, a model must not just be defensible on classical diagnostic grounds. It must also be able to predict the results based on the opposing model. If one of the two views fails this encompassing test, the one that passes is to be preferred.

On a priori grounds, we have reason to believe it will be more difficult for the growth rate specification to encompass the level specification than for the level specification to encompass the growth rate specification. To understand why, it is necessary to think through several types of specification errors.

From the level perspective, first differencing is a mistake. To see this, recall that there are two steps involved in estimating the dynamic response of a variable to a technology shock. First, we must estimate the technology shock itself. Second we must estimate the dynamic response of the variable to that shock. The first step involves an instrumental variables regression, while the second step involves an estimated vector autoregression.

Consider the instrumental variables regression first. Suppose the econometrician proceeds under the mistaken assumption that hours worked is a difference stationary variable. Also, for simplicity suppose that the only variable in X_t is hours worked. The econometrician would difference X_t a second time, *i.e.*, work with $(1-L)\Delta X_t$, and then estimate μ and the coefficients in the finite-ordered polynomials, $\beta(L)$ and $\alpha(L)$, in the following system:

$$\Delta f_t = \mu + \beta(L)\Delta f_{t-1} + \alpha(L)(1-L)\Delta X_t + \varepsilon_t^z.$$

Suppose that in truth X_t is stationary with a spectral density that is different from zero at frequency zero. Then, in the true relationship the term involving X_t is actually $\bar{\alpha}(L)\Delta X_t$, where $\bar{\alpha}(L)$ is a finite ordered polynomial. In this case, the econometrician commits a specification error because the parameter space does not include the true parameter values. To see this, note that the only way $\alpha(L)(1-L)$ could ever be equal to $\bar{\alpha}(L)$ is if $\alpha(L)$ has a unit pole, i.e., if $\alpha(L) = \bar{\alpha}(L)/(1-L)$. But, this is impossible, since there is no finite lag

5% significance level, the null of stationarity over the longer sample period.

[13] The maximal eigenvalue of the estimated level specification VAR is equal to 0.972. We also estimated univariate AR(4) representations for hours worked using the synthetic data sets and calculated the maximal roots for the estimated univariate representations of hours worked. In no case did the maximal root exceed one. Furthermore, 95 percent of the simulations did not have a root greater than 0.982.

polynomial, $\alpha(L)$, with this property. So, imposing the restriction that X_t has a unit root when it does not entails specification error.

We now turn to the VAR used to estimate the response to a shock. An hours worked series that is first differenced when it should not be has a unit moving average root. It is well known that there does not exist a finite-lag vector autoregressive representation of such a process. So here too, the specification error is one that has consequences even in large samples.

These considerations are only suggestive about the possibility that the level specification may encompass the results obtained with the difference specification. For it to pass the encompassing test successfully, it must predict that the mistake made by falsely first-differencing hours worked specifically leads to the finding that hours fall after a positive technology shock. We explore this more in the next subsection.

Now consider the first difference perspective. What are the consequences failing to assume a unit root in hours worked, when there in fact is one? An econometrician who proceeds under this assumption specifies the following IV regression:

$$\Delta f_t = \mu + \beta(L)\Delta f_{t-1} + \alpha(L)\Delta X_t + \varepsilon_t^z,$$

and neglects to the restriction, $\alpha(1) = 0$, when it is in fact true. This is not a specification error, however, because the econometrician's parameter space does not rule out the true parameters. A similar point is true for the VAR: if there is a unit root in hours worked, and it is not imposed, this can be imposed after all if the VAR coefficients themselves contain the appropriate unit root (see Sims, Stock and Watson (1990)).

So, technically an econometrician who fails to treat hours worked as a difference-stationary series when in fact it is, does not commit a specification error. This suggests the possibility that the difference specification predicts, counterfactually, that an analyst working with the level specification will conclude that hours worked fall after a positive technology shock. That is, it suggests that the difference specification may not encompass the level result. This conclusion would be premature, however. The analyst who applies the econometric procedure implemented here, who treats hours worked as a stationary variable uses the lagged level of hours as an instrument for the change in hours worked. When hours worked contain a unit root, we expect this to be a poor instrument: the change in hours worked is driven by the relatively recent history of shocks while the level of hours worked is heavily influenced by shocks that occurred long ago, so that the correlation between $(1-L)h_t$ and h_{t-1} is likely to be small. The resulting 'weak instrument problem', together with other small sample issues might well have the implication that the difference specification encompasses the level specification after all. A full analysis, which takes into account small sample considerations, appears two subsections below.

Still, the considerations raised in this section lead us to suspect it is be easier for the level perspective to account for results based on first differencing than it is for the difference perspective to explain the level results. If the level perspective is true, then first differencing induces specification error. If the first difference result is true, working with levels does not entail any fundamental specification error. These considerations neglect a variety of small sample issues, which are taken up in the simulation experiments in the following two subsections.

4.3 Does the Level Specification Encompass the Growth Rate Results?

When we proceed using the growth rate specification, our point estimates imply that hours worked declines after a positive technology shock. To investigate the plausibility of the view that this decline reflects distortions due to over-differencing, we proceeded as follows. First, we generated two groups of one thousand artificial data sets from the estimated VAR in which the second element of Y_t is the log level of hours worked. In the first and second group, the VAR corresponds to the one estimated using the long and short sample period, respectively. So in each case the data generating mechanism corresponds to the estimated level specification. The number of observations in each artificial data set of the two groups is equal to the corresponding number of data points in the sample period.

In each artificial data sample, we proceeded under the (incorrect) assumption that the growth rate specification was true, estimated a bivariate VAR in which hours worked appears in growth rate form, and computed the impulse responses to a technology shock. The mean impulse responses appear as the thin line with circles in Figure 4. These correspond to the prediction of the level specification for the impulse responses that one would obtain with the (mispecified) growth rate specification. The lines with triangles are reproduced from Figure 3 and correspond to our point estimate of the relevant impulse response function generated from the growth rate specification. The gray area represents the 95% confidence interval of the simulated impulse response functions.[14]

From Figure 4 we see that, for both sample periods, the average the impulse response functions emerging from the 'mispecified' growth rate VAR are very close to the actual estimated impulse response generated using the growth rate specification. Notice in particular that hours worked are predicted to *fall* after a positive technology shock even though they *rise* in the actual data generating process. Evidently the specification error associated with imposing a unit root in hours worked is large enough to account for the estimated response of hours that emerges from the growth rate specification. That is, our level specification attributes the decline in hours in the estimated VAR with differenced hours to over-differencing. Note also that in all cases the estimated impulse response functions associated with the growth rate specification lie well within the 95% confidence interval of the simulated impulse response functions. We conclude that the level specification convincingly encompasses the growth rate specification.

4.4 Does the Growth Rate Specification Encompass the Level Results?

When we proceed using the level specification, our point estimates imply that hours worked rises after a positive technology shock. We now consider whether the growth rate specification can account for this finding. To do this we proceed as above except that we now take as the DGP the estimated VAR's in which hours appears in growth rates. Figure 5 reports the analog results to those displayed in Figure 4. The thick, solid lines, reproduced from Figure

[14] Confidence intervals were computed point wise as the average simulated response plus or minus 1.96 times the standard deviation of the simulated responses.

2, are the impulse response associated with the estimated level specification. The thin line with the triangles are reproduced from Figure 3 and correspond to our point estimate of the relevant impulse response function generated from the growth rate specification. The gray area represents the 95% confidence interval of the simulated impulse response functions.

The thin line in Figure 5 with circles are the mean impulse response functions associated with the mispecified VAR in which hours appears in levels. These correspond to the prediction of the growth rate specification for the impulse responses that one would obtain with the level specification. The gray area is the associated confidence interval. Notice that the thin lines with the triangles and the circles are very similar. Evidently the distortions associated with not imposing a unit root in hours worked are not very large. In particular, average hours are predicted to fall for a substantial period of time after a positive technology shock, just as they do in the data generating process, where hours worked do have a unit root. So the distortion from not imposing this unit root is not large enough to account, on average, for the level specification result that hours worked rises (see the thick black line). At the same time, there is a wide confidence interval about the thin line, which includes the thick, solid line. So, the growth rate specification could account for the level specification results as reflecting the effects of sampling uncertainty.

4.5 Quantifying the Relative Plausibility of the Two Specifications

The results of the previous two subsections indicate that the level specification can easily account for the qualitative features of the growth rate specification results, while the opposite is not true. However given the large sampling uncertainty associated with the growth rate specifications, definitive conclusions cannot be drawn based on those results. One way to quantify the relative plausibility of the two specifications is to compute the type of posterior odds ratio considered in Christiano and Ljungqvist (1988) for a similar situation where differences and levels of data lead to very different inferences.[15] The basic idea is that the more plausible of the two VAR's is the one that has the easiest time explaining the facts: the level specification implies that hours worked rise after a technology shock while the growth rate specification implies that hours worked falls.

To proceed, it is convenient to summarize the findings for hours worked in the form of a scalar statistic. We work with the average percentage change in hours in the first six periods after a technology shock. The level specification estimates, imply this change, μ_h,is equal to 0.89 and 0.55, for the long and short sample period, respectively. The analog statistic, $\mu_{\Delta h}$, in the two sample periods, for the growth rate specification, is -0.13 and -0.17,respectively.

To evaluate the relative ability of the level and growth specification to simultaneously account for μ_h and $\mu_{\Delta h}$ we proceed as follows. We simulated $1,000$ artificial data sets using each of our two estimated VARs as the data generating mechanism. In each data set, we calculated $(\mu_h, \mu_{\Delta h})$ using the same method used to compute these statistics in the actual data. The results based on the long and short sample periods are reported in Figure 6 and

[15] Eichenbaum and Singleton (1988) found, in a VAR analysis, that when they worked with first differences of variables, there was little evidence that monetary policy plays an important role in business cycles. However, when they worked with a trend stationary specification, monetary policy seems to play an important role in business cycles. Christiano and Ljungqvist argued that the preponderance of the evidence supported the trend stationary specification.

7, respectively. In all the figures, the horizontal axis corresponds to μ_h, while the vertical axis corresponds to $\mu_{\Delta h}$. Each dot represents a realization of $(\mu_h, \mu_{\Delta h})$ in the artificial data generated by the VAR indicated in the figure header. The point denoted by X indicates the empirical estimate of $(\mu_h, \mu_{\Delta h})$. The elliptical thin black line denotes a 90% confidence interval around these estimates.[16]

Notice that, for both sample periods, the level specification assigns more density to the bottom right hand quadrant than the growth rate specification. This is the quadrant that contains the actual estimated values of μ_h and $\mu_{\Delta h}$. To quantify the relative ability of the two specifications to account for the estimated values of $(\mu_h, \mu_{\Delta h})$, we computed the frequency of the joint event, $\mu_h > 0$ and $\mu_{\Delta h} < 0$. For the long sample period, the level and growth rate specifications imply that this frequency is 65.2 and 34.2, respectively. That is,

$$\begin{aligned} P(Q|A) &= 0.65 \\ P(Q|B) &= 0.34, \end{aligned}$$

where Q denotes the event, $\mu_h > 0$ and $\mu_{\Delta h} < 0$, A indicates the level specification, B indicates the growth rate specification and P denotes the percent of the impulse response functions in the artificial data sets in which $\mu_h > 0$ and $\mu_{\Delta h} < 0$. Suppose that our priors over A and B are equal: $P(A) = P(B) = 1/2$. The unconditional probability of Q, $P(Q)$, is $0.65 \times 0.5 + 0.34 \times 0.5 = 0.495$. The probability of the two models, conditional on having observed Q, is:

$$\begin{aligned} P(A|Q) &= \frac{P(A,Q)}{P(Q)} = \frac{P(Q|A)P(A)}{P(Q)} = 0.657 \\ P(B|Q) &= \frac{P(B,Q)}{P(Q)} = \frac{P(Q|B)P(B)}{P(Q)} = 0.343. \end{aligned}$$

So, we conclude that given these observations, the odds in favor of the level specification relative to the growth rate specification are 1.9 to 1.

Similar results emerge for the short sample period. Here the percent of impulse response functions in the bottom right hand quadrant is 52.4 in the artificial data generated by the level specification, while it is 25.6 for the growth rate specification. The implied values of $P(Q/A)$ and $P(Q/B)$ are 0.672 and .328. So, the odds in favor of the level specification relative to the growth rate specification are slightly larger than two to one.

The result that for both sample periods the odds favor the level specification by roughly two to one fundamentally reflects the fact that this model has an easier time explaining the estimated values of μ_h and $\mu_{\Delta h}$ than does the other model. On these purely statistical grounds we argue that the level specification model and its implications are more 'plausible'

[16] We computed the confidence set as follows. For a given VAR as the DGP, we have a 1,000 pairs of simulated responses. Denote them by $(\mu_h, \mu_{\Delta h})_j$ for j from 1,to 1000..Using these simulated pairs as a sample, it is straigthforward to calculate a sample average $\bar{\mu} = \frac{1}{1000}\sum_{j=1}^{1000}(\mu_h, \mu_{\Delta h})_j$ and a sample variance covariance matrix V $= \frac{1}{1000}\sum_{j=1}^{1000}\left((\mu_h, \mu_{\Delta h})_j - \bar{\mu}\right)'\left((\mu_h, \mu_{\Delta h})_j - \bar{\mu}\right)$. Using these values, one can then construct a 90% confidence set C as given by

$$C = \left\{ \begin{pmatrix} \mu_h \\ \mu_{\Delta h} \end{pmatrix} : \left(\begin{pmatrix} \mu_h \\ \mu_{\Delta h} \end{pmatrix} - \bar{\mu}' \right)' V^{-1} \left(\begin{pmatrix} \mu_h \\ \mu_{\Delta h} \end{pmatrix} - \bar{\mu}' \right) \leq 4.6 \right\}$$

where 4.6 is the 90% critical value of a chi square distribution with two degrees of freedom.

than those of the other VAR. Of course the odds in favor of the level specification would be even higher if we assigned more prior weight to the level specification. For reasons discussed in the introduction this seems quite natural to us. Our own prior is that the growth specification simply cannot be true because per capita hours worked are bounded.

5 Moving Beyond Bivariate Systems

In the previous two sections we analyzed the effects of a permanent technology shock using a bivariate system. In this section we extend our analysis to allow for a richer set of variables. We do so for two reasons. First, the response of the other variables that we consider are interesting in their own right. Second, there is no a priori reason to expect that the answers generated from small bivariate systems will survive in larger dimensional systems. If variables other than hours worked belong in the basic relationship governing the growth rate of productivity, and these are omitted from (1), then simple bivariate analysis will not generally yield consistent estimates of innovations to technology.

Our extended system adds four basic macroeconomic variables to the bivariate system analyzed above. These are the federal funds rate, the rate of inflation, the log of the ratio of nominal consumption expenditures and nominal GDP, and the log of the ratio of nominal investment expenditures to nominal GDP.[17] The last two variables correspond to the ratio of real investment and consumption, measured in units of output, to total real output. Standard models, including those that allow for investment specific technical change, imply these two variables are covariance stationary.[18] Data on our six variables are displayed in Figure 1.

