NBER WORKING PAPER SERIES

LEARNING BY DOING AND AGGREGATE FLUCTUATIONS

Russell Cooper
Alok Johri

Working Paper 6898
http://www.nber.org/papers/w6898

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 1999

Early versions of this paper were presented at the 1997 NBER Summer Institute Meeting of the Macroeconomic Complementarities Group, the 1997 Canadian Macro Study Group, the 1998 North American meetings of the Econometric Society, and the 1998 NBER Summer Institute meeting of the Impulse and Propagation Group as well as the University of British Columbia, University of Toronto, Queens University, SUNY Buffalo and York. We thank participants at these conferences and seminars as well as V.V. Chari, L. Christiano, M. Eichenbaum, P. Kuhn, and L. Magee for useful comments. We are grateful to Jon Willis for outstanding research assistance on this project. Cooper acknowledges financial support from the National Science Foundation. Johri thanks the Social Science and Humanities Research Council and the Arts Research Board. This research was partially conducted at the Boston Research Data Center. This paper has been screened to insure that no confidential information is released. The views expressed here are those of the author and do not reflect those of the National Bureau of Economic Research or of the U.S. Bureau of Census.

Learning by Doing and Aggregate Fluctuations
Russell Cooper and Alok Johri
NBER Working Paper No. 6898
January 1999

ABSTRACT

A major unresolved issue in business cycle theory is the construction of an endogenous propagation mechanism capable of capturing the amount of persistence displayed in the data. In this paper we explore the quantitative implications of one propagation mechanism: learning by doing. Estimation of the parameters characterizing learning by doing is based both on aggregate 2-digit data and plant level observations in the US. The estimated learning by doing function is then integrated into a stochastic growth model in which fluctuations are driven by technology shocks. We conclude that learning by doing can be a powerful mechanism for generating endogenous persistence.

Russell Cooper
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
rcooper@acs.bu.edu

Alok Johri
Department of Economics
McMaster University
1280 Main Street West
Hamilton, Ontario
Canada L8S 4M4
johria@mcmaster.ca

I. Motivation

One of the major unresolved issues in business cycle theory is the construction of an endogenous propagation mechanism capable of capturing the amount of persistence observed in the data. The fact that the standard real business cycle (RBC) model has weak internal propagation mechanism is evident from early work on these models, such as King, Plosser and Rebelo [1988]. This point has been made again in recent papers by Cogley and Nason [1995] and Rotemberg-Woodford [1996] in the context of output growth: the data indicate that U.S. output growth is positively autocorrelated in contrast to the predictions of the standard RBC model.

In this paper we study the role of internal learning by doing in propagating shocks. This is done by supplementing an otherwise standard representative agent stochastic growth model by introducing internal learning by doing.[1] As we shall see, this creates a new state variable, which we label organizational capital, that reflects past levels of activity at the plant or firm. The idea is simply to capture the fact that the production process creates information about the organization of the production facility that improves productivity in the future. As discussed at some length in Section II, there is abundant empirical support for learning by doing. We see our notion of organizational capital as including the accumulation of experience by workers and managers.

Section III presents our analysis of the stochastic growth model with organizational

[1] Cooper-Johri [1997] investigates the role of dynamic complementarities in the production function for the propagation of temporary shocks to productivity and tastes and find the propagation effects can be quite strong : an iid technology shock creates serial correlation in output of about .95 using their estimated complementarities. In the present paper these complementarities are ignored in order to isolate the effects of internal learning by doing.

capital. As discussed in Section IV, we find empirical support for these learning effects in both aggregate and plant level data. Finally, the parameterized model is simulated to understand the effects of productivity shocks, as described in Section V.

Overall, we find some interesting implications of introducing internal learning by doing into the stochastic growth model. In particular, learning by doing can be a powerful propagation mechanism. Using our aggregate estimates the first two autocorrelation coefficients of the simulated output series can be as high as .53 and .42 even when the model is subject to **iid** technology shocks. In the presence of serially correlated technology shocks, we are able to generate the hump-shaped impulse response functions documented in Cogley-Nason [1995]. Further, for the case of stochastic trends, when technology follows a random walk, we find that our model produces an autocorrelation function for real output growth with positive coefficients for at least two periods. This is much closer to the autocorrelation function for output growth in U.S. data and contrasts sharply with the predictions of the RBC model without learning by doing.

Introducing learning by doing into the standard model creates another state variable whose movement shifts both labor supply and labor demand. This has a number of interesting implications that we explore as well. First, the model predicts a lower correlation between productivity and employment, in keeping with the finding of Christiano-Eichenbaum [1992]. Second, this same device is helpful in understanding the source of "taste shocks" and "unobserved effort" in aggregate fluctuations.

II. Previous Studies of Learning by Doing

The study of learning by doing (LBD) dates back to the turn of the century (see references in Jovanovic and Nyarko [1995]). Our specification, presented in detail in the next section, has two key elements:

- increases in output leads to the accumulation of organizational capital
- organizational capital depreciates.

We relate these components to the existing literature before proceeding with our analysis.

Since Wright's [1936] work, the typical form has involved estimating how costs of production fall as experience rises.[2] Generally, studies of learning by doing use cumulative output as a measure of experience which enters the cost function as a separate input. These studies use output price as a proxy for marginal cost.

In general these studies find considerable evidence for learning by doing in that costs tend to fall with cumulative output. For example, Irwin and Klenow [1994] report learning rates of 20% in the semi-conductor industry. Jarmin [1994], measures LBD in the early US rayon industry without having to use price as a proxy for marginal cost. Taking into account strategic interactions among producers, Jarmin reports negative first and positive second order learning effects on the cost function, i.e., costs fall as cumulative output increases but at a decreasing rate.

Another widely cited micro study of LBD by Bahk and Gort [1993] introduces experience into the production function as a factor that influences productivity. Bahk and Gort construct a dataset of new manufacturing plants and this allows them to construct two measures of the stock of experience: cumulative output since birth and time since birth. While the authors do not allow experience to depreciate over time, they do allow for learning to decline to zero over time.

Since they are interested in studying the effects of LBD separately by production factor,

---

[2] References to other early work can also be found in Bahk and Gort [1993].

they are careful to decompose productivity enhancements into two parts: those that can really be attributed to a change in inputs if measured in efficiency units and those that result from accumulation of experience. To capture the former effects they introduce human capital (measured by average wages) and the average vintage of capital as inputs that are separate from raw labor and capital. The latter effects are captured in two ways: one formulation introduces the stock of experience as a separate input and the other proceeds by allowing the Cobb Douglas input coefficients to change over time i.e., the coefficients are functions of time since the birth of the plant. Notice that these specifications are only able to get at LBD that is specific to the firm, any learning that is captured by the employee in the form of skills is lumped into the human capital measure.

When experience was proxied by cumulative output per unit of labor, a 1 percent change in cumulative output lead to a .08 per cent change in output. Using the other specification, Bahk and Gort find that capital learning continues for five to six years, labor or organizational learning appears to continue for ten years but results were relatively unstable.

