

Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries*

Kathryn Koenders
Arizona State University

Richard Rogerson
Arizona State University

October 10, 2004

Abstract

This paper proposes a new explanation for the apparent slow growth in employment during the last two recoveries. Our explanation emphasizes dynamics within growing organizations and the intertemporal substitution of organizational restructuring. A key implication of our analysis is that recoveries from recessions following long expansions will have slower employment growth. Empirical analysis shows that the recovery which began in 1970 also exhibited slow employment growth, consistent with this prediction of the analysis.

*This paper was prepared for the 2004 St. Louis Federal Reserve Bank Economic Policy Conference on "Productivity, Labor and the Business Cycle". Rogerson acknowledges support from the NSF.

1 Introduction

Since the work of Burns and Mitchell (1946), economists who study business cycle fluctuations typically refer to the “business cycle facts” without need to reference a particular episode in a particular country. One of the accepted stylized facts of business cycle movements is that employment and output are strongly positively correlated, although employment lags output by about one quarter. The apparent slow growth of employment in the recoveries following the last two US recessions (i.e., the so-called “jobless recovery” phenomenon) runs counter to this stylized fact. This paper suggests a possible explanation for this apparently anomalous behavior.

The two most recent recessions in the US share a common property: both followed unusually long expansions. Motivated by this observation we propose an economic mechanism that links the speed with which employment increases during the recovery from a recession to the length of the expansion preceding the recession. The mechanism stresses the manner in which organizations seek to eliminate unneeded labor and is easily described. We assume that inefficiencies regarding the use of labor emerge over time within an organization. Eliminating these inefficiencies (a process we refer to as reorganizing) requires scarce organizational resources that must be diverted away from current production. This tradeoff generates opportunities for intertemporal substitution and we show that reorganization will be postponed to periods in which production is relatively low. It follows that after a long expansion, many more organizations have postponed reorganization. Because reorganization leads to the shedding of unnecessary labor and takes time, this gives rise to an extended period in which the economy sheds labor, thereby delaying the date at which aggregate employment begins to increase during the recovery.

The first part of this paper presents one formulation of a model of organizations that generates these effects. The core model should be seen as an extension of the Lucas (1978) span of control model and the Hopenhayn (1992) industry equilibrium model to allow for a richer set of dynamics within an organization. Although we focus on the implications for business cycles, we believe this model may also prove useful for examining plant and firm level dynamics more generally. The model is purposefully simplified in order to highlight the key economic tradeoffs, and the analysis focuses on the qualitative nature of the interactions. We leave the task of building a model suitable for quantitative analysis of the forces for future work.

Any theory that links the speed of the recovery in employment to the length of the preceding expansion would seem to be a viable candidate for explaining the anomalous behavior of employment following the last two recessions. However, since the recession of 1969-1970 also followed an unusually long expansion, an obvious implication of any such theory is that a slow recovery of employment should also have been observed following this recession. The second part of the paper turns to this issue and argues that when viewed from the perspective of our model, the behavior of employment in the 1970 recovery is in fact very similar to the behavior of employment in the recoveries of 1991 and 2001, and is

qualitatively different from the behavior of employment in the other post WWII recoveries.

One interpretation of our explanation is that the recent recessions do not represent counterexamples to the standard set of business cycle facts, but rather that the business cycle facts need to be modified somewhat to acknowledge that recoveries following long expansions exhibit somewhat different dynamics. There are of course, other types of explanations that one might appeal to. For example, if one thought that it is only the two most recent recessions that appear different, one might consider the possibility that the business cycle facts are evolving and seek to understand what features of the economy are changing that would lead to the change in business cycle dynamics of employment. Schreft and Singh (2003) pursue this tack, arguing that increased flexibility in personnel policies are responsible for the change. A second general class of explanations is to posit that the anomalous behavior of employment in any particular cycle is due to an additional shock or policy change that happens to coincide with the recovery. The work of Cole and Ohanian (2004) on employment during the second half of the Great Depression is exactly this type of explanation—they argued that employment did not recover as one would have expected because of the adoption of the New Deal Policies and how they affected real wage rates. A third and related class of explanations stresses that trends in the economy change over time and that cyclical properties may be affected by the nature of the trend changes in the background. Along these lines a commonly heard explanation for the recent slow recovery of employment was that the world had become more uncertain, leading firms to postpone increases in employment. Groshen and Potter (2003) is an example of this class. They argue that the underlying amount of reallocation of labor across firms is currently greater and that this affects the pace of employment growth during the recovery. We do not carry out any comparison of our explanation versus these others, but do note that the extent to which the 1970 recovery is viewed as being similar to the two most recent recoveries would cast some doubt on the theory that there is evolution in the business cycle facts.

The mechanism that we describe is related to others that have appeared in the literature, and it is of interest to note the similarities and differences. An old idea in the business cycle literature is that recessions are periods of restructuring. However, modern formulations of this, such as Lilien (1982), are based on the notion that the key element of restructuring is across organizations, in particular, that resources need to be reallocated across sectors. In contrast, our model does not stress the reallocation of resources from one organization to another but rather the restructuring that takes place within organizations that leads to the elimination of wasteful employment. Hall (1991) argued that recessions should be thought of as reorganizations, but the reorganizations that he stresses were really reallocations of labor across activities. As noted above, Groshen and Potter (2003) have stressed that the amount of reallocation needed may vary over time and be higher in some business cycle episodes than in others.

A second related literature is associated with the work of Caballero and Engel (1999) and Krusell and Smith (1998). Both papers consider the possi-

bility that the distribution of individual state variables can influence how the economy responds to shocks. Caballero and Engel argue that the distribution of the difference between plant-level capital stocks and their ideal points was quantitatively significant for the response of the economy to shocks. Krusell and Smith examine how the distribution of asset holdings across consumers affects propagation of shocks. Qualitatively, our model emphasizes a similar channel, since we argue that the distribution of efficiencies of organizations matters for how the economy responds to shocks. However, despite the similarity the mechanics are quite different. In particular, if carried over to the labor setting, the Caballero and Engel model cannot explain why aggregate demand for labor would continue to decrease in the face of positive aggregate shocks, a result that can emerge in our model and is central to accounting for the delayed increase in employment that accompanied the last two recoveries.

An outline of the paper follows. In Section 2 we describe our benchmark model, which captures the evolution of an individual organization over its life-cycle. In Section 3 we consider how aggregate temporary shocks interact with the decisions of the organization, and derive our key result, which is that organizations will concentrate reorganizations during periods with low aggregate shocks. Section 4 discusses the implications of this finding for the cyclical properties of labor demand. Section 5 carries out an empirical analysis of postwar business cycles and shows that the recovery following the 1969-1970 recession exhibits patterns for employment that are quantitatively very similar to those found in the two most recent recoveries. Section 6 concludes.

2 Benchmark Model: A Life Cycle Model of Organizational Dynamics

In this section we formulate a model of the life cycle of an organization. In the next section we will use this benchmark model to investigate how the resulting pattern of organizational dynamics may be affected by shocks that we interpret to be business cycle shocks. The goal of these two sections is to highlight a particular interaction between organizational dynamics and business cycle shocks.¹ With this in mind, we purposefully work in a very simple setting in order to best highlight this interaction. We leave the development of a model that would be useful for a quantitative assessment of these interactions to future work.

The essence of benchmark model is as follows. We view an organization as producing a differentiated product and therefore facing a downward sloping demand curve for its product. Our model captures the following stylized evolution of an organization over its lifecycle—when an organization is first created it faces a relatively low demand for its product. But, over time this situation may change. Some organizations fail and disappear, while some others experience

¹Our model without aggregate shocks can also be viewed as extending the model of Hopenhayn (1992) to consider richer decision making processes within organizations.

large increases in the demand for their product and hence grow. However, even those organizations that grow and become large will at some point experience decreases in their demand and eventually cease to exist.

The model that we describe below is a model of an individual organization that faces stochastic demand for its product, but takes all input prices as given, where these prices remain constant over time. We assume that the organization maximizes the present discounted value of profits using a discount rate of β , which one can think of as one divided by one plus the interest rate. We now describe the specifics of the model in more detail.

2.1 Demand

Let $P(y_t, \varepsilon_t)$ be the inverse demand faced by the organization in period t , where y_t is the amount of output produced in the current period and ε_t is a stochastic shock to the demand for the output of the organization. For a given value of ε_t we assume that this function is twice continuously differentiable and strictly decreasing in y , that the product $P(y, \varepsilon)y$ is strictly concave in y , and satisfies the boundary conditions:

$$\lim_{y \rightarrow 0} \frac{d}{dy} P(y, \varepsilon)y = \infty, \text{ and } \lim_{y \rightarrow \infty} \frac{d}{dy} P(y, \varepsilon)y < 0.$$

The first condition will ensure that an organization that remains in existence will always want to produce a positive amount of output, while the second condition says that output can effectively be viewed as bounded for any given ε , since the organization will never produce beyond the point where revenues are decreasing in output.

We assume a very simple form of the stochastic process on ε . In particular, we assume that ε_t takes on only one of two values, ε^s and ε^l , where $\varepsilon^s < \varepsilon^l$, with the interpretation that ε^s is the demand state that will give rise to a small organization, while ε^l is the demand state that will give rise to a large organization. To generate the standard life cycle profile of organization size, we assume that when an organization is first created it will have a value of ε equal to ε^s , and that over time the state of demand may increase to ε^l . We simplify this process by assuming that the probability that ε increases from ε^s to ε^l is given by the value π^l , which is assumed to be constant over time. We also assume that the process never transits from ε^l back to ε^s . To capture the notion that state ε^l is better than state ε^s , we assume that $P(y, \varepsilon^l) > P(y, \varepsilon^s)$ for all positive values of y and that $\frac{d}{dy} P(y, \varepsilon^l)y > \frac{d}{dy} P(y, \varepsilon^s)y$ for all positive y as well.

If the stochastic evolution of ε described above was a complete description of the uncertainty facing a given organization, then all new organizations would eventually become “large” and remain that way forever. We incorporate the fact that organizations do not last forever by assuming that organizations face an exogenous probability of death, and that this probability varies with the state of their demand, ε . In particular we assume that a organization with demand state ε^i faces a probability λ^i of death. This assumption implies not only that

organizations do not last forever, but also that not all new organizations will necessarily become “large”.² We assume that the realizations of the random variables are independent.

Our assumption about timing is as follows. An organization that produced in period t with demand state ε^i finds out the realizations of the demand and death shocks at the beginning of period $t + 1$ but before any decisions are made in period $t + 1$.

2.2 Production

We assume that labor is the only factor of production. The production technology in our model has several key features, which we detail in several steps.

2.2.1 Scale Effects

First, we assume that there are different ways to organize production, and that the optimal way to organize production is dependent upon the scale of production. In general one could imagine a large set of possible ways to organize production, but given that we are restricting attention to a model with two demand states we also assume that there are only two ways in which production can be organized. We refer to each different way to organize production as a distinct technology. We let $h^i(y)$ denote the labor necessary to produce output y using technology i . Because we will implicitly restrict parameter values so that an organization with demand state ε^s will always use technology 1 while an organization with demand state ε^l will always use technology 2, we will also use s and l to index the two technologies. The idea is that technology one is better if producing at a low scale, while technology two is better at producing at a large scale. We adopt the following functional forms for the two technologies. We assume that $h^s(y) = a_s y$ for all positive values of y , where $a_s > 0$. The other technology is only operable above some minimum threshold scale \bar{y} , and for $y \geq \bar{y}$ we assume that $h^l(y)$ takes the form $h^l(y) = a_l y$, where $a_l < a_s$. More generally we could specify that this technology can operate at scale less than \bar{y} but that the average product of labor is sufficiently low that no one would choose to operate it below the scale \bar{y} . We assume free disposal of output, implying that an organization could choose to operate the large scale technology but only sell a fraction of the output, though in our analysis we will implicitly assume that this never happens.

