

Jobless Recoveries*

David Andolfatto
Simon Fraser University

Glenn MacDonald
Washington University in St. Louis

December 30, 2004

Abstract

Historically, when an economy emerges from recession, employment grows with, or soon after, the resumption of GDP growth. However, following the two most recent recessions in the United States, employment growth has lagged the recovery in GDP by several quarters, a phenomenon that has been termed the “jobless recovery.”

To many, a jobless recovery defies explanation since it violates both historical patterns and the predictions of traditional macro theory. We show that a recession followed by a jobless recovery is precisely what neoclassical theory predicts when new technology impacts different sectors of the economy unevenly and is slow to diffuse, and sectoral adjustments in the labor market take time to unfold.

1 Introduction

One of the familiar business cycle facts is that when an economy emerges from recession, employment tends to grow either contemporaneously with, or soon after, the recovery in GDP growth; indeed, this is part of the motivation for Lucas’ (1982) famous declaration that “business cycles are all alike.” But this key fact appears to be changing. Figure 1 plots the time path for real GDP (per adult) and employment (per adult) for the United States over the last thirty years (1972.4 – 2003.2). The sample period covers five periods of declining GDP and employment. While employment invariably begins to fall along with GDP, in the first three recoveries employment growth

*We would like to thank the conference participants at the 2004 SED meetings in Florence and the 2004 Canadian Macro Study Group in Montreal. In particular, we would like to thank Jonas Fisher for his many helpful comments and criticisms. Andolfatto acknowledges the financial support of SSHRCC.

resumes almost immediately, i.e., within one or two quarters after GDP growth turns positive. However, during the two most recent recessions, employment growth has lagged the economic recovery by several quarters, leaving observers to puzzle over a phenomenon that has been labeled the “jobless recovery.” (The jobless recoveries are shaded in the Figure.)

The phenomenon of an extended jobless recovery seems, to some, hard to understand since it violates both historical patterns and the predictions of familiar macroeconomic wisdom. In particular, a conventional view assumes that “aggregate demand” drives growth during a recovery. Since a major component of demand stems from the consumer, and since a major component of household income comprises labor earnings, it is plausible that the prospect of employment growth fuels the economic recovery. Lack of employment growth during a recovery is, therefore, problematic from this perspective. Some are even led to question whether a recovery in GDP is sustainable in the absence of employment growth. For example, according to Weller (2004): “Economic growth is more broad based than just a few quarters ago, with all economic industries contributing, which helps to stabilize and solidify the recovery. To maintain this momentum, though, the labor market has to improve. Otherwise, consumption, which comprises the vast majority of the economy, will not be able to grow at a strong pace, possibly putting a damper on growth in the medium-term.” Naturally, this way of looking at the jobless recovery influences discussions of economic policy. For example, Bernstein (2003) writes “The jobs picture is so serious that steps to stimulate the economy and generate job growth are urgently needed. Any stimulus proposal should be evaluated primarily on its impact on job creation and its ability to reverse the current trend of weakening wage growth.”

A jobless recovery, as a matter of arithmetic, obviously implies growth in labor productivity, which is generally viewed as a positive development. But some take the view that increases in labor productivity may have undesirable consequences in the short-run. For example, according to Governor Ben Bernanke (2003): “Strong productivity growth provides major benefits to the economy in the longer term, including higher real incomes and more efficient and competitive industries. But in the past couple of years, given erratic growth in final demand, it has also enabled firms to meet the demand for their output without hiring new workers. Thus, in the short run, productivity gains, coupled with growth in aggregate demand that has been insufficient to match the expansion in aggregate supply, have contributed to the slowness of the recovery of the labor market.”

We see several problems with explanations like those offered by Bernanke. First, the explanation takes as given an “insufficient and erratic growth in final demand.” Even assuming that final demand is a well-defined concept, why should its growth be insufficient and erratic? Leaving aside the issue of its erratic nature, one possible explanation for the “insufficiency” of aggregate demand is “sticky” product prices. If prices are sticky, then a period of rapid productivity growth increases profit margins and allows firms to meet available demand with fewer workers. But this explanation is implausible since it requires a degree of price stickiness (i.e., several quarters) that is inconsistent with the data; e.g., see Bils and Klenow (2002).

Our explanation for the jobless recovery is based on the following three assumptions: (1) Technological innovations vary in their size and scope, where “scope” refers to the breadth of potential applicability of the new technology across various firms; (2) New technology takes time and resources to diffuse; and (3) Labor cannot be instantaneously reallocated across different firms.¹ These assumptions are neither new nor, do we think, controversial. Greenwood, MacDonald and Zhang (1996) develop a real business cycle model that embeds assumptions (1) and (3), with firms interpreted as conventionally defined sectors. As with any framework that appeals primarily to a costly labor market adjustment mechanism, output and employment in that model share similar transition dynamics, ruling out a jobless recovery. In Andolfatto and MacDonald (1998) we investigated the properties of a neoclassical growth model that embeds assumption (2), and showed how the arrival and slow diffusion of new technology can generate booms and slowdowns in output that are consistent with the data. However, since intertemporal substitution is the only source of employment dynamics in that model, it has limited success in accounting for employment dynamics, and shows little promise of explaining a jobless recovery. No one, to our knowledge, has investigated the implications of these assumptions taken jointly within the context of a business cycle model. When we do this, we find that not only is neoclassical theory consistent with the phenomenon of a jobless recovery, but it also offers an explanation as to how technological advances

¹We use the terminology “firms” to refer to a variety of economic entities that might be affected by new technology. The reader can substitute “occupation”, “industry”, or “sector” if that seems more appropriate. But whatever the terminology, it is important to recognize that the scope of new technology may not match, e.g., the BLS definition of any particular occupation, or correspond to any particular SIC code. For example, an innovation based on bioluminescence will have impact on many occupations and industries conventionally defined, but there is no aggregate, e.g., manufacturing, that corresponds to these occupations and industries.

may lead to the recession that precedes the jobless recovery. As will become apparent, the model does not have to be “unusual”, or its parameter values “extreme”, for it to produce a quantitatively significant jobless recovery. We also show that (1) - (3) must all be present if a jobless recovery is to occur.

