Explaining Unemployment: Sectoral vs. Aggregate Shocks

Prakash Loungani
and Bharat Trehan

Federal Reserve Board and Federal Reserve Bank of San Francisco, respectively. The authors would like to thank Chan Huh, Mark Levonian, Rob Valletta and Joyce Zickler for helpful comments. We are also grateful to workshop participants at the Council of Economic Advisors for many constructive suggestions. Benjamin Bridgman provided excellent research assistance.

We include a stock market-based measure of sectoral shocks in a small VAR to examine the role played by these shocks in explaining the behavior of the unemployment rate. Sectoral shocks explain a significant proportion of the variation in the unemployment rate—especially the long-duration unemployment rate—even though other kinds of shocks (such as shocks to monetary policy, defense expenditures, and oil prices) are allowed to affect the unemployment rate. A historical decomposition reveals that sectoral shocks were most important during the 1974–75 recession, and they explain only a modest part of the rise in unemployment over the 1990 recession.

"A leading question—perhaps *the* leading question—in macroeconomics since the publication in 1982 of David Lilien's paper, 'Sectoral Shifts and Cyclical Unemployment,' is whether sectoral, rather than aggregate, shocks are the key factor responsible for fluctuations in the unemployment rate."

Yellen (1989)

"In an average week, between 350,000 and 400,000 jobs are destroyed. On average, a bit more than that are created. The flow of workers out of the old jobs and into the new ones is not seamless. The period of transition between jobs depends on many factors, including . . . the match between skills possessed and those needed . . . A large pool of unemployed workers might exist in a particular region even if most labor markets are viewed as 'tight.'"

Lindsey (1996)

In a controversial paper, Lilien (1982) suggested that frictions associated with the reallocation of labor across sectors of the economy accounted for as much as half of all fluctuations in unemployment. Though Lilien's paper inspired a significant amount of follow-up work,[1] the debate over the relative importance of sectoral shifts and aggregate shocks in unemployment fluctuations remains unresolved. We revisit Lilien's hypothesis in this paper. We are motivated in part by the lack of agreement on what causes business cycles that has been highlighted in some recent work. For instance, after an exhaustive review of the evidence, Cochrane (1996) concludes that "we haven't found large identifiable exogenous shocks to account for the bulk of output fluctuations" (though he suggests that "oil plus reallocation" may be a promising avenue). It is also telling that at the 1993 American Economics Association session entitled "What caused the recession of 1990–91?" Hall (1993) considered the relative importance of eight possible causes of the recession suggested by contemporary macro theories, but concluded that "established models are unhelpful in understanding this recession, and probably most of its predecessors." The failure of traditional models suggests that the sectoral shifts hypothesis may deserve another look.

---

1. Two examples are Davis (1987) and Campbell and Kuttner (1996).

Another recent development that helps motivate our study is the fact that the average duration of unemployment has been surprisingly high recently; for instance, in 1994 the average duration was nearly 20 weeks, roughly the same level as in the 1981–82 recession. This increase in duration appears to be related to the growing importance of permanent job loss relative to temporary layoffs, a phenomenon which was highlighted by Perry and Schultze (1993) and Hall (1995).[2] As we discuss below, sectoral shocks are a plausible candidate for explaining these changes.

To conduct our investigation we follow a suggestion by Black (1987), who conjectured that periods of greater dispersion in stock returns should be followed by increases in unemployment. The reason is that the stock market dispersion measure gives an "early signal of shocks that affect sectors differently, and puts more weight on shocks that investors expect to be permanent" (Black 1995). This latter point is important because it is presumably permanent shocks that motivate reallocation of labor across industries, thus significantly raising unemployment.

Two previous studies, Loungani, Rush, and Tave (1990) and Brainard and Cutler (1993), have provided evidence in favor of Black's conjecture. This paper extends their work in a number of ways; in particular, we are more careful about the measurement of aggregate shocks in our model as well as the kinds of shocks we include. For instance, since many observers, such as Romer and Romer (1989), consider shifts in monetary policy as the dominant source of recessions, it is important to control adequately for such shocks when trying to judge the importance of sectoral shifts. Both studies mentioned above used unanticipated money growth as a measure of monetary policy; Brainard and Cutler used M2 growth, for example. However, it is not obvious that the broad monetary aggregates provide a good measure of policy. For example, over the period 1979 to 1982, M2 growth was relatively robust, even though this period is generally thought of as one of restrictive monetary policy. Using money growth may therefore give a misleading picture of the relative importance of monetary policy and sectoral shifts over this period. Our solution is to employ the funds rate, since a lot of recent work (such as Bernanke and Blinder 1992) suggests that innovations in the federal funds rate are a better indicator of the stance of monetary policy.

