REAL BUSINESS CYCLE THEORY:
WISDOM OR WHIMSY?

Martin Eichenbaum

REAL BUSINESS THEORY:
WISDOM OR WHIMSY?

ABSTRACT

This paper assesses the empirical plausibility of the view that
aggregate productivity shocks account for most of the variability in post
World War II US output. We argue that the type of evidence forwarded by
proponents of this proposition is too fragile to be believable. First, our
confidence in the evidence is fundamentally affected once we abandon the
fiction that we actually know the true values of the structural parameters of
standard Real Business Cycle (RBC) models. What the data are telling us is
that, while productivity shocks play some role in generating the business
cycle, there is simply an enormous amount of uncertainty about just what
percent of aggregate fluctuations they actually do account for. The answer
could be 70% as Kydland and Prescott (1989) claim, but the data contain
almost no evidence against either the view that the answer is really 5% or
that the answer is really 200%. Second, we show that point estimates of the
importance of technology shocks are extremely sensitive to small
perturbations in the theory. Allowing for labor hoarding behavior in an
otherwise standard RBC model reduces the ability of technology shocks to
account for aggregate fluctuations by 50%. This finding provides some
support for the view that many of the movements in the Solow residual which
are labelled as productivity shocks may be an artifact of labor hoarding type
phenomena.

Martin Eichenbaum
Department of Economics
Northwestern University
Evanston, IL 60208

This paper discusses two simple questions which are of fundamental importance to macroeconomists. First — which impulses have been the primary sources of fluctuations in postwar US aggregate output? Just how important *have* aggregate technology shocks been? Second, how reliable are existing answers are to the first question?

The answers to these questions obviously matter from the perspective of optimal public policy. But just as important, the *perceived* answers also matter because they influence the research agenda of macroeconomists. Around 1977 it seemed just as obvious to the representative graduate student as it was to Milton Friedman or Robert Lucas that monetary instability is a critical determinant of aggregate output fluctuations. Granted there was substantial disagreement about the nature of the relationship between monetary and real phenomena. But the critical point is that those years were marked by enormous amounts of research aimed at understanding the propagation mechanisms by which monetary policy affects aggregate economic activity. That this was a critical item for business cycle research was, by and large, simply taken for granted.

The situation has clearly changed. Since Kydland and Prescott's (1982) apparent demonstration that productivity shocks can account for *all* output variability in the post War US, the need for an adequate theory of monetary and fiscal sources of instability has come to seem much less pressing. Not surprisingly, the amount of research devoted to these topics has declined precipitously.

Does the evidence *in fact* provide such overwhelming support in favor of the basic claim of existing Real Business Cycle (RBC) theories so as to rationalize this fundamental shift in our view of the business cycle? In my view it does not. This is because the evidence in favor of the proposition that productivity shocks can account for most of the variability in post World War II US output is simply too fragile to be believable.

(I) Small perturbations to the theory alter the conclusion in a basic way.

(II) Small changes in the statistical methods used alter the conclusion in a basic

1

way.

(III) Small changes in the sample period alter the conclusion in a basic way.

(IV) And most importantly, our confidence in the conclusion is fundamentally affected once we abandon the convenient fiction that we actually *know* the true values of the structural parameters of standard RBC models.

Indeed, once we quantify the uncertainty in model predictions arising from uncertainty about model parameter values, calibrated or otherwise, our view of what the data is telling us is affected in a first order way. Even if we do not perturb the standard theory and even if we implement existing formulations of that theory on the standard postwar sample period and even if we use the stationary inducing transformation of the data that has become standard in RBC studies — even then the strong conclusions which mark this literature are unwarranted. What the data are actually telling us is that, while technology shocks almost certainly play some role in generating the business cycle, there is simply an enormous amount of uncertainty about just what percent of aggregate fluctuations they actually do account for. The answer could be 70% as Kydland and Prescott (1989) claim, but the data contain almost no evidence against either the view that the answer is really 5% or that the answer is really 200%.

Under these circumstances, the decision to drastically de-emphasize the importance of traditional impulses like monetary and fiscal shocks in business cycle research ought to be viewed as whimsical, in the sense that Leamer (1983) uses that term. An inference is just not believable if it is fragile. And a decision based on a fragile inference is whimsical.

In this paper I discuss the first and fourth of the aforementioned contentions.[1] To do this it is useful to consider the quantitative implications of one widely used RBC theory —

---

[1]For a discussion of points II and III see Christiano and Eichenbaum (1989) and Burnside, Eichenbaum and Rebelo (1990).

the indivisible labor model associated with Gary Hansen (1985) and Richard Rogerson (1988). According to this model, the time series on the beginning of period t capital stock, k_t, time t consumption, c_t, and time t hours worked, n_t, correspond to the solution of a social planning problem which can be decentralized as a Pareto optimal competitive equilibrium. The planner ranks streams of consumption services, and leisure, $T-n_t$, according to the criterion function:

$$(1) \quad E_0 \sum_{t=0}^{\infty} \beta^t \{\ln(c_t) + \theta(T-n_t)\}.$$

Here T denotes the representative agent's time endowment, E_0 is the time 0 conditional expectations operator, θ is a positive scalar, and β is the subjective discount rate, $0 < \beta < 1$.