We assume that the organization hires labor in a competitive market and hence takes the wage rate as given. In what follows we normalize the wage rate to one.

²This specification amounts to assuming that there are three demand states, with the third state being an absorbing state in which the organization cannot sell any positive amount of output and still receive a positive price. We identify this third state with death of the organization.

2.2.2 Organizational Waste

An important goal for any organization is to use its resources efficiently. The large differences in measured productivity across organizations suggests that organizations vary in the degree to which they accomplish this. Inefficient use of resources may take several forms. We incorporate one particular form of inefficiency, which we refer to as waste. What we have in mind is that in any organization there is potentially some duplication of effort or unnecessary tasks being performed which affect labor productivity in an inframarginal way. In particular, in the context of the technologies described in the previous section we assume that an organization with waste ϕ has labor requirement $h(y) + \phi$ rather than $h(y)$. The key feature of this waste is that it affects average labor productivity but not marginal labor productivity. One could obviously consider inefficiency that also serves to alter the slope of $h(y)$. Inefficiencies of this form are certainly plausible. We assume inefficiency only of the form as characterized by the parameter ϕ because it is this type of inefficiency that will be central to our analysis. Alternative types of inefficiencies can be added to the analysis. We note as a matter of interpretation that we interpret this inefficiency as reflecting inefficiency in the organizational design and not inefficiency due to a worker that is shirking, for example.

For our purposes there are two key issues associated with these inefficiencies. The first is where they come from, and the second is how organizations can eliminate them. Again, our formulation will be somewhat specialized in order to isolate a particular effect. In general one could imagine that inefficiencies stochastically occur within any organization, and that it takes organizational resources to get rid of them. One could also assume that organizations devote resources to these activities *ex ante* in order to reduce the likelihood that they arise. Our formulation relies on the notion that changes in organizational scale are likely to be associated with the appearance of inefficiencies since the organization is less likely to know how to best use resources as it moves to a new organizational structure.³ Motivated by this idea, we assume that all organizations operating the small scale technology do so efficiently, but that whenever an organization switches from the small scale technology to the large scale technology it will necessarily move to a positive level of inefficiency that we denote by the parameter $\bar{\phi}$. That is, if an organization used technology s in period $t-1$ and switches to technology l in period t , then their labor requirement function will be $a_l y + \bar{\phi}$ for $y \geq \bar{y}$. We note that it would be straightforward to also assume that a new organization that is operating the low scale technology for the first time also begins with an inefficiency, but we abstract from this possibility for simplicity. We assume that the level of inefficiency is known to the organization. While one could consider interesting issues that arise from organizations not having complete information about the state of their efficiency and needing to learn over time about them, we abstract from them here.

Having described how inefficiencies arise, we now turn to the issue of how

³Bertschek and Kaiser (2001) present evidence from a data set of German firms that changes in productivity are linked with changes in organizational structure.

an organization can get rid of them. We adopt a simple and straightforward formulation. In particular, in any given period an organization makes a discrete decision about whether to try to eliminate inefficiency. Having done so, with probability π^e the organization will decrease its inefficiency to zero in the following period, while with probability $1 - \pi^e$ the organization will experience no change in its level of inefficiency. Assuming that there is no improvement in efficiency the organization can continue to choose to try to get rid of the inefficiency in each subsequent period.

Our formulation assumes that a given organization makes stochastic transitions between two levels of efficiency. If we had a large number of organizations all with inefficiency $\bar{\phi}$ and they all continued to choose to try and eliminate this inefficiency until successful, then the average level of inefficiency among this group of organizations would decrease monotonically and approach zero asymptotically. As we will see later in the analysis, this is the pattern that we want to generate. Of course, this pattern could also be generated by having each individual organization experience a monotonically decreasing level of inefficiency rather than the all or none form that we specify. We have chosen the all-or-nothing form of improvements to simplify the analysis of the model.

In the next subsection we describe in more detail the cost to the organization of trying to reduce its level of inefficiency. Note that we assume that there is no direct cost associated with changing from one technology to another. It would be straightforward to add such a cost but it is not central to the effects that we stress below. Lastly, our model is related to models of costly adjustment and models of organizational capital, so it is of interest to remark on these relationships. At a general level, the inefficiency associated with change of scale can obviously be interpreted as a form of adjustment cost. We note however, that our specification differs from most specifications of adjustment costs because the cost does not necessarily disappear in the periods following the adjustment. That is, most models of adjustment costs assume a one time cost associated with the adjustment, but here the cost is permanent unless the organization takes some actions. The restructuring that takes place within organizations in our model can also be interpreted as a form of investment in organizational capital. However, our model differs from many formulations of organizational capital in that we implicitly assume that this investment in organization capital is a substitute for labor input since it leads to a reduction in labor, while most analyses assume that increases in organizational capital lead to increases in the marginal product of labor.

2.2.3 The Role of the Manager

We assume that each organization has one manager and potentially many workers. The labor requirement functions described in the previous subsection should be thought of as specifying the required amount of non-managerial labor input. In this regard our model is similar to the standard span of control model of Lucas (1979). However, we deviate from that model by allowing a manager to choose between two primary uses of their time. In particular, we assume that

a manager can either devote their time to facilitating production, or to trying to reduce inefficiency. We shall cast this choice as having the manager choose between “producing” or “reorganizing”, which we will denote as $m = p$ and $m = r$, respectively. The cost of having the manager devote their time to reorganizing is that they are not able to focus on production. We model this cost as a decrease in the efficiency of labor used in the organization. In particular, we assume that if a manager of an organization using a large scale technology devotes their time to reorganization, then the labor requirement function becomes $(1 + \bar{\eta})h^l(y)$ where $\bar{\eta} > 0$ is the efficiency loss associated with having the manager not focus their attention on production. The benefit of having the manager focus on reorganizing is that it makes it possible for the organization to be more efficient in the future. As this description makes clear, a key tradeoff that a manager faces when making their time allocation decision is between current efficiency and future expected efficiency. As we will see in the next section, it is this tradeoff and how it interacts with business cycle shocks that is at the heart of our analysis.

We assume a competitive market for (homogeneous) managers. The organization will therefore also take the managerial wage as given, which we also assume to be constant over time. Because an organization cannot function without a manager, managerial compensation is effectively a fixed per period cost for the organization. We denote this wage by w_m . The only way that an organization can avoid having to hire a manager is to cease to exist. We assume that if an organization chooses this option that it cannot return in the future.

2.3 The Organizational Life Cycle

It is straightforward to formulate the optimization problem of the organization just depicted. We do it recursively. The state vector for the organization that remains alive is denoted by $s = (\varepsilon, \phi)$ where ε is the state of demand for its product and ϕ is its level of inefficiency if it chooses to operate the large scale technology. An organization is always born into the state $(\varepsilon^s, \bar{\phi})$, which is to say that a new organization begins with demand in the low state and an inefficiency level of $\bar{\phi}$. Note that the level of inefficiency begins at $\bar{\phi}$ because if a new organization were to use to large scale technology it would be faced with the inefficiency. But, as noted earlier, as long as it chooses to operate the low scale technology it can do so without experiencing any inefficiency. In each period, after observing its current state variable, if the organization remains alive it faces three choices: which technology to use ($i = s$ or l), how much output to produce (y), and how to allocate the manager’s time ($m = p$ or r). Given the organization’s state vector and choices for each of these decisions, we can determine the current revenues net of payments to nonmanagerial labor that would accrue to the organization, which we will denote by $R(\varepsilon, \phi, i, y, m)$. This function takes the following form:

$$R(\varepsilon, \phi, i, y, m) = P(y, \varepsilon)y - (1 + \bar{\eta}I_{m=r})h^i(y) - I_{i=l}\phi - w_m \quad (1)$$

where $I_{m=r}$ is the indicator function for $m = r$ (i.e., the manager reorganizes), and $I_{i=l}$ is the indicator function for using the large scale technology.

It is now easy to write the Bellman's equation for the maximization problem faced by the organization:

$$V(\varepsilon, \phi) = \max\{0, \max_{i,y,m}[R(\varepsilon, \phi, i, y, m) + \beta(1 - \lambda^i)EV(\varepsilon', \phi')]\} \quad (2)$$

where the outer max reflects the decision of whether or not to remain active, and the inner max reflects the optimal choices assuming that the organization remains active. We have assumed that if the organization ceases to exist, either through choice in the current period, or because of a death shock in the next period that all future returns will be zero. The expectation operator E incorporates two elements. First, it incorporates the dynamics in the demand state ε , and second, it incorporates the dynamics in the level of inefficiency if the organization chooses to have the manager devote their time to reorganization.

Given our assumptions thus far, we cannot rule out some rather extreme or degenerate outcomes that are of little interest. We describe some of these now. In what follows we do not offer any specific conditions on the model specification to rule out these outcomes, but do note intuitively what parameters would be relevant in ruling out certain outcomes.

It is possible that a newly created organization cannot earn positive expected lifetime profits and hence will choose to shut down. In particular, as noted previously the managerial wage acts like a fixed cost of being in operation, and it is well known that in a model with a fixed cost, it is not enough to guarantee positive net revenues from the variable factors, as our earlier assumption on P does. Of course, if a newly created organization is choosing to shut down and there is some cost associated with creating an organization in the first place, then this would imply that new organizations are never created. In an equilibrium context in which consumption is infinitely valued at the margin when consumption is zero, such an outcome could not be an equilibrium outcome. In view of this, it is natural to assume that wages are sufficiently low relative to the price of output that this is not the outcome. Given our assumptions on the price function P , it follows that if the expected present discounted value of profits are positive for a newly created organization, then they are positive for any feasible state vector. This does not necessarily imply that the organization will have positive current period profits in all states—it is possible that a new organization only remains active because of the possibility of transiting to the higher demand state and that the higher demand state is the only state that is profitable in a static sense. However, given our assumptions on P it is true that conditional on remaining active, an organization will always choose to produce a positive amount of output, even if current period profits were to be negative.

The model has been constructed to focus on the change in scale of production and the associated change in organization structure that occurs as organizations successfully mature. However, given a demand function P , if the value of \bar{y} is sufficiently large then no organization will ever choose to operate the large scale technology, and if \bar{y} is too small then all organizations will choose to operate

the large scale technology. In the context of our model, neither of these cases is particularly interesting. So, in what follows we assume that it is optimal for an organization in demand state ε^s to operate the low scale technology and for an organization in the demand state ε^l to operate the large scale technology. We note, as a feature of our specification, that it is not possible to eliminate future inefficiency while currently operating the small scale technology. It follows that even if an organization decides to reorganize and thereby experiences the current loss of efficiency associated with η , that the organization will still choose to operate the large scale technology.

Lastly, it is also possible that the values of ϕ and $\bar{\eta}$ are such that no organization would ever choose to reorganize. This could happen if the level of inefficiency (i.e., $\bar{\phi}$) is sufficiently small relative to the foregone productivity (i.e., $\bar{\eta}$ is large) or the probability of failure in reorganization (i.e., $1 - \pi^e$). Conversely, if the size of the inefficiency is sufficiently large relative to the cost of eliminating it, then an organization would always choose to reorganize. Since the case of no reorganization is not very interesting in the context of our model, we assume in what follows that we are in a region of parameter space in which organizations do sometimes choose to reorganize.

