# Productivity Shocks and the Unemployment Rate*

Bharat Trehan

*Research Advisor*
*Federal Reserve Bank of San Francisco*

Productivity grew noticeably faster than usual in the late 1990s, while the unemployment rate fell to levels not seen for more than three decades. This inverse relationship between the two variables also can be seen on several other occasions in the postwar period and leads one to wonder whether there is a causal link between them. This paper focuses on technological change as the common factor, first reviewing some recent research on the effect of technological change on the unemployment rate and then presenting some empirical evidence on the issue. While theoretical models make conflicting predictions about the effects of a technology shock on the unemployment rate, the empirical evidence presented here shows that a positive technology shock leads to a reduction in the unemployment rate that persists for several years.

## 1. Introduction

The economy boomed in the second half of the 1990s, with output and productivity growing rapidly. At the same time, unemployment fell to levels that had not been seen since the 1960s, leading to concerns that inflationary pressures were likely to build and suggesting to some (especially early in the boom) that policymakers might need to take action to prevent an acceleration of inflation.[1]

One aspect of this boom that was hard to miss was the large role that was being played by changes in technology. For one thing, these changes made themselves felt in the explosive growth of things like computers, the Internet, and cellular phones. Accompanying this growth was a dramatic decline in the prices of these products, which underlined the role that technical change appeared to be playing during this period. These price declines may have helped to exert downward pressure on the aggregate inflation rate as well; in any case, there was little evidence to suggest that inflation was picking up even after several years of rapidly growing output and record low unemployment.

This combination of circumstances led other observers to suggest that the "new economy" had a new "speed limit." More specifically, these observers suggested that because of technical change the economy could grow faster or operate at lower unemployment rates without experiencing higher inflation. Thus, they argued, there was no need for monetary policymakers to tighten policy.

While the subsequent recession put an end to the debate about whether policy needed to be any tighter (as policymakers began to worry about rising unemployment and falling output), it did little to answer the underlying question of how policy should react to a technology-driven boom. The appropriate response obviously depends upon how the effects of technical change make themselves felt in the economy. For instance, one possibility is that rapid technological change leads to a reduction in inflation independent of the prevailing unemployment rate. In this case, a return to unemployment rates that prevailed in the late 1990s without a pickup in the pace of technical change seen during that period would tend to be accompanied by a pickup in inflation. Another possibility is that improvements in technology lead to a permanently lower unemployment rate, in which case we could see a return to low unemployment in the near future without any pickup in inflation. Yet another possibility, which is incongruent with the U.S. boom of the 1990s but which cannot be ruled out a priori, is that rapid technological change destroys the job skills of some types of workers and leads to long-term unemployment for them. This would tend to raise the unemployment rate and to change the relationship between the unemployment rate and inflation as well.

*I would like to thank Chad Jones, Milton Marquis, Glenn Rudebusch, and Dan Wilson for helpful comments. I am grateful to Brian Johnson and, especially, to Seamus Smythe for capable research assistance.

1. Such a relationship is implied by the Phillips curve, as discussed, for example, in Mankiw (1992). According to Mankiw, the inflation rate depends on (among other things) "the deviation of unemployment from the natural rate, called *cyclical* unemployment" (p. 303). Thus, the concern in the second half of the 1990s was that cyclical unemployment had fallen too low, which would cause the inflation rate to go up. Also see Ball and Mankiw (2002).

This paper attempts to shed some light on this issue by focusing on what effects productivity shocks are likely to have on the unemployment rate. It seems useful to start by looking at the historical relationship between productivity and unemployment, which is shown in Figure 1. As mentioned above, during the 1990s the unemployment rate fell to levels that had not been seen for three decades. At the same time, productivity grew faster than it had at any time since the 1960s. More importantly, the figure reveals that this is not the only time that one can see such a pattern. Thus, the 1960s were a period of high productivity growth and relatively low unemployment as well. The same relationship can be seen in the 1970s, except that unemployment was rising while productivity growth languished. While some of the observed correlation is obviously short-run in nature, the evidence suggests that there is a long-run relationship between the two series as well. Stock and Watson (2001), for instance, construct measures of the trend (or long-run) components of unemployment and productivity growth and show that they are negatively related.

These comovements between the two variables suggest that there may be a causal relationship between them. This paper focuses on the causation running from (exogenous) changes in productivity to unemployment.[2] It begins by reviewing a theory of unemployment in Section 2 and goes on in Section 3 to discuss recent research that has made use of this theory to explain how changes in productivity can affect the unemployment rate. In Section 4 it discusses some other models that use changes in productivity to explain changes in the unemployment rate. This is followed by a discussion in Section 5 of the relatively limited empirical evidence on the issue. The paper then goes on to present some new estimates of the effects of productivity shocks on both the economy-wide unemployment rate and the unemployment rate of workers with different education levels. Section 6 concludes.

## 2. The Search Theory of Unemployment

Since much of the recent research on the relationship between productivity growth and unemployment is based on the search theory of unemployment, it is useful to begin with a simple overview of this theory. It starts with the assumption that workers have different skills and jobs have different skill requirements. Workers need to find desirable jobs at the same time that firms need to find the most productive workers. Neither firms nor workers have all the information they need about the options available to them, so they must engage in search. Since search is costly and time-consuming, both firms and workers must use some of their resources to find a good match.

FIGURE 1
PRODUCTIVITY GROWTH AND UNEMPLOYMENT

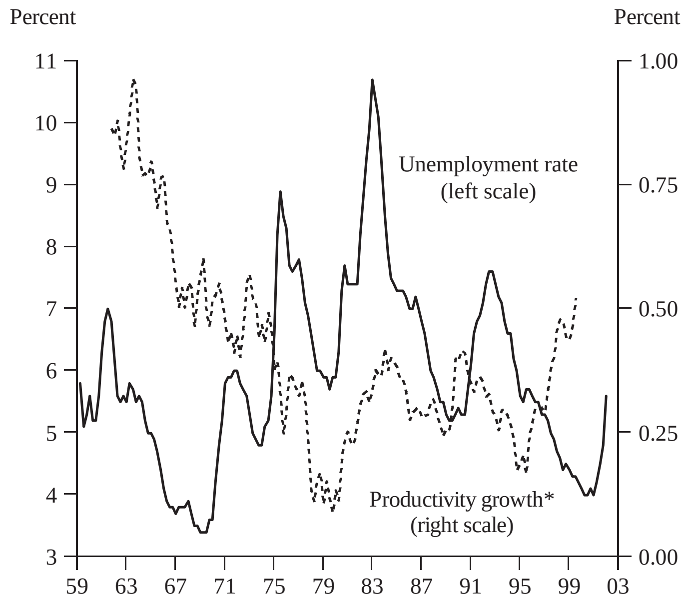


*Centered 5-year moving average.

Workers are assumed to search only when they are unemployed. They face an uncertain environment (just as firms do). When a worker gets a wage offer, for instance, she must decide whether to accept it or continue searching for a better offer. Accepting the offer means forgoing the chance of a higher wage offer later, while continuing the search means losing the wages she would have earned if she had accepted the offer and started working. The wage at which the worker is indifferent between continuing the search and accepting the current job is called the "reservation wage." The worker accepts all job offers above this wage and rejects all offers below it.

When a search is successful, that is, when there is a match between the needs of the worker and the firm, the worker leaves unemployment. However, existing matches sometimes fall apart, which leads to workers becoming unemployed. At the equilibrium unemployment rate, the number of workers leaving unemployment equals the number of workers becoming unemployed. The equilibrium unemployment rate moves around over time; it is often argued, for instance, that the equilibrium unemployment rate went up in the 1970s because a large number of workers with little or no experience entered the labor force.

2. Unemployment rates could have an effect on productivity as well. For instance, it has been argued that the unemployed tend to lose skills and become less productive. While this is a reasonable hypothesis, it is not obvious that it can be used to explain changes in economy-wide productivity levels over time.

The relative level of the reservation wage is obviously a crucial determinant of the level of unemployment in the economy. If the typical worker's reservation wage is significantly higher than the typical wage offer, she will tend to turn down more offers and spend more time searching for a job. Consequently, the unemployment rate will tend to be higher.

The wage offered by the firm is directly related to the worker's productivity. Assume, now, that there is an economy-wide increase in productivity that workers are not aware of. The higher productivity makes it more attractive for the firm to increase employment and allows it to do so by increasing the wage it offers to workers. This, in turn, increases the likelihood that the average worker will find an acceptable job offer and reduces the time she is likely to spend searching. Thus, the unemployment rate will decline in response to the increase in productivity.