There are at least two interpretations of the term involving leisure in (1). First, it may just reflect the assumption that individual utility functions are linear in leisure. The second interpretation builds on the assumption that there are indivisibilities in labor supply, so that individuals can either work some fixed positive number of hours or not at all. Assuming that agents' utility functions are separable across consumption and leisure, Rogerson (1988) shows that a market structure in which individuals choose the probability of being employed rather than actual hours worked will support the Pareto optimal allocation. Under these circumstances, criterion function (1) represents a reduced form preference ordering which can be used to derive the Pareto optimal allocation using a fictitious social planning problem. The parameter θ places no restrictions on the elasticity of labor supply at the micro level of the individual agent. At the macro level, the parameter θ serves only to pin down steady state per capita hours worked. For linear (ln or level) solutions of the type used in the literature, the value of θ has no impact on model

3

moments like the volatility of hours worked or the volatility of output.[2]

Output, y_t, is produced via the Cobb Douglas production function

$$(3) \quad y_t = A_t k_t^{1-\alpha} (\gamma^t h_t)^{\alpha}$$

where $0 < \alpha < 1$, γ^t is the constant unconditional growth rate of technology, and A_t is an aggregate shock to technology which has the time series representation

$$(4) \quad A_t = A_{t-1}^{\rho_a} \exp(\epsilon_t).$$

Here ϵ_t is a serially uncorrelated iid process with mean ϵ and standard error σ_ϵ, and ρ_a is a scalar satisfying $|\rho_a| < 1$.

The aggregate resource constraint is given by

$$(5) \quad c_t + k_{t+1} - (1-\delta)k_t \leq y_t.$$

The parameter δ, which governs the depreciation rate on capital, is a positive scalar satisfying $0 < \delta < 1$.

To discuss the quantitative implications of the theory, it is convenient to denote the model's structural parameters by the vector Ψ. Given a value for Ψ, it is straightforward to deduce the model's implications for a wide variety of moments which might be of interest. For example, the analyst might be interested in understanding the model's quantitative implications for an object like the variance of aggregate output. Existing RBC studies do this by conditioning on a particular value for Ψ and then compare the model's prediction for the variance of output with the corresponding moment in the data. When RBC analysts

---

[2]For a review of some of the solution procedures which have been used in the RBC literature, see Christiano (1990).

say that the model accounts for $\lambda\%$ of the variance of output, what they mean is that their model yields a value of λ given by

(6) $\lambda = \sigma_{ym}^2(\Psi)/\sigma_{yd}^2,$

Here the numerator denotes the variance of model output, calculated for a specific value of Ψ, and the denominator denotes the variance of actual US output. The claim that technology shocks account for most of the fluctuations in postwar US output corresponds to the claim that λ is a large number, with the current estimate being between .75 and 1.0, depending on exactly which RBC model is used (see for example Hansen (1988)).

To evaluate this claim, we abstract, for the moment, from issues like sensitivity to small perturbations in the theory. As decision makers, some obvious things we might want to know are:

How much information is there in the data about λ?
Is our calculation of λ sensitive to small perturbations in Ψ?
And just what is a small perturbation in Ψ?

Unfortunately, the existing RBC literature does not offer much help in answering these questions. Basically this is because that literature makes little use of formal econometric methods, either at the stage when model parameters values are selected, or at the stage when the fully parameterized model is compared to the data. Instead a variety of informal techniques, often referred to as "calibration" are used. Irrespective of what other virtues or defects calibration techniques may possess − one limitation is clear. By ignoring sampling uncertainty in the moments which underlie the values of Ψ that are ultimately adopted, calibration exercises do not lead, in any natural way, to a definition of what a small perturbation in Ψ is. This precludes the possibility of quantifying the sampling

5

uncertainty inherent in model predictions. Consequently, calibration exercises do not provide any information about how loudly the data speak on any given question.

In recent work Lawrence Christiano and I discuss one way to circumvent these problems, in a way that is similar in spirit to existing analyses of RBC models, but which uses formal econometric tools (see Christiano and Eichenbaum (1990)). The basic idea is to use a version of Hansen's (1982) Generalized Method of Moments procedure in which the estimation criterion is set up so that, in effect, the estimated parameter values succeed in equating model and sample *first moments of the data*. As it turns out these values are very similar to the values employed in existing RBC studies. For example, most RBC studies (see for example Prescott (1986)) assume that the quarterly depreciation rate, δ, and the share of capital in the aggregate production function, $(1-\alpha)$, equal .025 and .36, respectively. Our procedure yields point estimates of .021 and .35, respectively.