The basic economics is as follows. Suppose a new technology arrives, and, per assumption (1), it has more potential to impact some firms more than others. Because, per assumption (2), the new technology does not diffuse instantaneously, it is initially applied by very few firms. But there are many other firms that are eager to learn and implement it. To the extent that implementation is endogenous and costly, resources are diverted away from production toward general “learning activities.” This diversion itself, along the lines of our earlier work, may lead to some contraction of output in the impacted firms. At the same time, the prospect of future productivity growth in those firms that may potentially benefit from the new technology begins to attract workers from firms outside the scope of the new technology. But, per assumption (3), reallocating labor across firms is costly and entails a period of unemployment for some workers. Since technology is slow to diffuse, initially, overall productivity does not rise greatly. Thus, the diversion of labor away from productive activities (at firms within the scope who divert labor toward technology implementation activities) and the increase in unemployment (from firms outside the scope) leads to an initial contraction in GDP. As in our earlier work, the new technology subsequently spreads like a contagion. When this picks up, productivity begins to rise rapidly, and the labor market “restructuring” accelerates. The rapid growth in productivity more than makes up for the overall decline in employment, so that the economy experiences a “jobless recovery.” Eventually, the pace of labor reallocation begins to decline and employment begins to recover. Productivity growth remains positive, but slows as the new technology approaches full absorption. At this phase of the cycle, both output and employment are growing, i.e., the economy enters into a full expansion phase (with firms outside the scope continuing to suffer). The cycle completes its course as the new technology is fully implemented. At this stage, both output and employment growth approach their normal levels.

The balance of the paper is organized as follows. Section 2 sets out the basics of the economic environment. Section 3 details the assumptions we make concerning the economy’s adjustment technologies. Section 4 then analyzes the optimizing behavior of workers and firms, and characterizes the economy’s general equilibrium. Section 5 discusses the model’s parameterization. Section 6 reports the economy’s response to a technology shock

under the benchmark parameterization and shows that the model easily delivers quantitatively important jobless recoveries in response to technology change that impacts productive capabilities unevenly. Sensitivity analysis in Section 7 confirms that all of limited scope, slow diffusion, and nontrivial search are required for a jobless recovery, and also reveals that the extent and duration of such recoveries depends primarily on the scope of the technology shock. Section 8 concludes.

2 Basics

Time is discrete and the horizon infinite; $t = 0, 1, \dots$. The economy is populated by a unit mass of infinitely-lived individuals. Individuals have identical preferences defined over stochastic consumption profiles $\{c_t \mid t \geq 0\}$. Preferences are represented by the function:

$$E_0 \sum_{t=0}^{\infty} \beta^t c_t,$$

where $0 < \beta < 1$. We will focus on perfect foresight equilibrium, so the expectations operator E_0 reflects only the presence of agent-specific uncertainty. Each individual is endowed with one unit of time, which we will assume is devoted either to working (in either production or learning) or searching.

We will be concerned with technology change that impacts various parts of the economy differentially. To model this we assume a unit continuum of firms where each firm is endowed with $k > 0$ units of a firm-specific factor of production; e.g., immobile land or capital. (Keep in mind that what we label here as a firm may in fact be better thought of as a plant or occupation; i.e., in the data, a firm may consist of several plants and/or occupations.) Since k is distributed uniformly over the unit interval, k represents both the aggregate and firm-specific quantity of the fixed factor. For simplicity, we assume that each individual owns an equal share of the economy-wide stock k .²

Production at every firm is described by the neoclassical production technology $y = F(k, n)$, where y denotes (homogeneous and nonstorable) output, k denotes the firm-specific factor, n denotes the level of employment,

²The only substantive role k plays in the analysis is to distribute profits. The reader can assume, as we do for the numerical analysis, $k \equiv 1$ if desired.

and F is linear homogeneous and so has diminishing marginal product in each factor.

If labor is freely mobile across firms, the equilibrium wage at each date and firm is $w^0 = F_n(k, 1)$, with the population distributed uniformly across firms. Consumption and output for each individual is $c = y = F(k, 1)$.

2.1 Technological innovation

Suppose the allocation just described is the economy's "initial" position. Then, at some date, which we will label $t = 0$, a major new technology is discovered that, if implemented, improves production possibilities at some fraction μ ($0 < \mu \leq 1$) of firms. The parameter μ indexes the *scope* of the new technology. If implemented at some firm within the scope, new technology augments the productivity of k by the factor γ ($\gamma > 1$); production possibilities at these firms becomes $y = F(\gamma k, n)$. The parameter γ indexes the *size* of the technological improvement. Together, scope and size will determine the economy's long-run potential GDP. Call the set of firms where production possibilities have improved *sector 1*, and the rest *sector 2*.

We think of the arrival of major new technology, e.g., microelectronics or the Internet, as occurring infrequently, and, for simplicity, model this as the arrival being completely unanticipated. Expectations along the adjustment path subsequent to the arrival of a technological breakthrough, however, are assumed fully rational. As there is no aggregate uncertainty, the equilibrium dynamics subsequent to the technology shock will follow a perfect foresight path.

Let $n(j)$ denote the level of employment at a firm located in sector j , $j = 1, 2$. If the economy could adjust instantly to the new technology, then the new equilibrium would be characterized by the following conditions (asterisks denote equilibrium values):

$$\begin{aligned} w^* &= F_n(\gamma k, n^*(1)) \\ w^* &= F_n(k, n^*(2)) \\ 1 &= \mu n^*(1) + (1 - \mu)n^*(2) \\ y^* &= \mu F(\gamma k, n^*(1)) + (1 - \mu)F(k, n^*(2)) \end{aligned}$$

The first two conditions describe firms' optimizing choice of labor in each sector; the others describe market clearing in the labor and goods markets respectively. Clearly, $n^*(1) > n^*(2)$, so $w^* > w^0$. That is, while employment

in sector 1 expands at the expense of sector 2, since individuals are identical, equilibrium requires that they share equally in the higher wages induced by the new technology. Note that while labor productivity rises in both industries, it does so for very different reasons. In sector 1, a more efficient technology makes labor more productive. In sector 2, productivity rises because the capital-labor ratio increases.

3 Adjustment technologies

Realistically, it takes time and effort for firms to adopt major new ideas, and workers moving to more attractive jobs/occupations is not free either. To describe this process, we introduce two “adjustment technologies”.

3.1 Job search

We assume labor is perfectly mobile within a sector, but not across sectors. An individual who attempts to gain employment in a different sector *fails* to do so with probability $0 < \phi < 1$. Given failure, the individual foregoes his wage for one period. Then, the individual has the option of working in the sector in which he worked previously, or trying to switch sectors. Conditional on success, the individual immediately earns the competitive wage in his new sector.

To describe how the size of the workforce evolves in each sector, we anticipate some equilibrium behavior, viz., that no individual would choose to leave a firm within the scope of the technology shock. Thus, any firm in sector 1 will generally attract workers from outside the scope, so the workforce will typically be expanding in sector 1 and contracting in sector 2.