Level and Growth Rate Specification Results

To conserve on space we focus on the 1959 - 2001 sample period.[19] Figure 8 reports the impulse response functions corresponding to the level specification, i.e. the system in which the log of per capita hours worked enters in levels. As can be seen, the basic qualitative results from the bivariate analysis regarding hours worked and output are unaffected: both rise in hump-shape patterns after a positive shock to technology shock.[20] The rise in output

[17] Our measures of the growth rate of labor productivity and hours worked are the same as in the bivariate system. We measured inflation using the growth rate of the GDP deflator, measured as the ratio of nominal output to real output (GDP/GDPQ). Consumption is measured as consumption on nondurables and services and government expenditures: (GCN+GCS+GGE). Investment is measured as expenditures on consumer durables and private investment: (GCD+GPI). The federal funds series corresponds to FYFF. All mnemonics refer to DRI's BASIC economics database.

[18] See for example Altig, Christiano, Eichenbaum and Linde (2002). This paper posits that investment specific technical change is trend stationary. See also Fisher (2003), which assumes investment specific technical change is difference stationary. Both frameworks imply that the consumption and investment ratios discussed in the text are stationary.

[19] Data on the federal funds rate is available starting only in 1954. We focus on the post 1959 results so that we can compare results to the bivariate analysis. We found that our 6 variable results were not sensitive to using data that starts in 1954.

[20] The contemporaneous effect of the shock is to drive output and hours worked up by 0.51%, and 0.11%, respectively. The long run effect of the shock is to raise output by 0.97%. By construction the shock has no effect on hours worked in the long run.

is statistically significant for roughly two years after the shock while the rise in hours worked is statistically significant at horizons roughly two to eight quarters after the shock.

Turning to the other variables in the system, we see that the technology shock leads to a prolonged, statistically significant fall in inflation and a statistically insignificant rise in the federal funds rate. Both consumption and investment rise, with a long run impact that is, by construction, equal to the long run rise in output.[21] The rise in consumption is estimated with much more precision than the rise in investment.

Figure 9 reports the impulse response functions corresponding to the growth rate specification, i.e. the system in which the log of per capital enters in first differences. Here a permanent shock to technology induces a long lived decline in hours worked, and a rise in output.[22] In the long run, the shock induces a 0.55% rise in output and a 0.25% decline in hours worked. Turning to the other variables, we see that the shock induces a rise in consumption and declines in the inflation rate and the federal funds rate. Investment initially falls but then starts to rise. Perhaps the key thing to note is that there is a great deal of sampling uncertainty associated with the point estimates. For the horizons displayed, none of the changes in hours worked, output, consumption, investment or the federal funds rate are statistically significant. The only changes that are significant are the declines in the inflation rate. Evidently, if one insists on the growth rate specification, the data is simply uninformative about the effect of a permanent technology shock on hours worked or anything else except the inflation rate.

Encompassing Results

We now turn to the question of whether the level specification can encompass the growth rate specification results. As with the bivariate systems, we proceeded as follows. First, we generated one thousand artificial data sets from the estimated six variable level specification VAR. The number of observations in each artificial data set is equal to the number of data points in the sample period, 1959 - 2001.

In each artificial data sample, we estimated a six variable VAR in which hours worked appears in growth rates and computed the impulse responses to a technology shock. The mean impulse responses appear as the thin line with circles in Figure 10. These responses correspond to the impulse responses that would result from the growth rate specification VAR being estimated on data generated from the level specification VAR. The thin lines with triangles are reproduced from Figure 9 and correspond to our point estimate of the relevant impulse response function generated from the growth rate specification. The gray area represents the 95 confidence interval of the simulated impulse response functions.[23] The thick black line corresponds to the impulse response function from the estimated 6 variable level specification VAR.

[21] The contemporaneous effect of the shock is to drive consumption and investment up by 0.42% and 0.90%, respectively. The long run effect of the shock is to raise consumption and investment by 0.97%.

[22] The contemporaneous effect of the shock is to drive output up by 0.12% and hours worked down by -0.27%.

[23] These confidence intervals are computed as in the same manner as the intervals reported for the bivariate encompassing tests. The interval is the average simulated impulse response plus or minus 1.96 times the standard deviation of the simulated impulse responses.

The average impulse response function emerging from the 'misspecified' growth rate specification is very close to the actual estimated impulse response generated using the growth rate specification. As in the bivariate analysis, hours worked are predicted to *fall* after a positive technology shock even though they *rise* in the actual data generating process. Also, in all cases the estimated impulse response functions associated with the growth rate specification lie well within the 95% confidence interval of the simulated impulse response functions. So, as before, we conclude that the specification error associated with imposing a unit root in hours worked is large enough to account for the estimated response of hours that emerges from the growth rate specification.

We now consider whether the growth rate specification can encompass the level specification results. To do this we proceed as above except that we now take as the data generating process the estimated VARs in which hours appears in growth rates. Figure 11 reports the analog results to those displayed in Figure 10. The thick, solid lines, reproduced from Figure 8, are the impulse response functions associated with the estimated level specification. The thin line with the triangles are reproduced from Figure 9 and correspond to our point estimate of the impulse response function generated from the growth rate specification. The gray area represents the 95 confidence interval of the simulated impulse response functions.

The thin line in Figure 11 with circles is the mean impulse response function associated with estimating the level specification VAR on data generated using, as the DGP, the growth rate specification VAR. Notice that the lines with triangles and circles are very similar. So, focusing on point estimates alone, the difference specification is not able to account for the actual finding with our estimated level VAR that hours worked rise. Still, in the end the difference specification is compatible with our level results because it predicts so much sampling uncertainty.

The Relative Plausibility of the Two Specifications

As in the bivariate system, we quantify the relative plausibility of the level and growth rate specifications with a scalar statistic: the average percentage change in hours in the first six periods after a technology shock. The estimated level specification implies this change, μ_h, is equal to 0.31. The statistic for the growth rate specification, $\mu_{\Delta h}$, is -0.29.

We simulated 1,000 artificial data sets using each of our two estimated VARs as data generating mechanisms. In each data set, we calculated $(\mu_h, \mu_{\Delta h})$ using the same method used to compute these statistics in the actual data. The results are reported in the Figure 12. In each panel, the header indicates the underlying data generating mechanism. Each dot represents a realization of $(\mu_h, \mu_{\Delta h})$ in artificial data generated by the specification indicated in the figure header. The point denoted by X indicates the empirical estimate of $(\mu_h, \mu_{\Delta h})$, and elliptical thin black line denotes a 90% confidence interval around these estimates.

For both samples, we see that the level specification assigns more density to the bottom right hand quadrant than the growth rate specification. This is the quadrant containing the actual estimated values of μ_h and $\mu_{\Delta h}$. Using each of our two time series representations, we computed the frequency of the joint event, $\mu_h > 0$ and $\mu_{\Delta h} < 0$. This frequency is 66.7 across artificial data sets generated by the level specification, while it is 36.7 in the case of the growth rate specification. The implied odds in favor of the level specification over the growth rate specification are 1.8 to one. So as with our bivariate systems, we conclude on these purely statistical grounds that the level specification model and its implications are more 'plausible' than those of the growth rate specification.

6 Sensitivity Analysis

In this section we investigate the sensitivity of our analysis along three dimensions: the choice of variables to include in the analysis, allowing for deterministic trends and subsample stability.

6.1 Sensitivity to Choice of Variables

While the qualitative effects of a permanent shock to technology are robust across the bivariate and six variable systems, the quantitative effects are quite different. One way to see this is to compare the relevant impulse response functions (see Figures 2 and 8). A different way to do this is to assess the importance of technology shocks in accounting for aggregate fluctuations using the bivariate and six variables systems. In the next section, we show that technology shocks are much less important in the larger system.

To help us analyze the sources of this sensitivity, we now briefly report results from two four variable systems. In the first, the CI system, we add two variables to the benchmark bivariate system: the ratio of consumption expenditures to nominal GDP and the ratio of investment expenditures to nominal GDP. In the second, the $R\pi$ system, we add the Federal Funds rate and the inflation rate to the benchmark bivariate system.

Figure 13 reports the point estimates of the impulse response functions from the level specification six variable system (depicted by the thick line), the CI system (depicted by the line with '*') and the $R\pi$ system (depicted by the line with 'X'). Two results are worth noting. First, the six variable and the CI systems generate very similar results for the variables that are included in both. Second, the six variable and the $R\pi$ systems generate qualitatively different responses of hours worked. In both the six variable and the CI systems, the impact effect of a positive technology shock on hours worked is positive after which they continue to rise in a hump shaped pattern. But in the $R\pi$ system, hours worked falls for roughly 3 quarters after a positive technology shock.

The most natural interpretation of this result is specification error. Both the CI and $R\pi$ systems are mispecified relative to the six variable system. But the quantitative effect of the specification error associated with omitting consumption and investment from the analysis (the $R\pi$ system) is sufficiently large to affect qualitative inference about the effect of a technology shock on hours worked. Of course, if the six variable system is specified correctly, it should be able to rationalize the response of hours worked in the $R\pi$.

To see if this is the case, we proceeded as follows. First, we generated one thousand artificial data sets from the estimated 6 variable VAR. The number of observations in each artificial data set is equal to the number of data points in the short sample period. In each artificial data sample, we estimated a VAR for the four variable $R\pi$ system and computed the impulse responses to a technology shock. The mean impulse responses appear as the thin line with circles in Figure 14. These correspond to the prediction of the six variable VAR for the impulse responses one obtains using the $R\pi$ system VAR. The thin line with the 'X' are reproduced from Figure 13 and correspond to our point estimate of the relevant impulse response function generated from the $R\pi$ system. The gray area represents the 95 confidence interval of the simulated impulse response functions. The thick black line corresponds to the impulse response function from the estimated 6 variable VAR.

Note that the average impulse response functions emerging from the 'mispecified' $R\pi$ system are very close to the estimated impulse responses generated using the actual $R\pi$ system. So the specification error associated with omitting consumption and investment is large enough to account for the estimated response of hours that emerges from the $R\pi$ specification. In all cases the estimated impulse response functions associated with the mispecified $R\pi$ specification lie well within the 95% confidence interval of the simulated impulse response functions.[24]

We conclude that it is important to include at least C and I in our analysis. While it may be desirable to include R and π on a priori grounds, the results of central interest here seem to be less sensitive to omitting them.

6.2 Quadratic Trends

From Figure 1 we see that per capita hours worked seem to follow a U shaped pattern. This suggests the possibility that hours worked may be stationary around a quadratic trend. If so, then the systems considered above are mispecified and may generate misleading results. With this in mind, we investigate two issues. First, is the response of hours worked to a technology shock sensitive to imposing a quadratic trend in hours worked? Second, to the extent that the results are sensitive, which set of results are most plausible?

We begin by redoing our analysis of the six variable system with two types of quadratic trends. In case (i), we allow for a quadratic trend in all the variables of the VAR. This seems natural since other variables like inflation and the interest rate also exhibit U shaped behavior (see Figure 1). In case (ii), we allow for a quadratic trend only in per capita hours worked. Except for these trends the other variables enter the system as in the level specification. Figure 15 reports our results. The dark, thick lines correspond to the impulse response functions implied by the six variable level specification. The lines indicated with 0's and x's correspond to the impulse response functions generated from this system modified as described in (i) and (ii) above. The grey area is the 95 percent confidence interval associated with the lines indicated with x's. We report only this confidence interval, rather than all three, in order to give some sense of sampling uncertainty while keeping the figure relatively simple.

Three things are worth noting. First, if we allow for a quadratic trend in all of the variables in the VAR, after a small initial fall, hours worked rise as in the level specification in response to a positive technology shock. Second, if we allow for a quadratic trend only in hours worked, then hours worked do in fact fall in a persistent way after a positive shock to technology. Third, in either case, the impulse response function of hours worked is estimated

[24] For completeness, we repeated the analysis for the systems in which hours enter in growth rates. Again, the six variable and the CI system are more similar to each other than the $R\pi$ system. However, the response of consumption is much smaller in the CI system than in the six variable system. Finally, we computed the analog results to those in Figure 14 and again found that the 6 variable system can encompass the CI growth rate system.

with very little precision. One cannot reject the views that hours worked rise, fall or do not change. If one insists on allowing for quadratic trends, then there is simply very little information in the data about the response of hours worked to a technology shock.

Still, focusing on the point estimates alone, the estimated response of hours worked to a technology shock is sensitive to whether or not we include a quadratic trend in hours worked. We now turn to the question of which results are more plausible: those based on our 6-variable level model, or those based on the quadratic trend models.

We begin by performing a classical test of the null hypothesis that there is no trend in per capita hours worked. Specifically, we regress the log of per capita hours worked on a constant, time and time-squared. We then compute the t statistic for the time-squared term allowing for serial correlation in the error term of the regression using the standard Newey-West procedure.[25] The resulting t statistic is equal to 8.13. Under standard asymptotic distribution theory, this has probability value of essentially zero under the null hypothesis that the coefficient on the time-square term is zero. So, on the basis of this test, we would reject our level model. But, it is well-kown that the asymptotic theory for this t statistic is quite poor in small samples, especially when the error terms exhibit high degrees of serial correlation. This is exactly the situation we are in according to our level model, since its eigenvalues are quite large.[26] To address this concern, we adopt the following procedure. We simulate $1,000$ synthetic time series on per capita hours worked using our estimated level model. The disturbances used in these simulations were randomly drawn from the fitted residuals of our estimated level model. The length of each synthetic time series is equal to the length of our sample period. We found that 13.3 percent of these t statistics exceed 8.13. So, from the perspective of the level model, a t statistic of 8.13 is not particularly unusual. We conclude that our t test fails to reject the null hypothesis that the coefficient on the time-squared term is equal to zero.

This result may at first seem surprising in view of the U shape of the per capita hours worked data in Figure 1. Actually, such shapes are at all not unusual in a time series system with eigenvalues that are close to unity. This is why the apparent evidence of a U-shape trend in the hours data is not evidence against our level model.

Evidently classical methods cannot be used to convincingly discriminate between the level model and the quadratic trend model. We now turn to the encompassing and posterior odds approach.

Encompassing Results

Appendix A discusses our encompassing results. In discussing our results we refer to the two quadratic trend models as the Trend in All Equations and the Trend in Hours Only models. Our main results can be summarized as follows. The Level model easily accounts for the results obtained using the two quadratic trend models. This is true even if we focus on point estimates alone. In particular, the Level model successfully accounts for the fact that one quadratic trend model implies a fall in hours after a technology shock, while the other implies a rise. The encompassing result is even stronger when we take sampling uncertainty into account.

[25] We allow for serial correlation of order 12 in the Newey-West procedure.

[26] The largest two eigenvalues of the determinant of $[I - B(L)]$ in (4) are 0.9903 and 0.9126.

Focusing on the point estimates alone, the Trend in Hours Worked model is unable to encompass the results of either of the other two models. Specifically, it cannot account for the fact that hours worked rise in each of the other two models. However, once sampling uncertainty is taken into account, this encompassing test also does not reject the Trend in Hours Only model.

Two things are worth noting regarding the Trend in All Equations model. First, focusing on the point estimates alone, this model can encompass the results based on the Trend in Hours Only model. But, it does not encompass the results based on the Level model. In particular, the Trend in All Equations model predicts, counterfactually, that the Level model produces a fall in hours worked after a positive technology shock. Second, even when sampling uncertainty is taken into account, encompassing test rejects the Trend in All Equations model vis a vis the Level model.