This drop in the unemployment rate is unlikely to be permanent, however, even if there is no decrease in productivity. This is because workers will come to realize that all firms are offering higher wages than before and, consequently, their reservation wage will adjust gradually to the higher level of wage offers in the economy. As this occurs, the level of unemployment gradually will go back to the level that prevailed prior to the increase in productivity. Of course, the reservation wage could adjust slowly, and so it could take a while for the unemployment rate to return to its original level. Even so, the key implication is that a change in the level of productivity cannot have a permanent effect on the level of the unemployment rate. We now turn to a discussion of how recent research has used elements from this story to explain the recent behavior of the unemployment rate.

## 3. Models with Homogenous Labor

Ball and Moffitt (2001) present a variation on the process described above to explain the decline in unemployment in the second half of the 1990s.[3] Instead of reservation wages, they conduct their analysis in terms of "aspiration wages," which is defined as wages that workers consider to be fair. Research by Akerlof and Yellen (1990), for instance, shows that workers are likely to reduce the amount of effort they put in on the job if they perceive the wages they receive to be unfair. However, it is hard to come up with an operational definition of "fair wages," especially when one is dealing with aggregate data. Ball and Moffitt assume that workers use past wages to determine fair wages. As a consequence, an increase in the rate of productivity growth does not show up as an increase in the aspiration wage at once and so leads to an increase in employment and a fall in the unemployment rate. As the aspiration wage adjusts to the increase in productivity, the unemployment rate rises back up.

In Mankiw and Reis (2001), the target nominal wage depends upon the price level, productivity, and employment. One way to motivate this relationship is from the standpoint of a union that sets a target for the real wage depending upon productivity but also varies its demand in response to the level of employment in the economy. The crucial assumption in their model is that the target wage is set using less-than-perfect information. It takes time and effort to collect and process information, so individuals update their decisions gradually over time. It is important to distinguish this approach from another closely related approach that is more common in the literature, which assumes that it is costly to change wages (or prices). Under the Mankiw-Reis assumption, firms are free to change prices every period. Thus, faced with an annual inflation rate of 4 percent, they could decide to change wages by 1 percent every quarter (or as often as they wished). However, if the inflation rate unexpectedly drops to 3 percent, firms will take a while to detect this change and to alter the rate by which they change prices every period.

Consider what happens in this model when there is an increase in the rate of growth of productivity. Under the authors' assumptions, it will take a while for the target wage to catch up to this increase. Higher productivity implies higher output even if employment were to remain unchanged. However, employment rises as well because firms find that they can get more output per worker but do not have to pay workers any more than before. Prices fall as firms pass on the lower costs to customers. Of course, the target wage does adjust at some point. Note that this wage is subject to two influences: the increase in productivity will push it up, while the fall in prices will push it down. Under certain conditions, it is possible for these two influences to offset exactly so that the target wage stays at the level needed to maintain full employment.

While these papers focus on the effects of a slowly adjusting reservation wage, Pissarides (2000) emphasizes a different mechanism, one that causes the equilibrium unemployment rate to decline in response to an increase in the rate of technical change. He points out that a firm's decision to hire a worker involves balancing the costs of hiring that worker against the profits that will accrue once the worker is hired. The hiring costs are incurred now while the profits are realized over time. Other things equal, an in-

3. The ultimate aim of the Ball and Moffitt paper is to use the change in productivity growth to explain why inflation stayed low in the late 1990s even though the unemployment rate fell to unusually low levels. See Grubb, et al. (1982) for a similar explanation of the stagflation of the late 1970s.

crease in the trend rate of growth raises future profits and makes it attractive to increase hiring today. Thus, an increase in the trend growth rate will lead to a decrease in unemployment, while a decrease in the trend growth rate will lead to an increase in unemployment. As the increase in future profits implies an increase in the present value of the job, this is known as the "capitalization effect."

Aghion and Howitt (1998) point out that technological progress tends to destroy old jobs at the same time that it creates new ones. Thus, it creates the need for workers to move across jobs. An increase in the rate of technological progress leads to an increase in the pace at which worker skills as well as jobs embodying specific technologies become obsolete, which leads to an increase in the rate at which worker-firm matches are broken up. As a consequence, frictional unemployment goes up, in contrast to Pissarides.

Mortensen and Pissarides (1998) show that either of these results can be obtained depending upon what one assumes about the cost of adopting new technology. Firms are assumed to lock in the existing technology when they create a new job. Because of technical progress, the technology embodied in a particular job becomes obsolete over time. The firm then must choose whether to spend the money to update the technology in the existing job (which may involve retraining the worker) or to destroy the job. If updating costs are prohibitively high, the firm will choose to destroy the job; in this case, faster technical progress (which makes existing capital obsolete faster) leads to greater job destruction. Note that because job creation and destruction depend upon job updating costs, which are likely to vary by firm and by industry, the model does not provide an unambiguous prediction about the relationship between economy-wide productivity growth and unemployment in the data.

One way an improvement in technology would lead to an unambiguous reduction in the equilibrium unemployment rate is if it led to a permanent increase in the rate at which searching firms and workers "find" the right match. This is exactly what Gomme (1998) suggests that the Internet has done. Firms now routinely post vacancies on the Internet so that workers can look for jobs in multiple (perhaps remote) locations at almost no cost. Saving (2000) notes that several million resumes are now online and that the Internet is available to roughly half the U.S. population. At this point in time, though, the Internet has not been around long enough to allow economists to measure the magnitude of this effect.

Some other authors have focused on the effect of anticipated changes in technology on the unemployment rate. In Phelps (1999) and Phelps and Zoega (2001) an anticipated improvement in technology tends to lower the unemployment rate temporarily. In earlier work (see Phelps 1994), Phelps has argued that the value of firms' assets (which include physical capital as well as employees and customers) is the proximate force driving the demand for labor. In more recent work, he discusses how changes in productivity can affect firm valuations. After defining a structural boom as a wholly, or largely, temporary expansion in employment (which is not caused by aggregate demand), the argument proceeds as follows. The starting point of the boom is "the unanticipated arrival of the *prospect* of new opportunities for profitable use of capital beginning at some point in the medium-term future, perhaps several years ahead" (Phelps and Zoega 2001, p. 93). Entrepreneurs realize that the jump in productivity will increase the return on their assets in the future, which, in turn, raises the value of their assets today without raising the cost of acquiring them. Thus, firms engage in anticipatory investment including the hiring of more workers. When, and if, the anticipated productivity jump materializes, the cost of investment in workers and equipment goes up (for instance, the value of labor in other uses is raised) and investment and employment fall off. Thus, the news of an increase in productivity sometime in the future leads to a boom in the economy, which dissipates (perhaps gradually) once the productivity increase materializes. Since the stock market reflects changes in the valuation of the firm (though not, perhaps, the change in the valuation of a job, which is the key variable here), the empirical analysis the authors conduct focuses on the relationship between the unemployment rate and the stock market.

Specifically, Phelps and Zoega present results from a multicountry study of the relationship between unemployment and lagged values of the stock market. (Recall that the stock market is hypothesized to go up following news of a future increase in productivity.) Using data over the 1960–1999 period for the G-7 (excluding Japan) as well as six other OECD countries, they show that there is a negative relationship between the unemployment rate and share prices which are normalized by productivity.

By contrast, Manuelli (2000) argues that an anticipated improvement in technology is likely to lead to a long-lived (but not permanent) increase in the unemployment rate. In his model, an anticipated (but not yet realized) improvement in technology reduces the market value of existing firms, which causes firms to cut back on investment and job creation.[4] New firms are unwilling to enter at this time as well, because doing so would mean that they would have to adopt a technology that soon would become obso-

4. See Greenwood and Yorukoglu (1997) for an early statement of the hypothesis that the productivity slowdown in the 1970s and the acceleration during the 1990s were part of the same phenomenon.

lete. Thus, investment and job creation go down and the unemployment rate goes up. Once the new technology becomes available, firms begin to increase investment and create more jobs, causing the unemployment rate to fall. Manuelli argues that stock markets fell and unemployment rose in the mid-1970s partly because markets realized that new technologies were coming that would make existing ones obsolete. These new technologies (relating to computers and information technology) began to mature sometime in the 1980s, causing unemployment to fall and productivity to rise over time. Thus, his paper links both the increase in unemployment in the 1970s and the decrease in unemployment in the 1990s to an exogenous change in technology. His model does not predict a productivity slowdown in the 1970s, though others have proposed similar models that do.