Let $x_t(j)$ denote the workforce (i.e., those who might work or search) *per firm* in sector j . Since we anticipate that no worker in sector 1 will search, let u_t denote the number of individuals per sector 2 firm who elect to search. Then *total* unemployment is simply $(1 - \mu)u_t$, and there are $(1 - \mu)u_t/\mu$ searchers per sector 1 firm. It follows that the per firm workforce within each sector evolves according to:

$$\begin{aligned}
x_{t+1}(1) &= x_t(1) + (1 - \phi) \left(\frac{1 - \mu}{\mu} \right) u_t \\
x_{t+1}(2) &= x_t(2) - (1 - \phi) u_t.
\end{aligned} \tag{1}$$

Since, there is a unit mass of workers, we must have

$$\mu x_t(1) + (1 - \mu) x_t(2) = 1.$$

3.2 Technology diffusion

Following Andolfatto and MacDonald (1998), when news of the new technology arrives, firms in sector 1 learn they have the *potential* to benefit, but generally must undertake costly activities to implement the new technology. That is, there is a difference between understanding the availability of a technology and actually learning how to implement it. Let λ_t denote the fraction of sector 1 firms that have learned how to implement the new technology. Firms in sector 1 are then labelled either *high-tech* or *low-tech* depending on whether they have implemented the new technology. For simplicity, suppose that when the new technology arrives, some (small) fraction $\lambda_0 > 0$ of firms in industry 1 learn the new technology immediately and costlessly; we will treat λ_0 as a parameter.³

Assume that learning to implement the new technology takes time and resources, and is not fully predictable (at the firm level). Let ι_t denote imitation effort, specifically the number of workers employed in the learning process at a representative low-tech firm (high-tech firms will devote no resources to learning). These workers must be paid a competitive wage. Given ι_t , a low-tech firm successfully learns the technology, and so can use it in subsequent periods, with probability $\xi(\iota) \lambda_t$, where $\xi(0) = 0$, $0 \leq \xi < 1$ and $\xi' > 0 > \xi''$. The law of motion that describes the pattern of diffusion is then:

$$\lambda_{t+1} = \lambda_t + (1 - \lambda_t) \xi(\iota_t) \lambda_t. \tag{2}$$

³The assumption that an exogenous $\lambda_0 > 0$ firms learn is merely a convenience that allows us to model the diffusion of new technology exclusively as “imitation”, i.e., low tech firms learning from high tech, rather than as a blend of imitation and “innovation”, the latter meaning firms can learn independently of others. If $\lambda_0 = 0$, imitation-based diffusion cannot begin. Allowing $\lambda_0 = 0$ and innovation can also be accommodated along the lines of Jovanovic and MacDonald (1994).

Notice that the technology of learning is specified so that it becomes easier for a firm to adopt the new technology when others have already done so. The idea is that a new technology becomes progressively easier to learn the more widely it is in use because there is more commonly-known experience with learning how to implement the new technology. Accordingly, the diffusion of technology will follow the familiar S-shaped pattern.

4 Individual optimization and equilibrium

In this section we characterize optimal behavior for individuals and firms. Since we are modelling behavior following the arrival of the new technology, and individuals expect no further technology shocks, there is no aggregate uncertainty. Accordingly, when forming their decisions, individuals take as given a vector of deterministic sequences describing the evolution of real wages and the distribution of knowledge.

4.1 Firms

Let $w_t(i)$ denote the real wage in sector i at date t . We distinguish between high-tech and low-tech firms with superscripts H or L . Define:

$$\begin{aligned}\pi_t^H(1) &\equiv \max_{n_t^H(1)} \{F(\gamma k, n_t^H(1)) - w_t(1)n_t^H(1)\} & (3) \\ \pi_t^L(1) &\equiv \max_{n_t^L(1)} \{F(k, n_t^L(1)) - w_t(1)n_t^L(1)\} \\ \pi_t(2) &\equiv \max_{n_t(2)} \{F(k, n_t(2)) - w_t(2)n_t(2)\}\end{aligned}$$

For all firms, choosing employment devoted to production in each period coincides with optimal behavior. The only dynamic choice is faced by the low-tech firms who must also decide on the extent of effort to learn the new technology, ι_t .

Let $V_t^H(1)$ and $V_t^L(1)$ denote the capital value of optimizing high- and low-tech firms in sector 1 at date t . The sequence $\{V_t^H(1)\}_{t=0}^\infty$ satisfies, for all t :

$$V_t^H(1) = \pi_t^H(1) + \beta V_{t+1}^H(1), \quad (4)$$

and the sequence $\{V_t^L(1)\}_{t=0}^\infty$ satisfies, for all t :

$$V_t^L(1) = \max_{\iota_t} \left\{ \begin{array}{c} \pi_t^L(1) - w_t(1)\iota_t + \\ \beta [(1 - \xi(\iota_t)\lambda_t)V_{t+1}^L(1) + \xi(\iota_t)\lambda_t V_{t+1}^H(1)] \end{array} \right\}. \quad (5)$$

Anticipating $V_t^H(1) \geq V_t^L(1)$, the condition describing the optimal level of learning effort (assuming an interior solution) is:

$$-w_t(1) + \xi'(\iota_t)\lambda_t\beta [V_{t+1}^H(1) - V_{t+1}^L(1)] = 0. \quad (6)$$

Optimal learning effort is increasing in the expected discounted capital gain associated with success, $\beta [V_{t+1}^H(1) - V_{t+1}^L(1)]$, and in the current state of technology absorption λ_t . Likewise, optimal learning effort decreases with the cost, $w_t(1)$, of employing workers in such an activity.

4.2 Workers

Let $J_t(j)$ denote the capital value associated with an individual who is employed in sector j at date t . Since individuals who work in production earn the same wage as those employed in learning activities, we need not distinguish between the two. Similarly, let Q_t denote a searcher's (necessarily from sector 2) capital value. Anticipating that $J_t(1) \geq Q_t$ and $J_t(1) \geq J_t(2)$, these capital values must satisfy:

$$\begin{aligned} J_t(1) &= w_t(1) + \beta J_{t+1}(1) & (7) \\ J_t(2) &= w_t(2) + \beta \max\{J_{t+1}(2), Q_{t+1}\} \\ Q_t &= (1 - \phi)J_t(1) + \phi\beta \max\{J_{t+1}(2), Q_{t+1}\} \end{aligned}$$

The choice problem facing individuals in sector 2 is simple: if $J_t(2) \geq Q_t$, then search; otherwise, remain working in sector 2.

4.3 Equilibrium

Given our timing assumptions, the sector 1 labor market must always satisfy

$$\mu x_t(1) = \mu [\lambda_t n_t^H(1) + (1 - \lambda_t)(n_t^L(1) + \iota_t)] + (1 - \mu)(1 - \phi)u_t. \quad (8)$$

The final term in the expression above represents the flow of individuals who successfully make a transition from sector 2 into the sector 1 labor market. At the same time, the sector 2 labor market must also obey

$$(1 - \mu)x_t(2) = (1 - \mu)[n_t(2) + \phi u_t]. \quad (9)$$

Of course, (8) and (9) must agree with (3.1).