We conclude that the encompassing analysis allows us to exclude the Trend in All Equations model. However, it does not allow us to discriminate between the Level and the Trend in Hours Only model. With this motivation, we turn to the posterior odds ratio.

The Relative Plausibility of the Two Specifications

We quantify the relative plausibility of the three models with a scalar statistic: the average percentage change in hours in the first six periods after a technology shock. The estimated Level, Trend in All Equations, and Trend in Hours Only model imply this change is equal to $\mu_1 = 0.31$, $\mu_2 = -0.12$, and $\mu_3 = 0.16$, respectively.

We simulated 1,000 artificial data sets using each of our three estimated VARs as data generating mechanisms. In each artificial data set, we calculated (μ_1, μ_2, μ_3) using the same method used to compute these statistics in the actual data. For each data generating mechanism, we computed the frequency of the joint event, $\mu_1, \mu_2 > 0$, $\mu_3 < 0$. This frequency is 19.30, 3.50, 5.60 for the Level, Trend in All Equations, and Trend in Hours Only models, respectively. So the posterior odds in favor of the Level model relative to the Trend in All Equations and Trend in Hours Only model is roughly 5.5 and 3.4, respectively. On this basis, we conclude that the Level model and its implications are more 'plausible' than those of the two quadratic trend models.

6.3 Subsample Stability

In this subsection we briefly discuss the issue of subsample stability. Authors such as Gali, Lopez-Salido, and Valles (2002) among others have argued that monetary policy may have changed after 1982, and that this resulted in a structural change in VAR's. Throughout our analysis, we have assumed implicitly that there has been no structural change. This section reports a test of this null hypothesis. Consistent with the results in Christiano, Eichenbaum and Evans (1999), we find that we cannot reject it.

We proceed as follows. Define the following unrestricted VAR:

$$y_t = B(L)y_{t-1} + \gamma D(L)y_{t-1} + \varepsilon_t^u$$

where γ equals one only if t is greater than 1982Q4. Define the following restricted VAR:

$$y_t = B(L)y_{t-1} + \varepsilon_t^r.$$

Consider the following likelihood ratio statistic:

$$LR = T\left(\log(|\Sigma_u|) - \log(|\Sigma_r|)\right),$$

where Σ_u and Σ_r are the variance-covariance matrices of the fitted residuals in the unrestricted and restricted VAR's respectively. Under the null hypothesis, $\gamma = 0$, the asymptotic distribution of the LR statistic chi-squared with the degrees of freedom equal to the number of variables squared times the number of lags in the VAR. In addition to using this distribution theory, we also calculated the distribution of the test statistic using the Bootstrap procedure in which the restricted model was used as the data generating process. We did the analysis both for the level and difference VAR's. The results are summarized below:

	Test Statistic		95% Critical Value		
			Classical	Bootstrap	
Two Variable	Levels	Growth Rates	DF 16	Levels	Growth Rates
1948-2001	9.02	12.93	26.30	38.74	38.33
1959-2001	13.12	17.73	26.30	40.82	35.74
Six Variable			DF 144		
1959-2001	233.44	227.26	173.00	276.33	276.33

The key things to note here are as follows. First, in the bivariate systems, regardless of the distribution theory used, there is no evidence against the null hypothesis of structural stability, $\gamma = 0$. In the six variable system, there is evidence against this hypothesis if we use classical asymptotic distribution theory. However, this evidence is overturned when we apply the Bootstrap critical values. We conclude that the evidence against the null hypothesis is at best marginal.

7 How Important Are Permanent Technology Shocks for Aggregate Fluctuations?

In Section 4and Section 5, we argued that the weight of the evidence favors the level specification relative to the growth rate specification. According to the level specification, a permanent technology shock causes a hump shaped rise in output. So in principle these shocks can account for the strong cyclical positive comovement between hours worked and output. So one cannot dismiss technology shocks as the key impulses to the cycle using Gali's (1999) argument. Of course this doesn't mean that technology shocks actually are important at business cycle frequencies. In this section we investigate their importance in two ways. First, we examine the contribution of technology shocks to the cyclical variance of various aggregates. Second, we present a series of simulations that display how these aggregates would have evolved had technology shocks been the only source of aggregate fluctuations. Based on these exercises we conclude that (i) technology shocks are not particularly important at business cycle frequencies but they do play an important role at relatively low frequencies of the data, and (ii) inference based on bivariate systems greatly overstates the cyclical importance of technology shocks

7.1 Bivariate System Results

We begin by discussing the role of technology shocks in the variability of output and hours worked based on our level specification bivariate VAR. Table 1 reports the percent of forecast error variance due to technology shocks, at horizons of 1, 4, 8, 12, 20 and 50 quarters. By construction, permanent technology shocks account for all of forecast error variance of output at the infinite horizon. Notice that technology shocks account for an important fraction of the variance of output at all reported horizons. For example, they account for roughly 80% of the one step ahead forecast error variance in output. In contrast, they account for only a small percent of the one step forecast error variance in hours worked (4.5%). But they account for a larger percent of the forecast error variance in hours worked at longer horizons, with the percent exceeding 40% at horizons greater than 2 years.

The first row of Table 3, reports the percent of the variance in output and hours worked at business cycle frequencies due to technology shocks. This statistic was computed as follows. First we simulated the estimated level specification bivariate VAR driven only by the estimated technology shocks. Next we computed the variance of the simulated data after applying the Hodrick-Prescott (HP) filter. Finally we computed the variance of the actual HP filtered output and hours worked. For any given variable, the ratio of the two variances is our estimate of the fraction of business cycle variation in that variable due to technology shocks. The results in Table 3 indicates that technology shocks appear to play an significant role for both output and hours worked, accounting for roughly 64% and 33% of the cyclical variance in these two variables, respectively.

A different way to assess the role of technology shocks is presented in Figure 15. The thick line in this figure displays a simulation of the 'detrended' historical data. The detrending is achieved using the following procedure. First, we simulated the estimated reduced form representation (4) using the fitted disturbances, $\hat{u}_t$, but setting the constant term, α, and the initial conditions of Y_t to zero. In effect, this gives us a version of the data, Y_t, in which any dynamic effects from unusual initial conditions (relative to the VAR's stochastic steady state) have been removed, and in which the constant term has been removed. Second, the resulting 'detrended' historical observations on Y_t are then transformed appropriately to produce the variables reported in the top panel of Figure 15. The high degree of persistence observed in output reflects that our procedure for computing output makes it the realization of a random walk with no drift.

The procedure used to compute the thick line in Figure 15 was then repeated, with one change, to produce the thin line. Rather than using the historical reduced form shocks, $\hat{u}_t$, the simulations underlying the thin line use $C\hat{e}_t$, allowing only the first element of $\hat{e}_t$ to be non-zero. This first element of $\hat{e}_t$ is the estimated technology shock ε_t^z, obtained from 3 The results in the top panel of Figure 15 give a visual representation of what is evident in Table 1 and the first row of Table 3. Technology shocks appear to play a very important role in accounting for fluctuations in output and a smaller, but still substantial role with respect to hours worked.

We conclude this section by briefly noting the sensitivity of inference to whether we adopt the level or growth rate specification. The bottom panels of Tables 1 - 3 and the bottom panel of Figure 15 report the analog results to those in the top panels of Tables 1- 3 and the top panel of Figure 15 for the bivariate growth rate specification. Comparing across the Tables or the Figures the same picture emerges: with the growth rate specification, technology shocks play a much smaller role with respect to output and hours worked than they do in the levels specification. For example, the percent of the cyclical variance in output and hours worked accounted for by technology shocks drops from 64% and 33% in the level specification to 11% and 4% in the growth rate specification. So imposing a unit root in hours worked, not only affects qualitative inference about the effect of technology shocks, it also affects inference about their overall importance.

7.2 Results Based on the Larger VAR

We now consider the importance of technology shocks when we incorporate additional variables into our analysis. Table 2 reports the variance decomposition results for the six variable level specification system. Comparing the first two rows of Table 1 and 2, we see that technology shocks account for a much smaller percent of the forecast error variance in both hours and output in the six variable system. For example, in the bivariate system, technology shocks account for roughly 78% and 24% of the 4 quarter ahead forecast error variance in output and hours, respectively. In the six variable system these percentages fall to 40% and 15% respectively. Still technology shocks continue to play a major role in the variability of output, accounting for over 40% of the forecast error variance at horizons between four and twenty quarters. Also note that technology shocks do play an important role in accounting for the forecast error variance in hours worked at longer horizons, accounting for nearly 30% of this variance at horizons greater than 4 quarters, and more than 40% of the unconditional variance.

The decline in the importance of technology shocks is much more pronounced when we focus on cyclical frequencies. Recall from Table 3 that, based on the bivariate system, technology shocks account for roughly 64% and 33% of the cyclical variation in output and hours worked. In the six variable systems, these percentages plummet to ten and four, respectively. Interestingly, a similar result emerges from the four variable CI and $R\pi$ systems. For example in the latter system, technology shocks account for roughly 64% and 33% of the cyclical variation in output and hours worked.

Turning to the other variables, Table 3 indicates that technology shocks play a substantial role in inflation, accounting for over 60% of the one step ahead forecast error variance and almost 40% at even the 20 quarter horizon. Technology shocks also play a very important role in the variance of consumption, accounting for over 60% of the one step ahead forecast error variance and almost 90% of the unconditional variance. These shocks also play a substantial, if smaller role in accounting for variation in investment. These shocks, however, do not play an important role in the forecast error variance for the Federal Funds.

Turning to business cycle frequencies, two results stand out. First, technology shocks account for a very small percent of the cyclical variance in output, hours worked, investment and the Federal Funds rate, 10%, 4%, 1% and 7%,respectively. Second, technology shocks account for a moderately large percentage of the cyclical variation in consumption (16.7) and surprisingly large amount of the cyclical variation in inflation (32%).

Figure 16 presents the historical decompositions for the six variable level specification VAR. Note that technology shocks do relatively well at accounting for the data on output, hours, consumption, inflation and to some extent investment at the lower frequencies. A similar picture emerges from Figure 17, which reports the analog results for the six variable growth rate specification VAR.

8 Conclusions

This paper argues that hours worked rise after a positive permanent technology shock. We, however, also found that these technology shocks do not play a large role in business cycle fluctuations. It is possible that further refinements of our analysis could isolate other kinds of technology shocks that might play a larger role in business cycle fluctuations (see Fisher, 2003).

9 Appendix: Encompassing Analysis for Level and Quadratic Trend Models

In this appendix we briefly discuss the ability of each of the three models in section 6.2 to encompass the results of the other two. The results we focus on are the response of hours worked to a technology shock. The objective is to provide additional details to the general discussion about encompassing that appears in section 6.2.

As in the text, we refer to the three models as the 'Levels' model, the 'Trend in All Equations' model, and the 'Trend in Hours Only' model. Our diagnostic results are reported in Figure A. Each of Panels A, B and C report encompassing results for one particular model, the one indicated in the associated panel header. In each panel there are two columns. Each column focuses on the ability of the model to encompass the empirical results obtained using one of the other two models.

Consider Panel A first. This evaluates the Level model's ability to account for the results based on the Trend in All Equations model and the Trend in Hours Only model. To do this, we simulated 1,000 synthetic time series, each of length equal to our sample period. Using each of these time series, we estimated the Trend in All Equations model and the Trend in Hours Only model, and we computed the impulse response function of interest. The starred line in each column indicates the mean response across the 1,000 time series. The grey area indicates the associated 95% confidence interval. The dark, thick line indicates the estimated impulse response function based on the Level model. The line with circles represents the estimated impulse response function based on the Trend in All Equations model. The line with x's represents the estimated impulse response function of the Trend in Hours Only model.

The key result in Panel A is that the Level model encompasses the empirical results implied by the other two models. This is true if we focus on point estimates alone. In particular, the Level model successfully accounts for the fact that one quadratic trend model implies a fall in hours after a technology shock, while the other implies a rise. The encompassing result is even stronger when we take sampling uncertainty into account.

The quadratic trend models are not mispecified if the Level model is true, and so the result in the 1,1 panel, that the Level model correctly predicts fall in hours implied by the Trend in Hours only model, must be a small sample phenomenon. We verified that this is indeed true, by modifying the calculations in Figure A to allow synthetic data sets to be much longer. Doing this, we found that the dark thick line and starred line nearly coincide in the 1,1 panel (we found the same thing for panel 1,2).

Panel B evaluates the ability of the Trend in Hours Only model to account for the results based on the Level and the Trend in All Equations models. The labelling convention on the lines is the same as in Panel A. Focusing on the point estimates alone, the Trend in Hours Worked model is unable to encompass the results of either of the other two models. Specifically, it cannot account for the fact that hours worked rise in each of the other two models. However, once sampling uncertainty is taken into account, this encompassing test does not reject the Trend in Hours Only model.

Panel C evaluates the ability of the Trend in All Equations model to account for the results based on the Level and the Trend in Hours Only models. Again, the labelling convention on the lines is the same as in Panel A. Two things are worth noting here. First, focusing on the point estimates alone, the Trend in all Equations model can encompass the results based on the Trend in Hours Only model, but it does not encompass the results based on the Level model. In particular, the Trend in All Equations model predicts, counterfactually, that the Level model produces a fall in hours worked after a positive technology shock. Second, even when sampling uncertainty is taken into account, encompassing test rejects the Trend in All Equations model vis a vis the Level model.

References

[1] Altig, David, Lawrence J. Christiano, Martin Eichenbaum and Jesper Linde, 2002, 'An Estimated Dynamic, General Equilibrium Model for Monetary Policy Analysis,' Manuscript.

[2] Basu, Susanto, John G. Fernald, and Miles S. Kimball. 1999. 'Are Technology Improvements Contractionary?' Manuscript.

[3] Boldrin, Michele, Lawrence J. Christiano and Jonas Fisher, 2001, 'Asset Pricing Lessons for Modeling Business Cycles,' *American Economic Review*. 91, 149-66.

[4] Caner, Mehmet and Lutz Kilian. 1999. 'Size Distortions of Tests of the Null Hypothesis of Stationarity: Evidence and Implications for the PPP Debate,' University of Michigan Manuscript.

[5] Christiano, Lawrence J., and Lars Ljungqvist, 1988, 'Money Does Granger Cause Output in the Bivariate Money-Output Relation,' *Journal of Monetary Economics*. 22(2) pp. 217-35

[6] Christiano, Lawrence J. and Martin Eichenbaum, 1992, 'Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations,' *American Economic Review*. 82(3) pp. 430-50.

[7] Christiano, Lawrence J. and Richard M. Todd. 1996. 'Time to Plan and Aggregate Fluctuations'. *Federal Reserve Bank of Minneapolis Quarterly Review* Winter 14-27.

[8] Christiano, Lawrence J. and Terry Fitzgerald. 1999. The Band Pass Filter. *National Bureau of Economic Research Working Paper* 7257, and forthcoming. *International Economic Review*.

[9] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans, 1999, Monetary Policy Shocks: What Have We Learned, and to What End?, in Taylor and Woodford, *Handbook of Macroeconomics*.

[10] Christiano, Lawrence J., Martin Eichenbaum and Charles Evans, 2001, 'Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy', manuscript.

[11] Doan, Thomas 1992 *Rats Manual* Estima Evanston, IL.