Galí (1999) also presents a model in which positive productivity shocks lead to temporary declines in worker hours (and, by implication, in employment), although the negative effect here is much shorter lived than the one in Manuelli's paper. Prices are assumed to be sticky and aggregate demand depends upon the amount of real money balances in the economy. In addition, it is assumed that the monetary authority does not vary the money supply in response to technology shocks—so that aggregate demand is unchanged.[5] Since neither aggregate demand nor prices can be varied in response to the technology shock, firms respond by reducing the number of worker hours that they employ to produce the same amount of output as before. In the next period, when they are free to adjust prices, they reduce them; output and worker hours go up as a consequence. Galí also presents empirical evidence consistent with this hypothesis. Using a number of different specifications and data for the U.S. over the 1948–1994 period, he shows that a positive technology shock leads to an increase in productivity but a temporary decrease in worker hours.[6]

For our purposes the key question is whether the positive technology shock will cause unemployment to go up as well. In terms of the theory, the answer depends upon how workers are assumed to respond to the higher productivity. Wen (2001) shows that the positive income effect arising from a positive technology shock can lead to a decrease in labor supply (i.e., a withdrawal from the labor force) under fairly general conditions. A similar effect is present in one of the models in Francis and Ramey (2002). The net effect on unemployment then will depend upon the size of this effect relative to the decrease in labor demand emphasized by Galí. The issue of the empirical relationship is addressed below.

To sum up the discussion so far, positive productivity shocks tend to lower unemployment in the short run to the extent that the reservation wage tends to adjust slowly to changes in productivity (as argued by Ball and Moffitt as well as by Mankiw and Reis). But there are other channels in play as well, and here the effects are not so clear-cut. Phelps and Zoega argue that news of a future increase in productivity can cause investment and employment to increase and unemployment to fall in the short run; however, Manuelli argues the opposite. Galí argues that positive technology shocks lead to a temporary contraction in the demand for labor. The effect on the unemployment rate is not unambiguous; it could go up as a result but may not if labor supply contracts by more than demand, as pointed out by Wen. The long-term effects of changes in the rate of productivity growth appear to be even more ambiguous. As Mortensen and Pissarides point out, the net effect depends upon firm- or industry-specific variables such as the cost of updating a job. Of course, to the extent that technological change affects the process of search directly (as some have argued the Internet has done), predictions are easier to make: lower search frictions should lead to lower unemployment.

## 4. Models with Heterogeneous Labor

The models we have looked at in the previous section have involved a fairly high level of aggregation and, in particular, have ignored heterogeneity across the labor force. In models with different kinds of labor, technology (and other) shocks may be more likely to raise unemployment in the short run than in single sector models, often because in these models both firms and households must search harder to make the right match.

Acemoglu (1998) presents a model in which an increase in technology or the skill level of the labor force "splits" an economy with homogenous jobs into one with jobs that require different levels of skills and pay different levels of wages. In his model, firms make irreversible decisions about capacity (that is, about the amount of physical capital) before hiring labor; these decisions are meant to repre-

5. The assumption about monetary policy is not innocuous. Dotsey (1999) shows that if the monetary authority is assumed to follow the rule estimated in Clarida, et al. (2000), positive technology shocks lead to an increase in both productivity and employment in the same model in which Galí's assumption of exogenous money supply leads to a negative correlation between the two variables.

6. Galí's analysis has led to considerable debate about whether technological shocks do, in fact, reduce labor input. Basu, et al. (1999) show that technology shocks reduce input use and, in particular, labor hours upon impact, using annual data and a methodology that is very different from that used by Galí. By contrast, Shea (1998) finds that technology shocks tend to increase input use in the short run and lower it in the long run. Using Galí's identification condition, Francis and Ramey (2002) find that technology shocks do reduce labor hours, while Altig, et al. (2002) find that they do not.

sent the choices firms make about the type of business to run and the types of jobs to create. There are two kinds of workers: high-skilled and low-skilled. The search process is random; for instance, low-capacity and high-capacity firms are equally likely to meet high-skilled workers. Upon meeting, firm-worker pairs decide whether to enter into an employment agreement or to continue the search for a new partner.

Different types of equilibria are possible in this model. In a "pooling" equilibrium both high-skilled and low-skilled workers do the same kind of job (in the model this means that they work with the same amount of capital). It is profitable for the firm to offer the same kind of job to all workers if the productivity differential between the two kinds of labor is small. In this equilibrium, low-skilled workers work with a greater amount of capital relative to the level of their skills than do high-skilled workers, and wage differentials are small. In a "separating" equilibrium firms create separate jobs for high-skilled and low-skilled workers. High-skilled workers work with more capital than low-skilled workers, so the ratios of physical capital to human capital in the two kinds of jobs are equalized. High-skilled workers earn more than they would in the pooling equilibrium, while low-skilled workers earn less. Thus, the dispersion of wages is higher under the separating equilibrium. Unemployment is higher as well, because high-skilled workers prefer to keep searching until they find the job that is right for them while firms that have a large amount of physical capital refuse to hire low-skilled workers. There also exist equilibria in which the unemployment rates of low-skilled workers go up by more than those of high-skilled workers.

If technical progress is skill biased (which means that it raises the productivity of high-skilled workers relative to that of low-skilled workers), it can push the economy from the pooling equilibrium to the separating equilibrium. In other words, technical progress can be accompanied by rising wage dispersion and rising unemployment.[7] According to Acemoglu, his model provides an explanation for the labor market patterns observed in the U.S. in the 1970s and 1980s, which include falling wages for low-skilled workers, rising unemployment rates for all workers, and a change in the composition of jobs in the economy.[8]

7. An increase in the size of the high-skilled labor force can have the same effect, as it can make it more profitable for firms to create different kinds of jobs.

8. A related literature looks at how technical change may have affected the wages of high-skilled workers relative to those of low-skilled workers; see, for instance, Greenwood and Yorukoglu (1997) and Krusell, et al. (2000). These papers say little, if anything, about the unemployment rate.

Blanchard and Katz (1997) discuss another reason why technical progress may raise the unemployment rate in a world with heterogeneous labor. They postulate that low-skilled workers are paid a wage that is very close to their reservation wage, while high-skilled workers are paid a wage that is substantially above their reservation wage. Consider, now, what would happen if there were a shift in technology that raised the demand for high-skilled workers but reduced the demand for low-skilled workers. The supply curve for low-skilled labor is relatively flat, which means that the inward shift of the demand curve will lead to a relatively large increase in the unemployment rate for those workers. By contrast, the supply curve for high-skilled labor is relatively steep, which means that the increase in demand will lead to a relatively small decrease in the unemployment rate for those workers. As a consequence, the overall unemployment rate will go up.

Baumol and Wolff (1998) focus on the costs of learning and argue that an increase in the rate of technical progress implies that workers will have to be retrained and plants will have to be retooled more often. Workers will be unemployed while plants are being retooled. Further, because older workers are harder to train than younger ones, the hypothesized adverse effects of faster technical progress should be more evident in older cohorts.

Using multivariate regressions, the authors show that the mean duration of unemployment is positively related to productivity growth over the previous five years. In addition to a measure of productivity (and other variables), their regressions also include investment in office, computing, and accounting equipment. This variable is positive and statistically significant in all regressions as well. In other words, the mean duration of unemployment goes up when investment in high-tech equipment goes up. They go on to repeat their regressions by age group and gender and show that the coefficients associated with both productivity growth and investment in high-tech equipment become larger as the age of the worker increases. The effects of these variables do not depend upon the gender of the worker. Overall, then, their evidence is consistent with the hypothesis that older workers will be more (adversely) affected by technical progress because they are harder to retrain.[9]

Mincer and Danninger (2000) conduct some tests on the effects of technological change on unemployment that explicitly account for differences in worker skill levels. Using annual data over the 1970–1995 period, they show

9. Changes in the relevant surveys over time have led to changes in the measures of duration used by the authors. See Abraham and Shimer (2001) for a discussion of these changes, as well as for a discussion of other factors affecting duration.

that technical progress leads to a contemporaneous increase in the unemployment rate of low-skilled workers *relative* to that of high-skilled workers.[10] This finding is robust to the use of four alternate measures of technical change: total factor productivity growth, the number of computers per worker, research and development expenditures per worker, and computers as a fraction of total capital equipment. Interestingly, they find that this effect is reversed after five years; according to them, the reversal represents a labor supply response to the initial change in the relative unemployment rates.[11]

Turning to the aggregate unemployment rate, Mincer and Danninger find that an increase in any of their four technology variables leads to a synchronous decrease in unemployment, though the effect is not significant for two of the four cases. The negative effects turn out to be stronger in the long run (after five years), when three of the four variables they consider become significant.