All of these studies take a very specific view of LBD which is much narrower than that envisaged by us. This view is that learning occurs when a new plant is set up, or when a new or old industry introduces new technology. Our view is that the stock of organizational capital fluctuates over high frequencies due to learning and dislearning which occur in many less dramatic everyday situations: when production teams are re-organized; every time workers are hired or fired; when employees are promoted or redeployed to new tasks; when new capital or software is installed; when new management, supervision, or bookkeeping practices are introduced and so on. The list is potentially endless-so learning should be widespread, but often will last for a limited duration. Note also that since a lot of this learning is very specific to a task or a match there is probably a considerable loss of organizational capital. So, if previous experience does not readily carry over to new environments (as suggested by the work of Irwin and Klenow for example), then depreciation of experience is high.

The study that comes closest to our view of learning as a widespread phenomena is Jovanovic and Nyarko [1995]. Unlike us, Jovanovic and Nyarko take an information theoretic approach to LBD, i.e., the agent learns about the best decision to take by observing the consequences of his past actions which slowly reveal in a Bayesian manner the true value of an unknown parameter which itself represents the "right" action to take. As the agent comes closer to this value, productivity increases and eventually stops when the parameter is attained. The authors fit their Bayesian learning process to a number of micro situations using both the cumulative number of times a task is performed (this could be production at a plant as well as an individual performing a specific task) as well as experience measured in terms of time elapsed at the job. These situations all appear to display a lot of concavity in the learning process as well as a level of experience after which productivity growth stops altogether.

In terms of microeconomic evidence on the depreciation of organizational capital, Benkard [1997] studies learning by doing in the commercial aircraft building industry using a closely related specification of technology. In contrast to many studies, Benkard allows for depreciation of experience. He finds that even though labor requirements fall as cumulative output rises, this process is not monotonic: there are several periods in which labor requirements actually increase before they start falling again. Using a generalized method of moments procedure, Benkard shows that the model with depreciation of experience (which he refers to as organizational forgetting) is better able to account for the data than the traditional learning model. Based on a careful analysis of the specific features of aircraft production technology and the nature of union contracts in the industry he concludes that a large part of the estimated depreciation of experience may be explained by labor turnover and redeployment of existing workers to new tasks within the firm.

On the issue of the generality of Benkard's results regarding the depreciation of experience we first refer the reader to Benkard [1997] that lists other studies recording organizational forgetting in industries as diverse as ship production and pizza production. Secondly, to the extent that depreciation of experience is the result of not being able to re-hire previously laid-off workers we can look for evidence on this issue. According to Katz and Meyer [1990], for the 1979-1982 period only about 42 % of laid-off workers were recalled to their old jobs over the period of about a year.[3] The data was collected mainly on blue collar workers in non-agricultural sectors but included some from trade, services and administration. Thirdly, in the estimation section we provide our own microeconomic evidence from automobile assembly plants, though our specification differs a bit from that used by Benkard.

Overall, there appears to be substantial support in the literature for learning by doing. Given the wide array of studies, it is difficult to settle on estimates based on a single microeconomic study. For our purposes, we take this existing literature as supportive of our interest in understanding the aggregate effects of learning by doing. To study this quantitatively, we provide our own estimates of these effects, motivated, of course, by the extant literature.

[3] Their reported numbers vary depending on the states included, the specific sub period used and for different sectors, though the highest recall rate they reported was 57% from a dataset suspected of being biased towards seasonal layoffs.

III. The Model

It is convenient to represent the choice problem of the representative agent through a stochastic dynamic programming problem. We first present a general version of the model and then consider a specific example.

A. General Specification

Here we consider the representative household as having access to a production technology that converts its inputs of capital (K), labor (N) and organizational capital (H) into output (Y).[4] This technology is given by Af(K,N,H) where the total factor productivity shock is denoted by A.

The household has preferences over consumption (C) and leisure (L) denoted by u(C,L) where this function is increasing in both arguments and is quasi-concave. Assume that the household has a unit of time each period to allocate between work and leisure: 1=L+N. The household allocates current output between consumption and investment (I).

There are two stocks of capital for the household. The first is physical capital (K). The accumulation equation for physical capital is traditional and is given by:

$$K' = K(1-\delta_K) + I. \tag{1}$$

In this expression, δ_K measures the rate of physical depreciation of the capital stock. The second stock is organizational capital which is accumulated indirectly through the process of production. The evolution of this stock is given by:

---

[4]Using this integrated worker/firm model of we are thus agnostic about the question of whether the organizational capital is firm or worker specific. As indicated by our discussion of evidence, there are arguments in favor of both firm and worker specific accumulation. This distinction does not strike us as critical for our investigation.

$$H' = \varphi(H,Y) \tag{2}$$

where $\varphi(\cdot)$ is increasing in both of its arguments. In (2) we have assumed that the accumulation of organizational capital is influenced by current output rather than current employment. As discussed above in Section II, this is apparently the traditional approach.

For our analysis, it is convenient to substitute the production function for Y in (2) and then to solve for the number of hours worked in order to accumulate a stock of organizational capital of H′ in the next period given the two stocks (H,K) today and the productivity shock, A. This function is defined as:

$$N = \mathbf{\Phi}(K,H,H',A). \tag{3}$$

Given that the inputs (K,N,H) are productive in the creation of output and the assumption that organizational capital tomorrow is increasing in output today, $\mathbf{\Phi}$ will be a decreasing function of both K and H and an increasing function of H′.

The dynamic programming problem for the representative household is then given by

$$V(A,K,H) = \max_{K',H'} u(Af(K,H,N) + (1-\delta)K - K', N) + \beta E_{A'} V(A',K',H') \tag{4}$$

where we use (3) to substitute out for N. The existence of a value function satisfying (4) is standard as long as the problem is bounded.

The necessary conditions for an optimal solution are:

$$[u_c(c,1-N)Af_N - u_L(c,1-N)]\,\mathbf{\Phi}_{H'} + \beta E V_{H'}(A',K',H') = 0, \tag{5}$$

and

$$u_c = \beta E V_K(A',K',H') \ . \tag{6}$$

These two conditions, along with the transversality conditions, will characterize the optimal solution.

Equation (5) is the analogue of the standard first order condition on labor supply though in this more complex economy, it includes the effects of current labor input on the future organizational capital stock. Thus the accumulation of organizational capital is one of the benefits of additional work (i.e. $V_H>0$ for all points in the state space) leading to a labor supply condition in which the current marginal utility of consumption less the disutility of work is negative. We term this "excessive labor supply" in the discussion that follows.

To develop this condition further, we use (4) to find that

$$V_H(A,K,H) = u_c A f_H + [u_c A f_N - u_L]\mathbf{\Phi}_H. \tag{7}$$

So, giving the agent some additional organizational capital will directly increase utility through the extra consumption generated by this additional input into the production process. This is captured by the first term in (7). Second, given the higher level of H today and assuming that $\mathbf{\Phi}_H<0$, the agent can reduce labor supply in the current period which, given the excessive labor supply, is desirable. Thus the second term in (7) is the current utility gain from reducing labor supply times the reduction in employment created by the additional H. Of course, this condition is used in (5) once it is updated to the following period.