Optimization by firms gives:

$$\begin{aligned} w_t(1) &= F_n(\gamma, n_t^H(1)) = F_n(1, n_t^L(1)); \\ w_t(2) &= F_n(1, n_t(2)); \end{aligned}$$

in addition to condition (6).

Along the equilibrium path, unemployment will either be strictly positive or equal to zero. For strictly positive unemployment rates, individuals in sector 2 must be just indifferent between working sector 2 or searching for employment in sector 1; otherwise, working in sector 2 must dominate search. That is

$$u_t > 0 \text{ implies } J_t(2) = Q_t$$

and

$$u_t = 0 \text{ implies } J_t(2) \geq Q_t.$$

An *equilibrium* for this economy is a set of sequences:

$$\{x_t(1), x_t(2), u_t, \lambda_t, \iota_t, w_t(1), w_t(2), n_t^H(1), n_t^L(1), n_t(2), V_t^H(1), V_t^L(1), J_t(1), J_t(2), Q_t\}_{t=0}^{\infty}$$

and an initial condition $(x_0(1), x_0(2), \lambda_0)$ satisfying:

1. *Individual optimization*: given $\{w_t(1), w_t(2)\}_{t=0}^{\infty}$, the sequences $\{J_t(1), J_t(2), Q_t\}_{t=0}^{\infty}$ satisfy equation (7);
2. *Firm optimization*: given $\{w_t(1), w_t(2), \lambda_t\}_{t=0}^{\infty}$, the sequences $\{n_t^H(1), n_t^L(1), n_t(2), \iota_t, \}_{t=0}^{\infty}$ $\{V_t^H(1), V_t^L(1)\}_{t=0}^{\infty}$ satisfy equations (3), (4), (5) and (6);
3. *Labor market clearing*: (1), (3.1), (8) and (9) hold at every t ;
4. *Search evolution*: (4.3) and (4.3) hold at every t ; and
5. *Technology evolution*: (2) holds at every t .

5 Parameterization

We assume the following functional forms for the production and learning technologies:

$$F(\gamma^{\chi}k, n) = (\gamma^{\chi}k)^{\alpha}n^{1-\alpha}$$

and

$$\xi(\iota) = 1 - e^{-\eta\iota},$$

where $\chi = 1$ if the new technology has been implemented; $\chi = 0$ otherwise.

Taking a time interval to be one year, we assume a commonly-employed value for the discount factor, $\beta = 0.96$. We set $\alpha = 0.36$, implying a long run labor share of income equal to .64. The search failure probability is $\phi = 0.50$, so that workers who choose to switch occupations have a 50% chance of doing so successfully within one year; we explore the consequences of higher and lower values below. The parameter η governs the speed at which the new technology diffuses. Setting $\eta = 25$ implies that it takes several years for a new technology to diffuse fully; results will also be reported for different diffusion rates. Finally, we set $k = 1$.

The initial conditions are such that all firms share the same technology. Thus, when the new technology arrives, since search takes at least one period, all firms have the same initial workforce; i.e., $x_0(1) = x_0(2) = 1.0$. When the new technology arrives, we assume 1% of firms *within the scope* immediately understand how to implement it, $\lambda_0 = 0.01$.

The two remaining parameters describe size (γ) and scope (μ) of the new technology. Assuming that new technologies are eventually fully absorbed, these two parameters dictate the long-run increase in real per capita GDP. Below, we will consider different configurations of these parameters, with each configuration generating a long-run increase in real GDP equal to about 20%. In our benchmark parameterization, we set $\mu = 0.25$ and $\gamma = 3.75$.

6 Results

6.1 Benchmark parameterization

The initial situation features full employment, with sector 1 and 2 employment shares equal to .25 and .75, respectively. The initial wage rate is the same in each sectors, and equal to .64 (labor's share of initial output, which is normalized to unity).

The technology shock is narrow in scope, being (potentially) available to only 25% of firms. But for these firms, successful implementation of the new technology increases productivity by just over 60% (i.e., by a factor of $\gamma^\alpha = 3.75^{0.36}$). On impact, however, only 1% of firms in the favorably-affected sector are able to implement the new technology immediately. Thus,

the initial impact on economy-wide TFP is miniscule.

The top panel of Figure 2 displays the post-innovation time path of employment by sector, as well as unemployment. The bottom panel of Figure 2 displays the time path of wage rates and labor's share of income. At the outset, new technology does little to stimulate restructuring; unemployment remains at its normal level for two periods following the shock. Wages paid by firms within the scope, however, begin to rise almost immediately, albeit at a modest rate. The resulting "wage gap" across firms within and outside the scope eventually makes it attractive for some workers to invest in search.

The pace of restructuring begins to increase substantially four or five periods following impact. At this stage, the new technology is beginning to diffuse very quickly across firms within the scope. The rapid increase in productivity among the favorably affected firms increases their demand for labor, which continues to put upward pressure on the real wage. Workers currently located at firms outside the scope are attracted by the high wages being paid elsewhere, which is what accounts for the significant rise in unemployment. Notice that the drain of workers from firms outside the scope compels these firms to accept wage increases that are consistent with some workers continuing to work for these firms; thus, profits at these establishments decline. ("New" and "old" economy examples come to mind.)

Despite the fact that workers are able to switch occupations with moderate ease, the dynamics of this process are drawn out for several periods owing to the slow diffusion of technology. Unemployment peaks a full nine periods following the arrival of the technology shock. Eventually, the restructuring process completes its course after fourteen periods have passed. At this stage, those firms within the scope, which comprises 25% of all firms in the economy, employ over 50% of the workforce.

Figure 3 displays the time paths of GDP growth and employment growth. Notice that the arrival of the new technology initially generates a brief/mild recession (or, at least a growth slowdown). This initial decline in output occurs as some firms within the scope divert labor away from production toward learning activities. Later, output declines (despite rising productivity) as learning continues and workers from outside the scope become unemployed.

GDP growth turns positive four-five periods after the initial technology shock. At this stage, the number of firms that have learned the new technology reaches a critical mass that makes subsequent adoption much easier for laggards. As a result, the new technology begins to diffuse quickly. The

rapid adoption of technology leads to a surge in productivity growth. This rapid spread of new technology makes investment in search attractive and so stimulates the pace of sectoral readjustment, and employment continues to decline even more rapidly (as unemployment peaks). During this phase of the adjustment process, the economy experiences a jobless recovery. As the new technology approaches full absorption, the rate of diffusion must necessarily slow. In the final phase of the cycle, both GDP and employment growth peak and eventually decline to their normal growth rates (zero, in this model). This is the full expansion phase of the cycle.

6.2 Sensitivity

In this section, we examine the sensitivity of our results to various parameter perturbations. In the first experiment, we vary the parameter that governs the ease with which workers are able to switch sectors. Figure 4 reports the results for $\phi \in \{0.0, 0.25, 0.50, 0.70\}$. Recall that ϕ measures the probability that search is *unsuccessful*.