The theoretical research reviewed in this section suggests that unemployment should worsen after a positive technology shock, and also that less-skilled workers should suffer more than high-skilled workers. The empirical findings are more mixed. Of the two studies discussed above, one finds that the aggregate unemployment rate falls in response to positive technology shocks. The results from both studies suggest that workers who find it harder to learn new skills will be hit harder by technological change, though the two studies tend to classify workers in different ways.

## 5. Empirical Estimates of the Effects of Technology Shocks on Unemployment

In the discussion above, we have seen that different papers emphasize different channels through which technology shocks can affect the unemployment rate, and these channels lead to different predictions about how the unemployment rate will respond. Given these often contradictory predictions, it is natural to ask which one dominates in the data. As we also have seen above, the limited empirical research available on this issue does not provide a clear-cut answer either. Accordingly, this section presents some new evidence on this issue using data on U.S. unemployment and productivity and some related variables. The evidence presented here will be obtained by estimating models in the spirit of Galí (1999), who identifies technology shocks by assuming that these are the only shocks that can have a permanent effect on the level of productivity. Before doing so, it is worth mentioning one other antecedent to the empirical analysis below. Blanchard and Quah (1989) estimate a two-variable model containing output and the unemployment rate in which the underlying structure is identified by assuming that certain shocks do not have a permanent effect on the economy. These shocks are labeled "demand shocks," while shocks that have a permanent effect on the level of output are labeled "supply shocks." Blanchard and Quah show that while a positive supply shock raises output permanently it also raises unemployment in the short run. (Demand shocks raise output and lower unemployment in the short run.) This result appears to echo Galí's, though it is useful to keep in mind that the Blanchard-Quah assumption does not really distinguish between different kinds of shocks that have a permanent effect on output. A permanent increase in labor supply, for instance, would also increase output permanently and could lead to a temporary increase in unemployment.

The measure of productivity included in the model is obtained by deflating output by total labor hours, as in Galí (1999). The sample period extends from 1959:Q1 to 2001:Q4, with the starting date dictated by the availability of the output data. Before estimating the model it is important to determine whether the two variables to be included are stationary. It turns out that one cannot reject the null that the productivity process contains a unit root (after taking logs), but one can easily reject the null that the first difference of this process contains a unit root. In the case of the log of the unemployment rate, the augmented Dickey-Fuller test statistic is –2.87, which in absolute terms is slightly smaller than the 5 percent significance level of –2.89.[12] While one cannot reject the null of a unit root in this case, it turns out that one also cannot reject the null that the unemployment rate is stationary. Specifically, the KPSS test yields a value of 0.16, which is well below the 5 percent critical value of 0.46.[13]

Based on these results, the analysis below will be carried out under the assumption that the unemployment rate is stationary. Note that this assumption is not innocuous since it automatically rules out any permanent effect of a productivity shock on the unemployment rate (such as that associated with the Internet, for example). Because of that, it is worth discussing what alternative assumptions

10. Education levels are used to proxy for skill levels; see the discussion below.

11. This interpretation is consistent with the analysis of Murphy and Topel (1997), who show that low-skilled workers have reacted to the prolonged decline in their wages by gradually withdrawing from the labor force. This allows the unemployment rate for these groups to come back down, even though employment does not recover fully.

12. See Maddala and Kim (1998), pp. 74–76, for a discussion of the tests and tables containing critical values.

13. See Maddala and Kim (1998), pp. 120–122, for a description of the test.

might imply. One possibility is to assume that the unemployment rate contains a unit root. A look at the data (in Figure 1) suggests that the unemployment rate does possess some of the characteristics of a unit root process. Like a unit root process, it shows a lot of persistence. It was relatively low in the 1960s and early 1970s but then seemed to be on an upward trend over the next decade or so. The unemployment rate tended to remain around 6 percent from the mid-1980s to the mid-1990s but then declined dramatically during the second half of the 1990s. Yet the assumption of a unit root appears too extreme; it implies, for instance, that the unemployment rate could move anywhere, given enough time. A perhaps more reasonable assumption is that the unemployment rate is mean-reverting, but that this mean shifts over time. For instance, it is possible that the mean of the unemployment process shifted up during the 1970s and then down again in the 1990s. Tests developed by Bai and Perron (1998) were used to examine this possibility. These tests allow for one or more shifts in the mean of a series without the user having to specify the date or dates at which these shifts may have occurred. Neither of the two tests recommended by them rejects the null of no breaks in the series, though one of the other tests finds evidence of five breaks (which is the maximum number of breaks that was allowed in the test). These findings may reflect the fact that the unemployment rate is a very persistent series, a situation in which the tests may have trouble detecting breaks according to Bai and Perron (2000).

Given these results, it seems best to stay with the assumption that the unemployment rate is stationary. As mentioned above, output per hour appears to contain a unit root, so it is differenced before being included in the vector autoregression (VAR). The identification condition employed here is taken from Galí (1999), who in turn employs a version of the long-run restriction developed by Blanchard and Quah (1989). Specifically, the assumption is that "only technology shocks can have a permanent effect on the level of labor productivity" (Galí 1999, p. 256). Note that under this identification scheme even though the second shock to the VAR is nominally identified, it really serves as a catchall since it contains some linear combination of shocks to labor supply, monetary policy, etc.

Figure 2 shows how the unemployment rate and productivity react to the productivity shock in the two-variable system. A positive, permanent shock to productivity lowers the unemployment rate on impact and continues to push it down further for about a year and a half (Panel A); thereafter, the unemployment rate returns gradually to its original level. While the effect is not statistically different from zero in the first few quarters, it does become significant for about three years after that.[14] Panel B shows that output per hour ultimately ends up a little bit higher than it was upon the initial impact of the productivity shock.

Before going further, it is useful to examine how sensitive this result is to changes in model specification. Consider, first, what happens when we change the identification assumption used here. Panel A of Figure 3 shows how the unemployment rate reacts to a productivity shock in a system in which identification is achieved by means of a Choleski decomposition. Specifically, we assume that productivity shocks have a contemporaneous effect on the unemployment rate but that shocks to the unemployment rate do not have a contemporaneous effect on productivity. The impulse response function is quite similar to that shown in Figure 2. In particular, a positive productivity shock tends to push the unemployment rate down for a while, and the estimated effect is statistically significant for the first few years.

Second we ask what happens to the estimated response when additional variables are included in the VAR. To answer this question we add the three variables used by Galí in the five-variable version of his VAR; the additional variables are the real interest rate, the quantity of real balances, and the inflation rate. The technology shock is identified using Galí's identification restriction (just as in Figure 2), and the effect of this shock on the unemployment rate is shown in Panel B of Figure 3. The unemployment rate actually rises a little bit in the first period in this system but falls below its original level one quarter later. The initial response is not really distinguishable from zero, however, and it is only around the two-year mark that the decline in the unemployment rate becomes statistically significant.

Overall, the results in Figures 2 and 3 suggest that the unemployment rate may fall a little bit immediately after a technology shock hits the economy, but that this effect is not very large. However, roughly a year after a positive technology shock the unemployment rate is lower than it was when the shock hit, and it stays significantly lower than the initial value for three to four years after that. Thus, this evidence offers little support for theories implying that a positive technology shock raises the unemployment rate and instead tends to favor theories that predict that the unemployment rate will decline in response to a positive technology shock.