The key difference between the procedures does not lie so much in the point estimates of Ψ. Rather the difference is that, by using formal econometrics, our procedure allows us to translate sampling uncertainty about the moments which define our estimator of Ψ into sampling uncertainty regarding Ψ itself. This information leads to a natural definition of what a small perturbation in Ψ is, which in turn, makes it possible to quantify uncertainty about the model's second moment implications. The net result is that it is possible to convey how much confidence we have in statements like:

The model accounts for $\lambda\%$ of the variability of output.

Before reporting the results of implementing this procedure for the model discussed above I must digress for one moment and discuss the way in which growth is handled. In practice empirical measures of objects like y_t display marked trends, so that *some* stationary inducing transformation of the data must be adopted. A variety of alternatives are available to the analyst. For example, according to the balanced growth model

6

described above, the data ought to be trend stationary, with the ln of real variables, excluding per capita hours worked, growing as a linear function of time. So one possibility would be to detrend the time series emerging from the model as well as the actual data assuming a liner time trend and calculate the moments of the linearly detrended series.

A different procedure involves detrending model time series and the data using the filter discussed in Hodrick and Prescott (1980). Although our point estimates of the vector Ψ were *not* obtained using transformed data, the second moments results were generated using this transformation of model time series and US data.

I do this for three reasons. First, many authors in the RBC literature report results based on the Hodrick Prescott (HP) filter (see for example Kydland and Prescott (1982), Hansen (1985), Prescott (1986), Kydland and Prescott (1988) and Backus, Kehoe and Kydland (1989)). In order to evaluate their claims, it seems desirable to minimize the differences between our procedures. Second, the HP filter is in fact a stationary inducing transformation for trend stationary processes (see King and Rebelo (1988)). So there is nothing logically wrong with using HP transformed data. Using it just amounts to the assertion that you find a particular set of second moments interesting as diagnostic devices. And third, all of the calculations reported in this paper were also done with linearly detrended data as well as growth rates. The qualitative results are very similar, while the quantitative results provide even stronger evidence in favor of the points I wish to make. So presenting results based on the HP filter seems like an appropriate conservative reporting strategy.

The first row of Table 1 reports results for the baseline indivisible labor model in which the only shocks to the environments are stochastic shifts in the aggregate production technology. Here, σ_a denotes the standard error of the linearly detrended Solow residuals, σ_n denotes the value of the standard error of the ln of per capita hours worked generated by the estimated model, while σ_y denotes the corresponding standard error of the ln of per capita output. The statistic λ_n denotes the ratio of the variance of the ln of per capita

7

hours worked implied by the model to the variance of the ln of actual per capita hours worked in the US. The variable λ_y denotes the ratio of the variance of the ln of per capita output implied by our model to the ratio of the variance of the ln of per capita post war real output. Numbers in parentheses denote the standard errors of the corresponding statistics. All uncertainty in model statistics reflects only uncertainty regarding the values of the structural parameters.[3]

Two key features of these results deserve comment. First, the standard error of σ_a is very large. This is true even though the standard error of our point estimate of the coefficient on capital in the production function is very small (see footnote 3). Second, our point estimate of λ_y equals 80%. This is consistent with claims that technology shocks explain a large percentage of the variability in postwar US output. But notice that the standard error of λ_y is very large. There is simply an *enormous* amount of uncertainty regarding what percent of the variability of output the model accounts for. As it turns out, this uncertainty almost completely reflects uncertainty regarding the law of motion of the Solow residual, ρ_a and σ_ϵ, and hardly at all uncertainty regarding the values of the other parameters of the model.[4]

A different way to summarize this uncertainty is to consider the graph of the confidence interval of λ_y, depicted in figure 1. Each point on the graph is generated by fixing λ at a specific value, λ^*, and then testing the hypothesis that $\sigma_{ym}^2 = \lambda^* \sigma_{yd}^2$. The vertical axis reports the probability value of our test statistic for the corresponding value of λ. To see just how little information this model and the data contain regarding λ, consider the question: What values of λ could we reject at the 5% significance level? The answer is:

---

[3]The data and econometric methodology underlying these estimates are discussed in Burnside, Eichenbaum and Rebelo (1990). Our point estimates of α, θ, δ, ρ_a, and σ_ϵ equal .655 (.006), 3.68 (.04), .021 (.0003), .986 (.002) and .0089 (.02). Numbers in parentheses denote standard errors. We compute σ_a using the formula $\sigma_a = [\sigma_\epsilon^2/(1-\rho_a^2)]^{.5}$.

[4]See Burnside, Eichenbaum and Rebelo (1990).

8

Not many. Even granting that our algorithm breaks down when calculating probability values for negative values of λ, we ought to be very comfortable believing that the model explains *anywhere* between 5% and 200% of the variance in per capita US output. Evidently, the model and the data, taken together, are almost completely uninformative about the role of technology shocks in generating fluctuations in US output.[5] Decisions based solely on the point estimate of λ_y are whimsical in the extreme. If you thought that monetary policy was the key impulse in the business cycle — there is virtually no evidence here to change your mind.