[12] Eichenbaum, Martin, and Kenneth J. Singleton 1986. 'Do Equilibrium Real Business Cycle Theories Explain Postwar U.S Business Cycles?' NBER Macroeconomics Annual 1986, pp. 91-135.

[13] Fisher, Jonas, 2003, 'Technology Shocks Matter,' manuscript.

[14] Francis, Neville, and Valerie A. Ramey, 2001, 'Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited,' manuscript, UCSD.

[15] Gali, Jordi, 1999, 'Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?' American Economic Review, vol. 89, no. 1, 249-271.

[16] Gali, Jordi, Mark Gertler and J. David Lopez-Salido, 2001, 'Markups, Gaps and the Welfare Costs of Business Fluctuations,' May.

[17] Gali, Jordi, J. David Lopez-Salido, and Javier Valles, 2002, 'Technology Shocks and Monetary Policy: Assessing the Fed's Performance', National Bureau of Economic Research Working Paper 8768.

[18] Kwiatkowski, D., Phillips, P.C.B., Schmidt, P., and Shin, Y. 1992, 'Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root,' *Journal of Econometrics*, 54, 159- 178.

[19] King, Robert, Charles Plosser, James Stock and Mark Watson, 1991, 'Stochastic Trends and Economic Fluctuations,' *American Economic Review*, 81, 819-840.

[20] Leybourne, Stephen J. and B.P.M McCabe, 1994, 'A Consistent Test for a Unit Root', *Journal of Business and Economic Statistics*, 12(2) pp. 157-66

[21] Shapiro, Matthew, and Mark Watson, 1988, 'Sources of Business Cycle Fluctuations,' NBER,*Macroeconomics Annual*, pp. 111-148.

[22] Shea, John 1998. 'What Do Technology Shocks Do?,' National Bureau of Economic Research Working Papers 6632

[23] Sims Christopher, James Stock and Mark Watson, 1990. 'Inference in Linear Time Series Models with Some Unit Roots,' *Econometrica* 58(1), pp. 113-144.

[24] Vigfusson, Robert J. 2002 'Why Does Employment Fall After A Positive Technology Shock,' manuscript.

Figure 1: Data Used in VAR

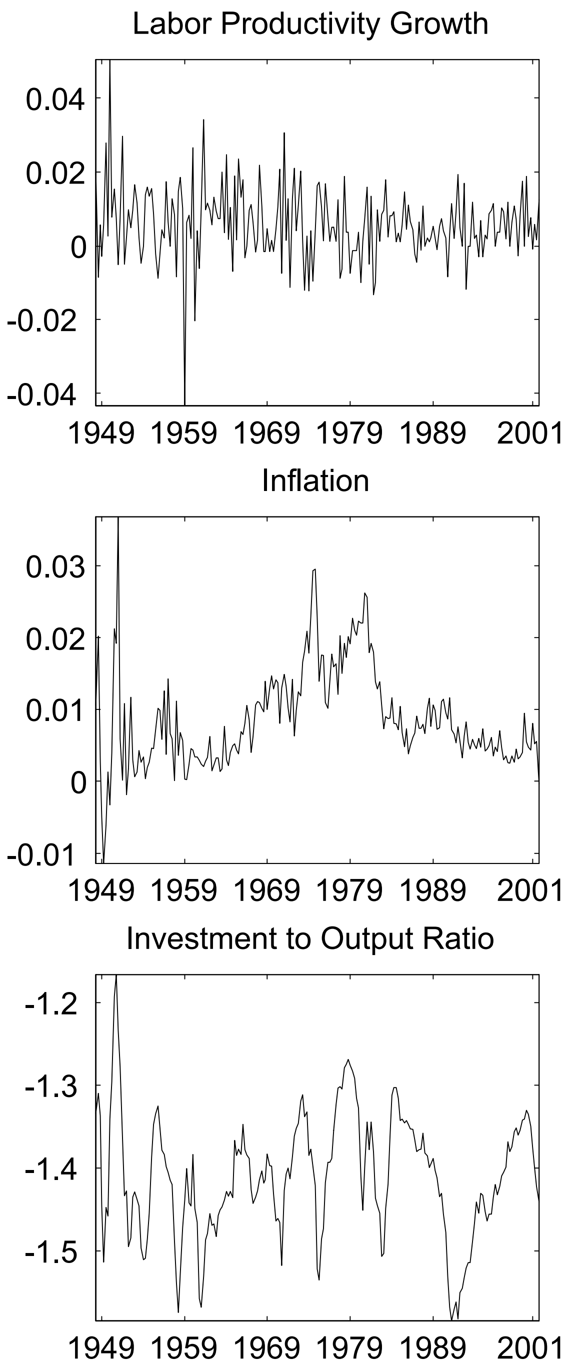

Figure 2: Response of log-output and Log-hours to a Positive Technology Shock (Specification with Level Hours)

Panel A: Sample Period, $1948Q1 - 2001Q4$

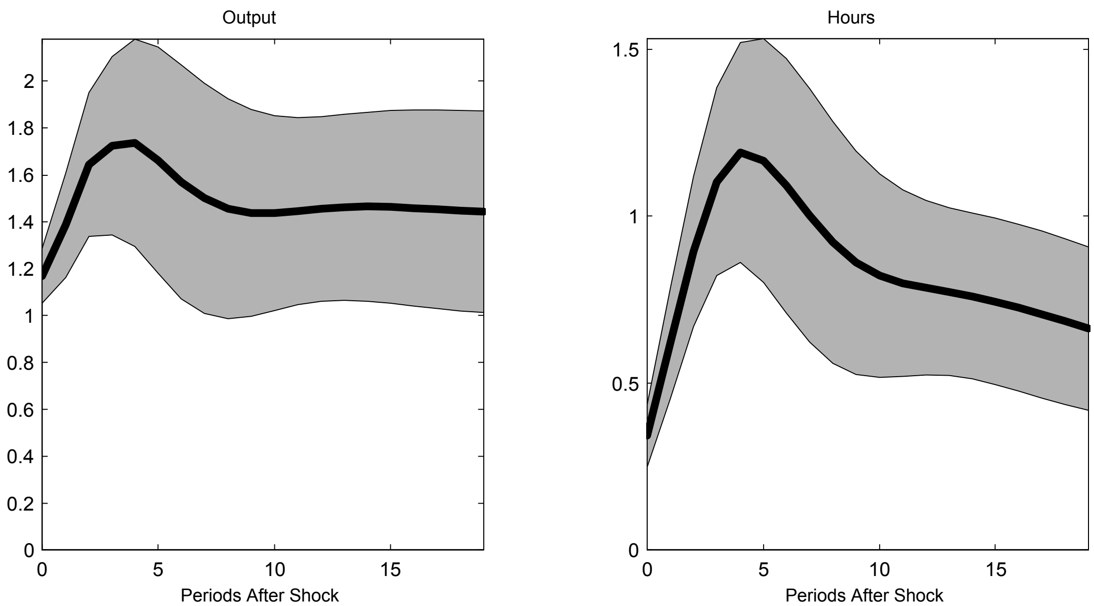

Panel B: Sample Period, $1959Q1 - 2001Q4$

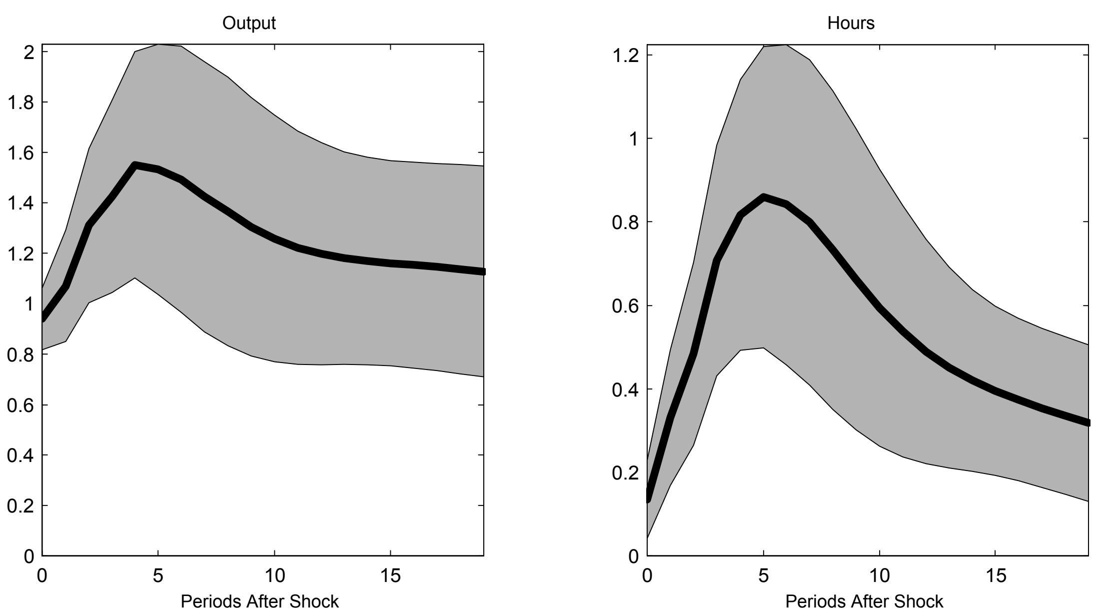

Figure 3: Response of log-output and Log-hours to a Positive Technology Shock
(Specification with Hours Growth)