Equation (6) is the Euler equation for the accumulation of physical capital: the marginal utility of consumption today is equated to the discounted value of more capital in the next period.

As with the labor supply decision, a gain to investment is the added output in the next period plus the accumulation of organizational capital that comes as a joint product. So, again using (4),

$$V_K(A,K,H) = u_c[Af_K + (1-\delta)] + [u_c Af_N - u_L]\Phi_K . \quad \textbf{(8)}$$

So, an additional unit of physical capital increases consumption directly, the standard result, and also allows the agent to work a bit less to offset the effects of the physical capital on the accumulation of organizational capital. As before, the updated version of (8) is used in (6) to complete the statement of the Euler equations.

In principle, one can characterize the policy functions through these necessary conditions. Alternative, the economy can be linearized around the steady state (assuming it exists) and then the linear system can be evaluated and the quantitative analysis undertaken. That is the approach taken here through a leading example presented in the next sub-section.

B. **A Leading Example**

Here we assume some specific functional forms for the analysis. These restrictions are used here to illustrate the model and then are imposed in some of our estimation/simulation exercises.

We assume that the production function displays constant returns to scale and is given by

$$Af(K,H,N) = AH^{\varepsilon}K^{\theta}N^{\alpha}. \quad \textbf{(9)}$$

where $\alpha+\theta+\varepsilon=1$. In this specification, ε parameterizes the effects of organizational capital on output.

The accumulation equation for organizational capital is specified as

$$\varphi(H,y) = H^{\gamma} y^{\eta}. \tag{10}$$

where γ captures the influence of current organizational capital on the accumulation of additional capital and η parameterizes the influence of current output on the accumulation of human capital. With the additional restriction that $\gamma + \eta = 1$, the model will display balanced growth.

Using the production function in the accumulation equation implies:

$$H' = H^{\gamma}(AH^{\varepsilon}K^{\theta}N^{\alpha})^{\eta} = H^{\gamma+\varepsilon\eta}(AK^{\theta}N^{\alpha})^{\eta} . \tag{11}$$

In this case, Φ(K,H,H',A) becomes

$$\left(\frac{H'}{H^{\gamma+\varepsilon\eta}(AK^{\theta})^{\eta}}\right)^{\frac{1}{\alpha\eta}} \tag{12}$$

So that N is increasing in H′ and decreasing in H, K and A as noted above.

Finally, assume the utility function is given by $u(c,N)=\ln(c)+\chi(1-N)$. Here χ parameterizes the contemporaneous marginal rate of substitution between consumption and leisure.

With these particular functional forms, the necessary conditions for an optimal solution become:

$$[\frac{\alpha Y}{CN} - \chi][\frac{N}{\alpha\eta H'}] + \beta E[[\frac{1}{C'}\ \frac{\varepsilon Y'}{H'}] - [\frac{\alpha Y'}{C'N'} - \chi][(\gamma+\varepsilon\eta)\frac{N'}{\alpha\eta H'})]] = 0. \tag{13}$$

and

$$\frac{1}{C} = \beta E[(\frac{1}{C'})(\frac{\theta Y'}{K'}+1-\delta) - (\frac{1}{C'}(\frac{\alpha Y'}{N'})-\chi)\ \frac{\theta N'}{\alpha K'}] \tag{14}$$

So, (13) and (14) are the analogues of (5) and (6) for this particular specification of functional forms. These conditions, along with the accumulation equations and resource constraint, fully characterize the equilibrium of the model.

IV. Estimation

The parameterization of the model utilizes both estimation and calibration techniques. The point of the estimation is to focus on the parameters of the production function and the accumulation technology. This procedure and our findings are discussed in some detail in this section. These parameters are estimated in two ways: in the first sub-section we directly estimate the technology using production function estimation techniques. The next sub-section discusses results from estimating the Euler equation. The remaining parameters, as in our earlier paper, King Plosser and Rebelo and many of the references therein, are calibrated from other evidence and are discussed in Section V.

A. **Production Function Estimates Using Sectoral Data**

We estimate our production technology and accumulation equation for organizational capital simultaneously using quarterly 2-digit manufacturing data for seventeen US manufacturing sectors. As is well known, the quarterly data display unit roots so all the following estimation exercises are done using data that has been rendered stationary using log first differences. The equivalent expression for (9) in log first differences is:

$$\Delta y_{it} = \alpha \Delta n_{it} + \theta \Delta k_{it} + \varepsilon \Delta h_{it-1} + \Delta a_{it} \tag{15}$$

where the lower case letters denote logs of variables and the subscript i refers to the i[th] 2-digit sector. The accumulation equation (10) may be written as:

$$\Delta h_t = \gamma \Delta h_{t-1} + \eta \Delta y_{t-1} = \frac{\eta}{(1-\gamma L)} \Delta y_{t-1} \tag{16}$$

where L is the lag operator and Δ denotes first differences. Replacing this expression in (15) and rearranging yields our first specification in Table 1, labeled SPEC 1, which corresponds to:

$$\begin{aligned} \Delta y_{it} = {} & \alpha \Delta n_{it} - \alpha \gamma \Delta n_{it-1} + \theta \Delta k_{it} - \theta \gamma \Delta k_{it-1} \\ & + (\gamma + (1-\alpha-\theta)(1-\gamma)) \Delta y_{it-1} + \Delta \hat{a}_{it} \end{aligned} \tag{17}$$

Note that we have imposed constant returns to scale on the estimation exercise. Since the parameters of interest are overidentified, we use a non-linear procedure to estimate them.

The first row of Table 1 reports the results for a non-linear system instrumental variable procedure where seventeen sectors are jointly estimated with the coefficients restricted to be the same across sectors. The variables used for estimation are quarterly data on gross output and hours in the US manufacturing sector at the 2-digit level from 72:1 -92:4, as in Burnside, Eichenbaum and Rebelo. The capital input is proxied by electricity consumption. The instrument list includes the second, third and fourth lag of innovations to the federal funds rate and to non borrowed reserves. The labor share, α, is estimated at .59, the share of physical capital θ is estimated at .33, the share of organizational capital is .08 while γ, which parameterizes the depreciation of organizational capital is estimated at .63. Throughout the paper, this is referred to as SPEC 1.

The second row of Table 1 corresponds to a specification in which we have imposed CRS in the production function but not in the accumulation equation. In this case, our model mimics the approach taken in many of the industrial organization studies where current output feeds directly into the accumulation of organizational capital. Throughout, we refer to this as SPEC 2.

