Macroeconomic Dynamics, **13** (2009), 366–389. Printed in the United States of America.
doi:10.1017/S1365100509080043

PROCYCLICAL SOLOW RESIDUALS WITHOUT TECHNOLOGY SHOCKS

ANDREW J. CLARKE
University of Melbourne

ALOK JOHRI
McMaster University

Most real business cycle models have a hard time jointly explaining the twin facts of strongly procyclical Solow residuals and extremely low correlations between wages and hours. We present a model that delivers both these results without using exogenous variation in total factor productivity (technology shocks). The key innovation of the paper is to add learning-by-doing to firms' technology. As a result, firms optimally vary their prices to control the amount of learning, which in turn influences future productivity. We show that exogenous variation in labor wedges (preference shocks) measured from aggregate data deliver around 50% of the standard deviation in the efficiency wedge (Solow residual) as well as realistic second moments for key aggregate variables, which is in sharp contrast to the model without learning-by-doing.

Keywords: Business Cycles, Learning-by-Doing, Productivity

1. INTRODUCTION

Although the neoclassical growth model has proved to be an extremely effective tool for capturing the basic features of business cycles, two major areas of weakness have been identified in the literature, in terms of the ability of the model to capture the dynamics of the aggregate data. These problems have been elegantly reinterpreted in terms of "wedges" in the equations of the model in a recent paper by Chari et al. (2007). In an accounting exercise, the authors show that the wedges embedded in the hours first-order condition and in the production technology are essential for "explaining" the bulk of the movement seen in aggregate U.S. data. In the business cycle literature, these wedges have been interpreted in the past as exogenous "shocks" that are the source of fluctuations. These interpretations of the two wedges as preference shocks and technology shocks, respectively, have been controversial, not least because they assume rather than explain a large fraction of the variation in the aggregate data. In this context, a number of authors have

The authors would like to thank SSHRCC for funding this research and Michelle Alexopolous for the data. Helpful comments of two referees, the associate editor and participants at the Canadian Economic Association 2006 meetings as well as Marc-Andre Letendre are greatfully acknowledged. Keqiang Hou provided stellar research assistance.
Address correspondence to: Alok Johri, Department of Economics, McMaster University, 1280 Main Street West, Hamilton, Ontario L8S 4M4, Canada; e-mail: johria@mcmaster.ca.

 1365-1005/09

attempted to develop models that produce fluctuations in the labor wedge and account for a part of the movement in the efficiency wedge.[1] However, much more work remains to be done in this area. As pointed out by Chari et al. (2007), "our results suggest that future theoretical work should focus on developing models which lead to fluctuations in efficiency and labor wedges. Many existing models produce fluctuations in labor wedges. The challenging task is to develop detailed models in which primitive shocks lead to fluctuations in efficiency wedges as well."[2] The goal of this paper is to take a step in this direction.

Unlike most business cycle models, there are no exogenous movements in technology in our model. Instead, we rely on movements in the labor wedge to induce fluctuations in aggregate variables, and the general equilibrium structure of the model to generate dynamic correlations between these variables. We show that when fluctuations in the labor wedge are fed into a prototypical real business cycle model, it is unable to account for many of the key features of the data. However, when they are fed into our model, it does quite well. In particular, the model does a good job of explaining fluctuations in the efficiency wedge (essentially the Solow residual), whereas the prototypical model would generate no endogenous fluctuations. These results are achieved with a simple modification to a standard real business cycle model: the presence of learning-by-doing effects in firms' technology that act as an endogenous propagation mechanism, converting shocks to the labor wedge into persistent movements in total factor productivity and output.

There is a large empirical literature documenting the pervasive influence of learning-by-doing in productive activities. Recent studies include those of Bahk and Gort (1993), Irwin and Klenow (1994), Jarmin (1994), Benkard (2000), Thompson (2001), Thornton & Thompson (2001), Cooper and Johri (2002), and Clarke (2007). These studies find that agents and organizations appear to become more productive as they gain experience at producing a particular product or service. A number of these studies also report spillover effects in learning across similar products, both within and across firms. Because there is evidence in favor of both internal and external learning-by-doing, and there is some controversy around which one is empirically more important, we look at both types of learning models and find not only that the results are qualitatively similar, but also that the quantitative differences are quite small at the aggregate level.

There exist a small number of papers that introduce learning-by-doing into dynamic general equilibrium business cycle models. Both Cooper and Johri (2002) and Chang et al. (2002) emphasize the ability of learning-by-doing to propagate shocks. In the former study, a representative consumer has access to a production technology with learning effects. The latter offers a specific decentralization in which workers learn and accumulate human capital. Our model offers a different decentralization in which firms accumulate production-related knowledge but workers do not.[3] Firms in our model operate in a market with monopolistic competition and therefore have the power to set prices. Nonetheless, our model is quite close to that of Cooper and Johri (2002) in that it directly borrows the way in which learning effects are introduced into the production technology.[4] As will be clear later, these changes have a number of novel implications. The pricing decision facing these firms becomes dynamic because raising current prices lowers current production and learning, which leads to future productivity decreases. Moreover, when firms are aware that they can learn and build up knowledge, price–cost markups can vary endogenously in an environment where they would be constant in the absence of learning-by-doing. In this sense, the model may be viewed as a new theory of time-varying markups. These features are absent in the earlier papers. In contrast to our model, those models imply that labor supply is dynamic. The implications of dynamic labor supply are discussed more fully in Johri and Letendre (2007). Although we believe that learning by workers is an important phenomenon, we choose to ignore it in this study in order to focus on the new implications of incorporating firm level learning in business cycle models.

Simulations from a linearized version of the model calibrated to U.S. data show that substantial improvements to the existing literature can be achieved in terms of matching the usual set of "key" second moments. For example, the model can deliver a small positive correlation between average labor productivity and hours, which has eluded most current business cycle models, either real or monetary. Moreover, the model generates a highly volatile and persistent efficiency wedge (Solow residual). Depending on the specification, we can account for around 40% to 50% of the actual volatility in the efficiency wedge and over 80% of the movement in aggregate output, hours, consumption, and investment. These results are obtained with fairly modest learning effects. In sharp contrast, the results from the model without learning-by-doing are quite unrealistic: average labor productivity is negatively correlated with both output and hours.[5] In addition, investment and hours are too volatile relative to output. Moreover the model explains none of the volatility in the efficiency wedge and only 67% of the volatility in output. Unfortunately, although the simulated markups are time-varying, the volatility is quantitatively too small to account for actual markup variations over the business cycle.

The next two sections of the paper presents the two models of learning-by-doing in order of decreasing complexity, starting with the internal learning model. Section 4 presents results from simulating linearized versions of the model calibrated to the U.S. economy. The final section concludes.

2. A DECENTRALIZATION OF THE INTERNAL LEARNING MODEL

The model presented below offers a decentralization of the Cooper and Johri (2002) model. A common feature of both models is that learning occurs as a by-product of production. The key difference occurs in who gets to learn. In Cooper and Johri (2002), learning occurs at the household level, whereas in the current model, it is intermediate goods firms that learn by accumulating production-related knowledge, which we will refer to as organizational capital. These firms operate in a market with monopolistic competition, which gives them the power to choose prices.[6] The firms control how much they wish to learn by varying the markup of price over marginal cost to maximize the present discounted value of profits. Because the behavior of consumers is standard, we discuss firms first.

2.1. Final and Intermediate Goods Producers

There are a large number of producers operating in a competitive industry that produce a final good Y_t using the following technology that converts intermediate goods $Q_t(i)$ indexed by $i \in [0, 1]$ into final goods:

$$Y_t = \left\{ \int_0^1 [\, Q_t(i)\,]^{1/\mu} \, di \right\}^{\mu}. \tag{1}$$

Each period these firms choose inputs $Q_t(i)$ for all $i \in [0, 1]$, and output Y_t to maximize profits. The conditional demand for each intermediate input $Q_t(i)$ that emerges from this exercise is

$$f\,(v_t(i), Y_t) = v_t(i)^{\mu/(1-\mu)}\, Y_t, \tag{2}$$

where $v_t(i)$ is the relative price charged by the ith intermediate goods producer. The price elasticity of demand, faced by intermediate goods producer i, will be given by $\mu/(1-\mu)$. A zero profit condition ties down the price charged by the final goods producer P_t as a function of intermediate good prices.