Panel A: Sample Period, $1948Q1 - 2001Q4$

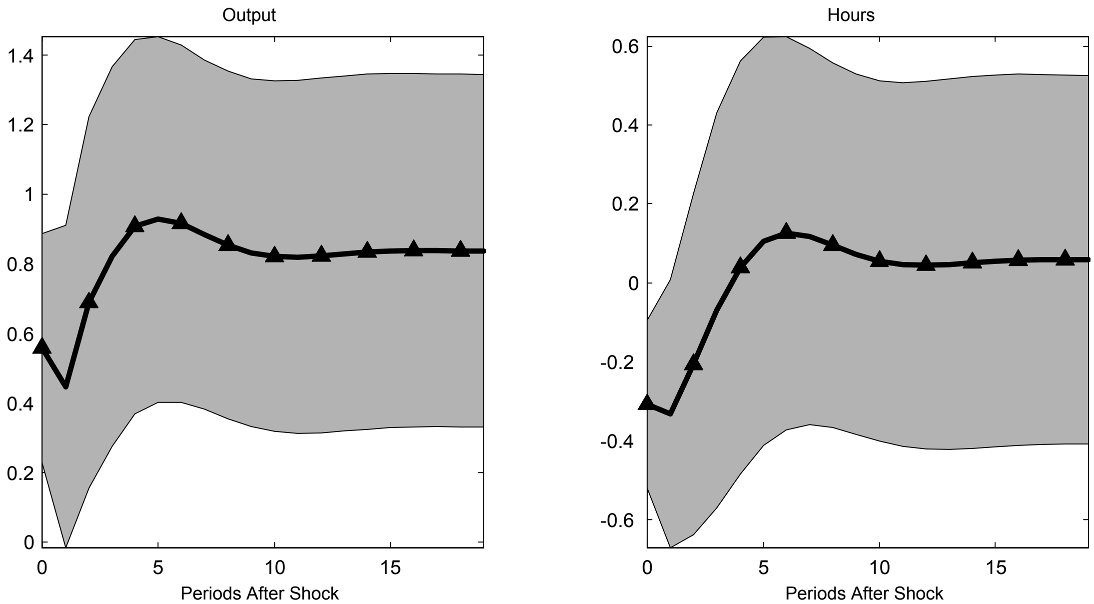

Panel B: Sample Period, $1959Q1 - 2001Q4$

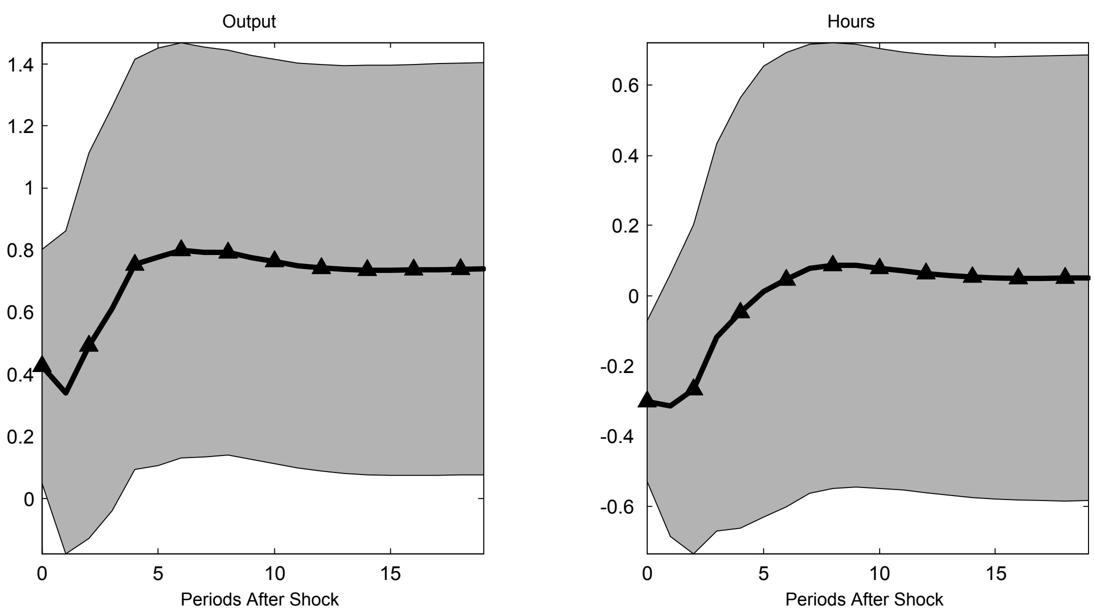

Figure 4: Encompassing with Level Specification as the DGP

Panel A: Sample Period, $1948Q1 - 2001Q4$

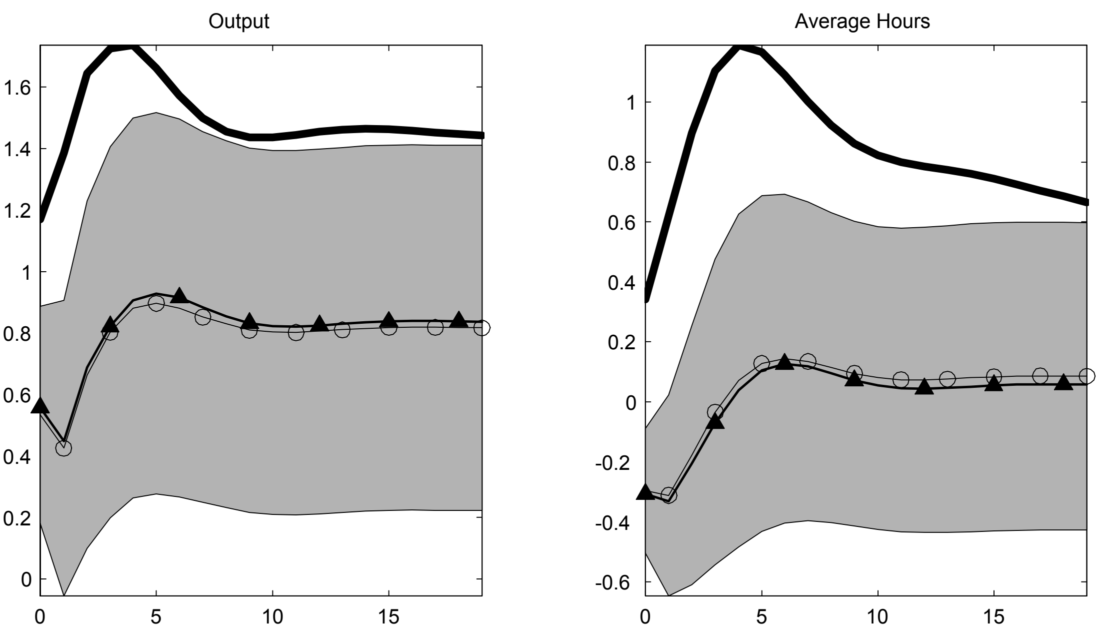

Panel B: Sample Period, $1959Q1 - 2001Q4$

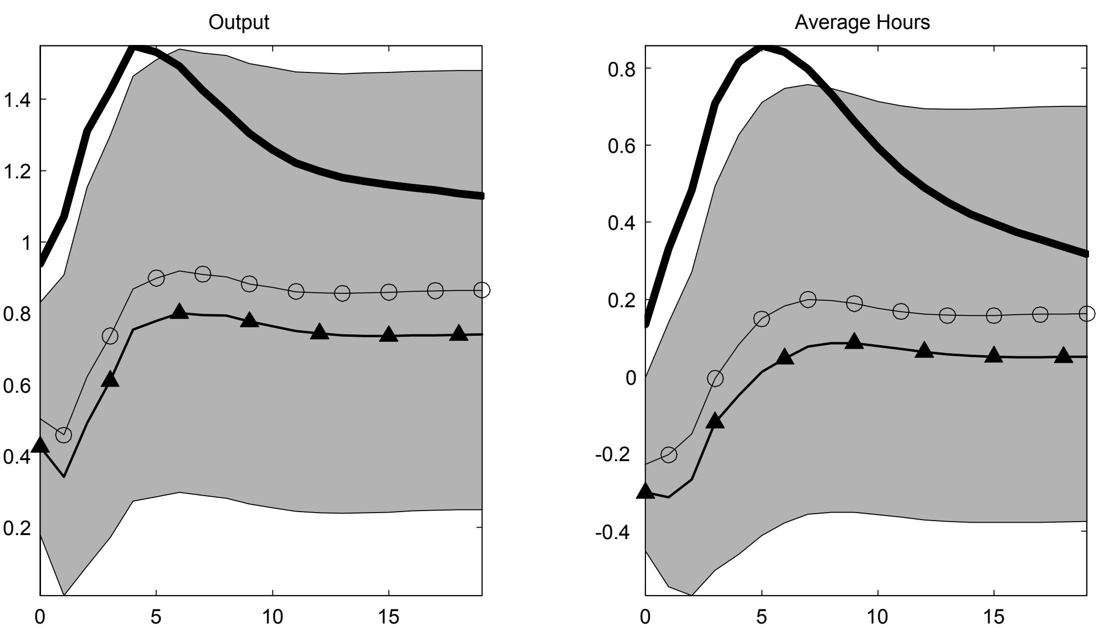

Figure 5: Encompassing with Growth Rate Specification as the DGP

Panel A: Sample Period, $1948Q1 - 2001Q4$

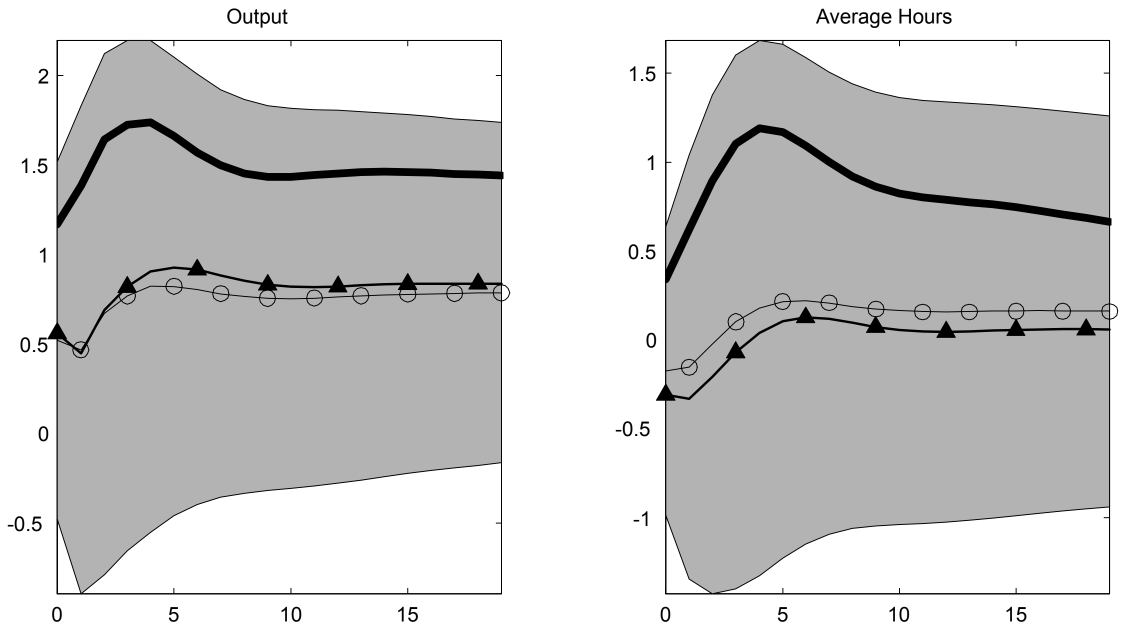

Panel B: Sample Period, $1959Q1 - 2001Q4$

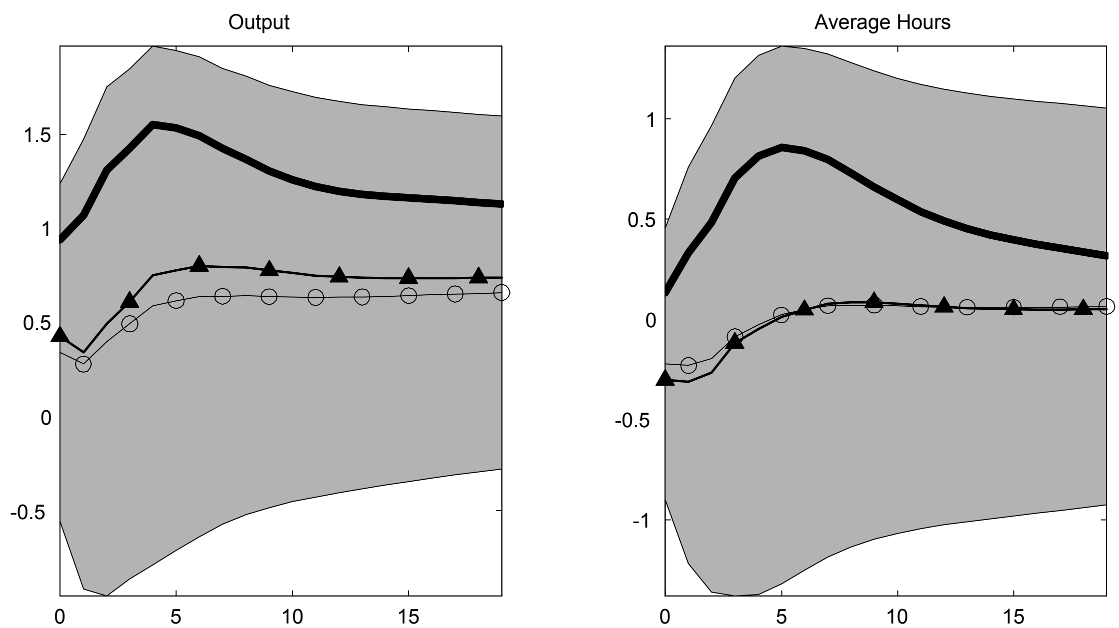

Figure 6: Average Response of Hours, Sample Period $1948Q1 - 2001Q4$

Hours: Level Specification | Hours: Growth Rate Specification

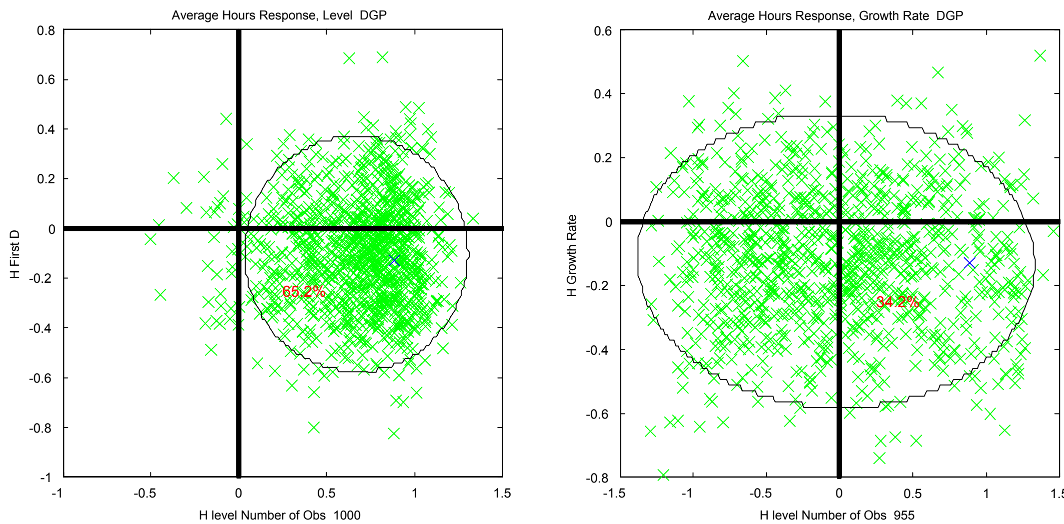

Figure 7: Average Response of Hours and Output - $1959Q1 - 2001Q4$

Hours: Level Specification | Hours: Growth Rate Specification

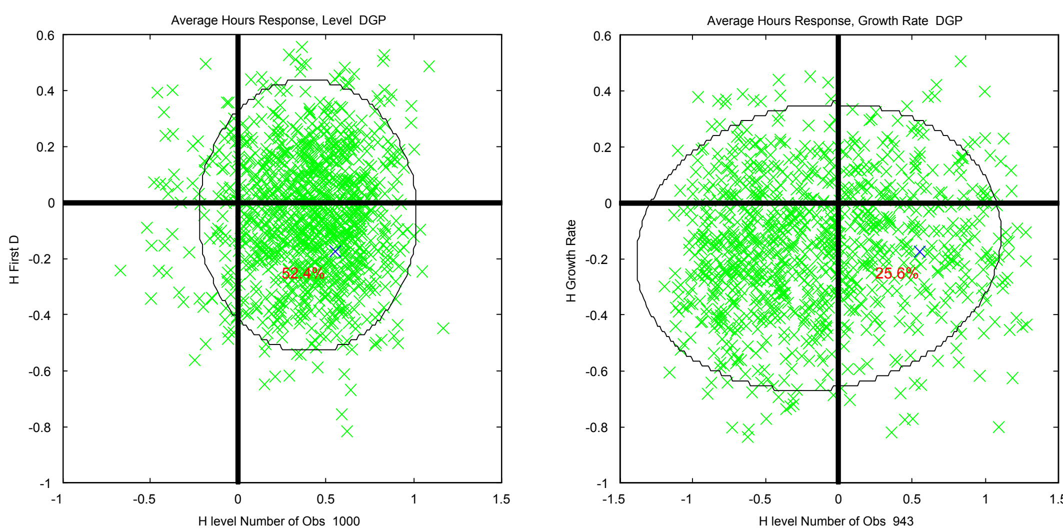

Figure 8: Six Variable System, Level Specification, Sample Period 1959-2001

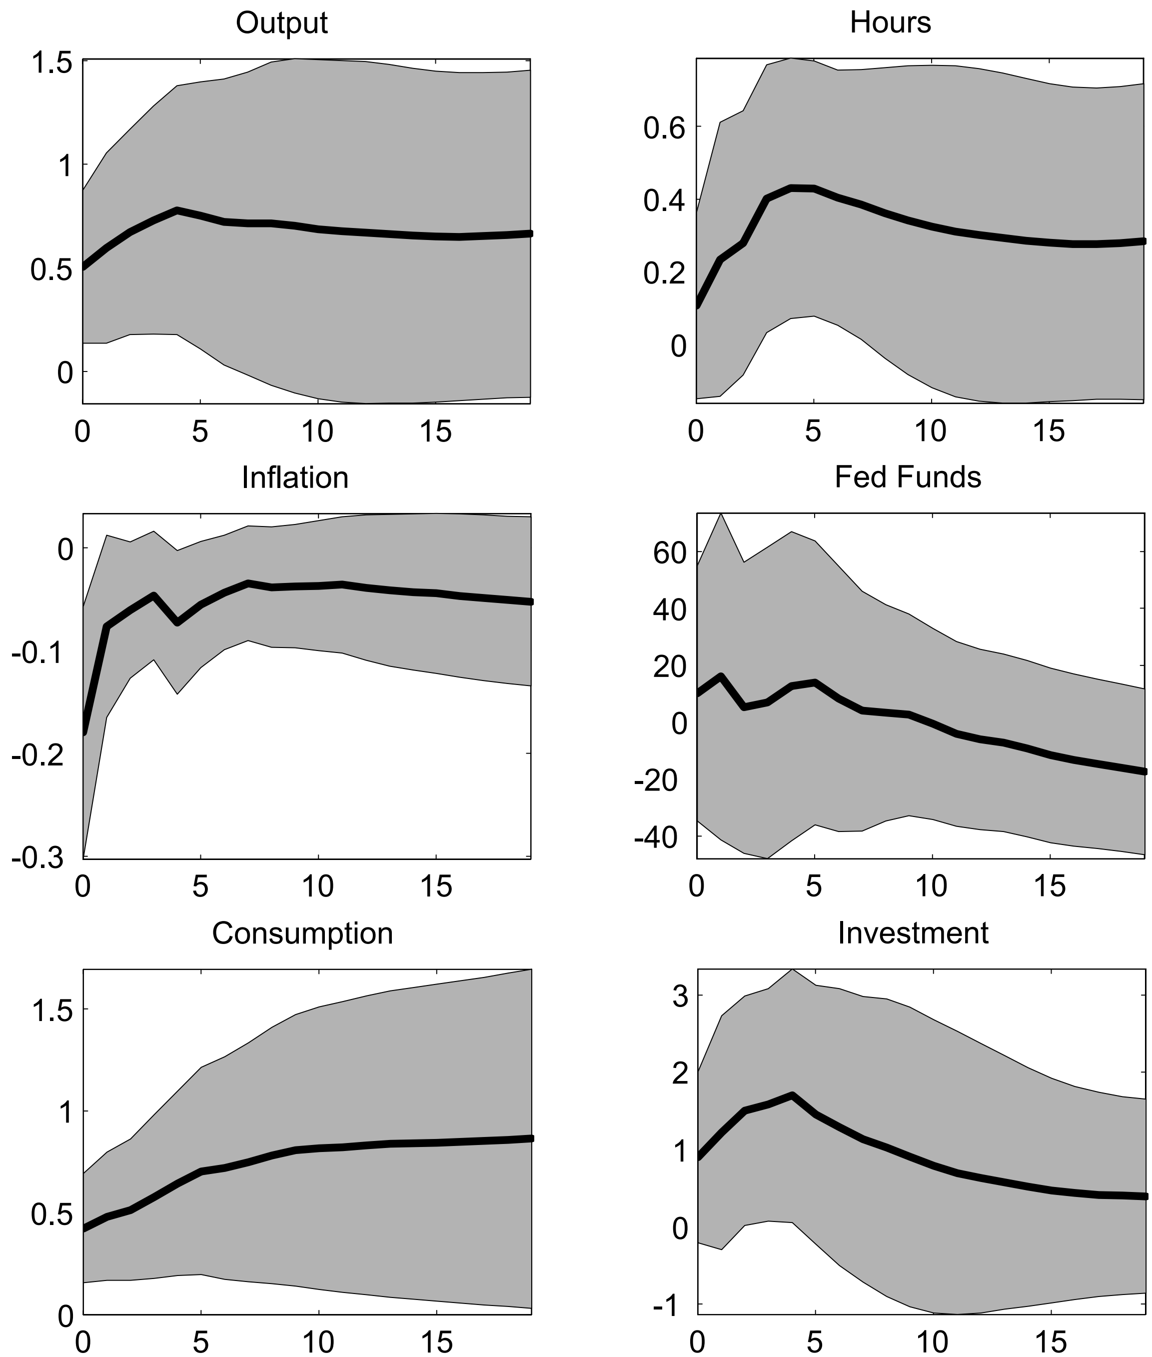

Figure 9: Six Variable System, Growth Rate Specification, Sample Period 1959-2001