When $\phi = 0$, workers are able to switch occupations easily. Accordingly, there is no jobless recovery since employment remains stable as the new technology diffuses. This result justifies our earlier claim that labor market frictions are a necessary (but not sufficient) condition for a jobless recovery. The qualitative properties of a jobless recovery appear to be robust for a wide range of parameter values. The effect of increasing the difficulty of switching sectors is to shorten the jobless recovery, while increasing the amplitude in the employment growth rate.

Now consider varying the scope of new technology, μ . Recall that μ measures the fraction of firms in the economy that may potentially benefit from the technology shock. Figure 5 reports the results for $\mu \in \{0.25, 0.50, 0.75, 0.90\}$. In each of these experiments, the size of the new technology, γ , is adjusted so that in each case the long-run GDP rises by approximately the same amount (20%). In particular, a technology shock with wider scope is smaller in size.

The top-left panel of Figure 5 records the benchmark parameterization. The effect of increasing the scope of new technology is to reduce the need for labor market restructuring. Accordingly, we see that the severity and length of the jobless recovery are reduced as the scope is widened. The increasing delay in widespread diffusion is accounted for by the fact that the size of the new technology is reduced as we increase its scope (so that the private incentive to adopt the new technology is not as great). While,

the qualitative properties of a jobless recovery remain intact for a wide range of parameter values, as μ becomes large, the jobless recovery, and unemployment generally, vanishes.⁴ This verifies our earlier claim that for a technology shock to induce a significant jobless recovery, a necessary (but not sufficient) condition is that the shock is not too broad in scope.

In the final experiment, we vary the speed at which the new technology diffuses. We do this in two ways. The first involves assuming different values for the initial fraction of firms (within the scope) who learn the new technology immediately upon its arrival, as governed by the parameter λ_0 . The idea here – somewhat outside the model – is to capture the notion that for some fraction of the scope, the new technology is very similar to what is already in use. For example, an improvement in the efficiency of bioluminescence technology may be easily implemented by firms already using it, but that this will be more difficult for firms who can eventually benefit from bioluminescence, but whose current production is not based on it, e.g., manufacturers of light bulbs.

Figure 6 reports the results for $\lambda_0 \in \{0.01, 0.10, 0.25, 1.0\}$. The top-left panel of Figure 6 is the benchmark parameterization; i.e., $\lambda_0 = 0.01$. Observe how the transition dynamics are altered significantly for even a moderate increase in λ_0 . For $\lambda_0 = 0.10$, the technology shock still induces a moderate recession on impact followed by a brief jobless recovery. However, when λ_0 is increased to 0.25, output remains virtually unchanged on impact, even though employment drops dramatically. In this case, the drop in employment is just compensated for by the increase in productivity. The bottom-right panel displays the limiting case where new technology is absorbed instantaneously (i.e., the standard real-business-cycle assumption). Observe that output and employment move in opposite directions only in the impact period of the shock. The subsequent transition dynamics are governed entirely by the adjustment costs in the labor market, with output and employment moving together throughout the transition period. Thus, slow technology diffusion is also a necessary (but not sufficient) condition for a technology shock to induce a jobless recovery.

The second way we vary the speed of diffusion of new technology is by assuming different values for η . An increase in η increases the probability that any given level of firm learning effort succeeds, in which case the firm

⁴Of course, such a shock may still induce fluctuations in the aggregate labor input, as demonstrated by Andolfatto and MacDonald (1998) in the context of a model that endogenizes the labor-leisure choice. However, the point here is that a broad scope technology shock is unlikely to induce a jobless recovery.

implements the new technology. Figure 7 reports the results of this experiment for $\eta \in \{7.5, 10, 25, 500\}$. Interestingly, even for very high η , e.g., 500, the jobless recovery exists. The reason is that the probability of a firm successfully imitating at t is $\lambda_t \xi(t)$, approximately λ_t for large η . Thus, high η implies that imitation is very likely to be successful once the technology is widely diffused, but not before that. So if η is low, e.g., $\eta = 7.5$, making imitation difficult, this slows diffusion and prolongs the jobless recovery. But even very large values, e.g., $\eta = 500$, do not remove the jobless recovery, i.e., finite η is not necessary. As we emphasized earlier, slow diffusion is necessary. Any factor that removes this feature will also eliminate a jobless recovery in our model.

7 Discussion

Our analysis suggests that a jobless recovery can be the result of a technology-driven recession (or growth slowdown). Our explanation hinges critically on three assumptions. First, technology advance is typically narrow in scope in the sense that not all technology benefits in the same way from innovations. Second, conditional on a shock that alters the relative returns to different economic activities, it takes time for resources to adjust to these new conditions. Lilien (1982) was one of the first to stress the idea that “sectoral shifts” in the structure of factor demands, together with costly readjustment, could explain a significant fraction of the increase in unemployment during recessions. One problem with Lilien’s hypothesis is that the act that there appears to be too much comovement in employment across conventionally defined sectors. But this may be more a problem with measurement than theory. In particular, it is conceivable that technology shocks that are narrow in scope alter the relative demands for different activities (e.g., occupations) within conventionally defined sectors or industries. While this idea sounds plausible, we are not aware of any data that can help us determine whether this is descriptive.

Finally, we assume that new technologies take time to diffuse. There is ample micro-level evidence that documents the well-known S-shaped diffusion pattern of new technologies; e.g., Griliches (1957). We think it is also reasonable to think this pattern also holds up for “general purpose technologies” (Bresnahan and Trajtenberg, 1996, and Jovanovic and Rousseau, 2003). At the macro level, Lippi and Reichlin (1994) argue that the stochastic process for GDP is more plausibly modeled as an ARIMA whose

impulse-response function follows an S-shaped pattern. Further empirical support for this idea is reported in Beaudry and Portier (2004), who investigate the joint behavior of stock prices and TFP movements. These authors find that business cycles appear to be driven largely by shocks that have little impact on TFP in the short run, but a big impact on TFP in the long-run, consistent with the way we have modeled new technology.

8 Conclusion

It is an empirical fact that the pattern of economic development in advanced economies is characterized by growth and fluctuations in GDP. The role of technological advance in generating growth is widely accepted. Real business cycle (RBC) theory asserts that since there is no *a priori* reason to expect the process of discovery to occur evenly over time, technology shocks may be largely responsible for both growth and fluctuations. In standard RBC environments, however, positive technological developments do not lead to recessions or jobless recoveries.


In this paper we explored the properties of an RBC model in which growth is driven by technological advances that improve factor productivity, that vary in the degree to which they affect the structure of the economy, that do not diffuse instantaneously, and that generate lasting labor market adjustments. The combination of technology advance with limited scope, less than instantaneous diffusion, and job search, yields income and employment dynamics that easily display recessions and jobless recoveries that are both quantitatively important and, in a general way, similar to the jobless recoveries whose emergence has proved so puzzling to many observers of aggregate economic activity.