Conditional on assuming that a newly created organization chooses to remain in existence, that organizations in the low (high) demand state operate the low (high) scale technology, and that organizations sometimes choose to eliminate inefficiency, it is fairly easy to characterize the life cycle dynamics that emerge. In particular, any newly created organization will operate the low scale technology and hire an amount of labor that we denote by h^s , and produces output denoted by y^s . Over time there are three things that may happen to this organization. It may receive a shock and cease to exist, it may remain in the same position and hence continue to hire h^s workers, or it may experience a shock that increases its demand to the high state.

If it is sometimes optimal to try and eliminate inefficiency, then because the organization's problem is recursive, it must be optimal to do it the first time the organization reaches the high demand state. While in the high state and reorganizing, the organization is employing the large scale technology and hires labor that we denote by h^r , producing an amount of output that we denote by y^r . Even though the manager is devoting time to reorganizing and this lowers the marginal product of labor, it still must be the case that $h^r > h^s$. To see this note the following. First, this organization must be producing at least \bar{y} units of output and this must exceed the amount of output produced in the small demand state, or else it would not have been optimal to use the low scale technology in the low demand state. Now, if it was possible to produce more output with less labor, then the organization could have chosen this combination in the previous state and chosen not to sell all of the output produced. It follows that h^r must exceed h^s . It follows that if an organization experiences an improvement in its demand state, it increases both its labor input and its output. Note that we cannot say anything about what happens to average labor productivity. This will depend on the magnitudes of the parameters $\bar{\eta}$ and $\bar{\phi}$.

An organization in the large demand state that is reorganizing can in turn

experience three different transitions. First, it may receive a bad shock and cease to exist. Second, it may be unsuccessful in eliminating inefficiency and remain in the same state, in which case it chooses the same actions again. Third, it may be successful in eliminating the inefficiency. We denote the levels of y and h that result in this case as h^l and y^l . How do the values of h^l , y^l , and y^l/h^l compare with the corresponding values from earlier in the life cycle? The first order condition for current period choice of output combined with our assumptions on P imply that output will definitely increase when the manager focuses on production rather than reorganization. This is because the term $(1 + \bar{\eta})$ goes away from the first order condition. However, it is ambiguous whether this leads to an increase or decrease in h , for two reasons. First, even if the only effect were the improved efficiency associated with the managerial time allocation, the effect on labor, as opposed to output, will depend upon the elasticity of the demand function. Second, the elimination of the inefficiency measured by $\bar{\phi}$ necessarily implies a decrease in labor in the amount of $\bar{\phi}$. However, although the effect on h is ambiguous, it is easy to see that independently of what happens to h , average labor productivity will necessarily increase.

Although with our implicit assumptions on parameter values that the organization will never choose to postpone reorganization once it reaches the large demand state, we can still ask what levels of h and y would be optimal if it chose to do so. Denote these levels by h^p and y^p . We ask how h^p and h^l compare. In making this comparison we are assuming that the current period marginal efficiencies are the same since in both cases the manager is focusing on production. Given our formulation, it follows that output will be the same in each case, i.e., $y^p = y^l$. However, because of the waste in the former case, we know that $h^p > h^l$. This will be of particular interest later on because it says that an inefficient firm that chooses to postpone reorganization will necessarily shed workers in the future.

As a final remark in this section, we note that our model emphasizes the restructuring that accompanies growth of an organization. It seems equally plausible that restructuring within shrinking organizations would also be of importance. For the implications that we stress we believe that similar results would emerge from this situation as well, so we have chosen to focus on growing organizations purely for simplicity.

3 The Model With Temporary Shocks

The previous model considered the decision making of an organization over its lifecycle. The only shocks in that model were highly persistent, and we interpreted them to be organization specific. In this section we add purely temporary shocks to the model and examine how these temporary shocks influence the life cycle dynamics that we studied earlier. Although our model is purely decision theoretic, we will interpret the iid shocks that we introduce as reflecting (aggregate) business cycle shocks.

Formally, we assume a second shock that influences demand, and for sim-

plicity we assume that this shock is multiplicative, so that we write the inverse demand function facing the organization as $\varepsilon_{2t}P(y_t, \varepsilon_{1t})$ where P is the same function as described earlier, ε_{1t} is the shock that we previously labelled as ε_t , and ε_{2t} is an iid shock that is drawn from a distribution with cdf $F(\varepsilon_2)$. We assume that realizations of ε_2 lie in the interval $[\varepsilon_{\min}, \varepsilon_{\max}]$ and that this interval contains 1 in its interior. Note that if $\varepsilon_{2t} = 1$ for all t then the model is identical to that considered previously. All other aspects of the environment are left unchanged.

Proceeding as before, we define the revenue associated with decisions in a particular state as $R(\varepsilon_1, \varepsilon_2, \phi, i, y, m)$. This function takes the following form:

$$R(\varepsilon_1, \varepsilon_2, \phi, i, y, m) = \varepsilon_2 P(y, \varepsilon_1) y - (1 + \bar{\eta} I_{m=r}) h^i(y) - I_{i=l} \phi - w_m \quad (3)$$

where $I_{m=r}$ is the indicator function for $m = r$ (i.e., the manager reorganizes), and $I_{i=l}$ is the indicator function for using the large scale technology. It is immediate that both R and R_y are increasing in ε_2 .

It is again easy to write the Bellman's equation for the maximization problem faced by the organization:

$$V(\varepsilon_1, \varepsilon_2, \phi) = \max\{0, \max_{i,y,m} [R(\varepsilon_1, \varepsilon_2, \phi, i, y, m) + \beta(1 - \lambda^i) EV(\varepsilon'_1, \varepsilon'_2, \phi')]\} \quad (4)$$

For future reference it is worthwhile to elaborate on the expected value term in more detail. As noted earlier, this expectation takes into account the evolution of the exogenous shocks as well as the evolution of the inefficiency variable in response to the decision about reorganization. If $\varepsilon_1 = \varepsilon^s$ then the only value of ϕ of interest is $\phi = \bar{\phi}$. Assuming this configuration plus an arbitrary value for ε_2 , the term for next period's value in the Bellman equation becomes:

$$\beta(1 - \lambda^i) EV(\varepsilon'_1, \varepsilon'_2, \phi') = \beta(1 - \lambda^s) [\pi^l \int V(\varepsilon^l, \varepsilon, \bar{\phi}) dF(\varepsilon) + (1 - \pi^l) \int V(\varepsilon^s, \varepsilon, \bar{\phi}) dF(\varepsilon)] \quad (5)$$

If the organization has $\varepsilon_1 = \varepsilon^l$, then the only efficiency value of interest is still $\phi = \bar{\phi}$, since otherwise there is no need for a decision about reorganization and the problem becomes static. For an arbitrary value of ε_2 if the organization decides to reorganize, then the future term in the Bellman equation becomes:

$$\beta(1 - \lambda^i) EV(\varepsilon'_1, \varepsilon'_2, \phi') = \beta(1 - \lambda^l) [\pi^e \int V(\varepsilon^l, \varepsilon, 0) dF(\varepsilon) + (1 - \pi^e) \int V(\varepsilon^l, \varepsilon, \bar{\phi}) dF(\varepsilon)] \quad (6)$$

whereas if it chooses not to reorganize then the same term becomes:

$$\beta(1 - \lambda^i) EV(\varepsilon'_1, \varepsilon'_2, \phi') = \beta(1 - \lambda^l) \int V(\varepsilon^l, \varepsilon, \bar{\phi}) dF(\varepsilon) \quad (7)$$

As was true in the previous subsection, depending upon parameter values there are various forms that the optimal decision rules may take. We modify

our previous assumptions marginally so we now assume that when $\varepsilon_1 = \varepsilon^s$ the organization will choose to remain active independently of the value of ε_2 .⁴ We furthermore assume that when an organization experiences an increase in its demand state from ε^s to ε^l that there is at least some interior value of ε_2 for which the organization would choose to reorganize.

Finally, we also place an implicit assumption on the size of the shocks to ε_1 and ε_2 . In particular we assume that life cycle shocks are much larger than business cycle shocks. The significance of this is that we assume that when $\varepsilon_1 = \varepsilon^s$ that the organization does not wish to operate the large scale technology independently of the realization of ε_2 . Similarly, we assume that when $\varepsilon_1 = \varepsilon^l$ that the organization never chooses to operate the small scale technology independently of the realization of ε_2 . This model is identical to the model of the previous section if we assume that $\varepsilon_{\min} = \varepsilon_{\max} = 1$. It follows that if the previous model implies technology i is only operated in demand state i , that this model will also generate this result if the range of ε_2 's is not too large.

We are now able to prove our main result, which is that the decision to reorganize when in state $(\varepsilon^l, \varepsilon_2, \bar{\phi})$ is characterized by a reservation value of ε_2 , with the property that it is optimal to reorganize if $\varepsilon_2 < \varepsilon_2^*$, and not to reorganize if $\varepsilon_2 > \varepsilon_2^*$. The intuition for the result is simple: it basically reflects intertemporal substitution of reorganization. Reorganization imposes a cost today in terms of foregone efficiency of labor, but offers a future gain in reducing waste. If ε_2 is iid, then future gains are the same independently of the current value of ε_2 . But, we will show that the current period cost of reorganizing is increasing in the amount of production desired in the event of not reorganizing, which in turn is increasing in ε_2 .

We now establish these results. Consider an organization in state $(\varepsilon^l, \varepsilon_2, \bar{\phi})$. Let y^p denote the optimal level of production if the organization were to choose not to reorganize, and let y^r denote the optimal level of production were the organization to choose to reorganize. Conditional upon deciding whether to reorganize, note that the resulting decision about the optimal choice of y is static, and can be represented as:

$$W(\varepsilon_2, \eta) = \max_y \{ \varepsilon_2 P(y, \varepsilon^l) y - (1 + \eta) h(y) \} \quad (8)$$

where η takes on the value 0 in the event of $m = p$, and $\eta = \bar{\eta}$ in the event that $m = r$. We denote the optimal choice of y as $y(\varepsilon_2, \eta)$. Note that the first order condition that defines this function is:

$$\varepsilon_2 [y P_1(y, \varepsilon^l) + P(y, \varepsilon^l)] = (1 + \eta) h'(y) \quad (9)$$

Given our assumptions it follows trivially that the optimal value of y is increasing in ε_2 and decreasing in η .