How important are technology shocks for the behavior of the unemployment rate on average? To answer this question, Table 1 shows the variance decomposition of the unemployment rate associated with the impulse responses shown in Figure 2 (which, in turn, are from a two-variable

14. The 95 percent confidence bands shown here and in all figures below are based on 500 simulations.

Figure 2
Effect of Technology Shocks: 1959:Q1–2001:Q4

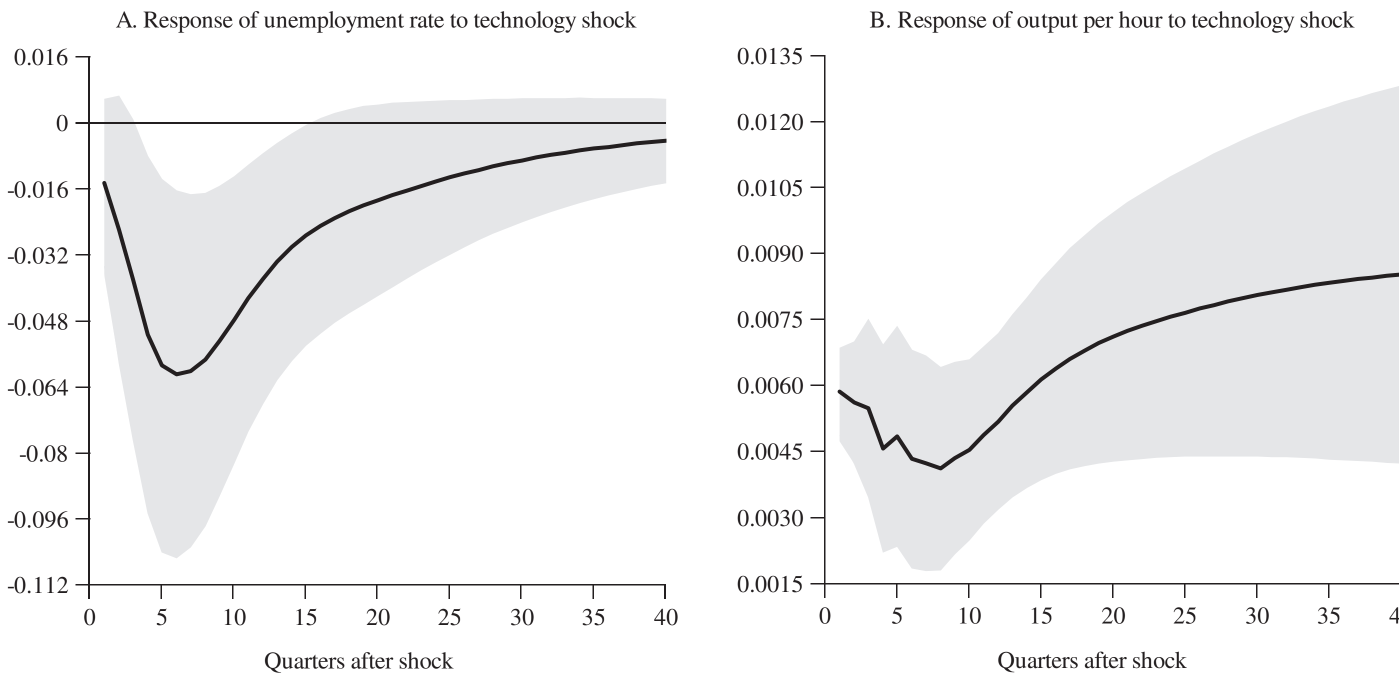


Note: 95 percent confidence bands are based on 500 simulations.

Figure 3
Effect of Technology Shocks on Unemployment under Alternative Specifications

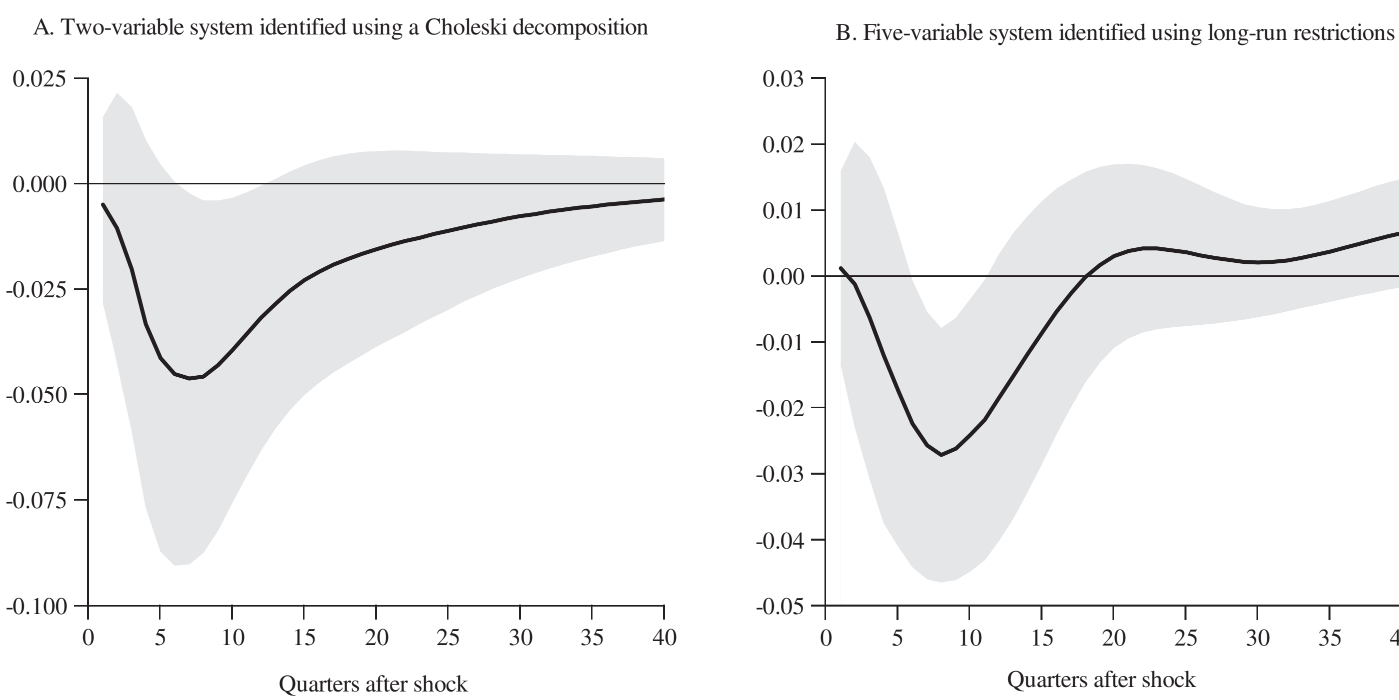


Note: 95 percent confidence bands are based on 500 simulations.

system that contains the unemployment rate and output per hour). Technology shocks do not have much of an impact on the unemployment rate initially, but their importance grows over time; they explain about a quarter of the forecast-error variance of the unemployment rate two years out and about one-third three years out. In the long run, their share of the variance settles to a little bit under 40 percent. For comparison purposes, it is worth noting that in the five-variable system the technology shock accounts for less than 1 percent of the variance of the unemployment rate after a year but for over 50 percent in the long run.

It is also interesting to compare the results in Figure 2 and Table 1 with the results for the Blanchard and Quah two-variable system, which contains real output and the unemployment rate. A key difference is that in their system (both in the original paper and when estimated over the sample period of this paper) supply shocks raise the unemployment rate for a relatively long time, while here technology shocks tend to lower the unemployment rate. Recall that in the Blanchard-Quah system, supply shocks are identified as having permanent effects on the level of output (and not just on productivity). So it is possible that the supply shocks identified by them also include the effects of shifts in labor supply, such as the increases in the labor force that took place when the baby boomers began to enter the labor force in the 1970s. The results from the variance decompositions are also consistent with this interpretation, as the supply shocks obtained under their identification account for a larger proportion of the variance of unemployment than the technology shocks identified here do. In their system (estimated over the sample period of this paper) supply shocks account for 26 percent of the variance of unemployment after four quarters and for 43 percent after two years. In the long run, the share of supply shocks rises to a little more than 50 percent.

To see how the estimated technology shocks have influenced the evolution of the unemployment rate over time, we construct a measure of how the unemployment rate would have evolved if the only shock hitting the economy were the technology shock and compare this measure to the actual unemployment rate over the 1961–2001 period. The results from this exercise are shown in Figure 4. The figure shows that the technology shock-driven unemployment rate moves fairly closely with the actual unemployment rate. Almost all of the decline through the late 1960s is explained by the supply component, as is much of the increase over the 1970s. Even so, technology shocks do not explain the peaks in the unemployment rate during the mid-1970s and early 1980s. The role of the technology shocks during the 1990s is especially interesting; the actual unemployment rate remains noticeably below the component that can be explained by technology shocks over this period, even though technology shocks were exerting constant downward pressure on the unemployment rate.