But what about the point estimate itself of λ_y? Just how sensitive is it to small perturbations in the theory?

One interesting perturbation is to consider the effects of labor hoarding on the analysis. Existing RBC studies interpret all movements in measured total factor productivity as being the result of technology shocks or to a much smaller extent as reflecting classical measurement error in hours worked. Various authors, ranging from Lawrence Summers (1986) to Robert Lucas (1989) have conjectured that many of the movements in the Solow residual which are labelled as productivity shocks are actually an artifact of labor hoarding type phenomenon.[6] To the extent that this is true, empirical work which identifies technology shocks with the Solow residual will systematically overstate their importance to the business cycle.

In fact, there is a substantial amount of evidence that the time series properties of

---

[5]Our method for estimating the model's structural parameters amounts to using an exactly identified version of Hansen's (1982) Generalized Method of Moments procedure. Presumably the confidence interval could be narrowed ny imposing more of the model's restrictions, say via a maximum likelihood estimation procedure or an over identified Generalized Method of Moments procedure. Using such procedures would result in substantially different estimates of Ψ, thus making comparisons with the existing RBC literature very difficult. See Christiano and Eichenbaum (1990) for a discussion of this point.

[6]For a more general critique of RBC models, see McCallum (1989).

9

Solow residuals are inconsistent with the notion that they represent exogenous technology shocks. For example Hall (1988) has argued that if Solow residuals represent exogenous technology shocks, then under perfect competition, they ought to be uncorrelated with different measures of fiscal and monetary policy. As it turns out this implication is counterfactual. Evans (1990) has pointed that the Solow residual is actually highly correlated with different measures of the money supply. Hall (1988) himself presents evidence they are also correlated with the growth rate of military expenditures. In interpreting his results as evidence of imperfect competition, Hall argues that labor hoarding alone will not produce significant procyclical behavior in the Solow residual, given perfect competition and flexible prices.

In ongoing research, Craig Burnside, Sergio Rebelo and I have tried to assess the sensitivity of inference based on Solow residual accounting to the Lucas/Summers critique. The model that we use incorporates a particular type of labor hoarding into a perfect competition, complete markets RBC model. Its purpose is to demonstrate, in a quantitative way, the fragility of existing claims about the cyclical role of technology shocks. Our basic findings can be summarized as follows:

(I)   RBC models can, in fact, be quite sensitive to the Lucas/Summers critique. Allowing for labor hoarding in our particular model reduces the ability of technology shocks to account for aggregate output fluctuations by over 50%.

(II)   We find that Hall's (1988) conjecture notwithstanding, labor hoarding with perfect competition and complete markets, is fully capable of accounting for the observed correlation between government consumption and the Solow residual.

Our model setup can be described as follows. Suppose, as in the standard indivisible labor model, that if an individual goes to work there is a fixed cost, ξ, denominated in terms of hours of foregone leisure. If a person does go to work, he stays there for a fixed

10

number of hours, h. The time t criterion of this person is given by

(7) $\ln(c_t^p) + \theta\ln(T - \xi - e_t h)$.

Here c_t^p denotes time t privately purchased consumption and e_t denotes the level of time t effort. The time t criterion function of a person who does not go to work is simply given by

(8) $\ln(c_t^p) + \theta\ln(T)$.

The aggregate production technology is given by

(9) $y_t = A_t k_t^{1-\alpha} (\gamma^t N_t e_t h)^{\alpha}$.

Here N_t denotes the total number of bodies going to work at time t.

Proceeding as in Rogerson (1988) it is easy to show that, since agents' criteria functions are separable across consumption and leisure, the social planner will equate the consumption of employed and unemployed individuals. The Pareto optimal competitive equilibrium corresponds to the solution of the social planning problem

Maximize

(10) $E_0 \sum_{t=0}^{\infty} \beta^t \{ \ln(c_t^p) + \theta N_t \ln(T - \psi - e_t h) + \theta(1 - N_t)\ln(T) \}$,

subject to the aggregate resource constraint

(11) $A_t k_t^{1-\alpha}(\gamma^t N_t e_t h)^{\alpha} = c_t^p + g_t + k_{t+1} - (1-\delta)k_t$.

11

Here g_t represents time t government consumption, which evolves according to

$$(12) \quad g_t = (\gamma^t) g_{t-1}^{\rho_g} \exp(\mu_t),$$

where μ_t is a serially uncorrelated iid process with mean ϵ and standard error σ_ϵ, while ρ_g is a scalar satisfying $|\rho_g| < 1$.[7]

If we assume that the social planner sees the time t realization of the technology shocks and government consumption *before* he chooses N_t and e_t, then this model is observationally equivalent to the standard indivisible labor model, modified to incorporate government consumption into the aggregate resource constraint. The second row of Table 1 reports the results of incorporating g_t alone into the analysis.[8] While the effect of this perturbation is very important for statistics like the correlation between hours worked and real wages (see Christiano and Eichenbaum (1990)), its effect on statistics like σ_y or λ_y is minimal.