Let $V^p(\varepsilon^l, \varepsilon_2, \bar{\phi})$ and $V^r(\varepsilon^l, \varepsilon_2, \bar{\phi})$ be the resulting optimal values obtained from choosing not to reorganize and to reorganize, respectively, assuming that

⁴More generally it would be straightforward to allow for the possibility that the organization will not remain in operation for all realizations of ε_2 . In this case there will be a reservation value of ε_2 that dictates whether the organization remains.

output is chosen optimally in each case. Now consider the difference between these two values. Using V^p and V^r to denote these functions, direct substitution gives:

$$\begin{aligned} V^p(\varepsilon^l, \varepsilon_2, \bar{\phi}) - V^r(\varepsilon^l, \varepsilon_2, \bar{\phi}) &= R^p - R^r + \beta(1 - \lambda^l) \left\{ \int V(\varepsilon^l, \varepsilon, \bar{\phi}) dF(\varepsilon) \right. \\ &\quad \left. - [\pi^e \int V(\varepsilon^l, \varepsilon, 0) dF(\varepsilon) + (1 - \pi^e) \int V(\varepsilon^l, \varepsilon, \bar{\phi}) dF(\varepsilon)] \right\} \end{aligned} \quad (10)$$

where we have written $R^p = R(\varepsilon^l, \varepsilon_2, \bar{\phi}, l, y^p, p)$ and $R^r = R(\varepsilon^l, \varepsilon_2, \bar{\phi}, l, y^r, r)$. Note that the value of the terms involving integrals are all independent of ε_2 . Denote these terms by the constant A . Moreover, the difference in the two revenues can be reduced to:

$$R^p - R^r = W(\varepsilon_2, 0) - W(\varepsilon_2, \bar{\eta}) \quad (11)$$

It follows that equation (10) can be written as:

$$V^p(\varepsilon^l, \varepsilon_2, \bar{\phi}) - V^r(\varepsilon^l, \varepsilon_2, \bar{\phi}) = [W(\varepsilon_2, 0) - W(\varepsilon_2, \bar{\eta})] + A \quad (12)$$

This equation is intuitive. The term in square brackets is the current period cost of reorganizing: it represents the loss in current revenue associated with having the manager devote their time to reorganizing. The term A is the future benefit to reorganizing. Given our assumption, this is simply a positive number that is independent of ε_2 .

We can now easily show that this difference is increasing in ε_2 . It is sufficient to show that the term in square brackets is increasing in ε_2 . Differentiation give:

$$\frac{d}{d\varepsilon} [W(\varepsilon_2, 0) - W(\varepsilon_2, \bar{\eta})] = W_1(\varepsilon_2, 0) - W_1(\varepsilon_2, \bar{\eta}), \quad (13)$$

so it is sufficient to show that $W_{12} < 0$. By definition,

$$W(\varepsilon_2, \eta) = \varepsilon_2 P(y(\varepsilon_2, \eta)) y(\varepsilon_2, \eta) - (1 + \eta) h(y(\varepsilon_2, \eta)) \quad (14)$$

Using the envelope condition, we have that

$$W_2(\varepsilon_2, \eta) = -h(y(\varepsilon_2, \eta)) < 0. \quad (15)$$

It then follows that

$$W_{12}(\varepsilon_2, \eta) = -h'(y(\varepsilon_2, \eta)) y_1(\varepsilon_2, \eta) \quad (16)$$

Since h is increasing and the solution for y is increasing in ε_2 , it follows that $W_{12} < 0$ and hence $V^p(\varepsilon^l, \varepsilon_2, \bar{\phi}) - V^r(\varepsilon^l, \varepsilon_2, \bar{\phi})$ is increasing in ε_2 . If the benefit to not reorganizing is monotone in ε_2 it follows that the optimal reorganization strategy is to employ a reservation value, as stated above.

If the reservation value is equal to ε_{\max} , i.e., the upper support of the distribution of the temporary shocks, then the organizational dynamics in this model

are qualitatively the same as in the previous subsection. That is, whenever a small organization gets an improvement in their idiosyncratic demand state ε_1 , they immediately choose to reorganize and continue to do so until the reorganization is successful. Fluctuations in ε_2 will lead to additional fluctuations in their labor input and output, but the lifecycle dynamics will be similar.

However, if the reservation value ε_2^* is interior to the interval $[\varepsilon_{\min}, \varepsilon_{\max}]$, then qualitatively different dynamics can emerge. In this scenario, if an organization in the small idiosyncratic demand state receives a shock that raises it to the large idiosyncratic demand state, the organization may or may not decide to reorganize at that point. In particular, if ε_2 is sufficiently high then the organization will choose to postpone the decision to reorganize in order to take advantage of the current temporarily high demand. As stated earlier, the organization will engage in intertemporal substitution of managerial actions. And, following from our discussion of the previous model, we know that an organization that chooses to postpone reorganization will necessarily employ less labor in the future when it does successfully reorganize, even holding the value of ε_2 constant.

3.1 Extension to the Case of Persistent Shocks

The previous analysis has assumed that the shock ε_2 is *iid* over time. Of course, if one wants to think of the ε_2 shock as proxying for business cycle movements in the demand faced by an individual organization, then the *iid* assumption is not very appealing. A well documented property of business cycles is that they are persistent, in the sense that if the economy is above trend today then we also expect it to be above trend next period. In view of this it is of interest to ask whether our result about the reservation value will extend to the case of persistent shocks. In fact, the argument is easily extended.

In the *iid* case, we argued that the current cost of reorganizing is increasing in the current value of ε_2 , and that the expected future benefit of reorganizing is independent of the current period value of ε_2 . The first statement is independent of whether the realizations of ε_2 are iid or not. However, a key observation about the structure of our model is that the benefit to successful reorganization is in fact independent of future realizations of ε_2 . The reason for this is that in our model successful reorganization does not impact on the marginal product of labor, and as a result an efficient organization and an inefficient organization will choose the same level of output conditional upon having the same managerial time allocation. The only effect on profit is from saving labor in the amount of $\bar{\phi}$ for each future period that the organization remains in existence, and this saving is independent of all future realizations of ε_2 .

The additional issue that needs to be addressed in the context of a model with persistent shocks to ε_2 is the following: An organization faced with a current realization of ε_2 can also consider the possibility of waiting a period to reorganize. If the benefit from waiting increases as ε_2 decreases, then the reservation property might not be preserved. In the *iid* case the benefit from waiting is actually increasing in ε_2 , since next period's expected one-period cost

of reorganizing is independent of ε_2 , so that a higher current value of ε_2 indicates lower expected costs in the future, whereas a low value of ε_2 indicates higher expected costs in the future.

With this in mind we can present the alternative characterization of the decision to reorganize. In particular, let $C(\varepsilon) = W(\varepsilon, 0) - W(\varepsilon, \bar{\eta})$ be the gain this period from not reorganizing. Let $G = \beta\pi^e\bar{\phi}/(1 - \beta(1 - \lambda^l))$ be the (expected) gain from choosing to reorganize today. Letting $H(\varepsilon)$ be the benefit of today's managerial choice relative to reorganizing today, we have that $H(\varepsilon)$ satisfies:

$$H(\varepsilon) = \max\{C(\varepsilon) + \beta \int H(\varepsilon')F(d\varepsilon', \varepsilon), 0\} \quad (17)$$

where the first term indicates the gain from not reorganizing today, and the second term indicates that if the manager chooses to reorganize then the gain is clearly zero. We know that $C(\varepsilon)$ is increasing in ε from our previous analysis. If we knew that the term $C(\varepsilon) + \beta \int H(\varepsilon')F(d\varepsilon', \varepsilon)$ is increasing in ε , the current realization of the shock, then we could easily conclude that the reservation property holds. Note as stated earlier that if ε is iid then the integral is independent of ε and the property holds based on the property of C . However, finding conditions under which the integral is increasing in ε is a standard problem in dynamic stochastic models. In particular, if we assume that

$$\int g(\varepsilon')F(d\varepsilon', \varepsilon) \quad (18)$$

is increasing in ε for any increasing function then we can easily show that the value function H will in fact be increasing and hence that the integral has the desired property.

Loosely speaking, a process for ε_2 that implied mean reversion would tend to satisfy this property. This is because a high value of ε_2 today implies that future values of ε_2 will be lower, implying that it is beneficial to wait to reorganize. On the other hand, a very low value of ε_2 today implies that future values of ε_2 will be higher, implying that there is greater incentive to reorganize today rather than in the future.

We conclude that the reservation value property for the optimal reorganization decision will also hold in the case of persistent shocks under reasonable conditions.

4 Implications for Business Cycle Dynamics

Our formal analysis has only considered the decision problem of an individual organization that takes demand for its product as given, assuming that wages for workers and managers and the real interest rate are constant over time. Such a model can be cast in an industry equilibrium setting such as Hopenhayn and Rogerson (1993) in which case the organizational dynamics that we describe can capture the steady state dynamics of general equilibrium model in which

all shocks are idiosyncratic, i.e., there are no aggregate shocks. If one introduces aggregate shocks into such a model, then one would need to take into account the effect that these shocks have on wage rates and the real interest rate. Veracierto (2002) and Thomas (2002) are examples of models in which these general equilibrium effects are considered. Hence, in its current form the model is really not appropriate to discuss how the economy responds to aggregate shocks. Nonetheless, in this section we want to discuss what our previous analysis suggests as some potential effects of interest for business cycle dynamics. We leave development of the appropriate framework and the associated formal analysis for future work.

Consider the following situation. We have a unit mass of entry of new organizations each period, each of which enters into the individual state $(\varepsilon^s, \bar{\phi})$. We consider the ε_2 shock to be common to all organizations while the ε_1 shock is idiosyncratic, and trace out the evolution of the economy assuming that wage rates and the interest rate remain constant. The first observation that we want to stress is that the aggregate state of this economy will be the realization of the aggregate shock ε_2 and the distribution of organizations across individual states. Entering a given period there are three types of organizations: those that are in the state $(\varepsilon^s, \bar{\phi})$, (which we call state 1), those that are in the state $(\varepsilon^l, \bar{\phi})$ (which we call state 2), and those that are in state $(\varepsilon^l, 0)$ (which we call state 3). We let μ_i be the mass of firms in each of the three states. The evolution of the μ_i 's is affected by the realization of the aggregate shock since it determines whether organizations in state 2 will be reorganizing, thereby influencing the probability that an organization transits to state 3. In particular, let μ_{it} be the distribution of active organizations in period t . Then, if $\varepsilon_{2t} > \varepsilon_2^*$, we have the following:

$$\begin{aligned}\mu_{1t+1} &= (1 - \lambda^s - \pi^l)\mu_{1t} + 1 \\ \mu_{2t+1} &= (1 - \lambda^l)\mu_{2t} + \pi^l\mu_{1t} \\ \mu_{3t+1} &= (1 - \lambda^l)\mu_{3t}\end{aligned}\tag{19}$$

If, on the other hand, we have $\varepsilon_{2t} < \varepsilon_2^*$ then the distribution will evolve according to:

$$\begin{aligned}\mu_{1t+1} &= (1 - \lambda^s - \pi^l)\mu_{1t} + 1 \\ \mu_{2t+1} &= (1 - \lambda^l - \pi^e)\mu_{2t} + \pi^l\mu_{1t} \\ \mu_{3t+1} &= (1 - \lambda^l)\mu_{3t} + \pi^e\mu_{2t}\end{aligned}\tag{20}$$

Two simple conclusions can be drawn from these laws of motion. First, note that the law of motion for μ_{1t} is independent of the realization of ε_2 . It follows that if entry is constant as we have assumed, that the value of μ_1 will approach a constant and will not be affected by realizations of ε_2 . The constant fraction of organizations in state 1 is easily computed to be $\lambda^l/(\lambda^l + \pi^l)$. The second point to note is that the remaining mass of organizations will be split between type 2 and type 3 organizations, and that this division will depend upon the history

of the ε_2 realizations. Specifically, if ε_2 remains above ε_2^* then there is a greater buildup of organizations in the second state at the expense of organizations in the third state. It follows that the longer the aggregate shock remains above the reservation value, the greater will be the buildup of organizations in the second state. In what follows, we illustrate the potential effects that this can have on how the economy responds to subsequent shocks.