There is a continuum of intermediate goods producers, indexed by the letter i, that operate in a monopolistically competitive economy. Each such producer produces a differentiated intermediate good $Q(i)$ according to the following production technology:

$$Q_t(i) = F\,[\,K_t(i), Z_t(i), H_t(i)] = K_t(i)^{\theta}\, H_t(i)^{\alpha}\, Z_t(i)^{\varepsilon} \qquad 0 < \alpha,\ \theta,\ \varepsilon < 1, \tag{3}$$

where organizational capital $Z(i)$ is combined with physical capital $K(i)$ and labor $H(i)$ to produce output $Q(i)$. The technology differs from a standard neoclassical production function because the firm carries a stock of organizational capital, which is an input in the production technology. Organizational capital refers to the information accumulated by the firm, through the process of past production, regarding how best to organize its production activities and deploy its inputs.[7] As a result, the higher the level of organizational capital, the more productive the firm.[8] Note that there are diminishing returns to accumulating organizational capital, a feature often found in empirical studies of learning-by-doing.

Although learning-by-doing is often associated with workers and modeled as the accumulation of human capital, a number of economists have argued that firms are also store houses of knowledge. Atkeson and Kehoe (2005) note that "At least as far back as Marshall (1930, bk.iv, chap. 13.I), economists have argued that organizations store and accumulate knowledge that affects their technology of production. This accumulated knowledge is a type of unmeasured capital distinct from the concepts of physical or human capital in the standard growth model." Similarly, Lev and Radharkrishnan (2003) write, "Organization capital is thus an agglomeration of technologies—business practices, processes and designs, including incentive and compensation systems—that enable some firms to consistently extract out of a given level of resources a higher level of product and at lower cost than other firms." There are at least two ways to think about what constitutes organizational capital. Some, such as Rosen (1972), think of it as a firm-specific capital good, whereas others focus on specific knowledge embodied in the matches between workers and tasks within the firm. Although these differences are important, especially in trying to measure the payments associated with various inputs, we believe they are not crucial to the issues at hand. As a result, we do not distinguish between the two.

Learning at the firm level is modeled through an accumulation equation for organizational capital, which is closely related to the empirical learning-by-doing literature in which each cumulative unit of past production contributes equally to the creation of knowledge. Our specification differs in that the contribution of production in any period to the current level of organizational capital is decreasing over time. Following Cooper and Johri (2002), who provide evidence on this specification, we write the accumulation technology as

$$Z_{t+1}(i) = Z_t(i)^{\eta}\, Q_t(i)^{\gamma} \tag{4}$$

or

$$\ln Z_t(i) = \eta^t \ln Z_0(i) + \gamma \sum_{k=0}^{t-2} \eta^k \ln Q_{t-1-k}(i),$$

where $Z_t(i)$ denotes the stock of organizational capital available to the production unit at time t, $Z_0(i)$ denotes the initial endowment of organizational capital, and $Q_t(i)$ denotes the current level of output. This initial stock of organizational capital constitutes the inherited knowledge of the organization, which might include a common component reflecting the prevailing best practice systems, structures, and procedures and an idiosyncratic component reflecting more specific knowledge imparted by the organization's founders that becomes a durable feature of the organization.

This modification of the traditional specification of learning has a number of advantages. First, it allows for the sensible idea that production knowledge may become less and less relevant over time as new techniques of production, new product lines, and new markets emerge. Second, it allows in a general way for the idea that some match-specific knowledge may be lost to the firm as workers leave or get reassigned to new tasks or teams within the firm. In addition, the knowledge accumulated through production experience will be a function of the current vintage of physical capital. The decision to replace physical capital will imply that the existing stock of organizational capital will be less relevant. Third, it allows for the existence of a steady state in which the stock of organizational capital is constant. In contrast, the traditional specification in the empirical learning-by-doing literature allows the stock of organizational capital to grow unboundedly. An alternative way to bound learning is to assume that productivity increases due to learning occur for a fixed number of periods. Although this may be appropriate for any one task or worker within the firm, we think of the internal context of firms as an environment with an ever-changing set of tasks, workers, teams, machines, and information. In this context, it may be better to model organizational capital as continually accumulating and depreciating.

The restriction $\eta < 1$ is consistent with the empirical evidence supporting the hypothesis of depreciation of organizational capital, often referred to as organizational forgetting. Argote et al. (1990) provide empirical evidence for this hypothesis of organizational forgetting associated with the construction of Liberty Ships during World War II. Similarly, Darr et al. (1995) provide evidence for this hypothesis for pizza franchises and Benkard (2000) provides evidence for organizational forgetting associated with the production of commercial aircraft. One difference between these studies and this paper is that the accumulation technology is log-linear rather than linear. Clarke (2006) shows that the additional curvature in this log-linear technology is unlikely to produce predictions for aggregate variables, in response to a technology shock, considerably different from those associated with a linear technology. It is the implied dynamic structure associated with the accumulation of organizational capital, rather than any functional form assumptions, that drives the results in Cooper and Johri (2002). We expect similar results to follow in the current context.[9]

It is useful to dissect the problem faced by an intermediate goods producer into two stages. In the first stage, the producer chooses the cost-minimizing quantities of labor and physical capital, for a given stock of organizational capital, to solve the following static cost minimization problem:

$$
\begin{aligned}
&C[w_t, r_t, Q_t(i), Z_t(i)] \\
&\quad = \min_{H_t(i),K_t(i)} \left\{ w_t H_t(i) + r_t K_t(i) \mid K_t(i)^\theta\, H_t(i)^\alpha\, Z_t(i)^\varepsilon \geq Q_t(i) \right\}. \qquad (5)
\end{aligned}
$$

Each intermediate goods producer is assumed to operate within perfectly competitive input markets such that the real rental price of physical capital r_t and the real wage w_t are taken as given. This cost minimization problem produces conditional factor demands that are a function of factor prices, the required level of output (Q_t), and the stock of organizational capital (Z_t). The cost function (5) will be a non-increasing function of the (given) stock of organizational capital.

In the second stage, the intermediate goods producer will solve a dynamic problem that selects the time path of output supply $Q_t(i)$, or the relative price $v_t(i)$, and the time path of organizational capital $Z_t(i)$ that maximizes the real expected discounted present value of profits, subject to the demand function (2), the accumulation technology for organizational capital (4), and the initial stock of organizational capital $Z_0(i)$.

Each intermediate goods' producer solves the following maximization problem:

$$\max_{v_t(i), Z_{t+1}(i)} E_0 \sum_{t=0}^{\infty} D_t \left(v_t(i) \cdot f\left[v_t(i), Y_t\right] - C\left\{w_t, r_t, f\left[v_t(i), Y_t\right], Z_t(i)\right\}\right)$$

subject to

$$Z_t(i)^{\eta}\left[f\left(v_t(i), Y_t\right)\right]^{\gamma} = Z_{t+1}(i),$$

given some initial stock of organizational capital $Z_0(i)$ and subject to an appropriate transversality condition on the stock of organizational capital. Because firms are owned by households, D_t is the appropriate endogenous discount rate for the firms.[10]

The solution to the (profit) maximization problem will satisfy the first-order conditions

$$v_t(i) f'\left[v_t(i)\right] + f\left[v_t(i)\right] - \mathrm{mc}_t(i)\, f'\left[v_t(i)\right] = -\lambda_t^f(i)\, \frac{\partial Z_{t+1}(i)}{\partial Q_t(i)}\, f'\left[v_t(i)\right] \quad \textbf{(6)}$$

and

$$\lambda_t^f(i) = E_t\left[\frac{D_{t+1}}{D_t}\left\{\lambda_{t+1}^f(i)\, \frac{\partial Z_{t+2}(i)}{\partial Z_{t+1}(i)} - \frac{\partial C_{t+1}(i)}{\partial Z_{t+1}(i)}\right\}\right], \quad \textbf{(7)}$$

where $\mathrm{mc}_t(i)$ denotes the marginal cost of producing output $Q(i)$. The term $\lambda_t^f(i)$ denotes the Lagrange multiplier associated with the maximization problem defined above and represents the discounted value of an additional unit of organizational capital, in terms of marginal real profits.