How can we perturb the model so as to capture labor hoarding type behavior? One particularly simple way to do this, which does not change the nonstochastic steady state of the model, is to just change the information structure facing agents when they make their work decisions. In particular, suppose that N_t must be chosen *before*, rather than *after*, time t government consumption and the level of technology is known. To provide a bound for the effects of labor hoarding in this setup, we maintain the assumption that the shift length, h, is constant.

The basic idea underlying this perturbation of the base line model is that it is costly for firms to vary the size of their work force. In the limit it is simply not feasible to change

---

[7]See Aiyagari, Christiano and Eichenbaum (1989) for a discussion of the effects of government purchases in the stochastic one sector growth model.

[8]Our point estimates of α, θ, δ, ρ_a, σ_ϵ, ρ_g, and σ_g equal .655 (.006), 3.68 (.04), .021 (.0003), .986 (.027), .0089 (.02), .979 (.021) and 0145 (.001). See Burnside, Eichenbaum and Rebelo (1990) for details.

12

employment in response to *every* bit of new information regarding the state of demand and technology. One way to capture this is to consider environments where firms must make their employment decisions conditional on their views about the future state of demand and technology, and then adjust, within a period of fixed time, to shocks along other dimensions. In our model this adjustment occurs by varying labor effort and is costly because workers care about *effective* hours of work. Consequently labor must be compensated for working harder. We need not be precise about the precise compensation scheme because the optimal decentralized allocation can be found by solving the appropriate social planning problem for our model economy.

Burnside, Eichenbaum and Rebelo (1990) show that, in this model, the ln of the Solow residual, S_t^*, the ln of the true technology shock, A_t^* and the ln of effort, e_t, are, in equilibrium, tied together via the relationship

$$(13) \quad S_t^* = A_t^* + \alpha e_t^*,$$

Here the superscript * denotes the deviation of the ln of a variable from its steady state value. The equilibrium law of motion for e_t^* is of the form

$$(14) \quad e_t = \pi_1 \dot{k}_t^* + \pi_2 \dot{N}_t^* + \pi_3 \dot{A}_t^* + \pi_4 \dot{g}_t^*$$

where the π's are nonlinear functions of the structural parameters of the model.

Given our estimates of the structural parameters, both π_3 and π_4 are positive.[9] This implies that, other things equal, it is optimal to work harder when faced with a positive innovation in government purchases or technology, i.e. *effort will be procyclical.*

---

[9]For this model our point estimates of α, θ, δ, ρ_a, σ_ϵ, ρ_g, and σ_g equal .655 (.006), 4.57 (.17), .021 (.0003), .981 (.027), .0062 (.02), .979 (.021) and .0145 (.001). See Burnside, Eichenbaum and Rebelo (1990) for details.

Consequently naive Solow residual accounting systematically overestimates the level of technology in booms, systematically underestimates the level of technology in recessions and systematically overestimates the variance of the true technology shock.

To understand the dynamic properties of the model, it is useful to consider the impulse response functions of the system, evaluated at our estimates of the model's structural parameters. Excluding the parameters which govern the law of motion of the *true* technology shock, these estimates are almost identical to those of the standard indivisible labor model (see footnotes 6 and 8). Figure 2 presents the response of the system to a 1% innovation in government consumption. By assumption employment cannot immediately respond to this shock. However, effort rises by over 15% in the first period and then reverts to its steady state level. Panel (a) shows the implied movement in the Solow residual. Since effort has gone up in the first period but total hours of work hasn't changed, the Solow residual increases by about .25%. This is true *even though there has been no technology shock whatsoever.* Naive Solow residual accounting falsely interprets the increase in average productivity to a shift in technology rather than an exogenous increase in government consumption. As panel (d) shows, labor productivity rises in the first period by .1% in response to the 1% innovation in government consumption. Like the mechanisms embedded in Lucas (1970) or Hansen and Sargent (1988), this simple perturbation of the model provides, at least in principle, an alternative to technology shocks as the sole explanation for the procyclical behavior of average productivity.

Figure 3 shows how the system responds to a 1% innovation in technology. Given agents willingness to intertemporally substitute effective leisure over time, they respond to the shock in the first period by increasing effort by about .4 of a percent. As a result the Solow residual rises by 1.3% in response to the 1% technology shock. Again naive Solow residual accounting exaggerates the true magnitude of the technology shock.

How do these errors translate into inference for λ_y? From the third row of Table 1 we see that the value of σ_y declines from .017 to .012. This translates into a 50% reduction

14

in λ_y which falls from .82 to .41. Evidently the point estimate of λ_y is quite sensitive to our perturbation of the theory. Notice also that the standard error of λ_y is reduced substantially, at least relative to its value in the standard model. Basically this reflects the fact that our point estimate of the standard error of the true technology shocks drops from .053 to .032.