One potential concern with these estimates is that the lagged terms may simply be picking up serial correlation in the error. This is addressed in row 3 of Table 1 which uses a non-linear IV procedure to jointly estimate ε, γ and ρ (the coefficient capturing serial correlation in the technology shock) in the following specification:

$$\begin{gathered}\Delta y_{it}=\alpha\Delta n_{it}-\alpha(\rho+\gamma)\Delta n_{it-1}+\\ \alpha\gamma\rho\Delta n_{it-2}+(\varepsilon+\rho+\gamma)\Delta y_{it-1}-\rho(\varepsilon+\gamma)\Delta y_{it-2}+\Delta\hat{a}_{it}\end{gathered} \tag{18}$$

Capital was dropped from the specification because the coefficients on the second lag of electricity were not significantly different from zero.[5]

With this specification, all coefficients except ρ are significantly different from zero. We found α=.44, ε = .31, γ = .39 and ρ = -.1. We also ran a specification which imposed $\gamma = 0$ on (18) in an attempt to see if serial correlation emerges. We found an estimate of ρ = .06 (other coefficients were similar) which was barely significant at the 5% level. This further strengthened the view that our estimates do not merely reflect serial correlation in the error term. Throughout, we refer to this as SPEC 3.

---

[5] Allowing capital affected the estimates of the coefficients a lot, generating increasing returns to scale (above 1.5) in labor and capital on one hand and insignificant estimates on ε, and ρ on the other hand. We have also not imposed constant returns to scale in the accumulation equation.

Overall, the evidence from aggregate 2-digit manufacturing data seems to us to strongly suggest the presence of learning by doing effects at the macro level. Further, these seem to reflect the effects of past output on current productivity rather than the serial correlation of underlying shocks.

B. **Euler Equation Estimation**

In this subsection we use a version of the labor supply Euler equation (13) to estimate the parameters of interest using aggregate data. For this purpose we work with the linear utility in leisure case where U(c, 1-N) = log(c) + (1-N). With these preferences (13) may be written as:

$$\frac{\alpha Y}{CN}+\beta E[(\gamma+\varepsilon(1-\gamma))\frac{N'}{N}-\alpha\gamma\frac{Y'}{C'N}]-1=0 \qquad \textbf{(19)}$$

The data used for this estimation exercise was the national income and product accounts quarterly series on real GDP, real non durable consumption and hours from 1964:1 to 1997:1.

The discount rate β was set to .98 and instruments used were the first seven lags of (C'/C) and (Y'/Y) which correspond to two years worth of information from the information set available at date t. These variables should be uncorrelated with expectational errors at date t. The estimated coefficients were α=.54 (2.7), ε=.5 (2.48) and γ=.8 (5.67) where the t-statistics are given in parentheses.

Clearly the estimated learning by doing from the Euler equation exercise is much larger than from the sectoral production functions. Note too that the evidence from the Euler equation implies that at least some of the learning by doing is internal to the firm.

C. **Plant-level Estimates**

Quite apart from the obvious advantages of estimating learning directly at the micro level, the plant level estimation also allows one to distinguish between learning that is internal to plants and external learning at the industry level. As in Cooper-Johri [1997], our plant level estimates come from the Longitudinal Research Database (LRD). In particular, we look at 49 continuing automobile assembly plants over the 1972-91 period. While this is certainly only a small subset of manufacturing plants, it is a group of plants that we have studied before and thus we have a benchmark for comparison.[6]

Our results are reported in Table 2. Here, all regressions include plant specific fixed effects.[7] The output measure is real value added at the plant level and the inputs are labor (specifically production worker hours, labeled ph) and physical capital (machinery and equipment at the plant level, labeled k).[8] The different specifications refer to the measure of organizational capital used in the estimation and the treatment of a time trend.

We consider two different measures of organizational capital. Lagged output at the plant is denoted lqv and corresponds to a specification in which γ=0. The variable cumv represents an alternative measure of experience which is a running sum of past output: i.e., it measures cumulative output and comes closest in spirit to the measure used by Bahk and Gort and most studies of LBD. This measure corresponds to the case of no depreciation of experience.[9] To

[6] In fact, this is a slightly different sample than used in Cooper-Johri [1997] due to the inclusion of one additional plant and two more years of data. Obviously one goal in this continuing research is to go beyond this group of plants.

[7] One issue to consider is the extent of the bias created by estimation with fixed effects and lagged dependent variables.

[8] This is in contrast to our previous study where we used electricity consumption as a proxy for capital utilization. Note that this measure of capital excludes the value of the plant.

[9] We are in the process of developing a non-linear estimation routine to allow for arbitrary depreciation.

estimate this model we start the experience accumulation process off at an arbitrary date. Obviously at that date different plants have different levels of experience so this level effect (of different levels of initial experience) is picked up in the fixed effects.

There are also three treatments of a trend that might come from technological advance. For some specifications the trend is ignored. For others, it is captured as a linear trend and finally we also allow for year dummies. In general, allowing for some time effects seem to matter though the distinction between a linear trend and time dummies doesn't seem to have important implications for other coefficients.

The results are generally supportive of some form of learning by doing at the plant level: the coefficients for the various output- based measures of organizational capital are statistically significant at the 5% level and range from .22 to .38 depending on the measure used and the trend.

V. Simulations

The simulations are based upon a linearized version of the model specified in Section III parameterized using the estimates from Section IV and other 'standard' parameters. In particular, we set β=.984 so that the real interest rate is 6.5% annually, δ=.1 and χ is chosen so that the average time spent working is .3. We report results for preferences given by:

$$u(c,1-N) = \log(c)+\chi(1-N).$$

As for the organizational capital part of the model, we consider a couple of alternative specifications to evaluate their implications for the behavior of the aggregate economy.

Since the main goal of the paper is to analyze the contribution of learning by doing to the propagation of shocks we study this issue using three diagnostic tools. First we study the model under iid technology shocks as this seems to highlight the ability of the model to propagate shocks. Second, using highly correlated technology shocks we compare the model generated data with key properties of US macro data in log-levels. The key diagnostic here will be the ability of the model to generate hump-shaped impulse response functions. Third, we analyze a version of our model assuming that technology shocks are permanent. With this stochastic trends formulation, we are better equipped to investigate the implications of the models for the autocorrelation of output growth. Here a key test of the model will be its ability to replicate the two positive autoregressive coefficients found in aggregate output data when measured in growth rates. Finally, in the last sub-section we conduct some counter-factual exercises based on the idea that when learning by doing effects are ignored, estimation exercises based on the Euler equation are likely to be mis-specified.

A. **Simulation Results from IID Technology Shocks**

To conduct the quantitative analysis, we consider a log linear approximation to the equilibrium conditions described by (13) and (14) plus the resource and accumulation conditions, around the steady state.[10] Using this system, our main question is how much propagation is created by internal learning by doing? We address this question by introducing temporary technology shocks into the model.

Table 3 summarizes our findings for the three different treatments we consider. Here we

[10]This is the same procedure followed in our earlier paper using a version of the economy specified in King, Plosser and Rebelo [1988]. For our parameterizations, the steady state was saddle path stable.

report statistics from the artificial economy for major macroeconomic variables: output (Y), consumption (C), investment (I), total hours (N) and average labor productivity, (W). For each, we present standard deviations of these variables relative to output, their contemporaneous correlations with output and the serial correlation of output. The first column of the table reports various specifications, which refer to the estimates reported in Table 1.