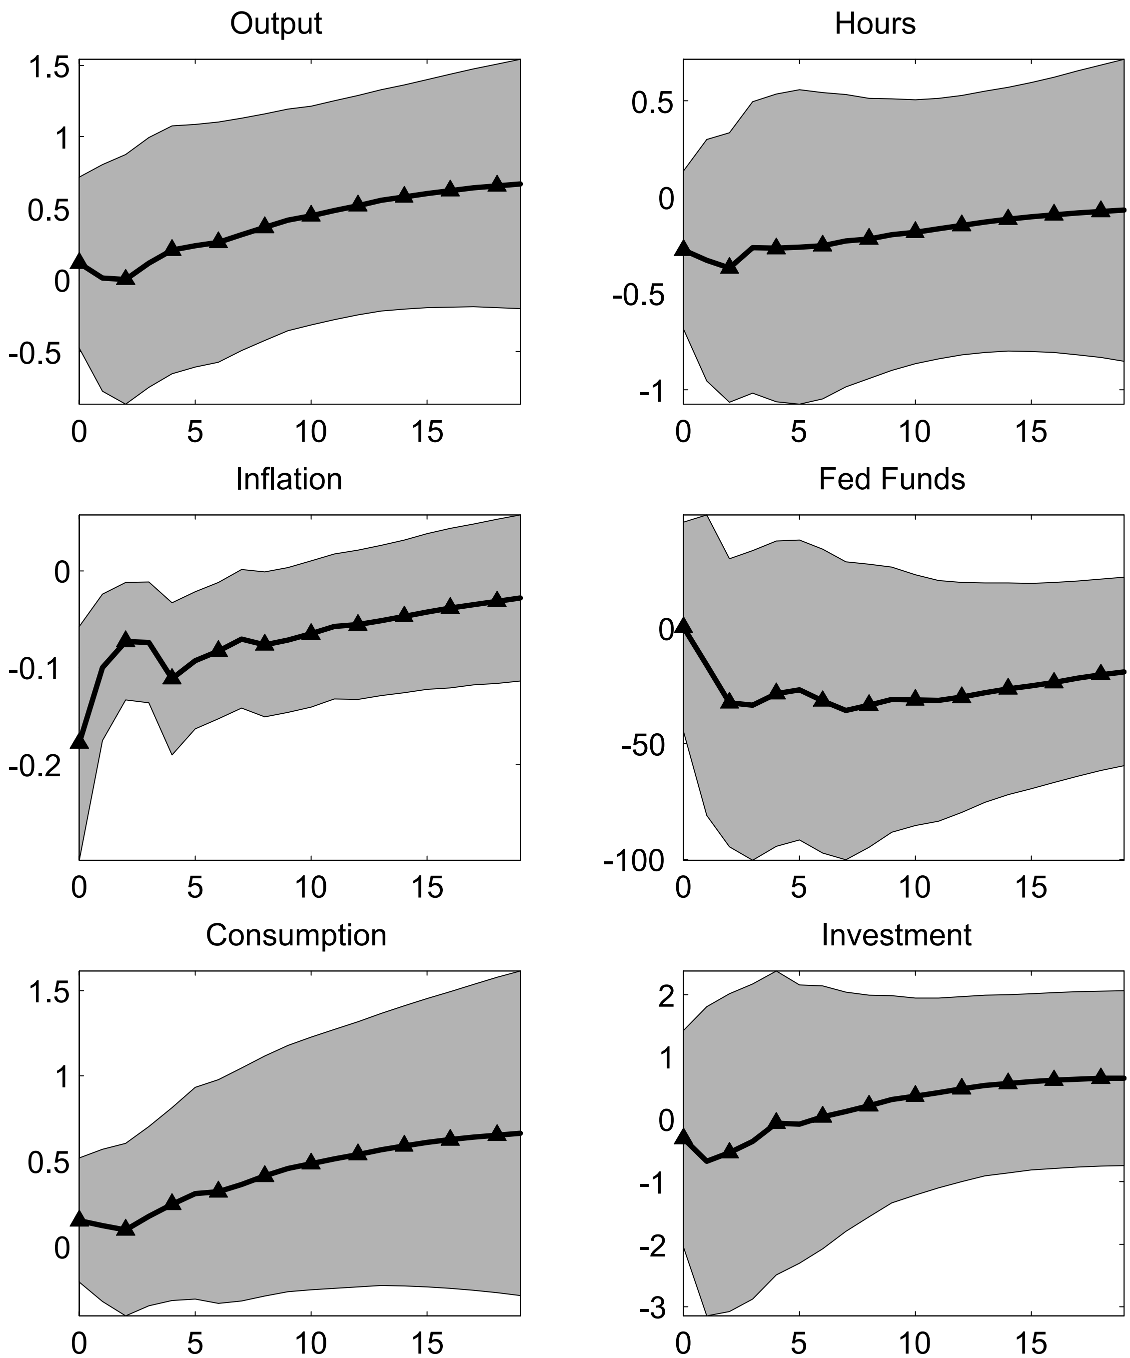

Figure 10: Encompassing Test with the Level Specification as the DGP, 1959-2001

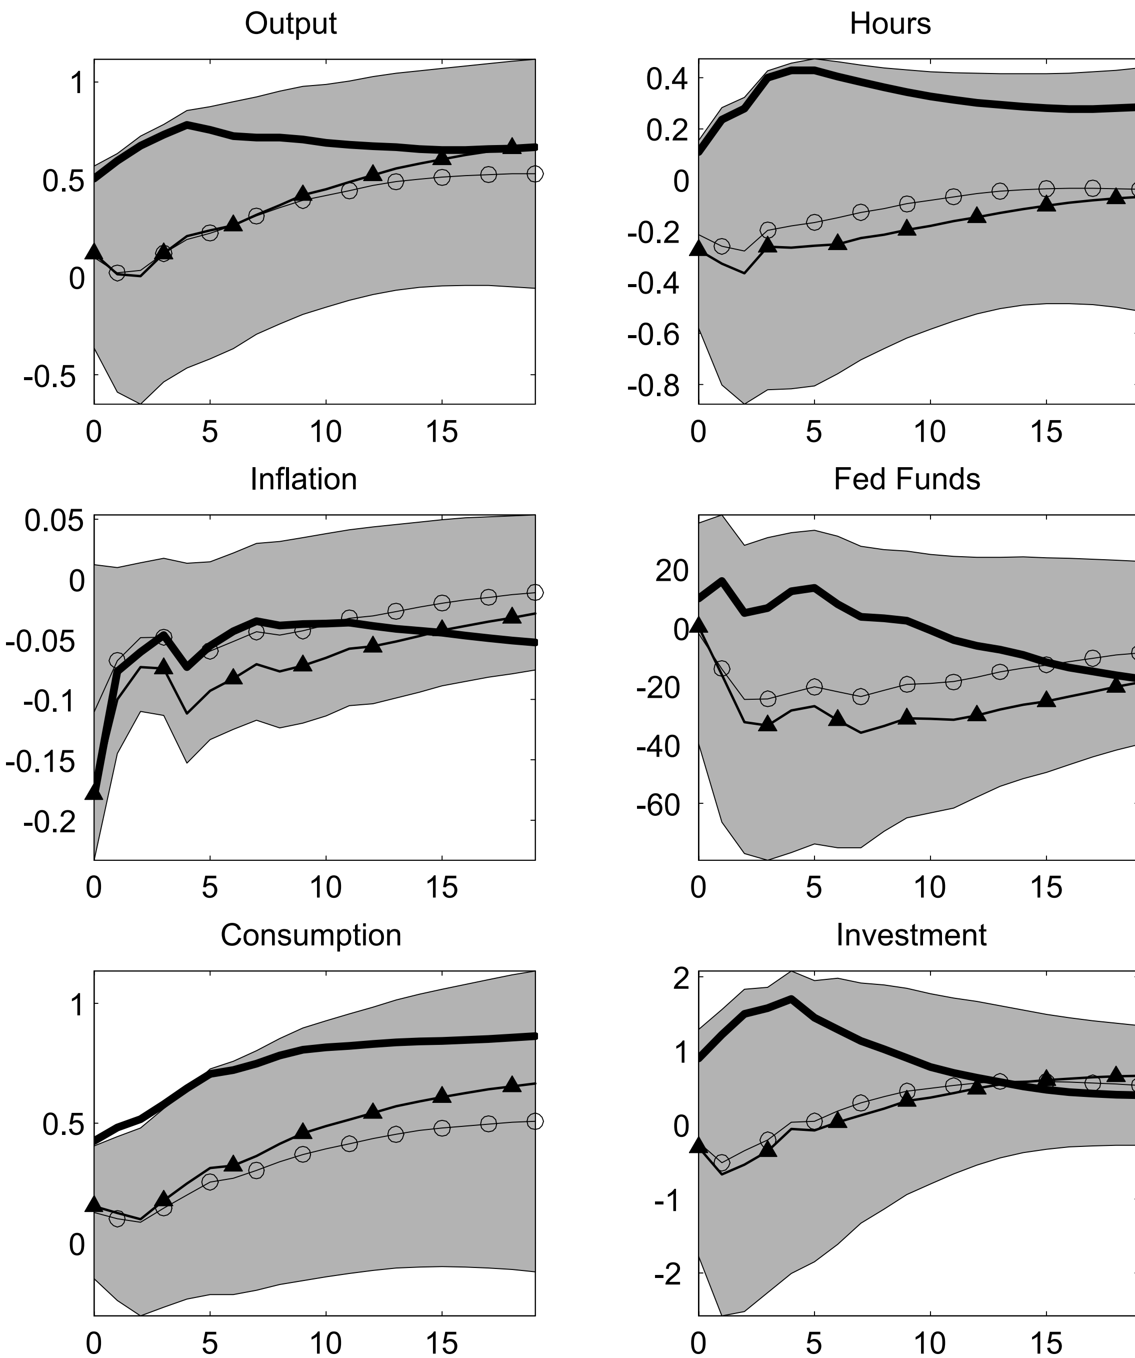

Figure 11: Encompassing Test with the Growth Rate Specification as the DGP, 1959-2001