Figure 4 plots the confidence intervals for λ_y implied by the three models which I have discussed. In all cases the maximal p value occurs at our point estimate of λ_y. Allowing for labor hoarding has two major effects. First, it shifts the whole distribution to the left — this reflects the fact that the point estimate of λ_y is now about .4 rather than about .8. Second, the whole graph becomes more centered around the peak. Now at the 5% significance we can reject values of λ_y that are less than 20% and those that exceed 80%.[10]

Finally before leaving my discussion of the labor hoarding model — let me point to one more bit of subsidiary evidence in favor of that model relative to existing RBC models. Suppose that we regress the growth rate of the Solow residual on the growth rate of government consumption. According to existing RBC models, this regression coefficient ought to equal to zero. In fact it equals .184 and is significantly different from zero.[11] Interestingly, our labor hoarding model implies that the probability limit of this regression coefficient is .104 with standard error of .026.[12] Taking sampling uncertainty into account one cannot reject, at conventional significance levels, the view that the model fully succeeds in accounting for the observed correlation between the Solow residual and government consumption.[13] Standard RBC models obviously cannot.[14]

---

[10] In part this increase in precision reflects the fact that we must impose more of the model's structure in order to disentangle changes in work effort from technology shocks.

[11] The standard error of this regression coefficient is .076.

[12] This standard error reflects sampling uncertainty on our estimates of the model's structural parameters.

[13] Hall (1989) argues that time varying effort is not a plausible explanation for explaining this correlation. To argue this, he first calculates the growth rate of effective labor input required to explain all of the observed movements in total factor productivity. From this measure he subtracts the growth rate of actual hours work to generate a time series on the growth rate in work effort. He argues that the implied movements in work effort are implausibly large. This calculation does not apply to our analysis because it

15

After all is said and done, what is my answer to the question advertised in the title to this talk: "Real Business Cycle Analysis: Wisdom or Whimsy?". My answer is — both. On the whimsy side, I have tried to convince you that the substantive claims in this literature regarding the cyclical role of technology shocks are exceedingly fragile. Decisions based on those claims ought to be viewed as whimsical.

On the wisdom side we *have* learned that dynamic stochastic general equilibrium models can be used to successfully organize our thoughts about the business cycle in a quantitative way. We *have* learned that technology shocks play *some* role in the business cycle. But we have not learned just how large that role is. Finally, to its great credit, work on quantitative Real Business Cycle models has reminded us that empirical work whose *sole* purpose is to answer the question: "Is the Model True" is not likely to very useful. Of course the model is not true. That much should have been obvious before we started. And it has been obvious to theorists all along. To take an obvious example — nobody objects to Lucas' (1972) model of the Phillips curve because old people aren't randomly whisked away in the middle of the night via unobservable helicopters. A formal statistical test which rejected the model because of that fact wouldn't be very useful or change anybody's mind about anything.

Convincing structural empirical work ought to address the question: Does the model succeed *quantitatively* in accounting for those features of the data it was designed to shed light on. But good empirical work also ought to tell us just how loudly the data speak in favor of a given hypothesis. And just as importantly it also ought to help us understand — at what cost did we succeed? What *didn't* we explain? What steps appear to be the most

---

presumes that there are *no* shocks to productivity, an assumption which is clearly at variance with our model.

[14]Burnside, Rebelo and I are currently pursuing our labor hoarding model to see whether it is quantitatively consistent with (a) the fact that average productivity leads the cycle, i.e. average productivity is positively correlated with future output and hours worked, and (b) the fact that average productivity tends to fall at the end of expansions (see Gordon (1979)). McCallum (1989) points out that existing RBC models fail to account for the dynamic correlations between average productivity and output.

promising in the inevitable and ongoing interaction between data and theory?

I conclude by trying to draw one final lesson from the way in which theorists proceed. Theorists are often told to be leery of econometricians bearing free parameters. They already know that they ought to be leery of qualitative conclusions which emerge *only* under highly specialized assumptions. Their response to this problem is to engage in *theoretical* fragility analyses. Indeed, Robert Lucas' (1989) paper on "The Effects of Monetary Shocks When Prices Are Set in Advance" provides an excellent example of this type of analysis. In motivating his paper Lucas writes:

> Models of monetary economies necessarily depend on the assumed conventions about the way in which business is conducted in the absence of complete markets, about who does what, when, and what information he has when he does it. Such conventions are necessarily highly specific, relative to the enormous variety of trading practices we observe. Do the various rigid price models have enough in common to have useful empirical or policy implications, or does everything hinge on the accuracy of assumptions in constructing each specific example?

Lucas concludes that the substantive implications which emerge from this class of models are, in fact, quite robust. Whether one agrees or not is not germane to this talk. What is germane is the effort to address the question.