Table 1
Variance Decomposition of the Unemployment Rate
(Sample: 1961:Q1 to 2001:Q4)

| Quarters Ahead | Technology Shocks (percent) | Other Shocks (percent) |
|---|---|---|
| 1 | 1.4 | 98.6 |
| 2 | 2.7 | 97.3 |
| 4 | 10.1 | 89.9 |
| 8 | 25.2 | 74.8 |
| 12 | 33.3 | 66.7 |
| 20 | 37.5 | 62.5 |
| 30 | 38.4 | 61.6 |
| 40 | 38.7 | 61.3 |

### 5.1. High-skilled versus Low-skilled Workers

As discussed above, a substantial body of research suggests that technology shocks should affect different kinds of labor in different ways. Blanchard and Katz (1997), for instance, argue that positive technology shocks lower the unemployment rate of high-skilled workers while raising that of low-skilled workers, while some of the arguments put forward by Acemoglu (1998) suggest that the

Figure 4
Actual and Simulated Unemployment Rate

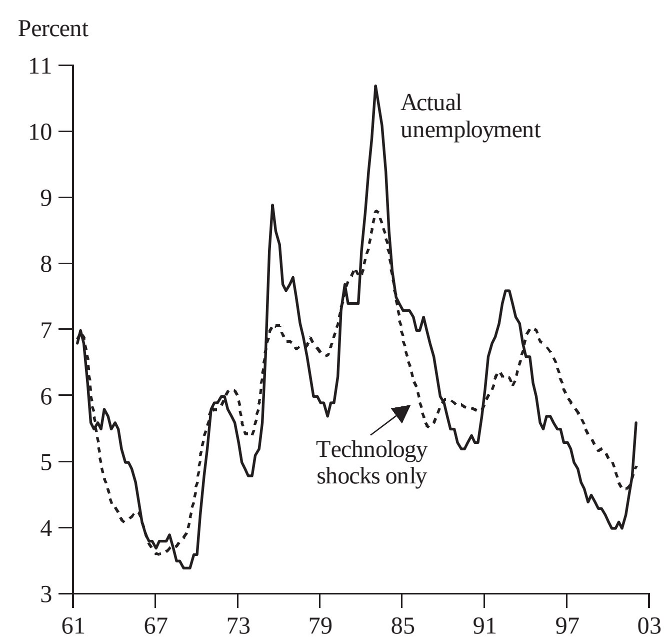

unemployment rate of low-skilled workers should go up more than that of high-skilled workers. This section attempts to determine what kind of evidence there is for such hypotheses.

Data on labor skills are extremely hard to get, particularly if one is looking for data on the entire economy over a reasonable length of time. The usual response is to use data on education as a proxy for labor skills, even though it is well understood that the two are not the same. It turns out that there are severe limitations on the availability of unemployment data by education levels as well. At a monthly frequency, data on unemployment rates by education levels are available since 1993 for workers 25 years and older. This sample is clearly too short to undertake the kind of analysis we have performed in the previous section. One can get annual unemployment data for workers with four different levels of education over the 1970–2000 period. However, these data are not consistent over time. For one thing, in the early 1990s the categories were redefined to focus on completed degrees rather than years in school. Other changes in the survey used to collect data also make it difficult to compare data after 1993 with earlier data. These data limitations mean that it will be difficult to get conclusive evidence on the impact of technology shocks on the unemployment rate across different kinds of workers using either data set. Nevertheless, in order to provide some sense of what the relationship might be like, this section will present the results from some simple analysis of both kinds of data.

Figure 5 shows monthly data on unemployment rates by education level since 1992. Broadly speaking, all four unemployment rates have fallen over most of this period. Unemployment rates were high at the beginning because the economy was emerging from a recession; they fell over the 1990s as the economy experienced a long expansion and rose towards the end as another recession hit. Note that in relative terms the recent recession has hit the most educated the hardest. For those with college degrees, unemployment rates are close to where they were in the early 1990s; this is not the case for those with lower education levels.

How are these changes related to changes in productivity over this period? The available sample of monthly data is clearly too small to allow the estimation of VARs similar to those in the previous section; for one thing, the sample is too small to impose a long-run restriction. Consequently,

Figure 5
Unemployment Rates by Education Level

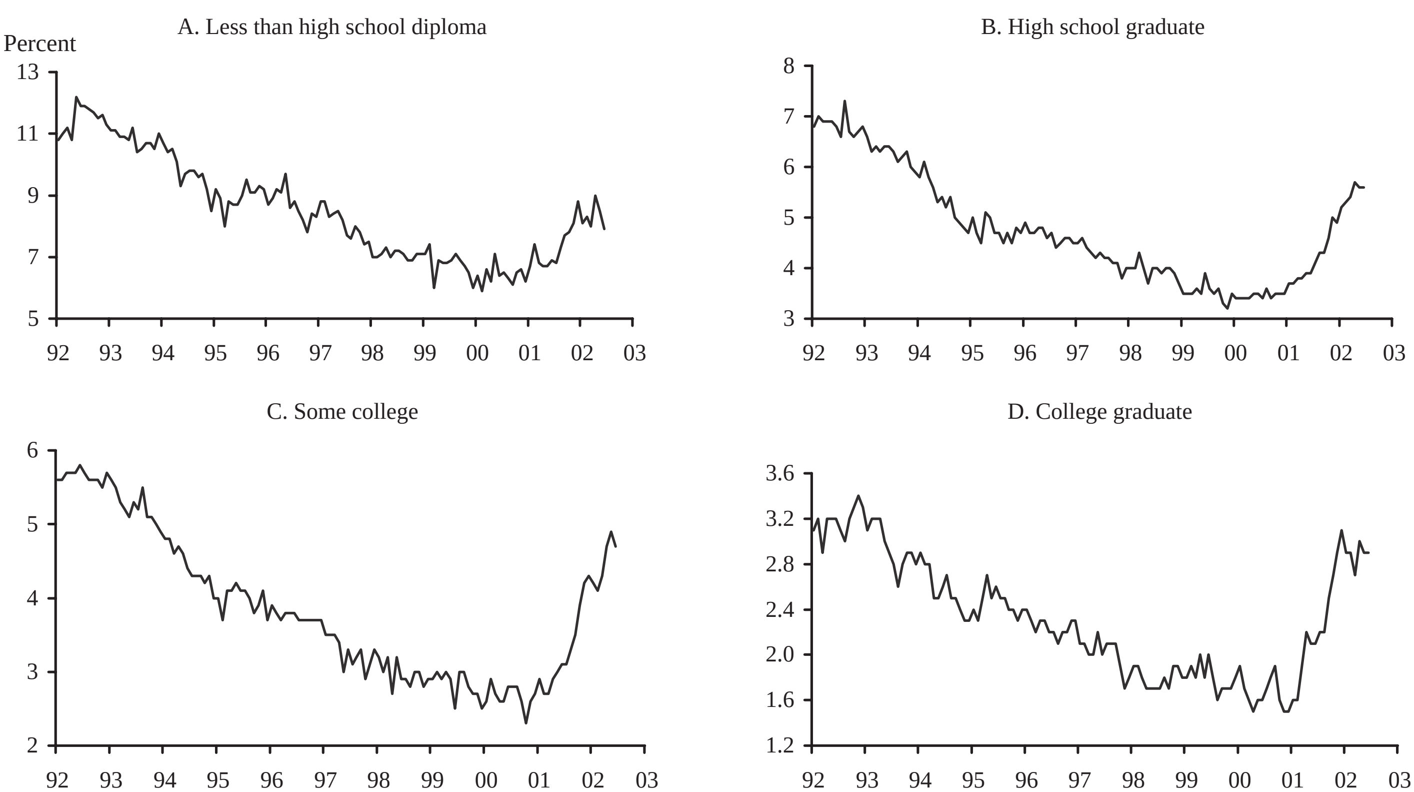

the analysis of the short sample here will take the technology shocks identified in the previous section as given and examine the effect that they have on the unemployment rates of different kinds of workers. An important benefit of this exercise, relative to the alternative of estimating separate VARs for each unemployment rate, is that the technology shock does not change across specifications.