The first-order condition (6) captures the nature of the dynamic trade-off that arises when intermediate goods producers face a downward-sloping demand curve. Fundamentally, these producers face a trade-off between maximizing current period profits and losing future productivity increases. The first-order condition (6) determines the optimal relative price v_t to be set by the producer of the intermediate input i. Since the intermediate goods producer faces a downward-sloping demand curve for its product, raising the output price by one unit causes demand for its product to fall. The first two terms in (6) capture the impact on current revenue of raising the relative price of output v_t. The third term represents the reduction in current period costs resulting from this corresponding lower level of output. The accumulation technology for organizational capital implies that a reduction in current output will lead to a reduction in the stock of organizational capital available in the next period. The term on the right-hand side of (6) represents the value of this reduced (future) stock of organizational capital. The term $f'(v_t(i))$ measures the reduction in current output due to the higher relative price and the term $\partial Z_{t+1}/\partial Q_t$ represents the reduction in the stock of organizational capital resulting from the reduction in current period output, evaluated at the marginal value of organizational capital to the intermediate goods producer. The first-order condition implies that the firm will set its optimal price by equating the current period benefit of increasing its relative price by one unit to the discounted future

https://doi.org/10.1017/S1365100509080043 Published online by Cambridge University Press

costs. This trade-off, captured by the (dynamic) term on the right-hand side of (6), will not appear in the standard model of monopolistic competition, without the accumulation of organizational capital.

The first-order condition (7) determines the value of an additional unit of organizational capital for use by the producer in period $t+1$. This additional unit of organizational capital has a (marginal) value, in terms of profits, of λ_t^f to the producer. Because the cost function is decreasing in the stock of organizational capital, an additional unit of organizational capital reduces the cost of producing output level Q_{t+1}. The accumulation equation for organizational capital implies that an additional unit of organizational capital will increase the stock of organizational capital available in period $t+2$. This higher stock of organizational capital has a value of λ_{t+1}^f to the producer. The condition (7) implies that organizational capital will be accumulated up to the point where the value of an additional unit of organizational capital today is equal to the discounted value of this organizational capital next period.

2.2. Households

The economy is populated by a continuum of identical, infinitely lived households. At time t, the representative household has preferences over consumption of final goods C_t and leisure $(T - H_t)$, where T represents the total time endowment. The representative household maximizes the expected discounted present value of utility by choosing consumption C_t, labor supply H_t, and investment in physical capital I_t, taking as given the real wage w_t and the real interest rate r_t. Consumers sell labor services and rent physical capital to the intermediate goods producers. As owners of the intermediate goods producers, they also receive the current profits of these producers. Physical capital is stored and accumulated by consumers according to a standard accumulation technology. The flow utility of the representative household is given by

$$U(C_t, T - H_t) = \ln C_t + B_t \mathbf{V}(T - H_t) \tag{8}$$

where

$$\mathbf{V}(T - H_t) = \begin{cases} \dfrac{1}{1-\nu}\,(T - H_t)^{1-\nu} & \text{for } \nu \geq 0 \text{ and } \nu \neq 1 \\ \ln\,(T - H_t) & \text{for } \nu = 1 \end{cases}$$

and the term B_t evolves according to

$$\ln B_t = (1 - \rho_b)\,\ln \bar{B} + \rho_b\,\ln B_{t-1} + \upsilon_{bt},$$

where υ_{bt} is an iid random variable with mean zero and standard deviation σ_b. B_t is the only exogenous source of fluctuations in our model. It represents the idea of the "labor wedge" explored in Chari et al. (2007). We chose this formulation

https://doi.org/10.1017/S1365100509080043 Published online by Cambridge University Press

of the labor wedge over representing it as a labor income tax because it is also consistent with the idea of preference shocks, which is increasingly being used in business cycle models. Early references include Parkin (1988) and Baxter and King (1991). More recently they can be found in Chang et al. (2002), which is a business cycle model with only real shocks, and Ireland (2004a, 2004b), which have both preference and money shocks.[11]

An obvious question that arises in this context is, What does the labor wedge represent? Our interpretation is that this is anything that exogenously shifts the aggregate labor supply at business cycle frequencies. Some possible sources of such shifts have been explored in the literature. For example, if the marginal utility of leisure is a function of omitted variables such as money balances or government spending, then variations in these might directly account for these shifts. See Parkin (1988) for an early discussion of this point. At a deeper level one might imagine these may be caused by changes in the amount of distortions in the economy caused by taxation and transfer payments, unionization, and other shifts in market power. See, for example, Mulligan (2002) and Chari et al. (2002) or Cole and Ohanian (2002, 2004) for a model in which shifts in government policy toward monopolies in a model with unions can account for these preference shifts in the context of the Great Depression. Rotemberg and Woodford (1991, 1999) consider models of imperfect competition in which distortions vary over the business cycle. Alternatively, a model with monetary shocks and sticky wages might account for them. Finally, Hall (1997) has suggested that these shifts might be accounted for in a labor search model where the number of aggregate hours spent searching for jobs varies with the cycle.[12]

Households own the imperfectly competitive firms and receive all (real) profits earned by firms. The household budget constraint then becomes

$$C_t + K_{t+1} - (1-\delta)\,K_t = w_t\,H_t + r_t\,K_t + \int_0^1 \pi_t(i)\,di. \tag{9}$$

The household problem may be represented by

$$\max_{C_t, H_t, K_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t\ [\ln C_t + B_t\,\mathbf{V}(T - H_t)]$$

subject to the budget constraint and the usual transversality condition given an initial stock of physical capital K_0. The first-order conditions associated with this problem are

$$B_t \mathbf{V}'(T - H_t) = \frac{w_t}{C_t} \tag{10}$$

and

$$\frac{1}{C_t} = \beta E_t \left\{ \frac{1}{C_{t+1}} [r_{t+1} + (1-\delta)] \right\}, \tag{11}$$

where δ represents the depreciation rate for physical capital. The interpretation of these first-order conditions is quite standard. Condition (10) requires that

https://doi.org/10.1017/S1365100509080043 Published online by Cambridge University Press

per-capita consumption and hours be chosen so that the marginal rate of substitution between consumption and leisure, for all t, is equal to the real wage rate. Condition (11) is the standard Euler equation for the accumulation of physical capital, which requires that, at the optimal solution, the utility cost of giving up one unit of consumption must be equal to the present value (in terms of utility) of this unit of consumption tomorrow.