4.1 A Reduced Form Example

For present purposes, the most effective way to illustrate the interaction between the distribution of organizations and the response of the economy to a given sequence of aggregate shocks is with a very specific reduced form example. Consider an economy at time 0 with unit mass of organizations that are distributed across types. We assume that the stochastic process for ε_2 has a sequence of realizations of the following form. In period 0 the economy is hit by a (very) negative value of ε_2 but subsequent to this experiences a constant and gradual increase back towards its unconditional mean value. We assume that it takes the economy 20 periods to reach this value, at which time it remains there. This type of realization could be thought of as tracing out the impulse response function in the presence of a mean-reverting process.

We assume the following reduced form properties in our example. First, we assume that reorganization is optimal for this entire range of realized ε_2 values. Second, we assume that when an organization reorganizes successfully, the effect on labor input is a decrease of e^f , holding ε_2 constant and is independent of the level of ε_2 . Third, we assume that each period in which ε_2 is increasing, the aggregate effect of this on labor input is an increase in labor input of e^h which we assume is distributed across organizations according to size.⁵ We assume that in period 0 a small organization employs 10 workers and that a large reorganized organization employs 100 workers. By assumption, a large organization in the process of reorganizing will employ $100 + e^f$ workers. The probabilities π^l , π^e , λ^s and λ^l are as before.

Our goal here is to illustrate the potential for the initial distribution of organizations to influence the resulting response of employment to a given sequence of shocks. With this in mind, we simulate the implied path of aggregate employment for several different initial conditions. As already noted above, with a constant rate of entry of new organizations, asymptotically there will be a constant mass of organizations and a constant fraction of them will be of type 1 (i.e., in the small demand state with inefficiency $\bar{\phi}$). In view of this we always consider the initial fraction of organizations that are of type 1 to correspond to this fraction. We normalize the total mass to equal one and hence assume that $\mu_{10} = \lambda^l / (\lambda^l + \pi^l)$. As noted earlier, however, at any given point the distribution of the remaining organizations between type 2 and type 3 will be influenced by the previous history of realizations of ε_2 , and we therefore consider different

⁵One could interpret this reduced form as reflecting a log linearization of the individual demand for labor functions.

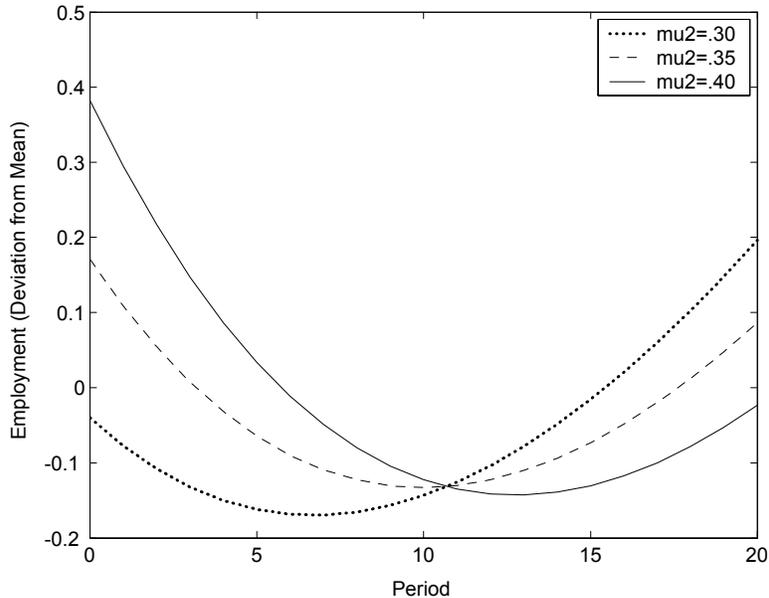


Figure 1: Effect of Initial Conditions on Employment

scenarios for how this remaining mass is allocated across type 2 and type 3 organizations.

For the paths shown in Figure 1 we have set $e^h = .1$, and $e^f = 10$. Setting $e^h = .1$ amounts to assuming that the accumulated increases in ε_2 over the 20 periods would increase aggregate employment by roughly 4% over the course of twenty periods, holding all else constant. We set $e^f = 10$, implying that successful reorganization leads to a reduction in employment of roughly 10%. We set $\lambda^l = \pi^l = .0025$, implying that over the course of the twenty periods the accumulated probability of failure for large organization or success for a small organization is roughly 5 percent. Finally, we set $\pi^e = .05$ which implies an expected duration of 20 periods for successful reorganization. While we have offered these quantitative guides to thinking about the parameter selections, we also emphasize that this example is intended purely as illustrative. We leave a rigorous quantitative assessment of the economic mechanisms described here for future work.

Figure 1 shows the employment paths that result for three different initial values of μ_2 .

In each case, the curve represents deviations from means for each path in order to focus attention on the implications for timing. The figure shows that as μ_{20} increases it takes longer to reach the turning point of employment. This result is intuitive. The greater the value of μ_{20} , the greater is the total amount of reorganization that needs to be done. As this reorganization takes place, suc-

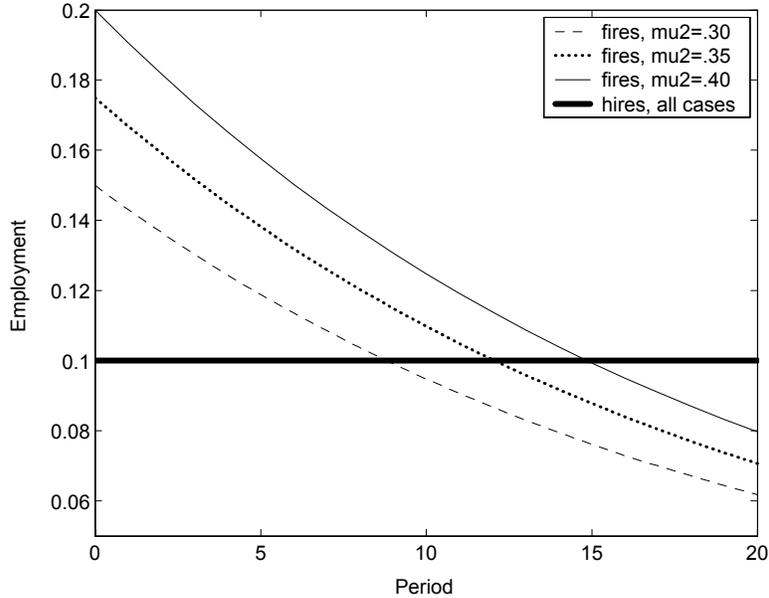


Figure 2: Initial Conditions and Hiring and Firing

cessful organizations will be shedding labor. This labor shedding is an opposing force to the increases in labor associated with the gradual improvement in the aggregate shock ε_2 . A key point to note is that in our model reorganization is potentially a long-lasting process in the sense that in any given period only a given fraction π^e of the remaining reorganization will be carried out. In fact, it is interesting to note the implications of the extreme case in which $\pi^e = 1$. In this case all of the “accumulated” reorganization will be carried out in the first period and there will be a large drop in employment, but after this we will see a continual increase in aggregate employment as ε_2 increases. Hence, a large amount of accumulated reorganization will simply lead to a very large one time drop in employment, but will not lead to a delayed turning point for aggregate employment.

To understand the dynamics of the opposing forces, the next figure shows the time paths of firing and hiring for each of the three scenarios considered in Figure 1.⁶

In all three cases the hiring associated with the improvement in ε_2 is constant over time and equal to .1. However, although each economy has the same fundamentals, the fact that the initial distribution of μ varies implies that the time path of fires associated with successful reorganization will be different. As

⁶An organization that successfully reorganizes may of course choose to fire fewer than 10 workers and not hire any new workers, so measured hiring and firing in the economy may not reproduce these curves.

the graph shows, the curves are effectively parallel shifts of each other, with the highest value of μ_{20} associated with the highest level of firing. Because successful reorganization takes time, the path of firing is fairly drawn out. The important feature to note is that aggregate employment will continue to drop as long as the firing curve lies above the hiring curve. And since a higher value of μ_{20} raises the firing curve but leaves the hiring curve unchanged it illustrates how restructuring can influence the point at which aggregate employment begins to increase.

We should emphasize that the dynamics that we have just traced out are obviously not definitive predictions of the model. As noted earlier, whether reorganization leads to labor shedding depends on parameter values. The main point that we want to emphasize is that the model suggests a mechanism which can produce these types of dynamics.

4.2 General Equilibrium Considerations

Although our analysis has only considered the decision problems of individual organizations and then aggregated holding wage and interest rates constant, it is worthwhile to discuss how general equilibrium considerations would possibly affect the types of outcomes that we were emphasizing. In particular, it is important to emphasize this issue in the context of the literature related to the work of Caballero and Engel (1999) referred to in the introduction. In that paper, they argued that changes in the distribution of individual firm state variables was important in influencing how the economy responds to shocks. However, the work of Veracierto (2002) and Thomas (2002) showed that in general equilibrium versions of the model the effects of interest rate movements basically offset the partial equilibrium effects.

There are two main issues that arise. The first concerns how changes in prices might impact the incentives for intertemporal substitution in our decision problem, as represented in our key analytic result showing the existence of a reservation value of ε_2 for the reorganization decision. With constant wages and interest rates, reorganization will be shifted away from times of high economic activity and toward periods of low economic activity. If real wages were procyclical, they could generate an opposing force to our intertemporal substitution. If wages are higher in periods of high economic activity and shocks are persistent then this produces a cost of not reorganizing that is procyclical. Simply put, the benefit of shedding labor is greater if wages are higher.

General equilibrium effects could also operate through changes in the real interest rate. However, if one views the case of procyclical real interest rates as the case of primary interest, then this effect will actually reinforce our result. If current real interest rates are high then current period costs are amplified and future benefits are attenuated, increasing the incentive to postpone reorganization in good times.

One could plausibly argue that some other margins that we have assumed to be constant over time would also exhibit variability over the cycle. For example, although we assumed that transition probabilities are constant over

time, one could argue that the probability of a small scale organization becoming successful increases in good times, i.e., that π^l is higher in good times. This by itself would tend to accentuate the effects that we have emphasized, since this will lead to a larger build up of type 2 organizations during good times. Similarly, if entry is higher in expansions then this will also tend to accentuate the build up of type 2 organizations.

A second issue that must be addressed in a more complete model is to understand why the workers that are being released due to restructuring do not find employment somewhere. One possible channel is standard intertemporal substitution effects. When an organization restructures, the shift of managerial time away from production and toward restructuring leads to a decrease in the marginal product of labor. A second channel that is not in our model but which may be important, is that it could be that the organizations with the greatest expected increases in employment in the future are those which have recently experienced large increases. If this were true, it could be that the organizations who decide to restructure are the same organizations that will eventually add the most workers. One could imagine a more detailed model in which an organization does not add to its existing labor force at the same time that it is trying to reorganize. Hence, the decision to reorganize is implicitly a decision to postpone new hires. Lastly, if one were to imbed our model into a model in which it takes time for workers to move from one organization to another, then a long lasting increase in separations would also lead to a long lasting decline in employment.