For the first exercise, we simply regress each of the unemployment rates on its own lags and lags of the technology shock variable. To avoid the problems raised by the redesign of the survey, we estimate these regressions over the 1994:Q1–2001:Q4 period. From each of these regressions we can compute how the relevant unemployment rate will behave over time in response to a technology shock. The four resulting impulse response functions are shown in Figure 6. Note, first, that all four unemployment rates decline immediately in response to a positive technology shock and that the unemployment rates continue to fall further for at least another year afterward. The impulse response functions are also similar enough that there is little hope of distinguishing them statistically from each other, especially in light of the small sample size. Even so, it is worth noting that the unemployment rate for workers without a high school degree falls least in response to a positive technology shock, followed by the unemployment rate for workers who have finished high school; in addition, the unemployment rates for workers with the two highest levels of schooling fall roughly twice as much as those for workers who did not finish high school. Finally, note that the impulse response function for workers with some college damps out very slowly; more likely than not, this is a reflection of the small sample size.

Figure 6
Response of Unemployment Rate to Technology Shocks: Quarterly Data from 1994:Q1 to 2001:Q4

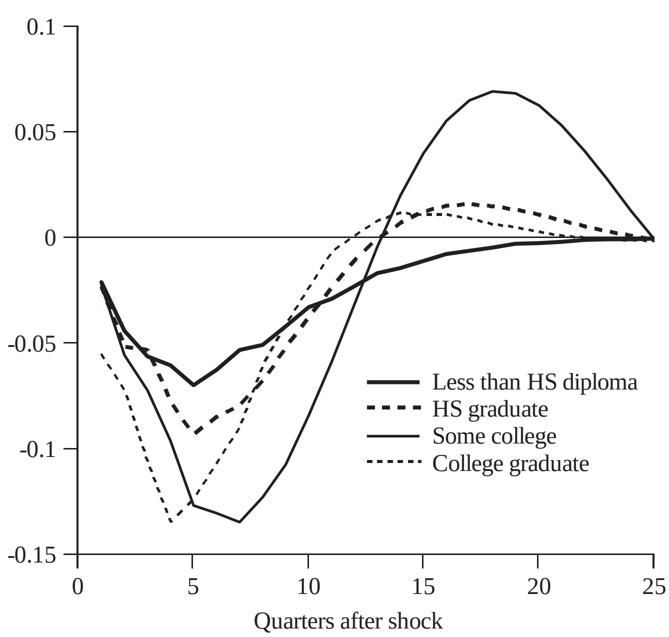


One way to judge whether the productivity shocks defined above affect different unemployment rates differently is to define a variable that measures the spread of the four different unemployment rates at each point in time and see if the technology shocks identified above make a difference to that variable. Variance-type measures come naturally to mind. One problem, however, is that such measures contain no information about the direction in which various unemployment rates change in response to a technology shock. For instance, a positive technology shock could lead to an increase in the dispersion of these four unemployment rates either because the unemployment rate of the workers with more education decreased more than that of workers with less education or because the former increased more than the latter. Consequently, we also construct a variable for the difference between the unemployment rate of the workers who have not finished high school and the unemployment rate of workers who have completed a college degree and see how this variable reacts to the technology shocks we have constructed. The results of these exercises are shown in Table 2.

The first column shows the results when the variance-based measure (called UDisperse) is regressed on the contemporaneous value of the difference of the log of real GDP (DGDP) and six quarterly lags of the technology shock variable from the exercise in Section 5 above. This specification was selected by starting with contemporaneous and lagged values of GDP and the technology shocks and eliminating insignificant lags, following the general-to-specific strategy recommended by Hendry (1995). Lagged values of UDisperse were included as well but were found to be insignificant. The results indicate that the dispersion of the unemployment rate goes up when real GDP growth picks up.[15] Dispersion also goes up in response to a technology shock, with some of this decrease being reversed in the long run.[16] The variables in the equation explain about half of the movement in the dispersion variable, and the chi-squared tests indicate that the technology shock variables are significant at the 1 percent level.

15. The standard errors shown in the table are robust to the existence of heteroskedasticity and autocorrelation.

16. In both equations shown here, lags 7 and 8 of the technology shock variable turn out to have negative coefficients when included in the equation; however, these lags are not significantly different from zero.

Table 2
Technology Shocks and the Dispersion of Unemployment (Sample 1994:Q1 to 2001:Q4)

| | UDisperse | URange |
|---|---|---|
| Constant | 0.20[1] | 1.23[1] |
| | (0.01) | (0.03) |
| $DGDP_t$ | 3.10[1] | 9.54[1] |
| | (0.82) | (2.59) |
| $Techshock_{t-1}$ | 0.96 | 2.92 |
| | (0.60) | (1.80) |
| $Techshock_{t-2}$ | 1.42[1] | 4.06[1] |
| | (0.49) | (1.47) |
| $Techshock_{t-3}$ | 1.11[5] | 3.37[10] |
| | (0.58) | (1.81) |
| $Techshock_{t-4}$ | 0.35 | 0.62 |
| | (0.37) | (1.09) |
| $Techshock_{t-5}$ | 1.06[5] | 2.59[5] |
| | (0.43) | (1.33) |
| $Techshock_{t-6}$ | –0.80[5] | –2.86[1] |
| | (0.33) | (1.05) |
| $R^2$ | 0.51 | 0.52 |
| $\chi^2(6)$[a] | 19.90[1] | 18.70[1] |

Note: Standard deviations are in parentheses.
[1] denotes significance at 1 percent.
[5] denotes significance at 5 percent.
[10] denotes significance at 10 percent.
[a] The null is that the techshock variable can be excluded from the equation.

The second column presents the results for the URange variable, which is defined as the difference between the unemployment rate of workers with the least education and the unemployment rate of workers with the most education. These results are similar to those in the first column. The advantage here is that one can see that (the increase in dispersion observed in the first column comes about because) the technology shock actually raises the unemployment rate of less-educated workers relative to those with more education. (The shortcoming of the second measure is that it ignores information about workers with intermediate skill levels.) Combined with the results in Figure 5, which shows us that the unemployment rates of both kinds of workers fall in response to a positive technology shock, the implication is that such a shock pushes down the unemployment rate of high-skilled workers by significantly more than it pushes down the unemployment rate of low-skilled workers.

As mentioned above, annual data on unemployment by education levels also are available over the 1970–2000 period. We carried out the same analysis using annual data with one important difference. Instead of using data on output per hour, we used data on multifactor productivity.[17] The advantage of using these data is that they are likely to be more closely related to technology shocks than are the data on average output per hour. These data were used, along with annual data on the aggregate unemployment rate, to derive a series of technology shocks, exactly as was done for the quarterly data.[18] Each of the four series of unemployment by education level then was regressed on the technology shock and its own lags. (Four lags were used.)

Figure 7 shows how each of the unemployment rates reacts to the technology shock. All four decline following a positive technology shock, though the initial response is positive for those who did not finish high school as well as for those who finished high school but did not attend college. The responses of the four unemployment rates are extremely close to each other and, given the small number of observations, there is not much point in trying to distinguish these responses from one another. An attempt was made to determine if these differences were statistically significant by constructing a spread variable similar to that for the quarterly data. However, the spread turns out to be

Figure 7
Response of Unemployment Rate to Technology Shocks: Annual Data from 1970 to 2001

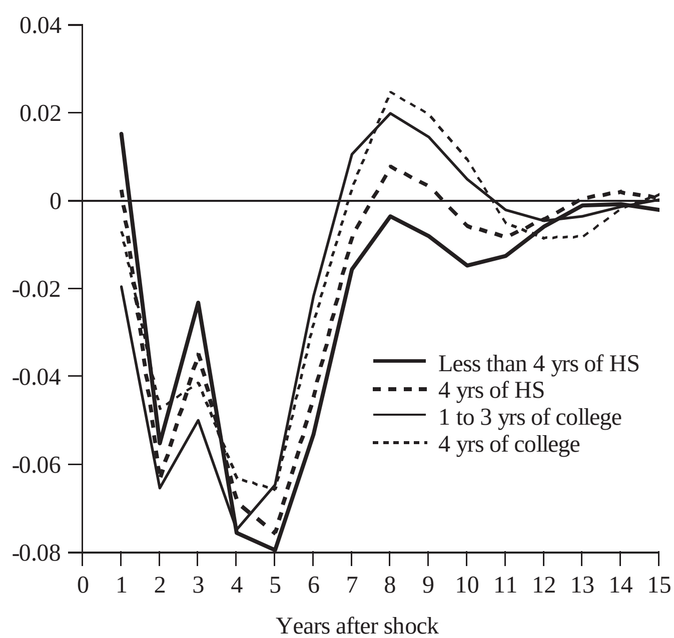


17. These data have been obtained from the Bureau of Labor Statistics website and are available only at an annual frequency.