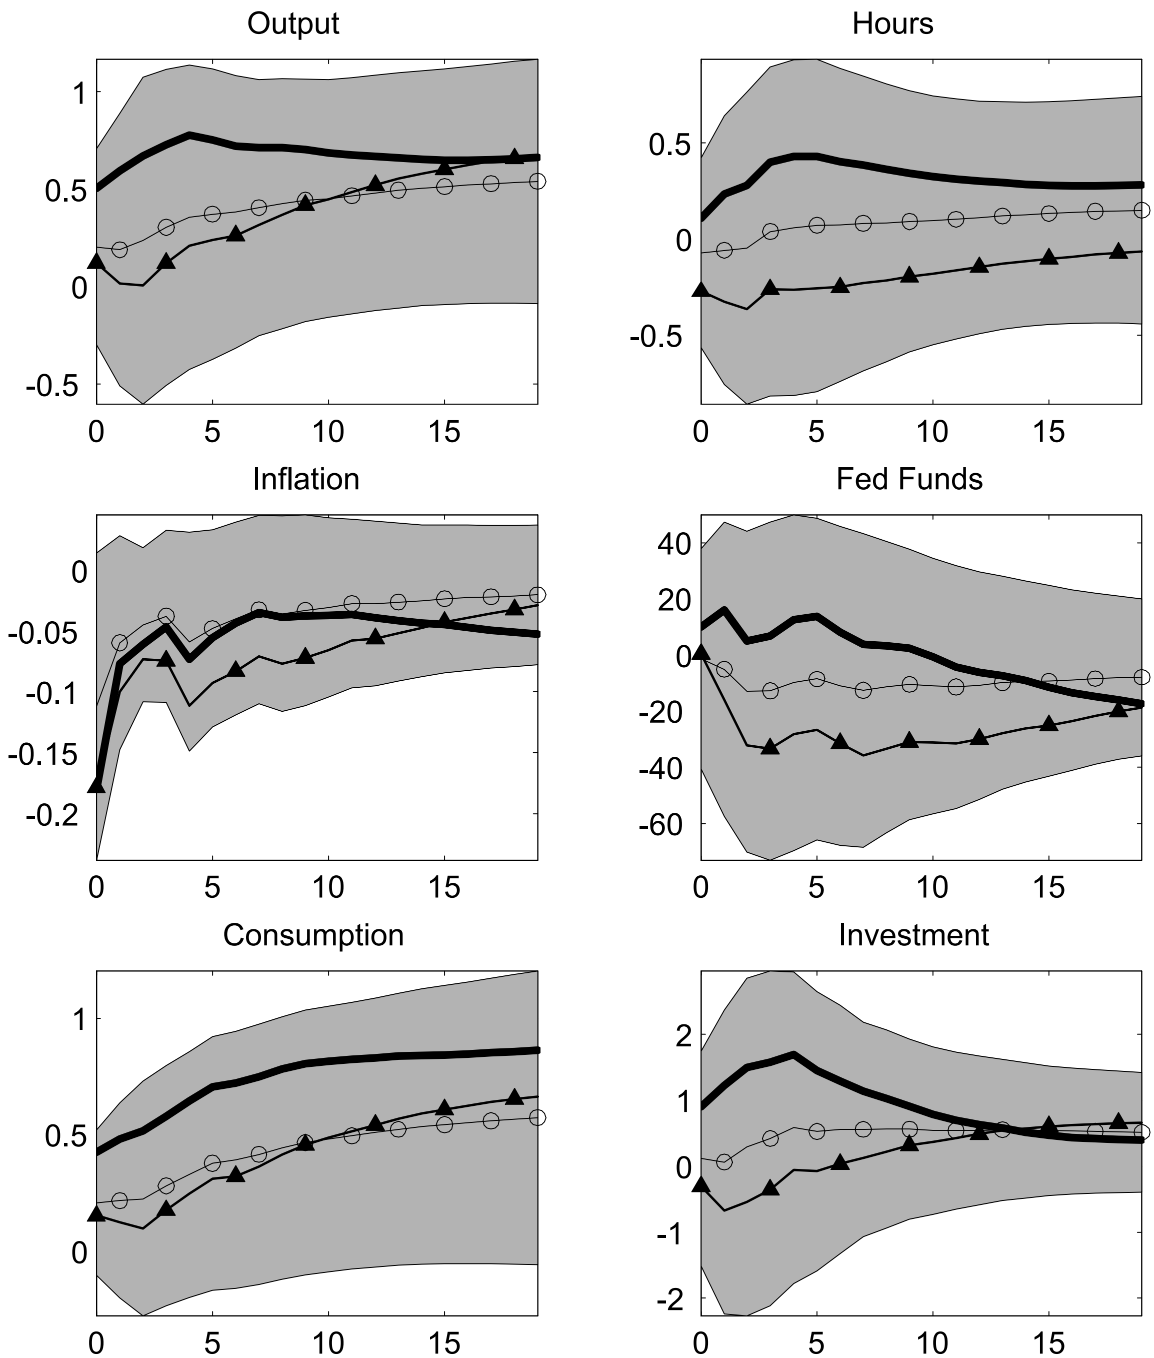

Figure 12: 6 Variables System Average Response of Hours - $1959Q1 - 2001Q4$

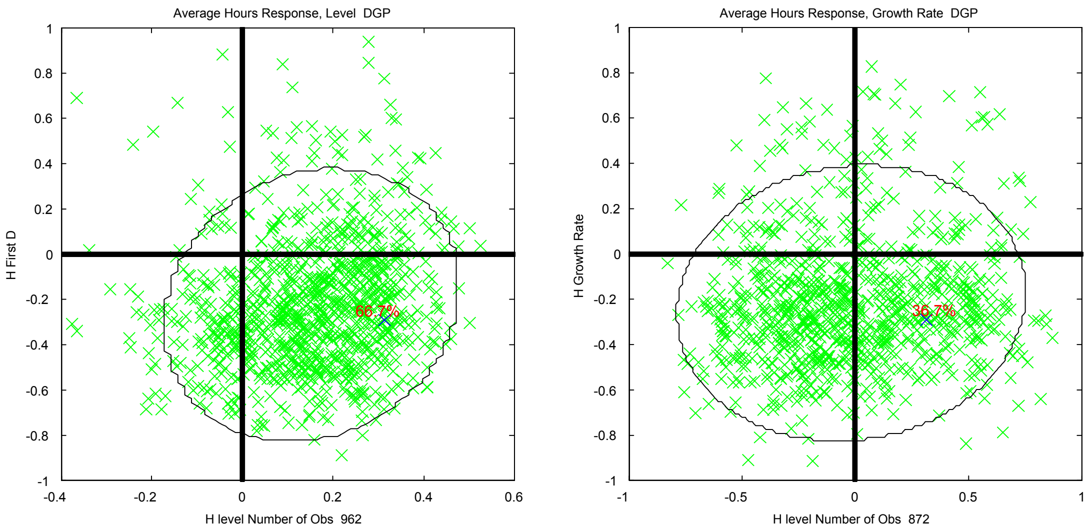

Figure 13: Comparing the Six Variable Specification to 2 different 4 Variable, Level Specification

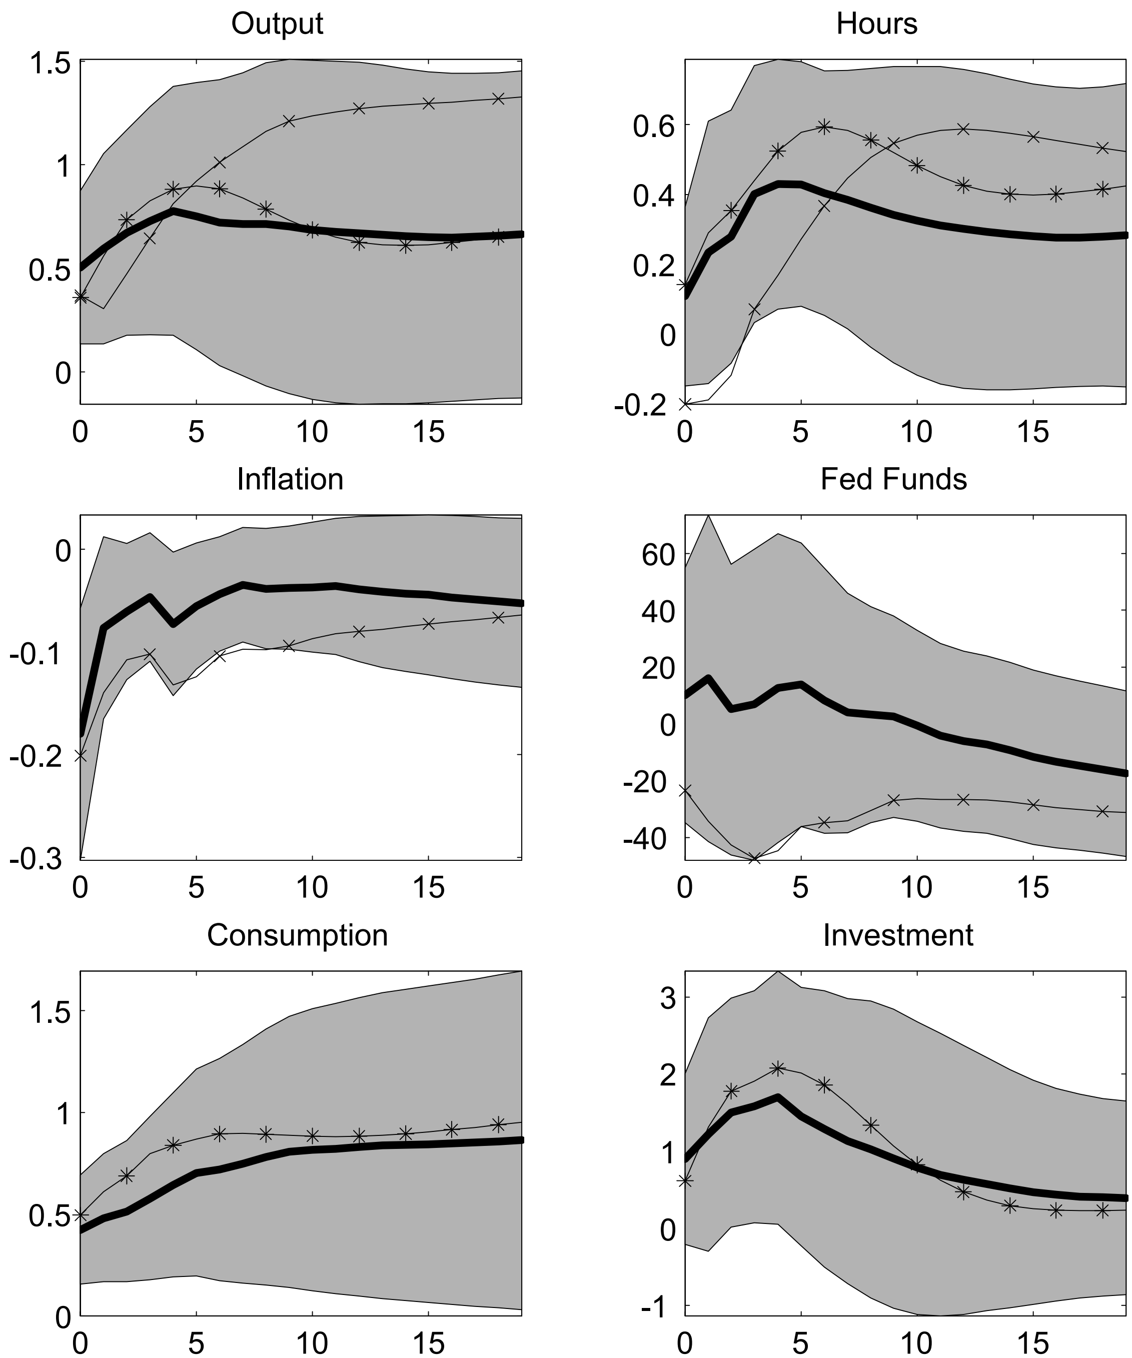

Thick Line has all six variables

'X' has hours, labor productivity, inflation and the federal funds rate.