5 A Closer Look at Jobless Recoveries

In this section, we argue that the insights derived from the preceding discussion may be relevant for understanding some features of business cycle dynamics, and that in particular they may be very relevant for the discussion of the phenomenon that has become known as the jobless recovery. As noted in the introduction, many individuals have coined the term “jobless recovery” to describe the apparent slow growth in employment following the troughs of the two most recent recessions in 1991 and 2001. Viewed in a broader perspective, the obvious implication of such a description is that not all business cycles are alike. This is an old and recurring theme in the business cycle literature. Burns and Mitchell (1946) were among the first to systematically measure the business cycle and argued that business cycles bear a remarkable similarity to each other along many dimensions. In particular, they developed the notion of a reference cycle to represent the “typical” business cycle. Influenced by this work, Lucas (1977) argued that a key stylized fact is that all business cycles are the same from the perspective of qualitative comovement of series. At the same time, there are many instances in which researchers have argued that some particular business cycle exhibits properties that distinguish it from its predecessors, while others

have argued that the business cycle phenomenon is slowly changing over time.⁷

The discussion of the previous section suggested the possibility that following the end of a long expansion it may take longer for employment to start to increase once again. This argument is consistent with the fact that each of the last two recessions has exhibited a relatively long period before employment began to increase, since each of the last two expansions has been extremely long by historical standards. However, there is another episode in the post war period which would seem to be relevant, and that is the recession of 1969-1970 which also followed a very long expansion. If the channel that we point to is quantitatively important, then this period should also have produced a “jobless recovery”. The goal of this section is to argue that the evidence is indeed consistent with this prediction. In particular, we will argue that there are three recessions in the post war period that stand out as distinct from the others in terms of the dynamics for employment in the subsequent recovery: the 1969 recession, the 1991 recession, and the 2001 recession. The material presented here draws on the results presented in Koenders (2004), which provides a much more thorough analysis.

5.1 A Review of Schreft and Singh

It is useful to begin with a summary of the analysis of Schreft and Singh (2003). They carry out the following calculation. They start with seasonally adjusted data for employment from the establishment survey. They then identify the level of employment at each of the NBER turning points that corresponds to the end of recessions. For each recovery, they plot the percentage change in employment from the turning point that occurs over the subsequent twelve months. Figure 3 below is equivalent to the picture that they produce except that we have included the two recessions from the 1950’s in our analysis, and we have time aggregated the employment data to quarterly frequency.

But our Figure 1 tells the same story as Chart 1 in their paper. Whereas the typical recovery shows steadily rising employment, with an increase of more than 3 percent in the first year of the recovery, the two most recent recessions show employment decreasing in each of the four quarters following the turning point. While this picture certainly suggests that the two recent recoveries are different than the average of the preceding ones, it obviously does not tell us if there are previous episodes that also resemble the two recent ones. For our purposes we are particularly interested in whether the recovery that began in 1970 also displays this pattern. Figure 4 below repeats the analysis of Figure 3 except we now consider three recoveries individually and compare them to the average of the remaining five recoveries. (Throughout this analysis we ignore the recovery in the early 1980’s since it was so short-lived.)

This figure suggests that the recovery that began in 1970 is much more similar to the average recovery than it is to the recoveries following the two

⁷A related but distinct issue is the extent to which business cycles have become less frequent.

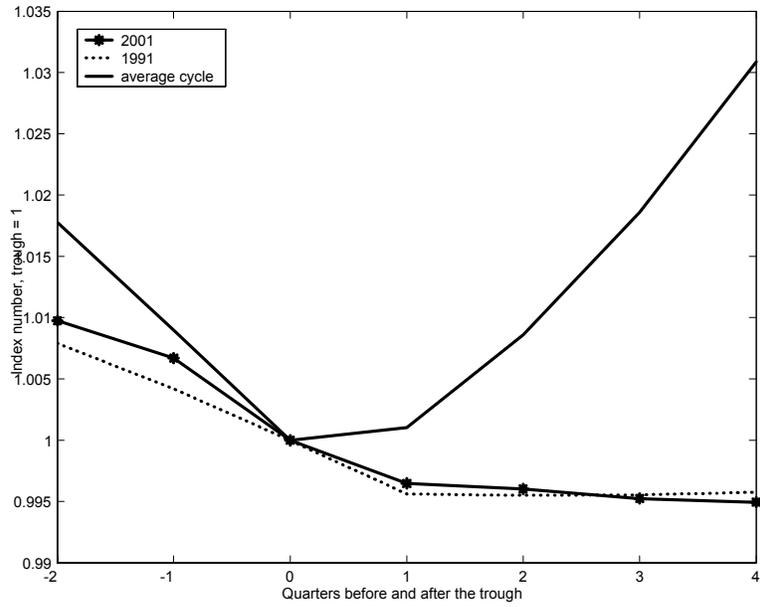


Figure 3: The Schreft-Singh Finding

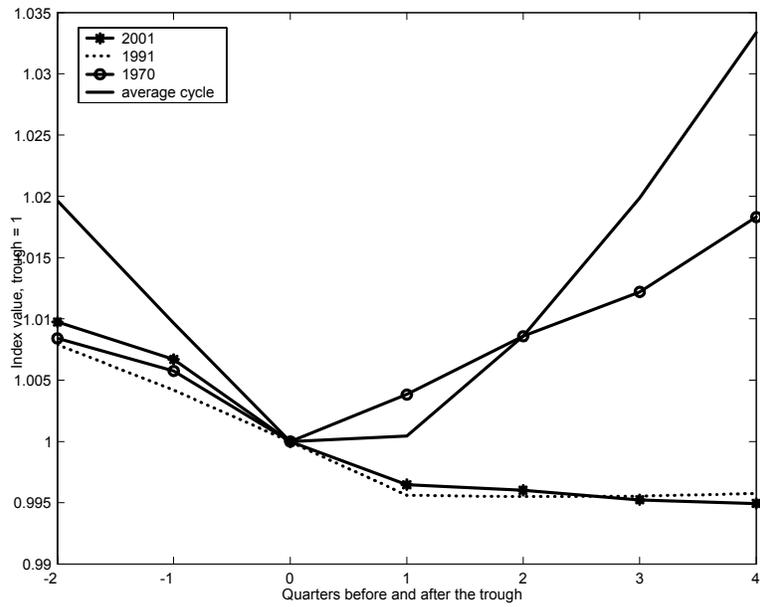


Figure 4: The 1970 Recovery (Schreft-Singh Method)

most recent recessions. One can repeat this analysis for the other recoveries as well, and one obtains a similar pattern in each case. Based on this analysis, one would be lead to conclude that it is only the two most recent recoveries that have had particularly distinctive employment dynamics.

However, there are several issues that we want to raise regarding the Schreft-Singh method of summarizing the data. The first issue is that the Schreft-Singh method is not consistent with modern views of the business cycle. Following Lucas (1977), modern business cycle analysis views the business cycle as deviations from a slowly changing trend. Properties of business cycles should be properties of the component of the time series that corresponds to these deviations from trend. The method of Schreft and Singh neglects this in two important regards. First, some recessions are more severe than others. To the extent that recessions are temporary departures from trend, a deeper recession would naturally be expected to be followed by higher subsequent growth in employment. The Schreft-Singh method does not incorporate this feature. Second, their method does not distinguish between movements in the trend and deviations from the trend. If the (raw) level of employment following the trough of a recession starts to increase, how are we to know to what extent we are moving closer to trend? If the trend always increased at the same rate, this issue would be irrelevant since it would affect all recoveries in the same fashion. However, a key feature of the post WWII labor market in the US is that trend employment growth has fluctuated substantially over time, due both to the entry of the baby boom into the labor market and the increased participation of women. To illustrate this Table 1 shows the decadal growth rates in employment for the US economy for the five decades following WWII.

Table One
Employment Growth by Decade

1950-1960	1960-1970	1970-1980	1980-1990	1990-2000
20.89	26.81	23.91	17.85	17.45

The table shows that the differences are large—the decadal growth rate in employment during the 1960’s is more than one and a half times as large as the decadal growth rates in the two most recent decades. It follows that sorting out relative movements in trend and deviation from trend may be an important consideration in documenting the differential pace of employment growth during recoveries.

Third, the method of Schreft and Singh compares the dynamics of recoveries by examining the behavior from the turning point forward. It is not clear that the turning point is the appropriate comparison point across cycles. In particular, if the downturns preceding the recoveries have been different, it is not clear that behavior should be the same from the turning point forward. In fact, we will argue later that our model suggests that the turning point should not be used as a common reference point.

Having raised some issues about the statistics that Schreft and Singh report, we now move on to describe the method that we use.

5.2 An Alternative Look at the Data

Our method is straightforward and is consistent with current practice in business cycle analysis in terms of documenting properties of cyclical fluctuations. In particular, let X_t be a quarterly series that is seasonally adjusted for which we have observations going from period 0 to period N . Define x_t to be the log of the series X_t . We define the trend component of x_t , denoted by x_t^T , by using the Hodrick-Prescott filter. In particular, x_t^T is the solution to the following optimization problem:

$$\min_{\{x_t^T\}} \sum_{t=1}^{N-1} \{[x_t - x_t^T]^2 + \lambda[(x_{t+1}^T - x_t^T) - (x_t^T - x_{t-1}^T)]^2\}$$

Following the literature, for quarterly data we use a value of $\lambda = 1600$.

The cyclical component of x_t , denoted by x_t^C is simply the deviation of x_t from its trend value: $x_t^C = x_t - x_t^T$. Because the series are measured in logs, the cyclical component reflects the percent deviation of the variable from its trend.

Figure 5 repeats the exercise of Schreft and Singh but uses the cyclical component defined above as opposed to the raw data. In particular, this graph shows the percent change in the cyclical component of employment in each of the four quarters following the NBER turning points.⁸

This figure tells a similar story to the one told by Figure 3, though we note that some details are different. In particular, whereas Figure 3 indicated that in a typical recovery employment begins to grow as soon as the turning point is reached, Figure 5 displays that well-known feature that employment lags GDP. In particular, this Figure shows that in a typical recovery, employment begins to increase one quarter after GDP begins to increase.⁹ (We note that the cyclical component of GDP defined as above has the property that the turning point for GDP following each recession coincides with the NBER turning point dates.)

Next, we again ask if the recovery that began in 1970 is similar to the two most recent recoveries. Figure 6 shows the results of carrying out the exercise analogous to that which generated Figure 4, i.e., we now consider three individual recoveries and compare them to the average of the other 5.

We see that Figure 6 offers a very different conclusion than does Figure 4. Based on analysis of the cyclical component of the employment series, the 1970 recovery also shows that even one year after the turning point employment remains below its level at the turning point. While this picture still indicates quantitative differences across the three recoveries, the qualitative behavior is in fact similar.

⁸It is important to note that whenever one detrends the data the behavior of the cyclical component at the very beginning and end of the sample is somewhat sensitive to the initial and terminal data points. This implies that the properties of the 2001 recovery may look somewhat different as more data becomes available.

⁹See, for example the cyclical properties as report in Cooley and Prescott (1991).