2.3. Equilibrium Prices and Quantities

A competitive equilibrium consists of

(1) allocations C_t, H_t, K_{t+1} that solve the consumer's problem, taking prices as given
(2) allocations K_{it}, H_{it}, $Z_{i,t+1}$ for $i \in [0, 1]$ that solve the firm's problem, taking all prices but its own output price as given
(3) allocations Y_t, Q_{it} that solve the final goods producer's problem, taking prices as given
(4) prices w_t, r_t, v_{it} for $i \in [0, 1]$

In addition, these allocations must satisfy the factor-market-clearing conditions and the aggregate resource constraint. Because the technology of the economy is assumed symmetric with respect to all intermediate goods producers, attention may be restricted to the symmetric equilibrium, where all producers in the intermediate goods sector charge the same price and produce the same output. In this case,

$$Y_t = \left\{ \int_0^1 [\, Q_t(i)\,]^{1/\mu}\, di \right\}^{\mu} = Q_t,$$

so that symmetry requires the relative price $v_t = 1$ for all firms. Because all intermediate goods producers have the same level of technology and have the same initial endowment of organizational capital, it will be true that $H_t(i) = H_t$; $K_t(i) = K_t$ and $Z_t(i) = Z_t$. In this case, the total demand for hours and the total demand for physical capital will be given by

$$\mathbf{H}_t^D = \int_0^1 H_t(i)\, di = H_t \qquad \text{and} \qquad \mathbf{K}_t^D = \int_0^1 K_t(i)\, di = K_t,$$

with the aggregate stock of organizational capital given by

$$\mathbf{Z}_t = \int_0^1 Z_t(i)\, di = Z_t.$$

Using the factor-market-clearing conditions, the household's budget constraint in the symmetric equilibrium will be given by

$$C_t + K_{t+1} - (1 - \delta) K_t = Y_t.$$

https://doi.org/10.1017/S1365100509080043 Published online by Cambridge University Press

3. AN AGGREGATE MODEL WITH LEARNING EXTERNALITIES

The most direct way to discuss the impact of learning externalities is to retain all features of the previous model except that firms no longer realize that production leads to the accumulation of organizational capital. Instead they take as given the economywide stock of organizational capital as an input in their production technologies. The obvious implication of this change in specification is that (7) is no longer part of the system of equations associated with equilibrium in the external model. Second, (6), the pricing equation, simplifies to the standard condition found in models of monopolistic competition—price is set as a fixed markup over marginal cost. In other words, the term on the right-hand side of equation (6) equals zero.

An alternative way to proceed is to abandon the decentralization involving monopolistic competition and solve a representative agent model economy with external learning effects built directly into the technology. An advantage of this approach is that the model can be directly compared to some earlier work, especially the closely related model presented in Cooper and Johri (1997, 2002). The two models then will differ in only two ways. First, whereas the agent in Cooper and Johri (2002) internalizes the learning effects built into the technology, here it will not. Second, the source of fluctuations in Cooper and Johri (2002) was technology shocks, whereas here it will be preference shocks.

The agent has preferences over consumption and leisure and has access to a Cobb–Douglas technology that produces a consumption/investment good (Y) using three inputs: hours (H), physical capital (K) and organizational capital (Z),

$$Y_t = K_t^{\theta} H_t^{\alpha} Z_t{}^{\varepsilon}. \tag{12}$$

The agent can use this good for consumption or to invest in physical capital. The specification for the accumulation of physical capital is

$$K_{t+1} = (1-\delta)K_t + I_t, \tag{13}$$

where δ is the depreciation rate. The stock of organizational capital depends on past production as well as past organizational capital, as follows:

$$Z_{t+1} = Z_t^{\gamma} Y_t^{\eta}. \tag{14}$$

The maximization problem of the representative agent is given by

$$\max_{C_t, H_t, K_{t+1}} E_0 \sum_{t=0}^{\infty} \beta^t \, [\ln C_t + B_t \mathbf{V}(T - H_t)],$$

subject to the aggregate resource constraint, the accumulation technology (14), the transversality condition $\lim_{t\to\infty} \Lambda_t^h K_t = 0$, and an initial stock of physical capital K_0.

https://doi.org/10.1017/S1365100509080043 Published online by Cambridge University Press

4. SOLUTION METHOD AND CALIBRATION

An approximate linear solution to the models outlined above is obtained using the method outlined in King and Watson (2002). This solution is derived by linearizing the equations characterizing the competitive equilibrium in the neighbourhood of the steady state. It is necessary to specify values for the discount rate β, the depreciation rate of physical capital δ, the production technology parameters α, θ, and ε, the accumulation technology parameters γ and η, the preference parameter ν, and the persistence ρ_b and the volatility σ_b of the preference shock. In addition, for the internal learning model, the demand parameter μ needs to be specified.

The following values for key parameters were used: $\delta = 0.02$, $\beta = 0.9926$, and $\theta = 0.34$. These values imply a steady state share of investment in output of 21.57% and a steady state (physical) capital-to-output ratio of 10.78. Previous attempts to measure the degree of returns to scale in the aggregate production function using instrumental variable techniques have yielded quite a range of estimates, often with large standard errors. Despite this, the consensus view [see the work of Basu (1996) and Basu and Fernald (1997)] is that the scale elasticity measured by $(\alpha + \theta)$ is close to unity. Consequently, given $\theta = 0.34$, the model is calibrated with $\alpha = (1 - \theta) = 0.66$.[13] Turning to the parameters associated with learning-by-doing, we draw on the results of Cooper and Johri (2002). To illustrate the impact of varying the parameters associated with learning, we report results for two different specifications of the learning technology that Cooper and Johri (2002) focused on. For our baseline case, we set the elasticity of output with respect to organizational capital at $\varepsilon = 0.264$. This value is close to manufacturing industry estimates reported in Cooper and Johri (2002) and corresponds to a learning rate of 20%, commonly estimated in microeconomic studies of learning-by-doing that do not allow past experience to become less valuable to the firm.[14]

Two alternative values for the parameters of the accumulation technology are considered. The first set ($\gamma = 0.5$ and $\eta = 0.5$) correspond to manufacturing level estimates reported in Cooper and Johri (2002). Sensitivity analysis is conducted with a second set ($\gamma = 0.2$ and $\eta = 0.8$) of values that are based on economywide estimates reported in Cooper and Johri (2002). Our value for μ varies across specifications to deliver a constant capital–output ratio and a steady state markup of 15%. We also report results for a much lower learning rate of 10%. This specification sets ($\gamma = 1$ and $\eta = 0.5$) as in Cooper and Johri (2002).

Although the calibration of the learning-by-doing parameters is based on estimates obtained in earlier studies, it may be useful to shed further light on the importance of organizational capital in the economy implied by these parameters. The level of organizational capital relative to output in the economy is governed by the accumulation equation (4). Because we have assumed (for the baseline case) constant returns to scale in this function, we end up with a steady state level of $Z/Q = 1$. By contrast, the ratio of physical capital to output is around 10 and Atkeson and Kehoe (2005) report that organizational capital is roughly 2/3 the size of physical capital, implying a Z/Q of roughly 6. Our justification for a lower

Z/Q is that we choose to focus on only one type of organizational capital, that which is accumulated as a by-product of production through learning-by-doing. The literature on organizational capital also emphasizes that firms intentionally invest conventional inputs in the creation of knowledge that may be captured in estimates of organizational capital. Finally, organizational capital may also include intangibles such as "work-culture" that may have an effect on productivity. To the extent that these features do not vary over the business cycle, their exclusion from the model should not strongly influence the results. In this context, note that if other aspects of organizational capital are not varying over the business cycle, then the accumulation equation can be modified to include a constant term:

$$Z_{t+1}(i) = Z_0 Z_t(i)^{\eta}\, Q_t(i)^{\gamma}. \tag{15}$$

In steady state this implies that $Z(i) = Z_0(i)^{1/(1-\gamma)}\, Q(i)$. In the current model $Z_0 = 1$. Given an estimate of the steady state ratio of organizational capital to output, Z_0 could be calibrated to deliver that number. However, the interesting thing to note about the level of steady state organizational capital is that it has little, if any, influence on the short-run dynamics of the model. A simple (albeit incomplete) way to see this is to examine the linearized version of the accumulation equation, where all lowercase variables are in deviations from steady state:

$$z_{t+1} = \gamma z_t + (1-\gamma) y_t. \tag{16}$$

Note that the level of organizational capital plays no role in its short-run dynamics and that Z_0 drops out of the equation.