We do not claim that all technological advances lead to jobless recoveries. According to our theory, whether a technological advance leads to a jobless recovery or not depends on the technology's scope and ease of implementation; these are parameters that are likely to vary across technological advances. And while advances in the technological frontier are initially characterized by recessions (or growth slowdowns) in our model, we also do not claim that all recessions are caused by technological advances. Our assertion is that technology advance may have been a contributing (and possibly primary) factor in some recent recessionary episodes. Other types of shocks, notably the oil price shocks of the early 1970s could, in the context of our model, generate a recession along with significant labor market re-

structuring. But because energy intensive sectors must absorb higher energy prices almost instantaneously, such a shock can be thought of as a reduction in TFP that diffuses quickly throughout the affected sectors. As we have demonstrated above, a shock that diffuses quickly will not generate a jobless recovery.

Of course, this leaves open the intriguing question of what factors determine the attributes of technological developments in terms of their size, scope, and ease of implementation. The fact that, according to our scope-based explanation of recently observed dynamics, new technology is systematically narrower in scope but greater in magnitude suggests the recent trend may not be serendipitous.

FIGURE 1
Jobless Recoveries
Real GDP per Adult
United States 1972.4 - 2003.2



Employment per Adult (%)
(Household Survey)
United States 1972.4 - 2003.2




FIGURE 2
 Employment, Wage, and Cost Dynamics
 Following a Technology Shock

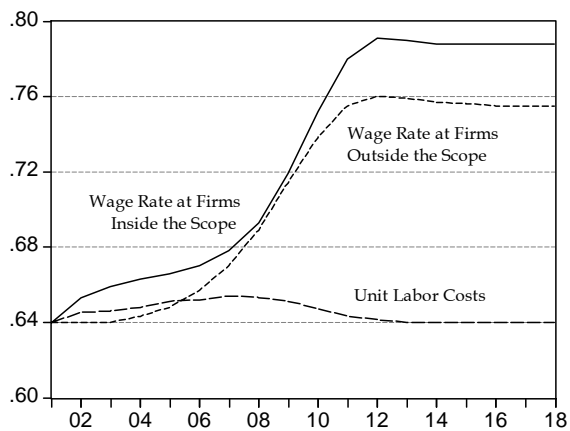
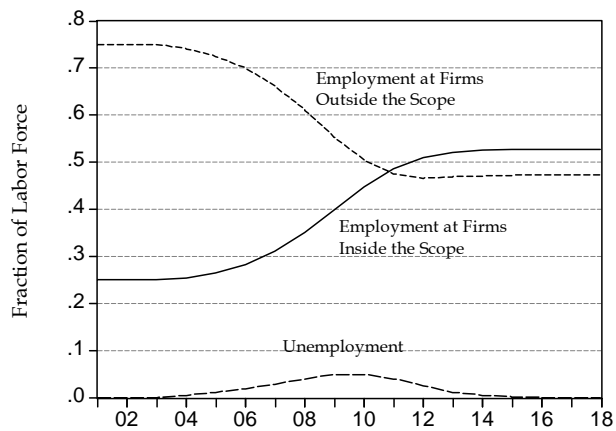


FIGURE 3
GDP and Employment Growth Rate Dynamics
Following a Technology Shock

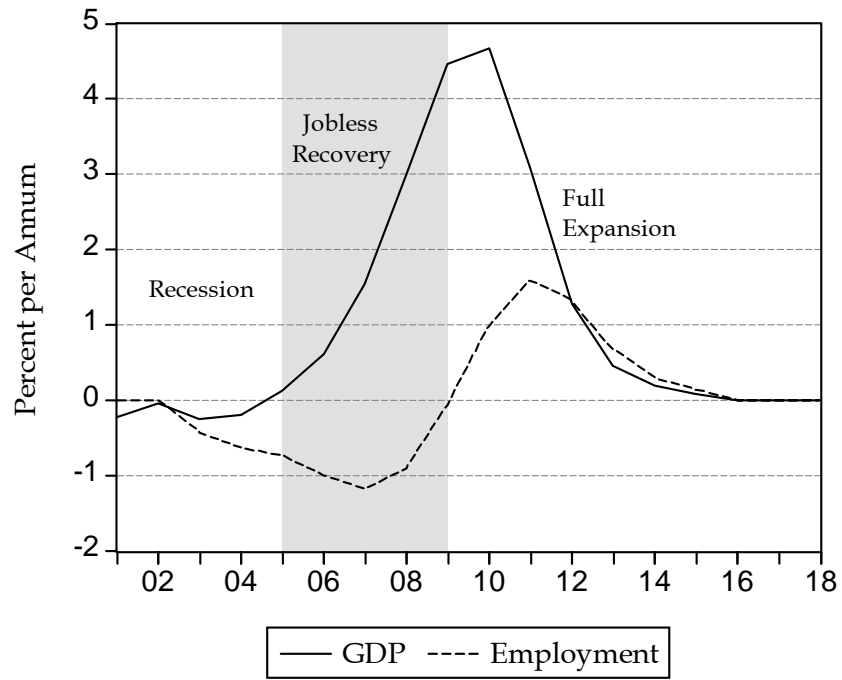


FIGURE 4
 GDP and Employment Growth Dynamics
 Following a Technology Shock for
 Different Job Finding Rates