Unfortunately, despite some important exceptions, notably Ed Leamer's (1984) work on international trade and recent work by Hansen, Sargent and Roberds (1990) on the time series implications of martingale models of consumption and taxes, it's hard to think of many analog examples in the empirical literature. The time has come for more empiricists to follow suit. Absent a greater willingness to engage in empirical fragility analysis, structural empirical work will simply cease to be relevant. We may continue to publish, but our influence will surely perish.
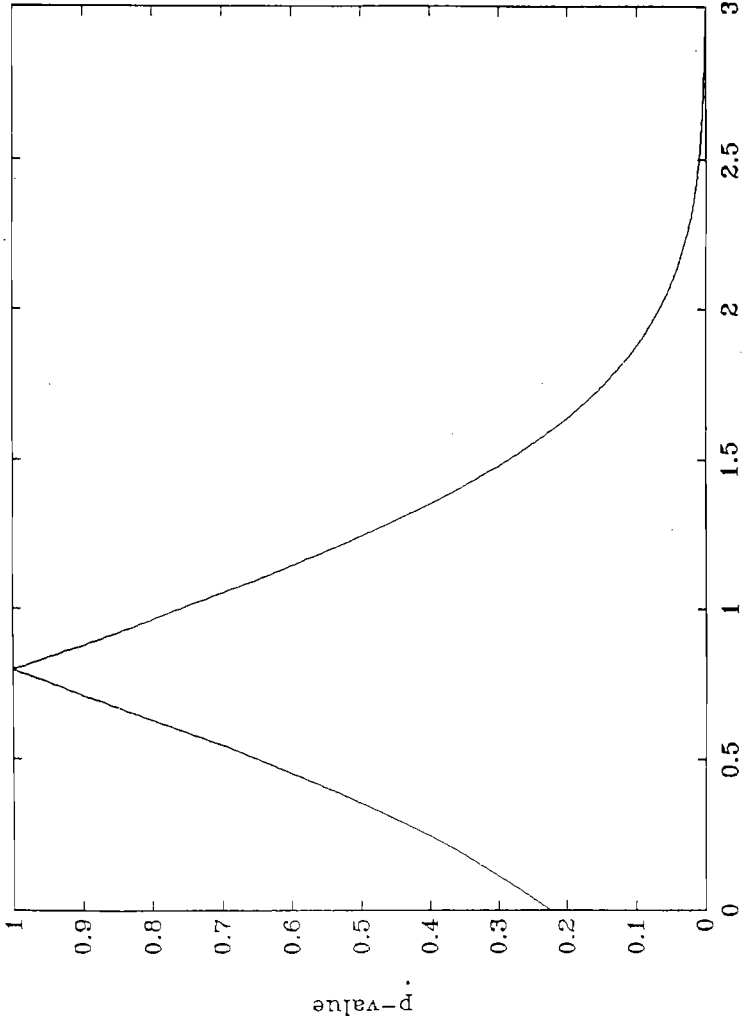
17

References

Burnside, Craig, Rebelo, Sergio T. and Eichenbaum, M., "Labor Hoarding and the Business Cycle," 1990, manuscript, Northwestern University.

Blanchard, Olivier J., "A Traditional Interpretation of Macroeconomic Fluctuations," *American Economic Review*, 1989, 79, 1146–1164.

Backus, David K., Kehoe, Patrick J. and Kydland, Finn E., "International Trade and Business Cycles," 1989, Federal Reserve Bank of Minneapolis Working Paper 425.

Christiano, Lawrence J. and Martin Eichenbaum, "Current Real Business Cycle Theories and Aggregate Labor Market Fluctuations", manuscript, Northwestern University, 1990.

Christiano, Lawrence J., Solving the Stochastic Growth Model by Linear–Quadratic Approximation and by Value Function Approximation, 1990, 8, 23–26.

Evans, Charles L.,"Productivity Shocks and Real Business Cycles," 1990, manuscript, University of South Carolina.

Gordon, Robert J., "The End–of–Expansion Phenomenon in Short–Run Productivity Behavior," *Brookings Papers on Economic Activity*, 1979 No. 2, pp. 447–61.

Hall, Robert E., "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, October 1988, 96, 921–47.

Hall, Robert E., "Invariance Properties of the Solow Residual", 1989, NBER Working Paper No. 3034.

Hansen, Gary D., "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, November 1985, 16, 309–28.

Hansen, Gary D., "Technical Progress and Aggregate Fluctuations," 1988, manuscript, University of California, Los Angeles.

Hansen, Gary D. and Sargent, Thomas J., "Straight Time and Over–Time in Equilibrium," *Journal of Monetary Economics*, March/May 1988, 21, 281–308

Hansen, Lars P., Sargent, Thomas J. and Roberds, William, "Implications of Expected Present Value Budget Balance: Application to Postwar Data," 1990, manuscript, University of Chicago.

Hansen, Lars P., "Large Sample Properties of Generalized Method of Moments Estimators," *Econometrica*, 1982, 50, 1029–54.

Hodrick, Robert J. and Prescott, Edward C., "Post–War U.S. Business Cycles: An Empirical Investigation, manuscript, Carnegie–Mellon University, 1980.

King, Robert G. and Rebelo, Sergio T., "Low Frequency Filtering and Real Business Cycles," manuscript, University of Rochester, 1988.