An alternative way to evaluate the impact of organizational capital on the economy is to compare the implied value for Tobin's q in the model with the data. Empirical estimates of average q vary quite widely. For example, Barnett and Sakellaris (1999) provide an estimate of $q = 1.77$ on average for the manufacturing sector. This accords well with the average estimate of $q = 1.7$ reported in Summers et al. (1981). A much lower estimate of $q = 1.11$ can be found in Lang and Schultz (1994). In our model for the baseline calibration, Tobin's q is equal to 1.643, which is in between the above estimates.[15]

One can also ask to what extent firms invest in organizational capital in steady state? The model provides an indirect answer to this question, because there are no direct investments in organizational capital by firms. However, because firms realize that producing more output yields additional future organizational capital, they choose to produce more than they would in the absence of learning-by-doing and more also than if they did not internalize the learning technology. How should we evaluate this additional "investment" in organizational capital? One way would be to calculate the costs associated with the extra output produced by firms that internalize learning effects. Recall, however, that the additional output also yields revenue, albeit at a lower price. Therefore it would appear that the best way to calculate the value of the investments in organizational capital is to calculate the difference in the steady state flow of profits between the internal and external learning models. Steady state profits are 68% larger in the external learning model as compared to the internal learning model. Attributing this difference in steady state profits to investments in learning, the steady investment-to-output ratio for organizational capital is approximately 0.0359. This number may be compared to the income share of investments in intangible capital that is broadly similar to organizational capital. Corrado et al. (2006) refer to these as investments in firm-specific resources. Their estimates suggest that such investments account for approximately one-third of all investments in intangible capital, providing an income share of investments in intangible firm-specific resources of 5%.

For the data under consideration, the average per capita total hours worked per quarter is given by approximately 300 hours per quarter.[16] With a calibrated total time endowment of 1,369 hours per quarter, this provides average per capita total hours, as a proportion of the total time endowment H/T, of 21%.[17]

The preferences of the representative household, given by (8), imply the following inter-temporal elasticity of labor supply (iels):

$$\text{iels} = \frac{1}{\nu} \frac{T - H_t}{H_t}.$$

The value of $\nu = 2.7226$ was picked to deliver an intertemporal labor supply elasticity of 1.3088 as estimated in Chang et al. (2002),[18] which is close to the upper end of estimates of the inter temporal labor supply elasticity using microeconomic data.[19]

The first-order conditions from the household's maximization problem imply the following relationship for the labor wedge:

$$\ln B_t = -\ln C_t + \nu (T - H_t) + \ln (W_t).$$

Consequently, an (implied) series for the exogenous process may be determined using data on consumption C_t, hours H_t, and wages W_t. This series B_t can then be used to estimate the second moments of the process according to

$$\ln B_t = (1 - \rho_b) \ln \bar{B} + \rho_b \ln B_{t-1} + \upsilon_{bt}.$$

Obviously, these moments will depend upon the value of the preference parameter ν. Based upon the data, the model with $\nu = 2.7226$ is calibrated with $\rho_b = 0.9536$ and $\sigma_b = 0.0107$.

5. SIMULATION RESULTS

This section presents quantitative results from simulating the calibrated linearized versions of the learning models along with a baseline real business cycle model without learning effects. We begin with a discussion of the implications for productivity and then briefly discuss the implied second moments usually presented in the business cycle literature.

TABLE 1. Productivity and wages

	σ_{sr}	ρ_{sr}	$\rho_{h,alp}$	$\rho_{h,w}$	$\rho_{w,y}$
U.S. data	0.0208	0.9457	0.0828	0.1900	0.7034
No learning	0	n/a	−0.4316	−0.4316	−0.1217
20% learning 1 ($\eta = 0.5$)	0.0094	0.9951	0.2453	0.2487	0.6730
20% learning 2 ($\eta = 0.8$)	0.0084	0.9981	0.1416	0.1273	0.6039
10% learning ($\eta = 0.5$)	0.0103	0.9954	0.2695	0.2752	0.7042

Table 1 presents some key results for productivity and wages. The first two columns contain the standard deviation and first-order autocorrelation coefficient of the Solow residual, whereas the next two columns present the contemporaneous correlations of hours with average labor productivity and wages. The final column presents the contemporaneous correlation between output and wages. Comparing row 2 of Table 1 with row 1, it is clear that the baseline business cycle model, without learning, generates results that are incompatible with U.S. data on wages and productivity. Not only is there no mechanism for endogenously generating movements in the Solow residual (efficiency wedge), but also average labor productivity and wages are both highly negatively correlated with hours, whereas the correlations in U.S. data are weakly positive. Similarly, wages are negatively correlated with output, whereas they are highly positively correlated in the data.

The next three rows present results for three variants of the model with internal learning-by-doing. Rows 3 and 4 correspond to a 20% learning rate ($\varepsilon = 0.26$), whereas the learning rate in row 5 is 10% with $\gamma = 1$, as is typical in the learning literature. Examination of these rows reveals that all three specifications generate considerable movement in the Solow residual. Depending on the specification, the model can account for 40% to 50% of the volatility of the Solow residual seen in the U.S. data.[20] The model with learning also does quite well in replicating the low positive correlation of hours with wages and average labor productivity as well as the high positive correlation between output and wages. The results from the external learning model are quite similar except that in general the model generates a little more endogenous movement in the Solow residual and slightly higher correlations of productivity and wages with hours and output. These results are presented in the Appendix with no discussion.

To understand why the learning models generate a small correlation between wages (average labor productivity) and hours, it is useful to think in terms of a labor demand–labor supply diagram. The shock causes agents to shift their labor supply curves outward, leading to an increase in hours worked. Absent the learning channel, this would cause a movement down the labor demand curve, leading to a sharp fall in wages and labor productivity, as in the model without learning. Once firms are allowed to learn from production, we get an offsetting second shift of the labor demand curve to the right, induced by the increase in future productivity due to the accumulation of organizational capital. The net result of the two shifts is a large movement in hours, even if the supply curve is not very responsive to

TABLE 2. Selected second moments

	$\sigma_{C/Y}$	$\sigma_{I/Y}$	$\sigma_{H/Y}$	$\rho_{C,Y}$	$\rho_{I,Y}$	$\rho_{H,Y}$	ρ_Y	σ_Y
U.S. data	0.79	1.76	0.72	0.92	0.83	0.78	0.97	0.0399
No learning	0.78	2.79	1.10	0.94	0.86	0.95	0.97	0.0267
20% learning 1 ($\eta = 0.5$)	0.79	2.15	0.77	0.94	0.89	0.88	0.99	0.0358
20% low learning 2 ($\eta = 0.8$)	0.81	2.05	0.80	0.95	0.89	0.87	0.99	0.0330
10% learning ($\eta = 0.5$)	0.80	2.12	0.74	0.94	0.89	0.88	0.99	0.0371

movements in the real wage, and a small movement in labor productivity. This analysis is reminiscent of Christiano and Eichenbaum (1992), but recall that in that model the two shifts occurred in response to two separate shocks; a government spending shock, and a technology shock, whereas here the only exogenous shift is in labor supply.

The analysis is also reminiscent of the discussion in Cooper and Johri (2002) of how the representative agent responds to a technology shock. There, as is usual, the technology shock causes a rightward shift in labor demand whereas the wealth effect of the accumulation of organizational capital is to shift the labor supply curve to the left. Once again, the model delivers a lower correlation of hours and average labor productivity than the baseline model, but it also delivers a smaller movement in hours, which is counterfactual. This occurs because the labor supply curve shifts left, countering some of the rightward movement of labor demand.

Given this analysis, it is useful to compare the moments reported in Table 1 with similar moments generated by total factor productivity (tfp) shocks instead of preference shocks.[21] These are reported in Table A.3.[22] As expected, the no-learning model generates too high correlations of wages (and productivity) with hours and output. The actual values are 0.54 and 0.97, respectively. As discussed above, the internal learning-by-doing models are able to lower the correlation of wages with hours, but the quantitative impact of this is small. As a result, we conclude that the performance of the learning-by-doing model with preference shocks is superior to its performance with technology shocks.

Figures 1 and 2 present impulse response plots for the high-learning model (stars) as well as the baseline model without learning effects (solid lines). The impact of learning-by-doing on endogenous productivity movements is clearly visible in the strong hump-shaped response of the Solow residual. As expected, absent any learning mechanism, average labor productivity falls sharply when hours rise in the impact period and remains below steady state levels for four years. With learning, organizational capital is built up due to above–steady state output. On its own, this effect increases productivity, thus counteracting the negative influence of the increase in hours worked. As a result, labor productivity is below steady state for only six quarters, as opposed to fourteen quarters.