The first row of the table provides results for the baseline real business cycle model in which technology shocks are iid and there are no learning by doing effects. This simple model produces many interesting features: procyclical productivity, consumption smoothing and investment that is more volatile than output. However, for this case there is essentially zero serial correlation in output. That is, the model does not contain an endogenous propagation mechanism.

The first treatment with learning by doing (row 2) uses our estimates from SPEC 1 of Table 1 in which we estimated labor share α=.59, capital share θ=.33, share of organizational capital ε=.08 and γ= .63. From the second row of the table, we see that all of the variables are positively correlated with output. There is again evidence of consumption smoothing. Notably, the first two autocorrelation coefficients of output are .07 and .05 which are more than ten times higher than in the baseline model. As can be seen from the next two rows learning by doing can generate much stronger propagation of temporary shocks over time for slightly different specifications. For example in row 3, a marginally higher ε raises the autocorrelation coefficient to .21 whereas the first two autocorrelation coefficients of output in row 4, corresponding to SPEC 3, are .53 and .42.

To understand the mechanism at work, we consider the impulse response functions for a temporary increase in total factor productivity using SPEC 3. These are shown in Figure 1. An increase in TFP causes an immediate increase in labor input to take advantage of this temporary shock. Consumption and investment both increase as well. The resulting increase in output leads to an increase in organizational capital in the subsequent period. Likewise, the burst of investment leads to an increase in the capital stock. Thus in the period after the burst of productivity, both stocks are higher so that output and employment remain above their steady state values. After this period, the stock of organizational capital slowly falls towards steady state for 20 quarters. Employment is above steady state for only 6 quarters while output is above steady state for about 20 quarters. Thus, a single transitory shock causes some richer dynamics relative to the standard model.

The source of the richer transition seems to be employment. In the traditional model, a high value of the capital stock causes employment to be below the steady state. In our economy, the shock causes organizational capital to be above its steady state after the initial period and employment is pulled above steady state during the initial part of the transition. But this effect diminishes quickly enough that the traditional dynamics dominate after four periods. This pattern of employment movement following a temporary shock is similar to the pattern reported in Cooper-Johri [1997] for the case of external learning by doing.

B. **Serially Correlated Technology Shocks**

Another way to see the impact of the internal propagation of the model is to study the behavior of the impulse response functions when the economy is hit by highly persistent technology shocks when measured in log -levels not in first differences. Cogley and Nason [1995] show that the data display characteristic hump shaped response functions in response to shocks whereas the benchmark models just replicate the behavior of the shock series.

There is no direct way to go from our empirical work in which shocks are modeled as having a unit root to this one in which shocks are highly persistent but still stationary in levels. There are however two indirect ways to use our estimates.

Row 3 of Table 1 jointly estimates the AR1 coefficient, ρ, in the shock process along with the other parameters of the technology and accumulation equations. In that specification we found some evidence of over-differencing: ρ=-.1 but not significant at the 5% level. Nonetheless, this estimate can be taken to suggest that the shock process in levels has an AR1 coefficient of .9. When we use technology parameters estimated from this specification to calibrate the model along with a persistent technology shock with an AR1 coefficient of .9, the learning by doing model displays hump shaped responses in output and other important variables. The impulse responses for this specification are displayed in Figure 2. The statistics for this specification are given in row 3 of Table 4.

An alternative approach is to use the estimates from row 2 (SPEC 1) of Table 1 and then choose the serial correlation in the technology shock so that the resultant serial correlation in output just matches that seen in the data. To match the observed serial correlation in US output of .96, we set the serial correlation in the shock process to .94. The impulse responses are similar to those in Figure 2 and are thus not reported. The quantitative implications of the model with serially correlated shocks are summarized in row 2, Table 4.

From Table 4 we see that the learning by doing model performs similarly to the benchmark model in terms of replicating the basic features of the business cycle. The final column of the table addresses an important issue: the contemporaneous correlation between hours and average labor productivity.

For our specifications, the correlation between average labor productivity and hours is smaller than in the baseline case though it is still higher than in the data.[11] Essentially, as in our discussion of taste shocks below, the movement in the stock of organizational capital acts as a "labor supply shifter" thus, as in the arguments of Christiano-Eichenbaum [1992], reducing the correlation between productivity and employment.

In the US data the dynamic correlation between average labor productivity and hours displays a distinct pattern: productivity leads hours. In fact the correlation between hours and lagged productivity is positive while the correlation between hours and leads of productivity is negative, though all the correlations are quite small when measured using linearly detrended data. For example the correlations between hours and the first two lags of productivity are .12 and .18. The corresponding leads of productivity are -.04 and -.12. So we see a pattern of the correlation declining in value as we go from the second lag of productivity to the second lead. The model also displays a similar dynamic pattern albeit at a higher level, ie., all the correlations are positive. However the correlations go from being the largest for the second lag of productivity to the smallest for the second lead of productivity. These correlations are reported in Table 5 using the SPEC 3 parameterization.

C. Stochastic Growth

As argued in the introduction, a major weakness of the standard real business cycle model is the lack of an internal propagation mechanism. This point was highlighted in Cogley-Nason

[11] Interestingly the correlation between average labor productivity and hours falls with the persistence of the shock process. For example if the serial correlation in the productivity shock is set at .95 to match the serial correlation in the Solow residual constructed from quarterly aggregate US data, this correlation falls to .1for spec 3.

[1995] by pointing out that the dynamics of output growth in that class of models to a very close approximation replicated the dynamics of the growth rate of the technology shock process that was built into the model. This is a serious shortcoming because typically, the technology shock process is estimated (using the Solow residual as a proxy) as a random walk whereas the growth rate of output displays at least two positive autoregressive coefficients. In this section we explore these issues, by reworking our model to accommodate stochastic trends.[12]

Our results are illustrated in Figure 3 which provides a plot of the autocorrelation function for output generated by the learning by doing model when technology shocks follow a random walk along with that found in U.S. data and that produced by the standard RBC model. Note that our model generates two autocorrelation coefficients which are positive and close to those seen in the data.

D. Some Counter-Factual Exercises

In the traditional RBC exercises, the Solow residual is viewed as a measure of technology shocks to the economy. By now it is widely recognized that movements in the Solow residual may not represent exogenous shocks to technology since the identifying assumptions used in early exercises may not hold. Cooper-Johri [1997] showed that the naively constructed Solow residual displayed a high degree of persistence even when the model was originally disturbed by iid technology shocks. Clearly similar results will be obtained here. While the "endogeneity" of the Solow residual has received a lot of attention, similar arguments can be made about a number of other unobservables that have been identified using Euler equations and first order conditions in the labor market. In this section we focus on two such series and discuss

[12]We are grateful to Jeff Fuhrer for supplying us with computer code.

the implications.

i. Taste Shocks

Baxter-King [1991] studied the quantitative implications of shifts in preferences as a driving force to explain economic fluctuations at the business cycle frequency. This is achieved by introducing a parameter into the utility function which varies the individuals marginal rate of substitution between consumption and leisure. This preference shift parameter is identified from the Euler equation that equates the marginal rate of substitution between consumption and leisure with the marginal product of labor. An important characteristic of that preference shock series is that it is extremely persistent. A possible interpretation of large preference shocks is that it measures the poor fit of the traditional Euler equation to the macro data. In an empirical study of various driving forces, Hall [1994] argues that preference shifts appear to be the most important driving force that "explains" aggregate fluctuations and that this should be viewed as a reason for focusing more on atemporal analysis as opposed to inter-temporal analysis.