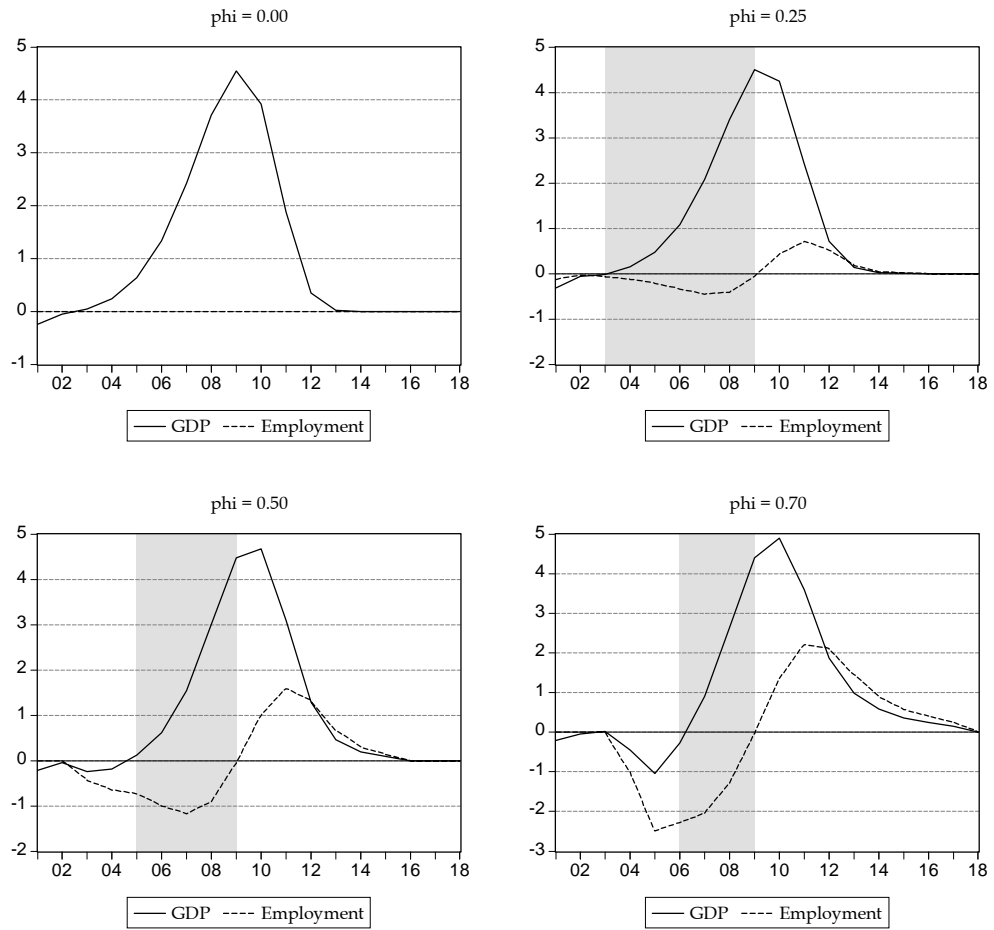


FIGURE 5
GDP and Employment Dynamics
Following a Technology Shock
of Varying Scope

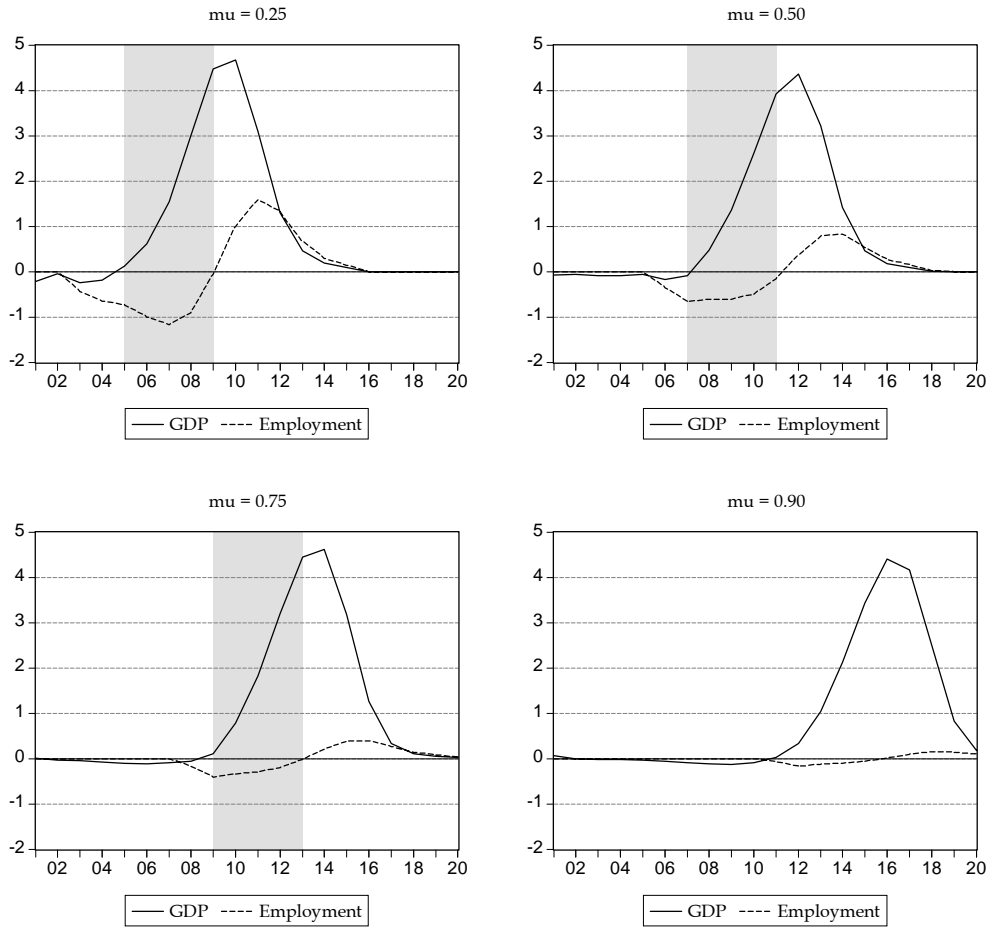


FIGURE 6
 GDP and Employment Dynamics
 Following a Technology Shock
 for Various Diffusion Rates