Table 2 presents selected second moments for the four models discussed above, as well as for linearly detrended aggregate U.S. data. Each row of the table is set up as before, with the first three columns reporting standard deviations relative to output. The next three columns report contemporaneous correlations with output,

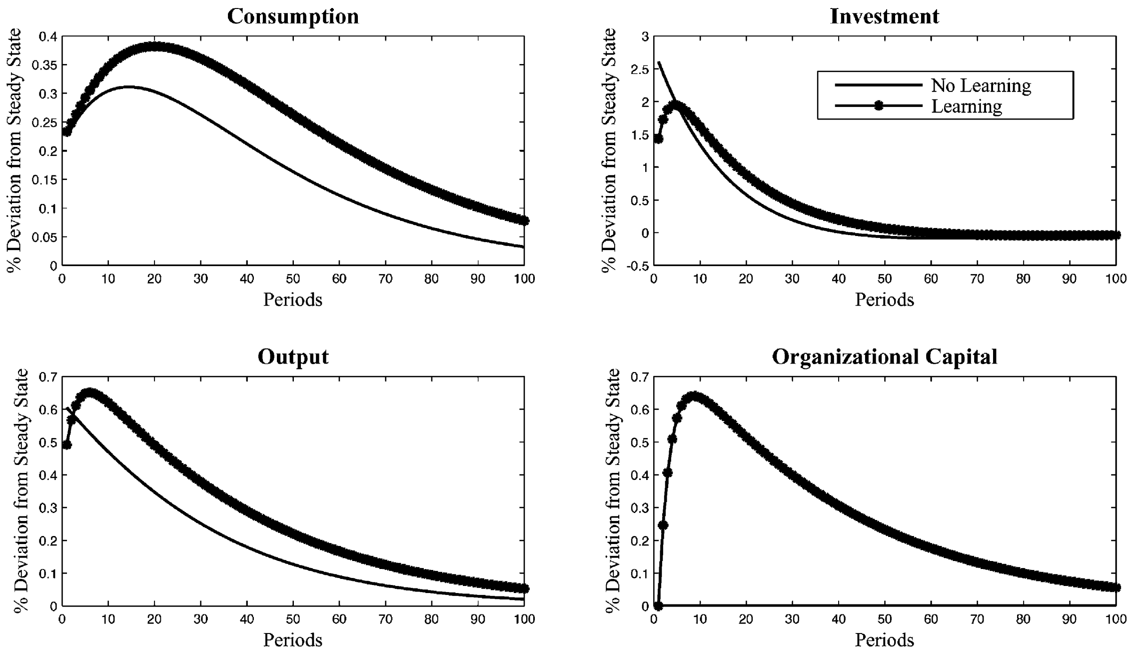

FIGURE 1. Internal learning-by-doing: selected impulse responses to a 1% shock to the labor wedge.

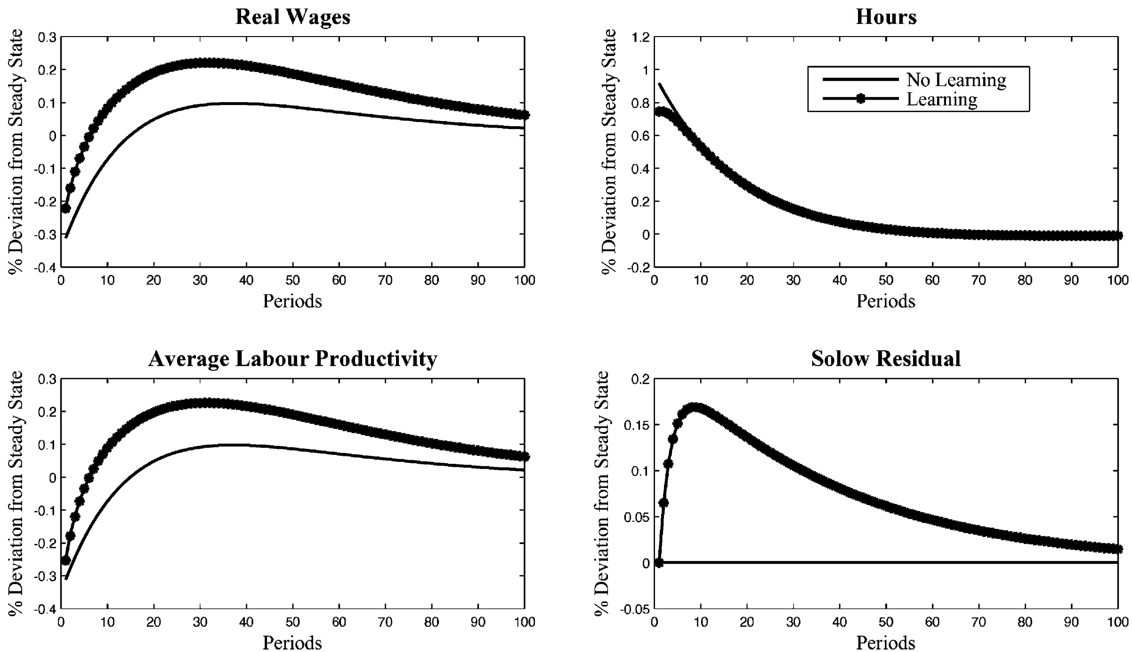

FIGURE 2. Internal learning-by-doing: selected impulse responses to a 1% shock to the labor wedge.

whereas the last two report the first-order autocorrelation and standard deviation of output, respectively. A glance at Table 2 suggests that all the models do a reasonable job of capturing the basic features of consumption and investment over the business cycle. As in the data, consumption, investment, and hours are positively correlated with output, and consumption is less volatile than output, whereas investment is more volatile. Looking down the second column, one notices that learning tends to temper the relative volatility of investment, which is too high compared to the data in the baseline model without learning. This moderating influence is clearly evident in Figure 1, where we see that investment rises by just over 2.5% in the impact period without learning, whereas the rise is just under 1.5% in the model with learning. A small hump-shaped pattern is also evident in this case.

The impact of learning-by-doing shows up more starkly when we study the behavior of aggregate hours. The baseline model generates a very high correlation between hours and output (0.95), which is reduced in all three specifications of the internal learning model. Similarly, hours are too volatile in the baseline model. Without learning effects, the relative volatility of hours with respect to output is 1.1, whereas this falls to around 0.7 with learning. This is much closer to the U.S. data, which exhibit a relative volatility of 0.72. The last two columns of Table 2 display the effectiveness of learning-by-doing as a propagation mechanism. Not only does learning-by-doing increase the persistence of output, but also it magnifies the effect of shocks, as is evident in the increase in the volatility of output from 0.0267 in the baseline model to roughly 0.0358 in the 20% learning specification, which is about 90% of the volatility of output seen in aggregate U.S. data. The extra propagation effect of learning-by-doing can also be seen in the strong hump-shaped response of output in Figure 1.

The external learning model produces results that are very similar to those reported above both qualitatively and quantitatively. Tables A.1 and A.2 in the Appendix provide the details. Here we quickly summarize the key differences. In general, given our specification for the learning parameters, the external model tends to magnify and propagate shocks more strongly. As a result, the volatility of productivity measures, output, and investment tends to be higher. For the same reason, the correlation of labor productivity and hours is somewhat higher as well.

Turning to the behaviour of the internal learning-by-doing model in the presence of technology shocks, we see from Table A.4 in the Appendix that the moments largely mirror the results reported above. As widely documented in the literature, the key area of underperformance for both the models with and without learning effects is the volatility of hours. With persistent tfp shocks, hours do not move enough.