Since the preference shocks are unobserved, they can be uncovered only under certain identifying assumptions. The preference shock series is calculated from the static first order condition on labor supply which is the term in parentheses in (5). However this ignores the effects of labor supply on the accumulation of future experience or organizational capital. Ignoring that element can then be viewed as a potential explanation of the poor fit of the standard Euler equation. This point is made by using the Baxter-King specification for calculating taste shocks on the simulated data from our model. Baxter-King assume that current utility is derived from the log of current consumption and leisure. In order to be consistent with their exercise we redo our model for log-log preferences. This gives rise to their specification for calculating the taste shock (denoted by ξ_t) in logs as:

$$\frac{\hat{\xi}_t}{c} = \hat{c}_t - \hat{w}_t + \frac{n}{1-n}\hat{n}_t$$

where c and n denote steady state values of consumption and hours.

We find that using the Baxter-King procedure on our simulated data uncovers a persistent series that appear as taste shocks even though the data was generated from a model with only technology shocks. The moments for our constructed preference shock series are reported in Table 6. When the model is parameterized using the estimates from SPEC 1 of Table 1 and a serially correlated (.94) technology shock , we see that the constructed taste shock series is very persistent (autocorrelation coefficient of .96) and highly correlated with output (.92). In comparison, the series generated by Baxter-King had an autoregressive coefficient of .97.

ii. Labor utilization

Recently a number of studies have argued that there is an unobserved component to the labor input, namely effort, which becomes a source of measurement error in the labor input series. Since effort will be procyclical this means that observed labor input series understate the contribution of labor to movements in output. Here we develop an expression for the effort series using a greatly simplified structure based on Burnside, Eichenbaum and Rebelo [1993].

Imagine that labor input can be varied by firms along two margins: hiring more labor hours and getting workers to put in more effort per hour. Assume that firms choose the number of hours to hire before they observe the technology shock and then adjust along the effort margin once the shock is realized. Then the unobserved series on effort can be uncovered from two conditions.

The first is the intratemporal first-order condition, obtained from a specification of preferences in which current utility depends on the log of consumption and the log of leisure. This is given by:

$$\frac{C_t}{1-N_tU_t}=MP_L$$

The second expression is a Cobb-Douglas technology, which determines the marginal product of labor:

$$Y_t=K_t^{1-\alpha}(N_tU_t)^{\alpha}A_t$$

where U measures unobserved effort. Using these conditions, the effort series can be uncovered from actual data.

When the above specification is used to calculate effort using simulated data from our model, we can simply back out the effort series from the production function since we observe the simulated technology shock. This procedure uncovers a highly cyclical and persistent series even though effort is constant in the data generating model. Some important moments of the 'naive' effort series are an autocorrelation coefficient of .98 and a contemporaneous correlation with output of .86.

VI. Conclusions

The point of this paper is relatively simple: what does learning by doing contribute to business cycles? As a theoretical proposition the answer is: quite a lot. From a quantitative perspective, it appears that the answer is the same. In particular, we estimate substantial and statistically significant learning by doing effects and these dynamic interactions can influence observed movements in real output. Thus, these links can serve as propagation devices.

To the extent that these interactions are excluded from standard models, they can lead to mis-specification errors. This was highlighted by our reinterpretation of results concerning the presence of taste shocks and unobservable labor effort. It appears that these phenomena can be "explained" by a richer stochastic growth model that incorporates learning by doing.

Finally, the model generates a hump-shaped response in the level of output to a serially correlated shock which highlights the internal propagation ability of the model. Hump shaped response functions were suggested as an important diagnostic tool by Cogley and Nason for this class of models.

This analysis hinges on the presence of technology shocks. The next step in this research will be to explore the effects of other disturbances in a model with learning by doing.

References

Bahk, B.H. and M. Gort, "Decomposing Learning by Doing in New Plants," Journal of Political Economy, 101 (1993), 561-83.

Basu, S. and J.G. Fernald, "Are Apparent Productive Spillovers a Figment of Specification Error?" Journal of Monetary Economics, 36 (1995), 165-88.

Baxter, M. and R. King, "Productive Externalities and Business Cycles," Institute for Empirical Macroeconomics, Federal Reserve Bank of Minneapolis, Discussion Paper #53, November 1991.

Beaudry, P. and M. Devereux, "Monopolistic Competition, Price Setting and the Effects of Real and Monetary Shocks," mimeo, 1993

__________________, "Towards an Endogenous Propagation Theory of Business Cycles", mimeo, University of British Columbia, Oct. 1995.

Benkard, C., "Learning and Forgetting: The Dynamics of Aircraft Production," mimeo, Yale University, November 1997.

Burnside, C. and M. Eichenbaum, "Factor-Hoarding the Propagation of Business-Cycle Shocks" American Economic Review, 86 (1996), 1154-74.

Burnside, C, M. Eichenbaum and S. Rebelo, "Labor Hoarding and the Business Cycle," Journal of Political Economy, 101 (1993), 245-73.

__________________, "Capital Utilization and Returns to Scale" in B. Bernanke and J. Rotemberg, eds., NBER Macroeconomics Annual, MIT Press: Cambridge, Ma., 1995, 67-123.

Christiano, L. and M. Eichenbaum, "Current Real-Business-Cycle Theory and Aggregate Labor-Market Fluctuations," American Economic Review, 82 (1992), 430-50.

Christiano, L., M. Eichenbaum and C. Evans, "The Effects of Monetary Policy Shocks: Evidence from the Flow of Funds," REStat, 78 (1996), 16-34.

Cogley, T. and J. Nason, "Output Dynamics in Real-Business-Cycle Models," American Economic Review, 85(1995), 492-511.

Cooper, R. and J. Haltiwanger, "Evidence on Macroeconomic Complementarities," REStat, 78 (1996), 78-93.

Cooper, R. and A. John, "Coordinating Coordination Failures in Keynesian Models," Quarterly Journal of Economics, 103 (1988), 441-63.

Cooper, R. and A. Johri, "Dynamic Complementarities: A Quantitative Analysis," Journal of Monetary Economics, 40 (1997), 97-119.

Hall, R. "On the Sources of Economic Fluctuation," mimeo, paper presented at the Economic Fluctuations Research Meeting, NBER, (1994), .

Irwin, D. and P. Klenow, (1994) "Learning-by-Doing Spillovers in the Semiconductor Industry," Journal of Political Economy, 102(6).

Jarmin, R., (1994) "Learning by Doing and Competition in the Early Rayon Industry," Rand Journal of Economics, 25(3).