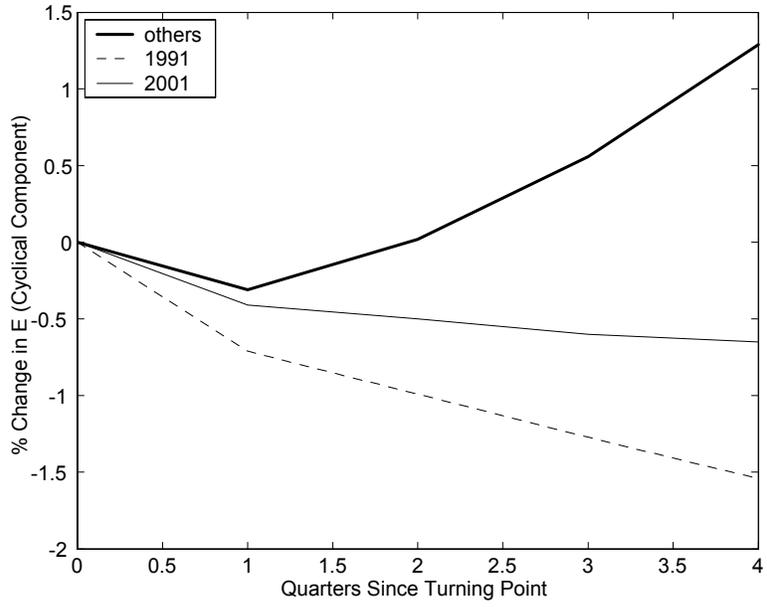


Figure 5: The 1991 and 2001 Recoveries (Detrended)

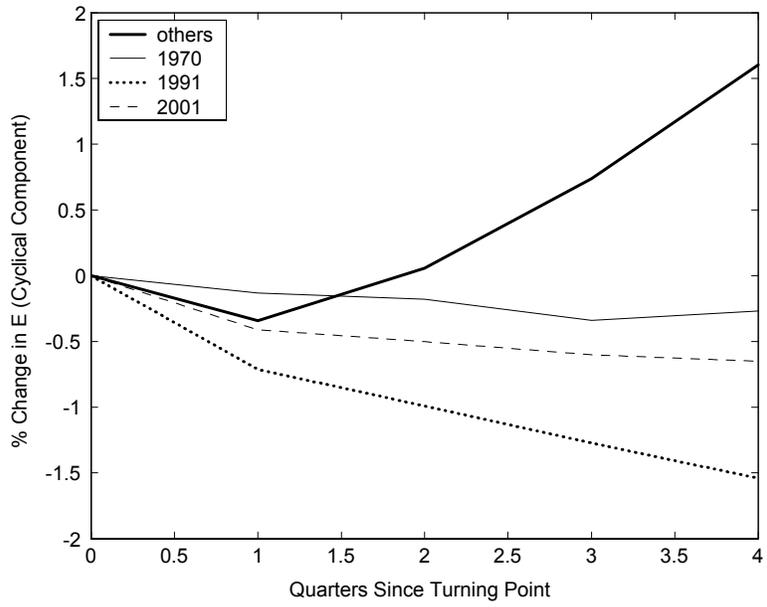


Figure 6: Recoveries of 1970, 1991 and 2001 (Detrended)

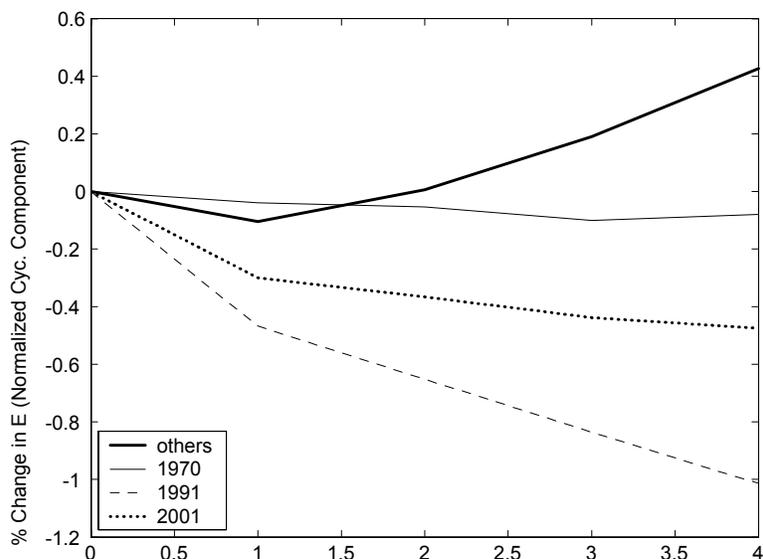


Figure 7: Cyclical Employment, Normalized by Magnitude of Trough

These figures still suffer from the first problem we mentioned earlier—they do not take into account the fact that the recessions vary quite substantially in their severity, and hence we do not know how much of the variation in growth simply reflects differences in distances from trend. There are many ways that we might normalize business cycles to account for the differing magnitudes. We employ a simple procedure here, which is to normalize by the magnitude of the recession as measured by the maximum percent deviation of output below trend. We then scale all of the employment deviations by dividing through by the absolute value of this number. This is shown in Figure 7.

Qualitatively this figure presents the same qualitative conclusions as the earlier figure, but we note that it does impact on the quantitative differences across episodes.

One way to summarize the properties of the above pictures are in terms of the extent to which the turning point for employment lags the turning point for GDP. The following table shows the values for each of the post 1950 recessions.

Table 2
Lag in Employment Turning Point Relative to GDP

1954	1958	1961	1970	1975	1982	1991	2001
1	0	1	3	1	1	6	7

This statistic confirms a property that we saw in the figures presented earlier: while a lag of 0 or 1 quarter is typical for the postwar period, there are three

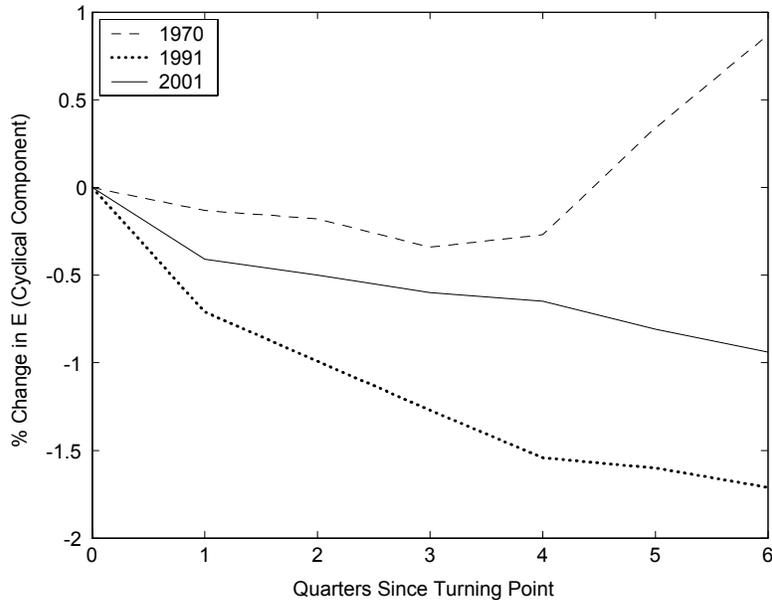


Figure 8: Comparing the 1970, 1991 and 2001 Recoveries

recoveries in which the lag is longer, and all three of these are recoveries from recessions that follow long expansions. This evidence supports our earlier claim about characterizing post war recoveries. At the same time, this table suggests that there is quite a significant difference between the 1970 recovery and the two most recent recoveries. The lags in the two most recent recoveries are in fact much longer than was experienced in the 1970 recovery. In particular, if we extend the analysis from the four quarters following the turning point to 7 quarters, then we obtain Figure 8.

Figure 8 would seem to suggest that the 1970 episode is not so similar to the two most recent ones. What we argue next, however, is that this difference is illusory. In particular, we argue that it is driven by the choice of initial point, and that when this initial point is chosen in a way that we believe is more consistent with the theory laid out earlier in the paper, the differences disappear.

To understand the issue, consider the discussion from the previous section. The key point there was that the time before employment begins to increase is tied to the time required for reorganization to diminish sufficiently. The key point to take away from this is that the turning point is not the point at which reorganization begins. Presumably reorganization starts to take place at some point during the recession preceding the turning point. This is significant because the duration of the recession preceding the turning point differs significantly across recessions. In particular, this period is either average or below average

for the two most recent recessions and is much above average for the 1969-1970 recession.

To implement this element, we need to have a method for picking out at what point reorganization begins. One possibility would be that it begins when the downturn first begins. However, when a recession first begins output is still quite high above trend and the analytic result that we proved earlier in the paper suggested that it is the level of output that is particularly important. With this in mind we identify the point at which reorganization begins to be the first quarter in which the cyclical component of output lies below trend. The following table shows how many periods this lies prior to the NBER turning point for each of the 8 recessions. Table 3 presents the data.

Table 3

Quarters Before Turning Point with GDP Below Trend

1954	1958	1961	1970	1975	1982	1991	2001
2	2	2	4	2	4	1	2

The table indicates that the typical number of quarters that GDP is below trend prior to reaching its turning point is 2, but that in 1970 this value was 4, and in 1991 it was only 1.

What we do next is to repeat our earlier analysis but instead of using the turning point as the initial condition for each recession, we use the period as indicated in the above table. Once again we continue to normalize using the magnitude of the drop in GDP as our scale factor. The next graph shows the average of the five recessions versus each of the three more prolonged recoveries.

It shows that the three recessions all stand out as different than the average of the others. Finally, we show a comparison of the three special recessions but for a much longer period, extending the analysis to 8 quarters. This allows us to include part of the recoveries in all three cases.

To this point we have focused on employment dynamics. Employment is simply one dimension along which labor input can vary. An important issue is the extent to which different behavior of employment across cyclical episodes also represent differences in labor input. Alternatively, it could be that the dominant difference across episodes is the compositional changes in labor input. More generally, one should also be concerned about measuring effective labor input rather than simply bodies or hours. We leave a more careful analysis of these issues for future work, but do want to present at least one piece of evidence to suggest that the differences in the behavior of employment that we have noted also extend to the behavior of aggregate hours. We only have hours data from 1964 on, so this analysis has only 5 cyclical episodes to compare. The next figure compares the three recoveries that we have focused on as compared to the average of the other two recoveries. Period 0 in this figure refers to the first period in which GDP drops below trend prior to the associated recovery. While the basic pattern in this figure is qualitatively similar to the one found in our graphs for employment, it is also true that the differences are much smaller in this graph than in the employment graphs. We infer from this that a more