6. CONCLUSIONS

This paper presents two business cycle models with learning-by-doing that are quite successful in explaining the broad features of the U.S. business cycle. The only shock to the economy comes from exogenous movement in the labor wedge, which we interpret as an unmodeled fluctuation in labor supply. Our models endogenously generate procyclical movements in labor productivity, wages, and the Solow residual that are highly persistent and volatile. The learning-by-doing models also do well in generating a realistically small correlation between average labor productivity (wages) and hours, an elusive goal for most business cycle models. We show that the learning effects built into these models are key aspects of this success. Absent learning, the baseline model performs quite poorly: as expected, labor productivity is strongly countercyclical and there is no movement in the Solow residual. Moreover, hours and investment are too volatile relative to output. Learning-by-doing is able to temper this excessive volatility. Whereas the benchmark no-learning model can account for around 67% of the volatility of aggregate output seen in U.S. data, the learning models can account for virtually all of it.

Chari et al. (2007) argue that wedges in the hours first-order condition (labor wedge) and in the production function (efficiency wedge) of a prototypical real business cycle model are key ingredients in accounting for modern business cycle fluctuations. In this paper we show that learning-by-doing can be an effective ingredient in developing an explanation of the efficiency wedge. Specifically, we show that exogenous movements in the labor wedge can lead to endogenous movements in the efficiency wedge. The model can account for about 50% of the observed movement in the wedge.

NOTES

1. See the discussion and references in King and Rebelo (1999).
2. See Chari et al. (2007, p. 828).
3. There is considerable evidence that organizations accumulate intangible capital that is specific to the organization. Some (though not all) aspects of this are related to learning-by-doing. Recent attempts to quantify different aspects of organizational capital under varying names can be found in Brynjolfsson and Hitt (2000), Hall (2000), Lev and Radharkrishnan (2003), and Atkeson and Kehoe (2005) and the references therein.
4. This paper also differs in that business cycles are not generated by technology shocks.
5. This countercyclical (and counterfactual) movement in average labor productivity has long troubled models of the business cycle driven by aggregate demand disturbances. See Rotemberg and Woodford (1991, 1999) for a discussion of the various attempts to overcome these problems.
6. The model outlined below takes the standard Blanchard and Kiyotaki (1987) model of monopolistic competition, commonly used in aggregate general equilibrium models, modified to allow for the accumulation of organizational capital. Neither Cooper and Johri (2002) nor Chang et al. (2002) includes this feature.
7. Atkeson and Kehoe (2005) model and estimate the size of organizational capital for the U.S. manufacturing sector and find that it has a value of roughly 66% of physical capital. Although their broad interpretation of organizational capital is similar to ours, they do not allow for depreciation of past learning.

8. Because we are focusing on movements at the business cycle frequency, we ignore technical progress of any form.

9. Another seemingly crucial feature of our model is that organizational capital is accumulated as a by-product of production. This ignores the considerable intentional investments made by firms in raising productivity. Hou and Johri (2007) show that allowing for intentional investments in organizational capital results in only small differences from the Cooper and Johri (2002) model.

10. Given our specification for preferences outlined in the next section,

$$D_{t+1} = \beta E\left[\frac{C_t}{C_{t+1}}\right].$$

11. See Uhlig (2004) for the use of preference shocks in the context of an interesting model of labor hoarding involving workplace leisure.

12. There is also a literature that suggests these shifts arise as a result of "aggregation error" in a model with heterogenous agents. See Maliar and Maliar (2004) and Chang and Kim (2007).

13. Papers in the organizational capital literature often work with decreasing returns in labor and capital. Typically the implied capital output ratios are much lower than found in aggregate U.S. data. Our results are not very sensitive to this assumption.

14. This value is also (broadly) consistent with industry-level and plant-level estimates in Clarke (2007), based upon structural estimation of the firm side of the internal learning model outlined above. Note that Benkard (2000) has argued that allowing for "organizational forgetting" leads to higher estimates of the learning rate. For example, in his work on aircraft production, the estimated learning rate rises to 39% once forgetting is allowed. Keeping this in mind, and the even higher estimate of $\varepsilon = 0.49$ for the aggregate economy in Cooper and Johri (2002), it would not be unreasonable to allow an $\varepsilon = 0.4$, which corresponds to a learning rate of roughly 32%. Results for this case were reported in an earlier version of this paper and are available from the authors.

15. Business cycle models are typically calibrated to a markup of 11.11%. The implied q is greater than 2 for a typical calibration.

16. We use aggregate U.S. data at the quarterly frequency covering the period 1955:Q1–1992:Q4. We thank Michelle Alexopoulos for providing these data to us.

17. See Christiano and Eichenbaum (1992), Burnside et al. (1993), or Burnside and Eichenbaum (1996).

18. We also explored the impact of setting $\nu = 1$ which, together with the steady state h/T, implies an intertemporal labor supply elasticity of 3.563. This value is commonly used in the business cycle literature. The results were similar and are available from the authors upon request.

19. See Browning et al. (1999) for a review of this literature.

20. Obviously the amount of variation in the Solow residual will depend on how productive organizational capital is in output production. In an earlier version of the paper we showed that even with a moderate learning rate of 32% the model could account for 80% of the variation in the Solow residual.

21. This requires some modification of the model outlined above. In the symmetric equilibrium, let the aggregate production technology be given by $Y_t = A_t K_t^{\theta} H_t^{\alpha} Z_t^{\varepsilon}$. The tfp process A_t evolves according to

$$\ln A_t = (1-\rho_a)\ln \bar{A} + \rho_a \ln A_{t-1} + \upsilon_{at}.$$

The process is calibrated to deliver the observed autocorrelation and variance of the Solow residual (see Table A.3) with $\rho_a = .9457$ and $\sigma_a = 0.0068$.

22. The value for μ varies across specifications to deliver a constant capital–output ratio and a steady state markup of 15%.

REFERENCES

Argote, L., S.L. Beckman, and D. Epple (1990) The persistence and transfer of learning in industrial settings. *Management Science* 36(2), 140–154.

Atkeson, A. and P.J. Kehoe (2005) Modeling and measuring organization capital. *Journal of Political Economy* 113(5), 1026–1053.

Bahk, B.-H. and M. Gort (1993) Decomposing learning by doing in new plants. *Journal of Political Economy* 101(4), 561–583.

Barnett, S.A. and P. Sakellaris (1999) A new look at firm market value, investment, and adjustment costs. *Review of Economics and Statistics* 81(2), 250–260.

Basu, S. (1996) Procyclical productivity: Increasing returns or cyclical utlilization? *Quarterly Journal of Economics* 111(3), 719–751.

Basu, S. and J.G. Fernald (1997) Returns to scale in U.S. production: Estimates and implications. *Journal of Political Economy* 105(2), 249–283.

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

Benkard, C.L. (2000) Learning and forgetting: the dynamics of aircraft production. *American Economic Review* 90(4), 1034–1054.

Blanchard, O. and N. Kiyotaki (1987) Monopolistic competition and aggregate demand disturbances. *American Economic Review* 77(4), 647–666.

Browning, M., L.P. Hansen, and J. Heckman (1999) Micro data and general equilibrium models. In J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics, Vol. IA*, pp. 543–633. Amsterdam: Elsevier Science.

Brynjolfsson, E. and L.M. Hitt (2000) Beyond computation: information technology, organizational transformation and business performance. *Journal of Economic Perspectives* 14(4), 23–48.

Burnside, C. and M. Eichenbaum (1996) Factor-hoarding and the propagation of business-cycle shocks. *American Economic Review* 86(5), 1154–1174.

Burnside, C., M. Eichenbaum, and S. Rebelo (1993) Labor hoarding and the business cycle. *Journal of Political Economy* 101, 245–273.

Chang, Y., J.F. Gomes, and F. Schorfheide (2002) Learning-by-doing as a propagation mechanism. *American Economic Review* 92(5), 1498–1520.

Chang, Y. and S.-B. Kim (2007) Heterogeneity and aggregation in the labor market: Implications for aggregate preference shifts. *American Economic Review* 97(5), 1939–1956.

Chari, V.V., P.J. Kehoe, and E.R. McGrattan (2002) Accounting for the great depression. *American Economic Review* 92(2), 22–27.

Chari, V.V., P.J. Kehoe, and E.R. McGrattan (2007) Business cycle accounting. *Econometrica* 75(3), 781–836.