'*' has hours, labor productivity, consumption and investment

Figure 14: Encompassing 4 variable R π system with 6 variable system

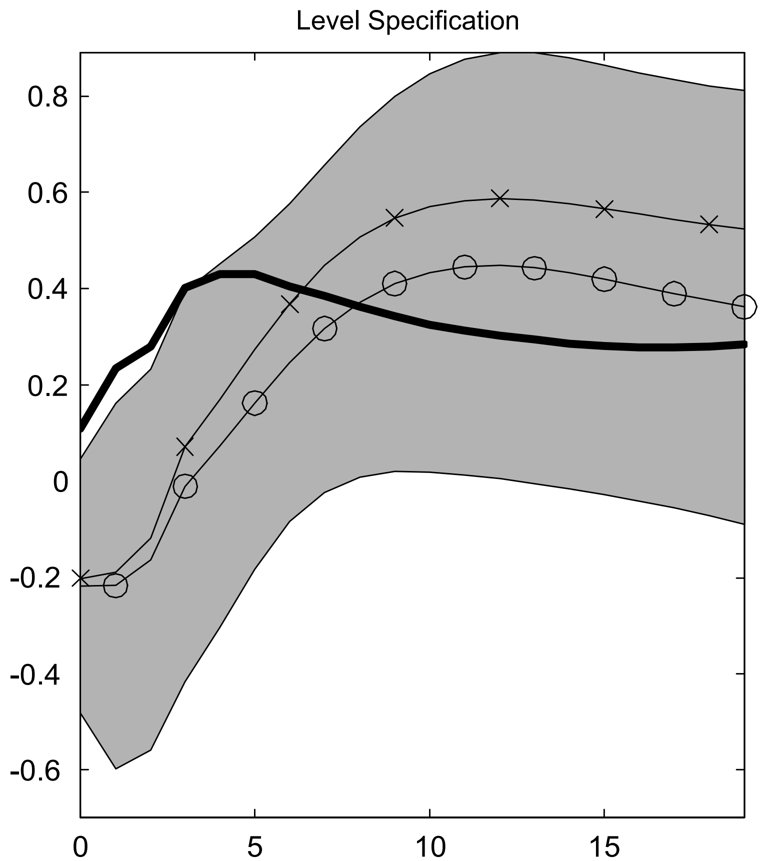

Figure 15: Historical Decomposition: Bivariate System,
Level Specification

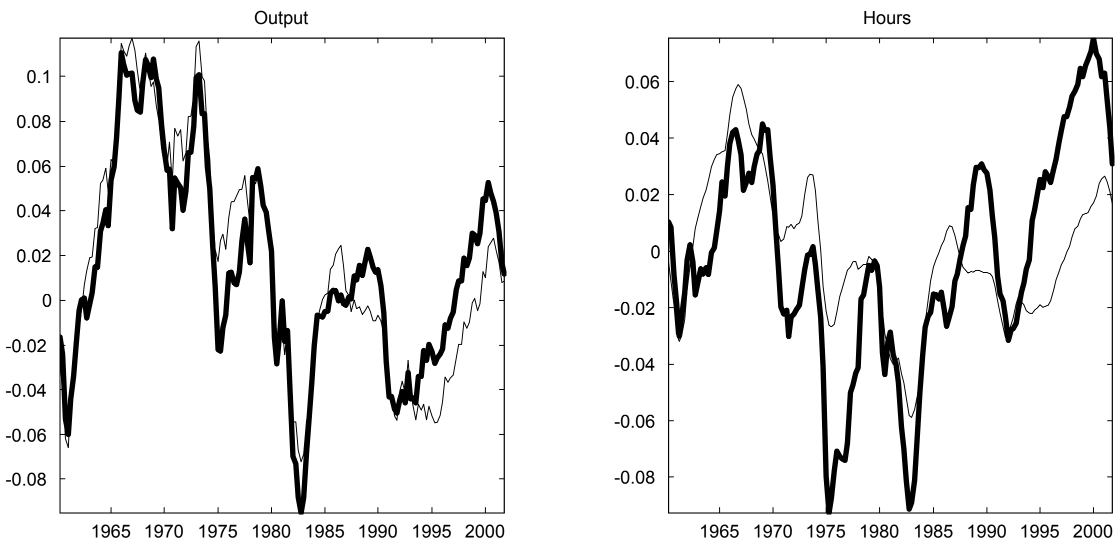

Growth Rate Specification

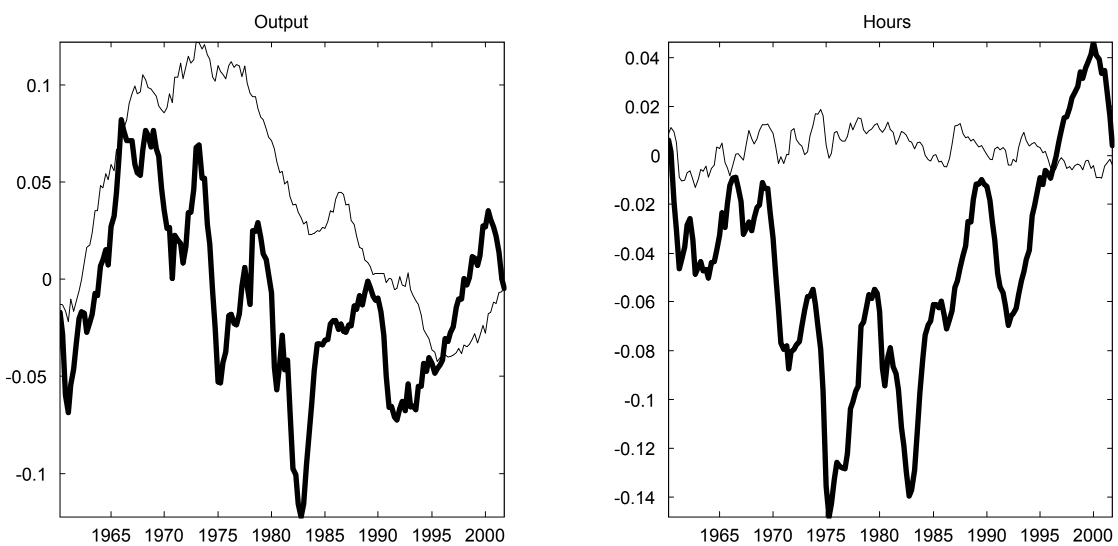

Figure 16: Historical Decomposition: 6 Variable System, Level Specification

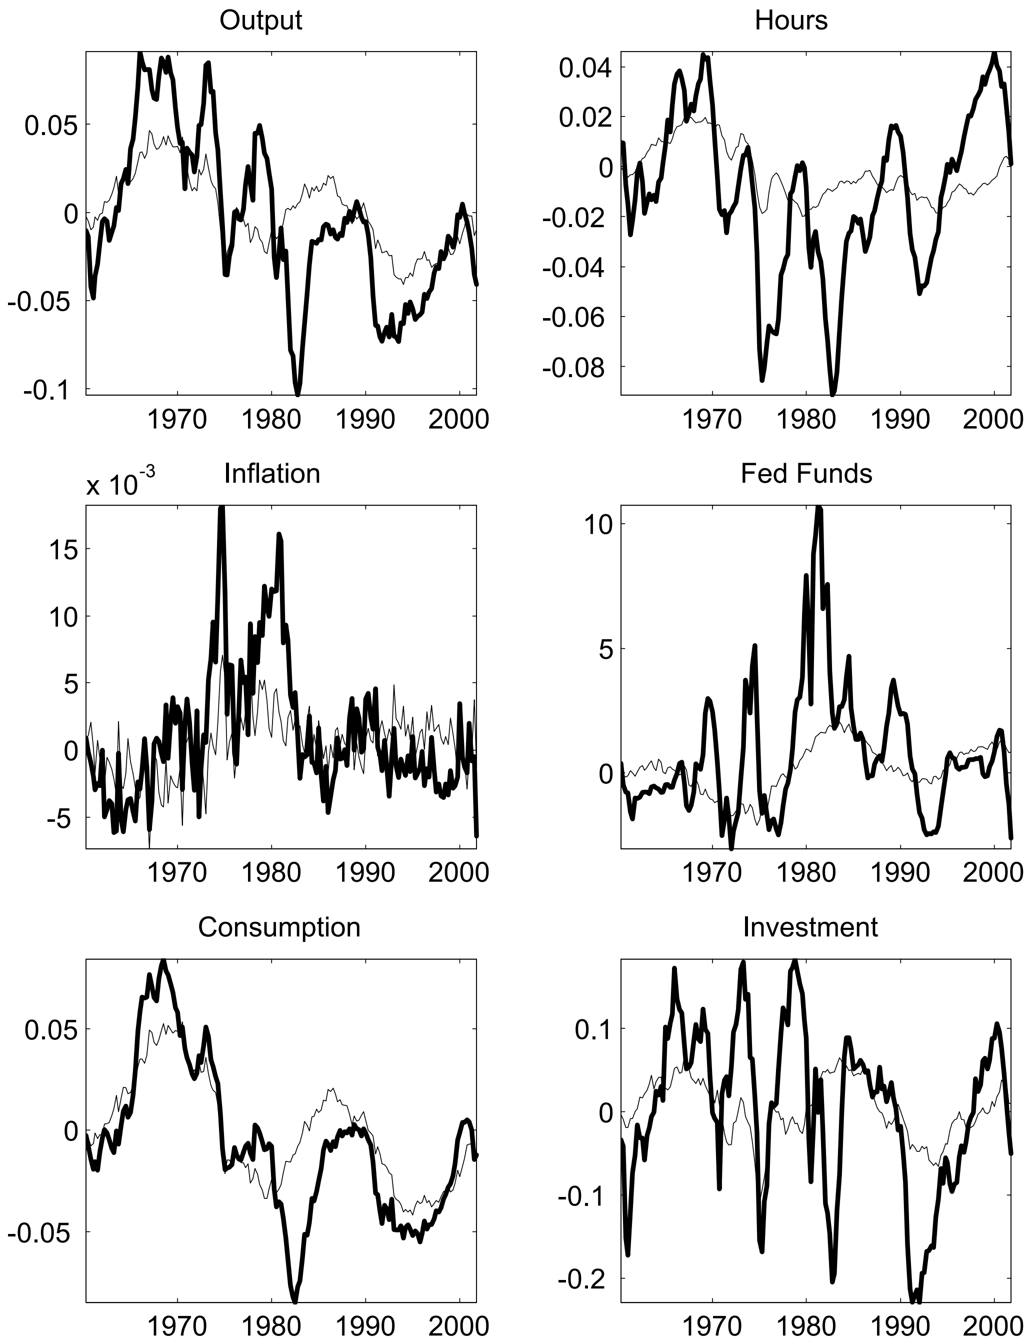

Figure 17: Historical Decomposition: 6 Variable System, Growth Rate Specification

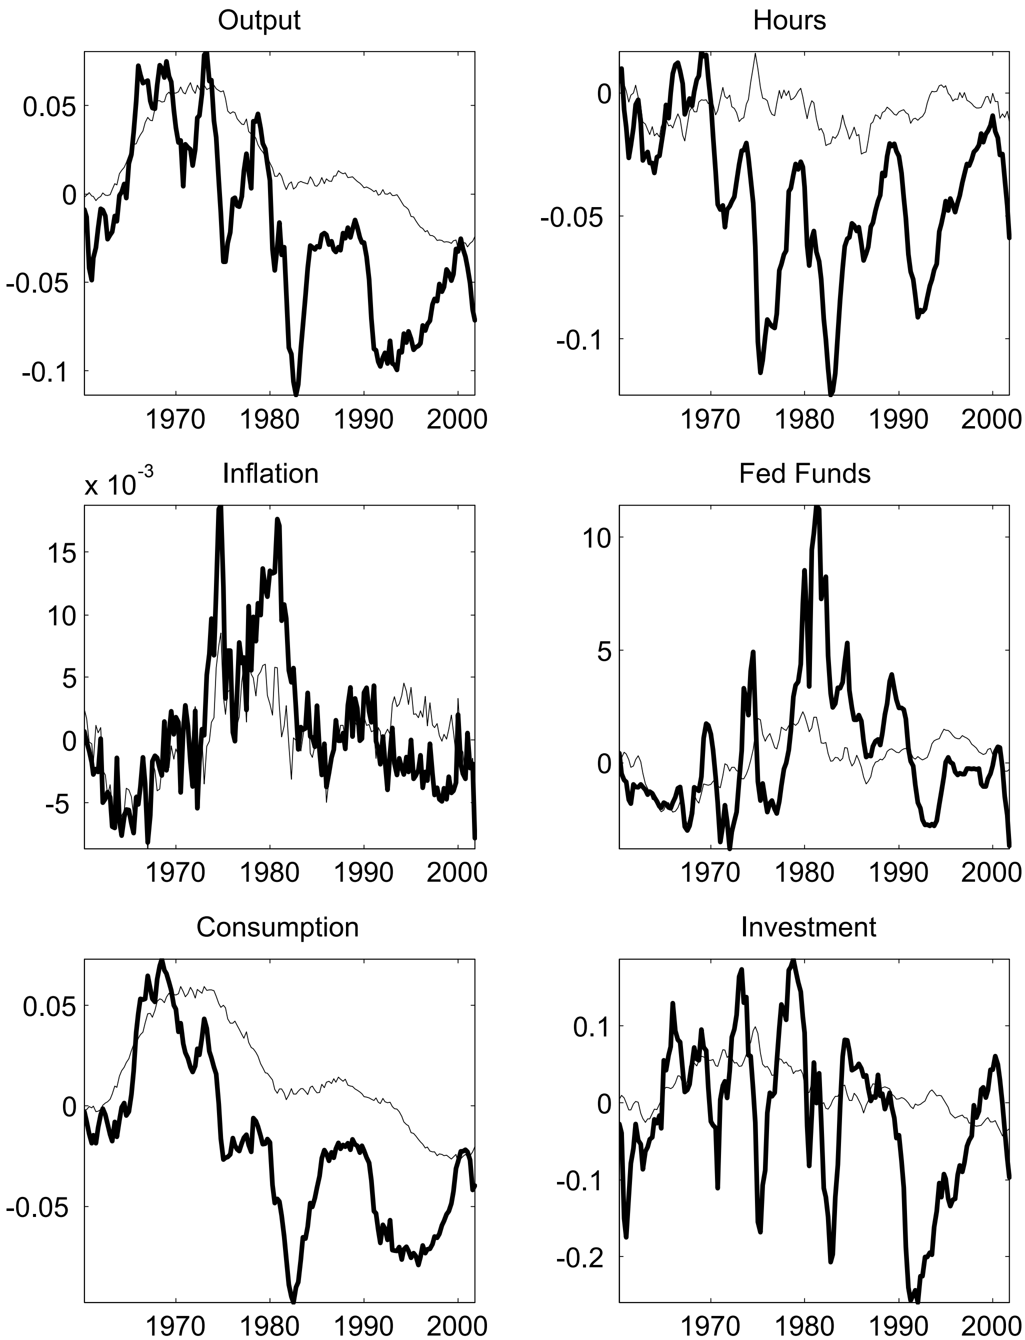

Figure A: Encompassing Analysis for Level and Quadratic Trend Models

Panel A: DGP - Levels

Trend in Hours Only

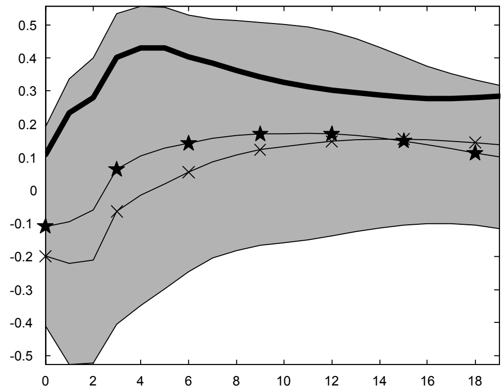

Trend in all Equations

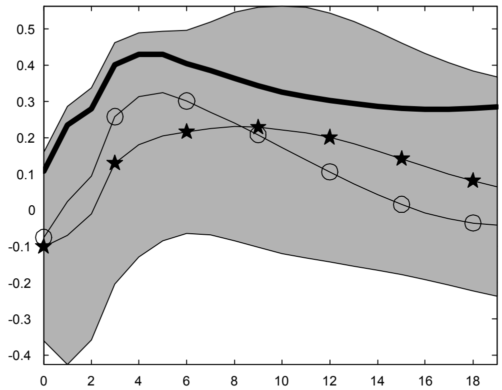

Panel B: DGP - Trend in Hours Only

Levels

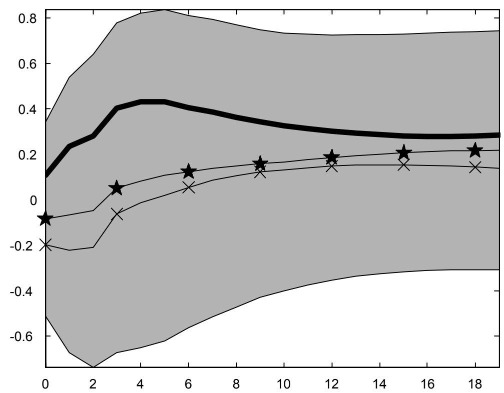

Trend in all Equations

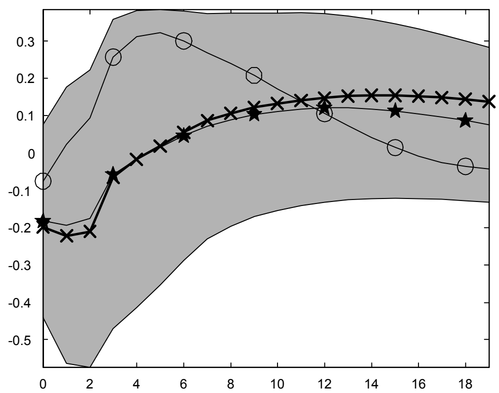

Panel C: DGP - Trend in all Equations

Levels

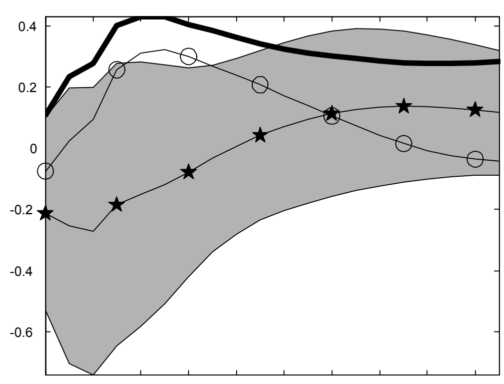

Trend in Hours Only

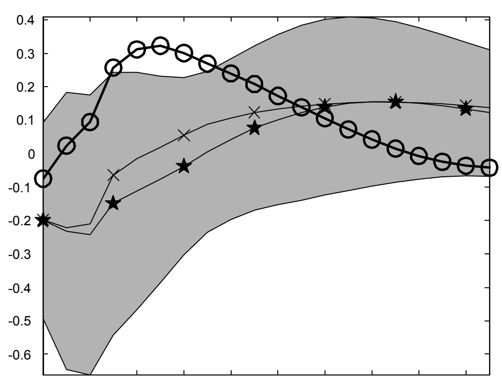

Notes: Thick Dark Line - Est. Levels Model, Stars - Predicted Mean Response, X's - Est. Trend in Hours Only, Circles - Est. Trend in all Equations.

Table 1: Contribution of Technology Shocks to Variance,Bivariate System

Level Specification

	Forecast Variance at Indicated Horizon					
Variable	1	4	8	12	20	50
Output	81.1	78.1	86.0	89.1	91.8	96
Hours	4.5	23.5	40.7	45.4	47.4	48.3

Growth Rate Specification

	Forecast Variance at Indicated Horizon					
Variable	1	4	8	12	20	50
Output	16.5	11.7	17.9	20.7	22.3	23.8
Hours	21.3	6.4	2.3	1.6	1.0	0.5

Table 2: Contribution of Technology Shocks to Variance, Six Variable System

Level Specification

	Forecast Variance at Indicated Horizon					
Variable	1	4	8	12	20	50
Output	31.2	40.3	44.6	41.5	44.8	70
Hours	3.6	15.4	28.8	28.4	28.8	43.9
Inflation	60.2	47.0	43.2	41.1	39.5	47.7
Fed Funds	1.6	1.4	1.7	1.7	3.7	23.3
Consumption	61.6	64.2	67.3	66.8	71.8	88.4
Investment	10.3	20.1	24.1	20.9	20.4	25.3

Growth Rate Specification

	Forecast Variance at Indicated Horizon					
Variable	1	4	8	12	20	50
Output	1.7	0.6	2.6	6.4	17.2	35.5
Hours	20.8	11.9	8.0	7.1	5.7	2.3
Inflation	58.5	54.7	55.6	52.4	47.4	33.8
Fed Funds	0.0	7.5	10.5	13.7	17.2	16.9
Consumption	7.9	4.1	8.7	14.3	25.3	34.3
Investment	1.1	2.0	1.1	1.3	3.7	13.8

Table 3: Contribution of Technology Shocks to Cyclical Variance (HP Filtered Results)

Level Specification

Variables in VAR	Output	Hours	Inflation	Federal Funds	Consumption	Investment
Y,H	63.8	33.4				
Y,H,$\Delta P, R$	17.8	17.9	53.2	11.2		
Y,H,C, I	19.9	18.5			20.1	20.7
Y,H,$\Delta P, R, C, I$	10.2	4.1	32.4	1.3	16.8	6.7

Growth Rate Specification

Variables in VAR	Output	Hours	Inflation	Federal Funds	Consumption	Investment
Y,ΔH	10.6	7.0				
Y,ΔH,$\Delta P, R$	6.8	8.5	48.4	8.1		
Y,ΔH,C, I	1.3	6.3			0.32	5.5
Y,ΔH,$\Delta P, R, C, I$	1.6	6.1	35.2	4.9	3.7	2.6