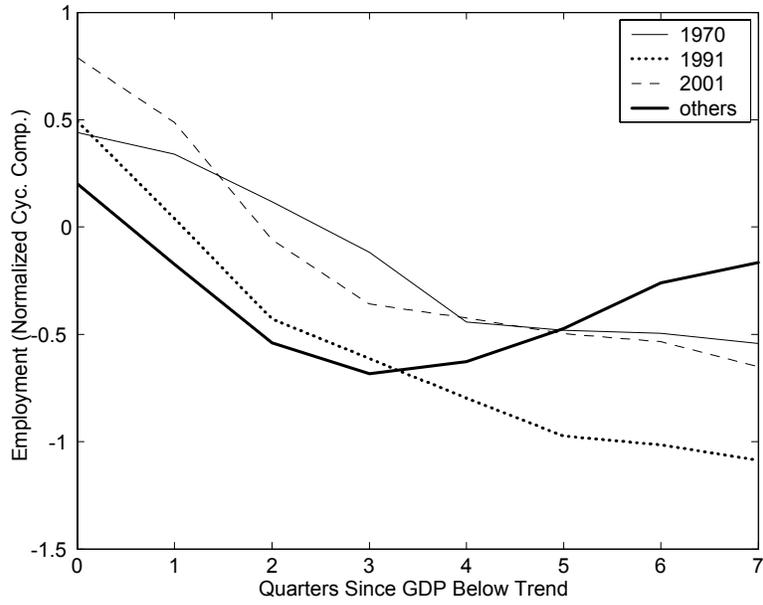


Figure 9: Recoveries with Alternative Initial Point

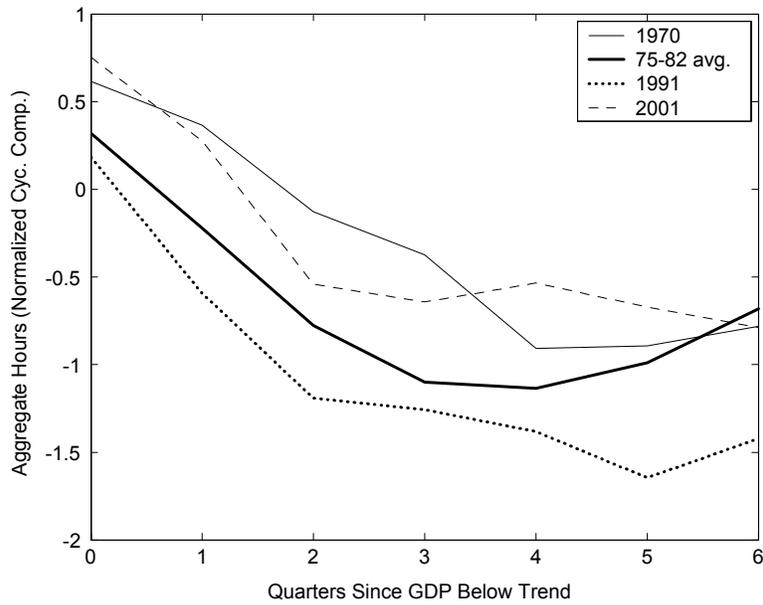


Figure 10: Behavior of Aggregate Hours

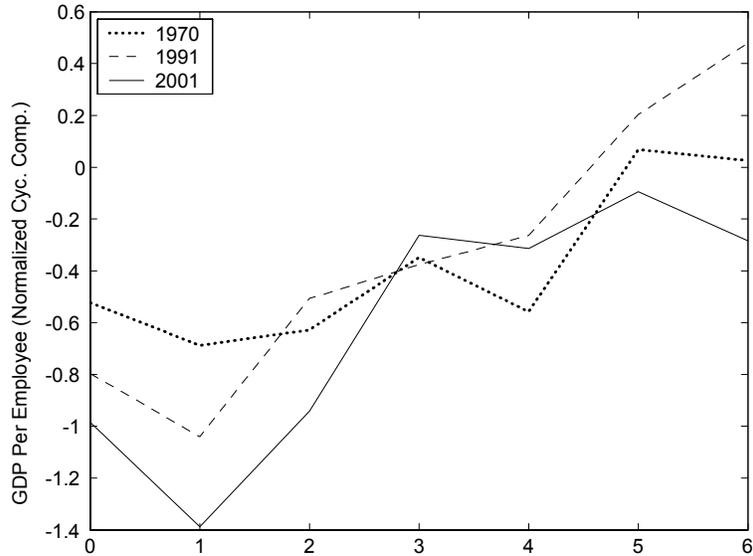


Figure 11: Productivity Per Employee in the 1970, 1991, and 2001 Recoveries

careful study of the behavior of labor input along the intensive and extensive margins is warranted.

Lastly, it is of interest to examine the behavior of productivity across the cyclical episodes. Our model predicts that when organizations switch from producing to reorganizing that they experience a decrease in productivity. However, as time passes and successful reorganization occurs, we should see increases in productivity. Having said this, we also feel that a large degree of caution need be taken with respect to assessing the implications for productivity. We have implicitly assumed that aggregate fluctuations in our model are driven by shocks to the demand for the output of each organization. While this is a convenient formulation for our analysis, it could be that the increase in demand is driven by improvements in product quality which in a more complete model would also show up as productivity changes.

Having offered this caveat we next turn to analyze the dynamics of productivity. Because hours data are available only since 1964, we use two different measures of productivity. First we compare output per worker in order to be able to use all eight recessions, and then we use output per hour to compare the five most recent recessions. Figure 11 shows how productivity per worker evolves in the three recoveries of interest.

Note that in all cases productivity drops in the initial period but then increases thereafter. While the magnitudes are somewhat different across the three episodes, the pattern is quite similar. At a very qualitative level this

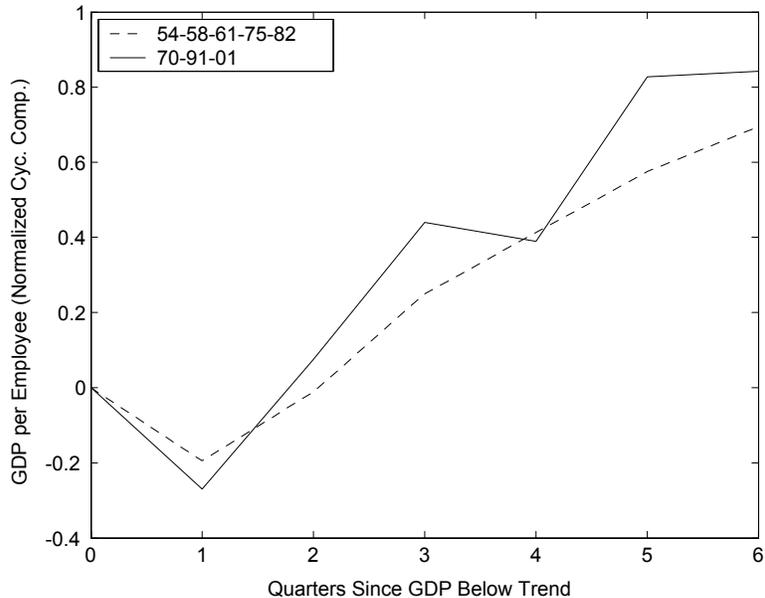


Figure 12: Comparison of Productivity per Employee Across Recoveries

seems to accord well with the implications of our model described above. Next we consider how productivity changes vary across the different types of recoveries. The next figure compares the average of these three recoveries with the average across the other five recoveries.

Both curves show the same qualitative behavior. The average of the three recoveries following long expansions does indeed have a slightly larger drop in the initial period and does show somewhat higher subsequent growth. Qualitatively these patterns are consistent with what one would expect from our model, though the quantitative differences do not seem that large.

Next we compare the behavior of productivity per hour for the post 1964 recoveries.

This graph shows the same two features: the three recoveries associated with recession following long expansions have a somewhat larger initial drop in productivity and subsequently experience somewhat higher growth.

6 Conclusion

We have highlighted a simple economic mechanism that we argue may be relevant for understanding the different behavior of labor market aggregates across business cycles. The model stresses two key effects. First, it argues that internal organizational dynamics are affected by aggregate shocks. Second, it stresses

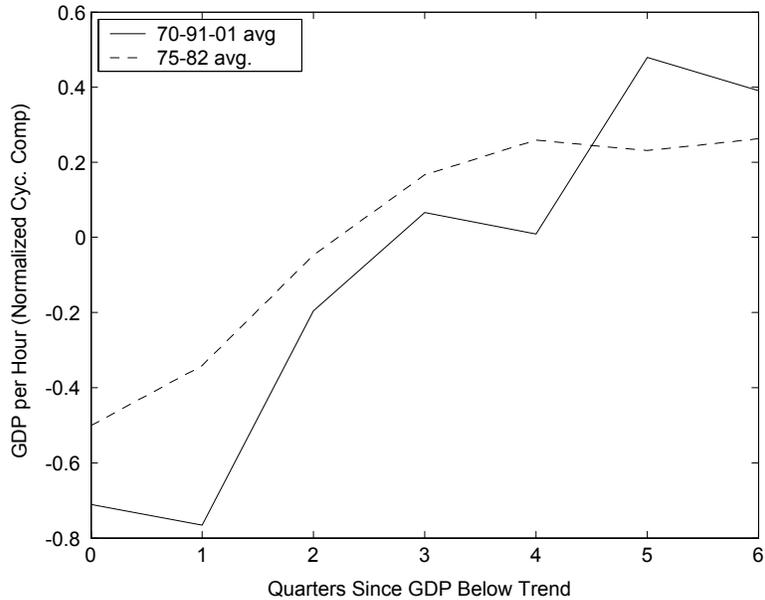


Figure 13: Comparison of Productivity Per Hour Across Recoveries

that the situations of organizations affect the manner in which the economy responds to aggregate shocks. In periods of high economic activity, organizations postpone structural changes to take advantage of current opportunities. But once an organization begins the process of restructuring itself, it is less likely to hire workers and more likely to release workers. These effects suggest the possibility that long expansions will be followed by recoveries in which employment starts to increase much later than output.

We then assess this link by studying eight US recessions in the post 1950 period. We argue that all three recoveries from recessions that followed long expansions exhibit the pattern of a long delay in the turning point for aggregate employment. This finding contrasts sharply with the characterization that it is the two most recent recoveries that are distinct.

While we think this work is suggestive, we must also emphasize that it is indeed only suggestive. A more rigorous quantitative assessment of the economic mechanism is called for, as is a more thorough analysis of the data.

References

- [1] Bertschek, I., and U. Kaiser, “Productivity Effects of Organizational Change: Microeconomic Evidence,” mimeo, 2001.

- [2] Burns, A. and Mitchell, W., *Measuring Business Cycles*, New York, NBER, 1946.
- [3] Caballero, R., and E. Engel, "Explaining Investment Dynamics in US Manufacturing: A Generalized (S,s) Approach, *Econometrica* 67 (1999), 783-826.
- [4] Cole, H., and L. Ohanian, "New Deal Policies and the Persistence of the Great Depression: A General Equilibrium Analysis," *Journal of Political Economy* 112 (2004), 779-816.
- [5] Cooley, T. and E. Prescott, "Economic Growth and Business Cycles," in *Frontiers of Business Cycle Research*, Princeton University Press, 1991, 1-38.
- [6] Groshen, E., and S. Potter, "Has Structural Change Contributed to a Jobless Recovery?," *Federal Reserve Bank of New York Current Issues in Economics and Finance*, Volume 9, no. 8, 2003.
- [7] Hall, R., "Labor Demand, Supply and Employment Volatility," *NBER Macroeconomics Annual*, edited by Olivier Blanchard and Stanley Fischer, MIT Press, 1991, 17-61.
- [8] Hopenhayn, H., "Entry, Exit and Firm Dynamics in the Long Run," *Econometrica* 60 (1992), 1127-1150.
- [9] Hopenhayn, H., and R. Rogerson, "Job Turnover and Policy Evaluation: A General Equilibrium Analysis," *Journal of Political Economy* 101 (1993), 915-938.
- [10] Koenders, K., "Long Expansions and Slow Recoveries: A Closer Look at Employment Fluctuations," mimeo, ASU, 2004.
- [11] Krusell, P., and T. Smith, "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy* 106 (1998), 867-896.
- [12] Lilien, D., "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy* 90 (1982), 777-793.
- [13] Lucas, R., "Understanding Business Cycles," *Carnegie-Rochester Series on Public Policy* 5 (1977), 7-29.
- [14] Lucas, R. "On the Size Distribution of Business Firms," *Bell Journal of Economics* 9 (1978), 508-523.
- [15] Schreft, S., and A. Singh, "A Closer Look at Jobless Recoveries," *Federal Reserve Bank of Kansas City Economic Review*, Second Quarter, 2003, 45-73.
- [16] Schweitzer, M., "Another Jobless Recovery?," *Federal Reserve Bank of Cleveland Economic Commentary*, March 1, 2003.

- [17] Thomas, J., "Is Lumpy Investment Relevant for the Business Cycle?," *Journal of Political Economy* 110 (2002), 508-534.
- [18] Veracierto, M., "Plant Level Irreversible Investment and the Business Cycle," *American Economic Review* 92 (2002), 181-197.