Christiano, L.J. and M. Eichenbaum (1992) Current real-business-cycle theories and aggregate labor-market fluctuations. *American Economic Review* 82(3), 430–450.

Clarke, A.J. (2006) Learning-by-doing and aggregate fluctuations: Does the form of the accumulation technology matter? *Economics Letters* 92, 434–439.

Clarke, A.J. (2007) Learning-by-Doing and Productivity Dynamics in Manufacturing Establishments. Manuscript, University of Melbourne.

Cole, H.L. and L.E. Ohanian (2002) The U.S. and U.K. great depressions through the lens of neoclassical growth theory. *American Economic Review* 92(2), 28–32.

Cole, H.L. and L.E. Ohanian (2004) New deal policies and the persistence of the great depression: A general equilibrium analysis. *Journal of Political Economy* 112(4), 779–816.

Cooper, R.W. and A. Johri (1997) Dynamic complementarities: A quantitative analysis. *Journal of Monetary Economics* 40, 97–119.

Cooper, R.W. and A. Johri (2002) Learning-by-doing and aggregate fluctuations. *Journal of Monetary Economics* 49, 1539–1566.

Corrado, C.A., C.R. Hulten, and D.E. Sichel (2006) Intangible capital and economic growth. NBER working paper 11948.

Darr, E.D., L. Argote, and D. Epple (1995) The acquisition, transfer, and depreciation of knowledge in service organizations: productivity in franchises. *Management Science* 41(11), 1750–1762.

Hall, R.E. (1997) Macroeconomic fluctuations and the allocation of time. *Journal of Labor Economics* 15(2), S223–S250.

Hall, R.E. (2000) E-capital: The link between the stock market and the labor market in the 1990's. *Brookings Papers on Economic Activity* 2, 73–102.

Hou, K. and A. Johri (2007) Costly Investments in Organizational Capital and Learning-by-Doing. Manuscript, McMaster University.

Ireland, P.N. (2004a) Money's role in the monetary business cycle. *Journal of Money, Credit and Banking* 36(6), 969–983.

Ireland, P.N. (2004b) Technology shocks in the new keynesian model. *Review of Economics and Statistics* 86(4), 923–936.

Irwin, D.A. and P.J. Klenow (1994) Learning-by-doing spillovers in the semiconductor industry. *Journal of Political Economy* 102(6), 1200–1227.

Jarmin, R.S. (1994) Learning by doing and competition in the early rayon industry. *RAND Journal of Economics* 25(3), 441–454.

Johri, A. and M.-A. Letendre (2007) What do "residuals" from first-order conditons reveal about dge models? *Journal of Economic Dynamics and Control* 31(8), 2744–2773.

King, R.G. and S.T. Rebelo (1999) Resuscitating real business cycles. In J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. IB, pp. 927–1007. Amsterdam: Elsevier Science.

King, R.G. and M.W. Watson (2002) System reduction and solution algorithms for singular linear difference systems under rational expectations. *Computational Economics* 20, 57–86.

Lang, L. H.P. and R.M. Schultz (1994) Tobin's q: Corporate diversification and firm performance. *Journal of Political Economy* 6, 1248–1280.

Lev, B. and S. Radharkrishnan (2003) The Measurement of Firm-Specific Organizational Captial. NBER working paper 9581.

Maliar, L. and S. Maliar (2004) Preference shocks from aggregation: time series evidence. *Canadian Journal of Economics* 37(3), 768–781.

Mulligan, C.B. (2002) A Century of Labor-Leisure Distortions. NBER working paper 8774.

Parkin, M. (1988) A method for determining whether parameters in aggregative models are structural. *Carnegie-Rochester Conference Sreies on Public Policy* 29, 215–252.

Rosen, S. (1972) Learning by experience as joint production. *Quarterly Journal of Economics* 86(3), 366–382.

Rotemberg, J.J. and M. Woodford (1991) Markups and the business cycle. *NBER Macroeconomics Annual* 63–129.

Rotemberg, J.J. and M. Woodford (1999) The cyclical behaviour of prices and costs. In J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. IA, pp. 1051–1135. Amsterdam: Elsevier Science.

Summers, L.H., B.P. Bosworth, J. Tobin, and P.M. White (1981) Taxation and corporate investment: A q-Theory Approach. *Brookings Papers on Economic Activity* 1, 67–140.

Thompson, P. (2001) How much did the liberty shipbuilders learn?: New evidence for an old case study. *Journal of Political Economy* 109(1), 103–137.

Thornton, R.A. and P. Thompson (2001) Learning from experience and learning from others: An exploration of learning and spillovers in wartime shipbuilding. *American Economic Review* 91(5), 1350–1369.

Uhlig, H. (2004) Do technology shocks lead to a fall in total hours worked? *Journal of the European Economics Association* 2, 361–371.

APPENDIX

SIMULATION RESULTS: EXTERNAL LEARNING-BY-DOING

TABLE A.1. External model: Productivity and wages

	σ_{sr}	ρ_{sr}	$\rho_{h,alp}$	$\rho_{h,w}$	$\rho_{w,y}$
U.S. data	0.0208	0.9457	0.0828	0.1900	0.7034
No learning	0	n/a	−0.4316	−0.5782	−0.1217
20% learning 1 ($\eta = 0.5$)	0.0114	0.9960	0.2886	0.2886	0.7258
20% learning 2 ($\eta = 0.8$)	0.0103	0.9984	0.1967	0.1967	0.6814
10% learning ($\eta = 0.5$)	0.0126	0.9963	0.3104	0.3104	0.7530

TABLE A.2. External model: Selected second moments

	$\sigma_{C/Y}$	$\sigma_{I/Y}$	$\sigma_{H/Y}$	$\rho_{C,Y}$	$\rho_{I,Y}$	$\rho_{H,Y}$	ρ_{Y}	σ_{Y}
U.S. data	0.79	1.76	0.72	0.92	0.83	0.78	0.97	0.0399
No learning	0.78	2.79	1.10	0.94	0.86	0.95	0.97	0.0267
20% learning 1 ($\eta = 0.5$)	0.84	2.31	0.72	0.97	0.87	0.87	0.99	0.0434
20% low learning 2 ($\eta = 0.8$)	0.85	2.18	0.75	0.97	0.88	0.85	0.99	0.0343
10% learning ($\eta = 0.5$)	0.84	2.28	0.69	0.97	0.87	0.86	1.00	0.0384

SIMULATION RESULTS: TECHNOLOGY SHOCKS

TABLE A.3. Productivity and wages

	σ_{sr}	ρ_{sr}	$\rho_{h,alp}$	$\rho_{h,w}$	$\rho_{w,y}$
U.S. data	0.0208	0.9457	0.0828	0.1900	0.7034
No learning	0.209	0.9457	0.5448	0.5448	0.9734
20% learning 1 ($\eta = 0.5$)	0.0339	0.9789	0.5312	0.5253	0.9822
20% learning 2 ($\eta = 0.8$)	0.0315	0.9763	0.4934	0.4872	0.9847
10% learning ($\eta = 0.5$)	0.0352	0.9804	0.5270	0.5203	0.9828

TABLE A.4. Selected second moments

	$\sigma_{C/Y}$	$\sigma_{I/Y}$	$\sigma_{H/Y}$	$\rho_{C,Y}$	$\rho_{I,Y}$	$\rho_{H,Y}$	ρ_{Y}	σ_{Y}
U.S. data	0.79	1.76	0.72	0.92	0.83	0.78	0.97	0.0399
No learning	0.76	2.97	0.27	0.93	0.86	0.72	0.96	0.0361
20% learning 1 ($\eta = 0.5$)	0.77	2.26	0.23	0.93	0.89	0.68	0.99	0.0549
20% learning 2 ($\eta = 0.8$)	0.79	2.15	0.21	0.94	0.89	0.64	0.98	0.0502
10% learning ($\eta = 0.5$)	0.78	2.23	0.22	0.93	0.89	0.67	0.99	0.0570