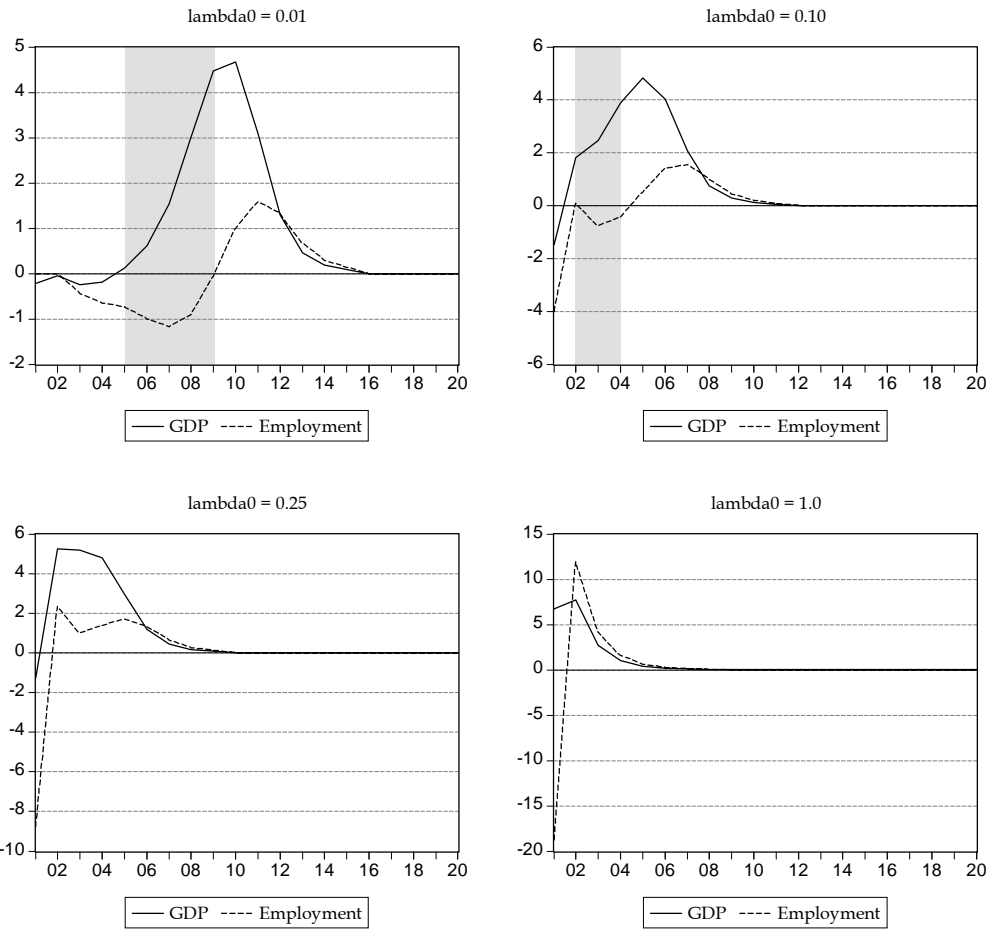
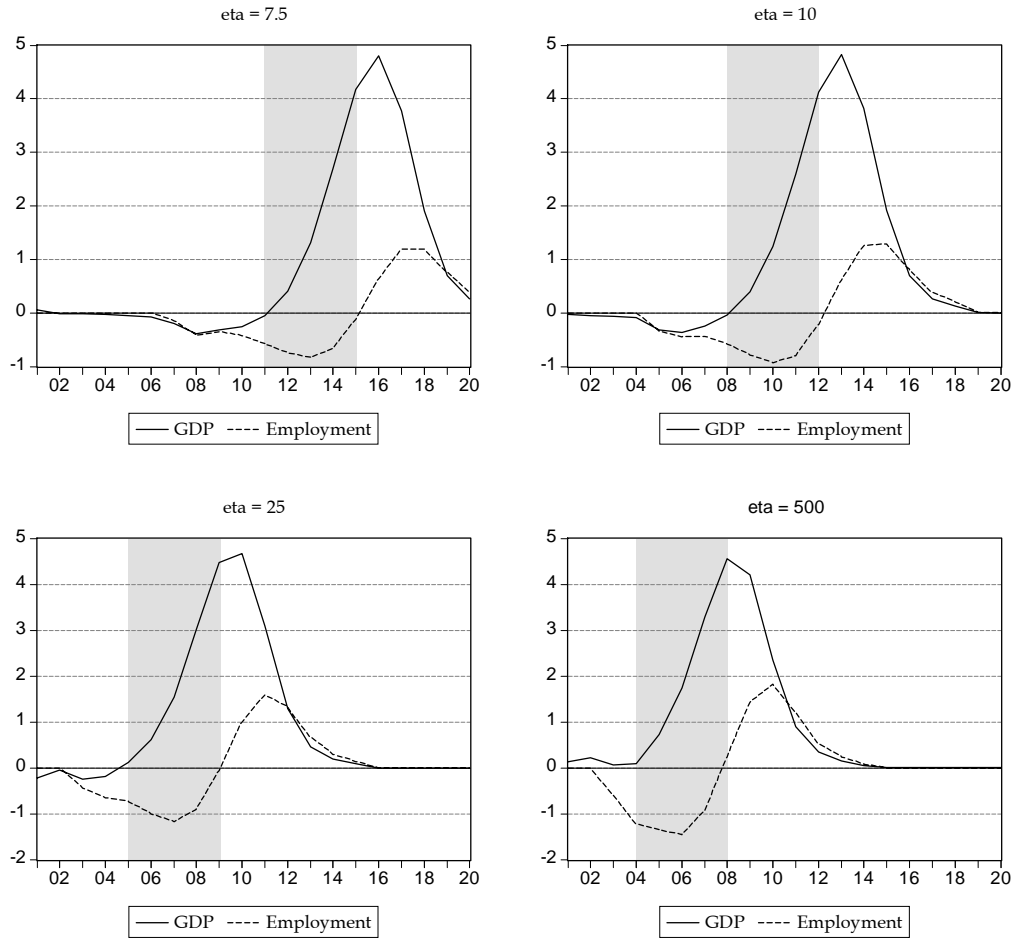


FIGURE 7
GDP and Employment Growth Dynamics
Following a Technology Shock as the
Ease of Imitation Varies



9 References

1. Andolfatto, David and Glenn M. MacDonald (1998). “Technology Diffusion and Aggregate Dynamics,” *Review of Economic Dynamics*, April, 1(2): 338-370.
2. Bernanke, Ben S. (2003). “The Jobless Recovery,” www.federalreserve.gov/boarddocs/speeches/2003/200311062/default.htm
3. Bernstein, Jared (2003). “The Jobless Recovery,” www.epinet.org/content.cfm/issuebriefs_ib186
4. Beaudry, Paul and Franck Portier (2004). “Stock Prices, News, and Economic Fluctuations,” Manuscript: University of British Columbia.
5. Bils, Mark and Peter J. Klenow (2002). “Some Evidence on the Importance of Sticky Prices,” Manuscript: Federal Reserve Bank of Minneapolis.
6. Bresnahan, T. F. and M. Trajtenberg (1996). “General Purpose Technologies: Engines of Growth?” *Journal of Econometrics*, Annals of Econometrics, 65: 83–108.
7. Greenwood, Jeremy, Glenn MacDonald and Guang Zhang (1996). “The Cyclical Behavior of Job Creation and Job Destruction: A Sectoral Model,” *Economic Theory*, 7:95–112.
8. Griliches, Zvi (1957). “Hybrid Corn: An Exploration of the Economics of Technological Change,” *Econometrica*, 25(4): 501–522.
9. Groshen, Erica L. and Simon Potter (2003). “Has Structural Change Contributed to the Jobless Recovery?” Federal Reserve Bank of New York, *Current Issues in Economics and Finance*, 9(8): 1–7.
10. Jovanovic, Boyan and Glenn MacDonald (1994). “Competitive Diffusion,” *Journal of Political Economy*, 102: 24-52.
11. Jovanovic, Boyan and Peter L. Rousseau (2003). “General Purpose Technologies,” Manuscript: New York University.

12. Lilien, David (1982). "Sectoral Shifts and Cyclical Unemployment." *Journal of Political Economy*, 90: 777–794.
13. Lippi, Marco and L. Reichlin (1994). "Diffusion of Technical Change and the Decomposition of Output into Trend and Cycle," *Review of Economic Studies*, 61(2): 19–30.
14. Lucas Jr., Robert E., (1982). "Understanding Business Cycles," in *Studies in Business Cycle Theory*, London, England, MIT Press.
15. Weller, Christian (2004). "America's Jobless Recovery," progressive-trail.org/articles/040130Weller.shtml