18. Neither slope nor intercept dummies that were meant to control for the effects of the changes in surveys discussed above turned out to be significant in the equations estimated on annual data.

nonstationary,[19] and after one accounts for this nonstationarity it is hard to find any role for the technology shocks.

To sum up, disaggregating the data by education levels shows that positive technology shocks tend to lower the unemployment rates of workers across all skill levels. This result seems to be true whether we use quarterly data since the early 1990s or annual data since 1970. Quarterly data for the last eight years also suggest that the unemployment rates of high-skilled workers tend to fall by more than those of low-skilled workers. However, the small sample size argues against putting a lot of weight on this finding. Further, annual data do not provide evidence of a significant difference across categories (even though the point estimates are consistent with the quarterly results).

## 6. Conclusions

This paper has looked at some recent research on the effects of changes in productivity growth on the unemployment rate. Models that postulate that the reservation wage adjusts sluggishly to changes in productivity (and that assume homogeneity of labor and ignore the increase in job destruction that is likely to come about as a result of a higher pace of technical progress) make an unambiguous prediction: high productivity growth implies that unemployment falls. A more complex picture emerges as some of these restrictions are relaxed. To the extent that rapid technical change leads to more job destruction, it raises frictional unemployment as more workers and firms must spend time looking for the right match. This effect will be amplified if technical change increases heterogeneity, since each individual must spend more time on search as well. There is also some ambiguity about how firms respond to news about higher productivity (or new technology) in the future. Some models also suggest that productivity shocks are likely to have different effects on workers with different skill levels; generally speaking, workers with relatively low skill levels are not likely to do as well as workers with high skill levels.

There has been relatively little empirical research on these issues. The results from two of the studies discussed above suggest that technology shocks lead to lower unemployment, while another study finds that the duration of unemployment goes up in response to technology shocks. The empirical results in this paper are relatively unambiguous. Specifically, positive shocks tend to lower unemployment, with effects that build up over several years before damping out. This effect appears to be robust to a classification of workers by skill levels, in that the unemployment rate of each of four groups of workers (differentiated by the Bureau of Labor Statistics on the basis of education levels) declines in response to a positive technology shock. There is some evidence that the unemployment rates of highly educated workers decline by more than those of workers with lower education levels, though available sample sizes are too small to place a lot of weight upon this finding.

These findings are consistent with the predictions of models that emphasize sluggish adjustment of the reservation wage and with models that predict an economic boom when news about improved technology arrives. And they are certainly consistent with the boom observed in the second half of the 1990s, when a surge in productivity growth was accompanied by a sharp decline in the unemployment rate. This does not mean that models which stress job destruction and worker as well as job heterogeneity are wrong in some way, but empirically the effects working through these channels appear to be dominated by the positive effects of technology shocks on the unemployment rate.

19. The unemployment rate of workers who did not finish high school goes up over the 1970s and stays relatively high throughout the sample, while the unemployment rates of other workers tend to fall back.

## References


Abraham, Katharine, and Robert Shimer. 2001. "Changes in Unemployment Duration and Labor Force Attachment." NBER Working Paper 8513 (October).

Acemoglu, Daron. 1998. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." NBER Working Paper 6658 (July).

Aghion, Phillipe, and Peter Howitt. 1998. *Endogenous Growth Theory.* Cambridge, MA: MIT Press.

Akerlof, George, and Janet Yellen. 1990. "The Fair Wage-Effort Hypothesis and Unemployment." *Quarterly Journal of Economics* 105(2) (May) pp. 255–283.

Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde. 2002. "Technological Shocks and Aggregate Fluctuations." Mimeo (June).

Bai, Jushan, and Pierre Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica* 66(1) (January) pp. 47–78.

Bai, Jushan, and Pierre Perron. 2000. "Computation and Analysis of Multiple Structural Change Models." Mimeo.

Ball, Laurence, and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives* 16(4) (Fall) pp. 115–136.

Ball, Laurence, and Robert Moffitt. 2001. "Productivity Growth and the Phillips Curve." NBER Working Paper 8421.

Basu, Susanto, John Fernald, and Miles Kimball. 1999. "Are Technology Improvements Contractionary?" Mimeo, University of Michigan.

Baumol, William J., and Edward N. Wolff. 1998. "Speed of Technical Progress and Length of the Average Interjob Period." Jerome Levy Institute Working Paper No. 237 (May).

Blanchard, Olivier, and Lawrence Katz. 1997. "What We Know and Do Not Know about the Natural Rate of Unemployment." *Journal of Economic Perspectives* 11(1) (Winter) pp. 51–72.

Blanchard, Olivier, and Danny Quah. 1989. "The Dynamic Effects of Aggregate Demand and Supply Disturbances." *American Economic Review* 79(4) (September) pp. 654–673.

Clarida, Richard, Jordi Galí, and Mark Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115(1) (February), pp. 147–180.

Dotsey, Michael. 1999. "Structure from Shocks," FRB Richmond Working Paper No. 99-6. http://www.rich.frb.org/pubs/wpapers/pdfs/wp99-6.pdf

Francis, Neville, and Valerie Ramey. 2002. "Is the Technology-Driven Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited." NBER Working Paper 8726 (January).

Galí, Jordi. 1999. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review* 89(1) (March) pp. 249–271.

Gomme, Paul. 1998. "What Labor Market Theory Tells Us about the 'New Economy.'" FRB Cleveland *Economic Review* QIII, pp. 16–24. http://www.clevelandfed.org/Research/review98/gomme.pdf

Greenwood, Jeremy, and Mehmet Yorukoglu. 1997. "1974." *Carnegie-Rochester Conference Series on Public Policy* 46, pp. 49–95.

Grubb, D., R. Jackman, and R. Layard. 1982. "Causes of the Current Stagflation." *Review of Economic Studies* 49, pp. 707–730.

Hendry, David. 1995. *"Dynamic Econometrics."* Oxford: Oxford University Press, pp. 270–271.

Krusell, Per, Lee Ohanian, Jose-Victor Rios-Rull, and Giovanni Violante. 2000. "Capital Skill Complementarity and Inequality: A Macroeconomic Analysis." *Econometrica* 68(5) (January) pp. 1,029–1,053.

Maddala, G.S., and In-Moo Kim. 1998. *Unit Roots, Cointegration, and Structural Change.* Cambridge: Cambridge University Press.

Mankiw, Gregory. 1992. *Macroeconomics*. New York: Worth Publishers.

Mankiw, N. Gregory, and Ricardo Reis. 2001. "Sticky Information: A Model of Monetary Nonneutrality and Structural Slumps." NBER Working Paper 8614.

Manuelli, Rodolfo E. 2000. "Technological Change, the Labor Market, and the Stock Market." NBER Working Paper 8022 (November).

Mincer, Jacob, and Stephan Danninger. 2000. "Technology, Unemployment, and Inflation." NBER Working Paper 7817 (July).

Mortensen, Dale T., and Christopher A. Pissarides. 1998. "Technological Progress, Job Creation, and Job Destruction." *Review of Economic Dynamics* 1(4) pp. 733–753.

Murphy, Kevin, and Robert H. Topel. 1997. "Unemployment and Nonemployment." *American Economic Review* 87(2) (May) pp. 295–300.

Phelps, Edmund S. 1994. *Structural Slumps: The Modern Equilibrium Theory of Unemployment, Interest, and Assets*. Cambridge, MA: Harvard University Press.

Phelps, Edmund S. 1999. "Behind This Structural Boom: The Role of Asset Valuations." *American Economic Review* 89(2) (May) pp. 63–68.

Phelps, Edmund S., and Gylfi Zoega. 2001. "Structural Booms." *Economic Policy* (April) pp. 83–126.

Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory.* Cambridge, MA: MIT Press.

Saving, Jason L. 2000. "The Effect of Welfare Reform and Technological Change on Unemployment." FRB Dallas *Economic and Financial Review* QII, pp. 26–34. http://www.dallasfed.org/htm/pubs/pdfs/efr/efr0002c.pdf

Shea, John. 1998. "What Do Technology Shocks Do?" *NBER Macroeconomics Annual,* pp. 275–309.

Stock, James, and Mark Watson. 2001. "Prices, Wages, and the U.S. NAIRU in the 1990s." NBER Working Paper 8320 (July).

Wen, Yi. 2001. "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations? Comment." Cornell University Department of Economics Working Paper 01-19.