Jovanovic, B. and Y. Nyarko, (1995) " A Bayesian Learning Model Fitted to a Variety of Empirical Learning Curves," Brookings Papers on Economic Activity.

Katz, L., and B. Meyer, "Unemployment Insurance, Recall Expectations, and Unemployment Outcomes," Quarterly Journal of Economics, (1990).

King, R., Plosser, C. and S. Rebelo, "Production, Growth and Business Cycles: I. The Basic Neoclassical Model," Journal of Monetary Economics, 21 (1988a), 195-232.

____________, "Production, Growth and Business Cycles: II. New Directions," Journal of Monetary Economics, 21 (1988b), 309-41.

Klenow, P."Not Learning by Not Doing," mimeo, University of Chicago, 1993.

Rotemberg, J. and M. Woodford, "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption," American Economic Review, 86 (1), (1996), 71-89.

Table 1

Estimation Using Aggregate 2-digit Quarterly Data*

(T-statistics given below coefficient estimates)

model**	α	ε	γ	ρ	φ
spec 1	.58 (8.7)	0.08	.63 (7.0)		.33 (2.5)
spec 2	.57 (8.9)	0.11	.55 (4.7)		.32 (3.4)
spec 3	.44 (7.7)	.31 (7.4)	.39 (5.3)	-.1 (1.13)	

* Instruments: The instrument list includes 2nd to 4th lags each of FF and NBR. The data set is gross output and hours and electricity consumption in the US manufacturing sector at the 2-digit level used by Burnside Eichenbaum and Rebello. 72:1 -92:4.

** Description of Specifications

Spec 1 corresponds to CRS. $Y_{it}=\alpha N_{it}-\alpha\gamma N_{it-1}+\theta E_{it}-\theta\gamma E_{it-1}+((1-\alpha-\theta)(1-\gamma)+\gamma)Y_{it-1}+A_{it}$.

Spec 2 corresponds to CRS in production function but not in accumulation equation..
$Y_{it}=\alpha N_{it}-\alpha\gamma N_{it-1}+\theta E_{it}-\theta\gamma E_{it-1}+((1-\alpha-\theta)+\gamma)Y_{it-1}+A_{it}$.

Spec 3 allows for serial correlation in the shock.

$$Y_{it}=\alpha N_{it}-\alpha(\rho+\gamma)N_{it-1}+\alpha\gamma\rho N_{it-2}+(\varepsilon+\rho+\gamma)Y_{it-1}-\varepsilon(\rho+\gamma)Y_{it-2}+A_{it}.$$

Table 2

Plant Level Estimates*

labor	k	year	lqv(-1)	cumv
0.97 (.04)	0.13 (.05)	0.04 (.003)		
0.980	0.140	Dums		
0.850	0.090	0.035	0.220	
0.9 (.04)	0.11 (.04)	-0.001 (.005)		0.29 (.03)
0.87 (.05)	0.11 (.04)	Dums		0.38 (.04)

Notes:

1. All coefficients are significantly different from zero at the 5% level.
2. Dums refers to treatments which include a year dummy.
3. Standard errors in parentheses.

* Dependent variable is the log of real output.

Table 3

IID Technology Shocks

Treatment	Corr. with Y Contemporaneous				Standard Deviation Relative to Y				Statistics for Y	
	C	Hr	In	W	C	Hr	In	W	sd	sc
BASELINE*	0.36	0.99	0.99	0.37	0.17	0.95	4.40	0.17	0.02	0.005
SPEC 1	0.38	0.98	0.99	0.66	0.19	0.87	4.07	0.23	0.02	0.07
SPEC 2	0.46	0.95	0.98	0.84	0.26	0.70	3.65	0.39	0.01	0.21
SPEC 3	0.59	0.86	0.97	.90	0.37	0.51	2.9	.62	0.01	0.53

* this specification imposes sn=.64, sk=.36, $\varepsilon=0$, $\gamma=0$.

† This specification imposes constant returns to scale on the production function but not on the accumulation equation. See row 2 of table 1.

Table 4

Persistent Technology Shocks

Treatment	Corr. with Y Contemporaneous				Standard Deviation Relative to Y				Statistics for Y		Corr. W and
	C	Hr	In	W	C	Hr	In	W	sd	sc	Hrs
BASELINE	0.88	0.73	0.89	0.88	0.72	0.49	2.43	0.72	0.06	0.95	0.33
SPEC 1	0.88	0.74	0.90	.88	0.70	0.5	2.5	.71	0.06	0.96	0.35
SPEC 3	0.88	0.62	0.92	.90	0.71	0.44	2.03	.80	0.06	0.98	0.22
U.S. Data log levels	0.89	0.71	0.60	0.77	0.69	0.52	1.30	1.10	0.04	0.96	0.100

Table 5

	Correlation of hours(t) and ALP(t+i)				
i	-2	-1	0	1	2
SPEC 3	.32	.28	.22	.14	.06
data	.18	.12	.10	-.04	-.12

Table 6

Counterfactual Experiments*

counterfactual	correlation with output	AR(1)
taste shock	.92	.96
unobserved effort	.86	.98

* The simulated data for these experiments was generated using SPEC 1 in table 1.

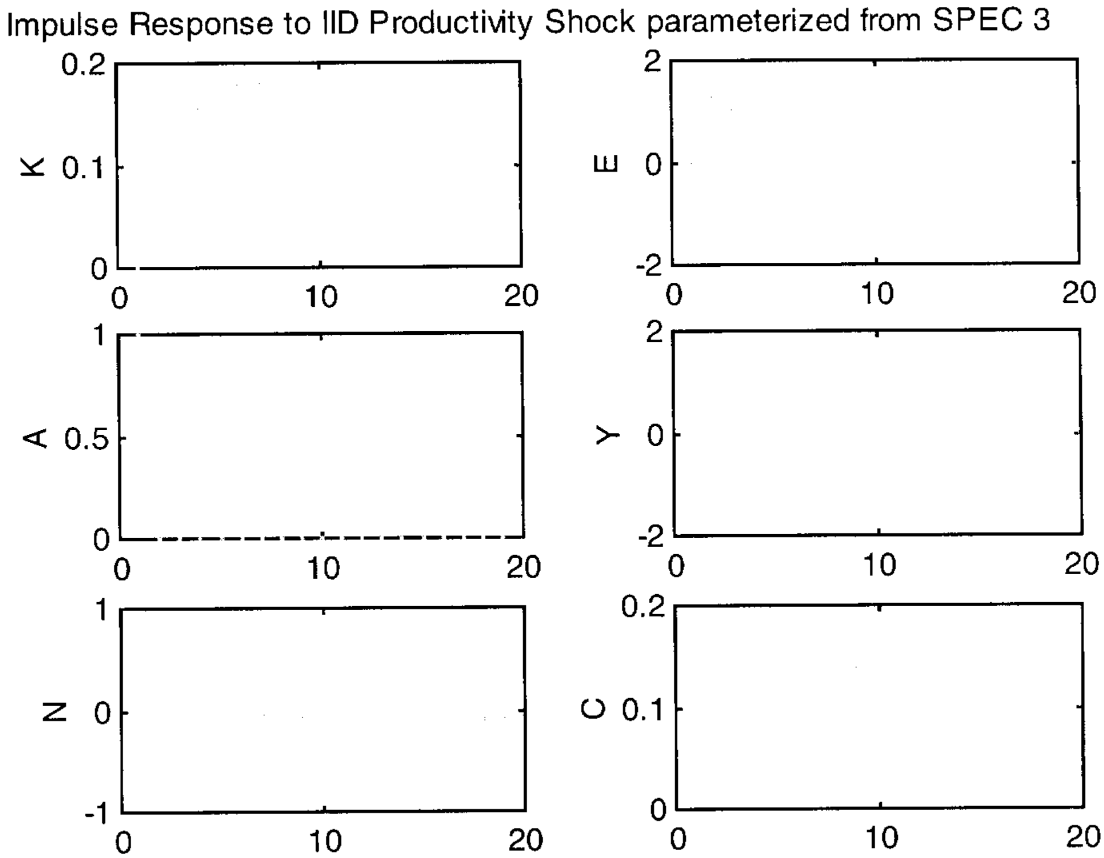
Impulse Response to IID Productivity Shock parameterized from SPEC 3
K
0.2
0.1
0
0
10
20
E
2
0
-2
0
10
20
A
1
0.5
0
0
10
20
Y
2
0
-2
0
10
20
N
1
0
-1
0
10
20
C
0.2
0.1
0
0
10
20