Kydland, Finn E. and Prescott, Edward C., "Time to Build and Aggregate Fluctuations," *Econometrica*, November 1982, 50, 1345–70.

Kydland, Finn E. and Prescott, Edward C., "The Work Week of Capital and Its Cyclical Implications," *Journal of Monetary Economics*, March/May 1988, 343–60.

Kydland, Finn E. and Prescott, Edward C., "Hours and Employment Variation in Business Cycle Theory," August 1989, Institute for Empirical Economics, Discussion Paper No. 17.

Leamer, E.E., "Let's Take the Con Out of Econometrics," *American Economic Review*, March 1983, 73, 31–43.

Leamer, E.E., "Sources of International Comparative Advantage: Theory and Evidence," 1984, MIT Press, Cambridge MA.

Lucas, Robert E. Jr., "Capacity, Overtime, and Empirical Production Functions," *American   Economic Review, Papers and Proceedings*, 1970, 6, 1345–1371.

Lucas, Robert E. Jr., "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 1972, 4, 103–124.

Lucas, Robert E. Jr., "The Effects of Monetary Shocks When Prices Are Set In Advance," November 1989, manuscript, University of Chicago.

McCallum, Bennett T., "Real Business Cycle Models," in Modern Business Cycle Theory, ed. by Robert J. Barro, 1989, 16–50, Harvard University Press.

Prescott, Edward, C., "Theory Ahead of Business Cycle Measurement," Federal Reserve Bank of Minneapolis *Quarterly Review*, Fall 1986, 10, 9–22.

Rogerson, R., "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, January 1988, 21, 3–17.

Summers, Lawrence J., "Some Skeptical Observations on Real Business Cycle Theory,"Federal Reserve Bank of Minneapolis *Quarterly Review*, Fall 1986, 10, 23–27.
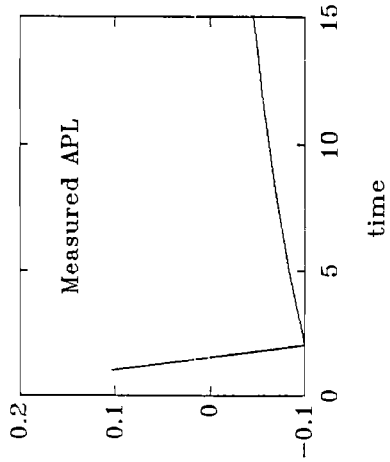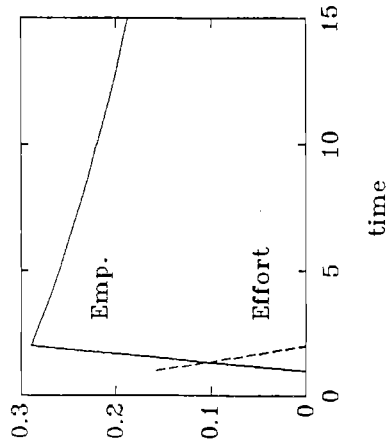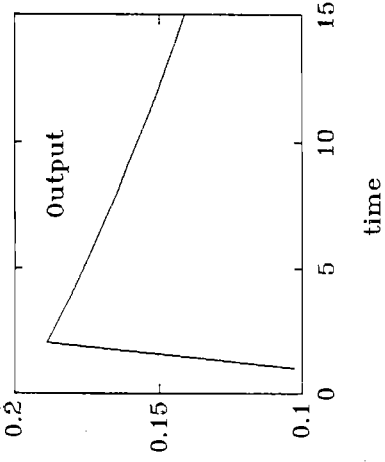
TABLE 1

SUMMARY STATISTICS

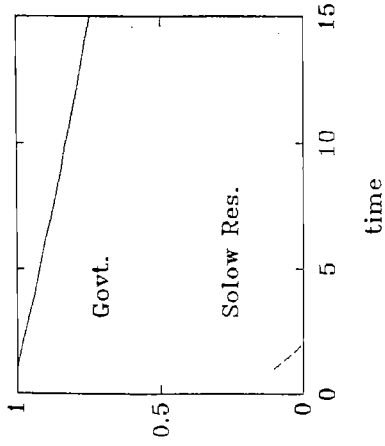|	σ_a	σ_n	σ_y	λ_n	λ_y
Standard Model (Constant Government)	.053 (.046)	.012 (.005)	.017 (.007)	.52 (.47)	.78 (.64)
Standard Model (Variable Government)	.053 (.046)	.013 (.005)	.017 (.007)	.64 (.47)	.82 (.64)
Labor Hoarding	.032 (.031)	.012 (.002)	.012 (.002)	.51 (.21)	.41 (.11)
US Data		.016 (.002)	.019 (.002)		

1

Standard Model (Constant Government)

p-value

Variance of Y (Model) / Variance of Y (Data)

SHOCK TO GOVERNMENT

SHOCK TO TECHNOLOGY

Output

Solow Res.

Tech. Shock

time

time

Measured APL

Emp.

Effort

time

time