Figure 1

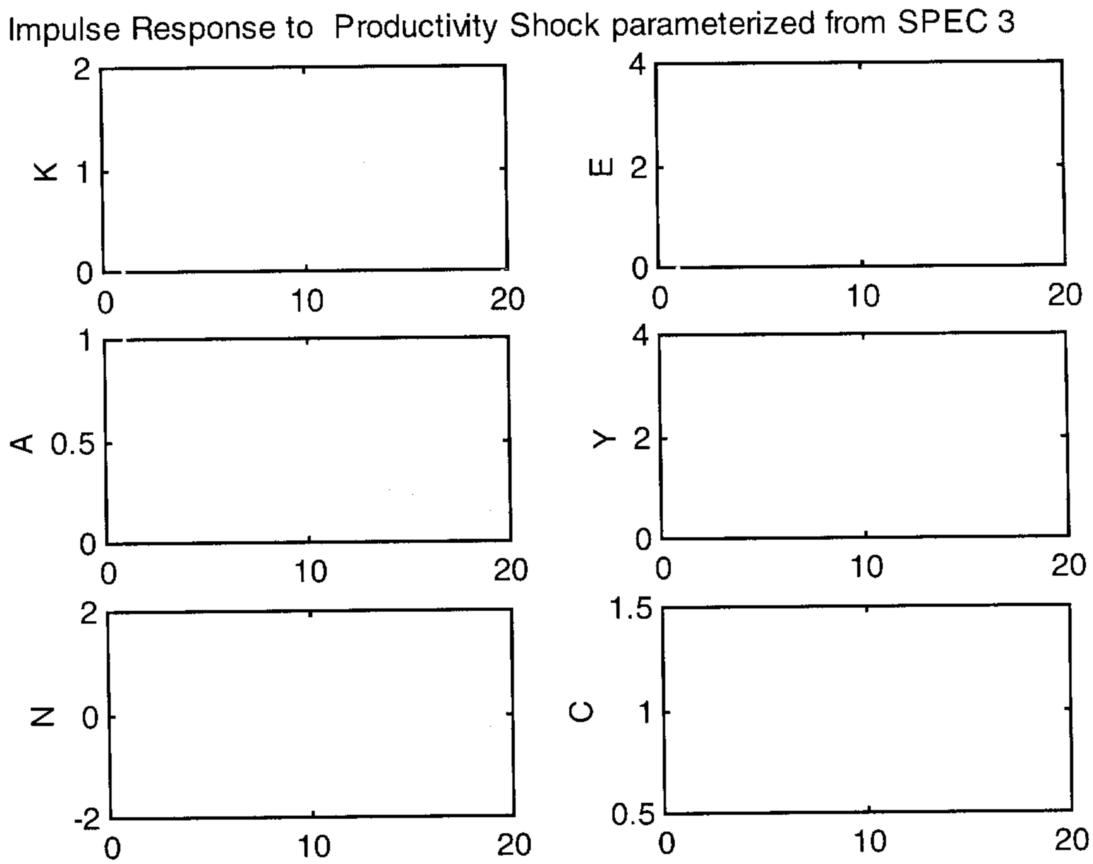
Impulse Response to Productivity Shock parameterized from SPEC 3
K
E
A
Y
N
C
0
10
20

Figure 2

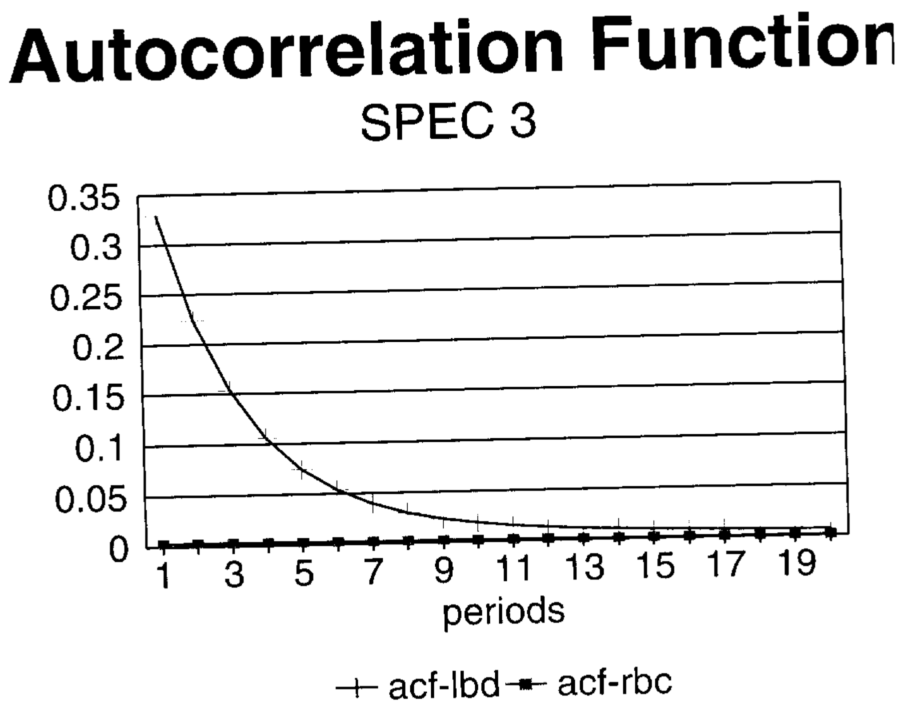
Autocorrelation Function
SPEC 3
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
1
3
5
7
9
11
13
15
17
19
periods
acf-lbd
acf-rbc

Figure 3