Second, in contrast to earlier studies, the system we estimate also includes real output. We believe that including real output is important for at least two reasons. For instance, "Okun's Law," which is a key component of Keynesian models, explains changes in unemployment in terms of the growth of output. More generally, inclusion of real output helps control for other shocks hitting the economy. Thus, in trying to determine how important sectoral shifts are likely to be in explaining unemployment, it seems desirable to account for the effects of changes in output.

Finally, our sample period extends to 1995, which allows us to attempt to explain the 1990–91 recession as well as the high duration of unemployment over the last few years.

The basic model we employ to estimate the relative importance of sectoral shifts and aggregate shocks is a Vector Autoregression (VAR) that contains the civilian unemployment rate[3] (plus other variables described below). We find that our measure of sectoral shifts accounts for roughly 30% of the fluctuations in the civilian unemployment rate at a horizon of three to five years. While this is not a small number, the funds rate appears to be even more important, accounting for roughly 40% to 50% of the fluctuations in the unemployment rate over this period.[4] To address issues concerning the average duration of unemployment, we also estimate VAR models for the long-duration unemployment rate (which is constructed using unemployment spells that are 27 weeks or more in length). The dispersion index plays a larger role in explaining long-duration unemployment than the funds rate does. At a three to five year horizon, for example, it accounts for something like 30% to 45% of the fluctuations in the long-duration unemployment rate, while the contribution of the funds rate is about 10% to 15% smaller.

The remainder of this paper is organized as follows. In Section I we motivate the empirical measure of sectoral shifts we use and present some evidence on how it performs relative to the measure introduced by Lilien (1982). In Section II, we add this measure to a standard macro VAR and examine how well we can explain movements in the aggregate unemployment rate, while in Section III we use our VAR to try to explain movements in long-duration unemployment. Section IV uses the VARs to examine the role played by various factors in the evolution of the unemployment rate over the 1971–1995 period, and Section V concludes.

2. Duration data are derived from the CPS survey, which was revised in 1994. According to Polivka and Miller (1995) ". . . the new methodology significantly increased the proportion of unemployed who had long spells of unemployment and significantly decreased the proportion of unemployed with spells of unemployment less than 5 weeks."

3. Figure 3 plots the behavior of the unemployment rate over the last 25 years, while the long-duration unemployment rate is shown in Figure 4.

4. These numbers are taken from the variance decompositions of the unemployment rate in a 5-variable VAR where the dispersion index is ordered last and the funds rate is placed in the middle.

I. Measuring Sectoral Shifts

Lilien (1982) and Black (1987) suggested that the amount of labor reallocation that an economy has to carry out can change significantly over time. Some periods may be marked by relatively homogeneous growth in labor demand across sectors, whereas others may be characterized by shifts in the composition of labor demand. While beneficial in the long run, the reallocation of labor in response to sectoral shifts imposes short-run costs in the form of increases in unemployment. The greater the divergence in the fortunes of different industries, the more resources must be moved, and the larger will be the resulting increase in unemployment.

While these ideas are fairly intuitive, constructing a satisfactory measure of sectoral shifts poses an empirical challenge for a couple of reasons. First, as stated by Barro (1986), shocks to the expected profitability of an industry can arrive from "many—mostly unobservable—disturbances to technology and preferences [that] motivate reallocations of resources across sectors." Second, Davis (1985) points out that "allocative disturbances *from any particular source* are likely to occur rather infrequently over available sample sizes," [italics ours] which makes it difficult to incorporate variables explicitly that capture the effects of sectoral shifts into an aggregate unemployment equation.

These considerations motivated Lilien's construction of a cross-industry *employment* dispersion index to proxy for the intersectoral flow of labor in response to allocative shocks. Many researchers, most notably Abraham and Katz (1986), have questioned Lilien's use of employment dispersion as a measure of labor reallocation. Their basic point is that movements in employment dispersion may simply be reflecting the well-known fact that the business cycle has non-neutral effects across industries. The increase in the dispersion of employment growth rates could reflect not increased labor reallocation, but simply the uneven impact of aggregate demand shocks on temporary layoffs in different industries. Under certain conditions—for instance, if cyclically responsive industries have low trend growth rates of employment—aggregate demand shocks also can lead to a positive correlation between the dispersion index and aggregate unemployment. Hence there is an observational equivalence between the predictions of the sectoral shifts hypothesis and the more traditional "aggregate demand hypothesis."

Loungani, Rush, and Tave (1990) and Brainard and Cutler (1993) attempt to circumvent these problems by constructing an index based on stock prices. Assuming that stock markets are efficient, so that shocks to the expected profitability of an industry are reflected in its stock market return, and assuming that these shocks are followed by changes in that industry's use of inputs such as labor, their hypothesis is that the *dispersion* of stock returns across industries can be used as a proxy for shocks to the desired allocation of labor, i.e., as a measure of sectoral shifts. For instance, the arrival of news regarding the relative profitability of industries is likely to be followed by an increase in stock price dispersion. It is likely that this news also will lead to a change in the output mix of the economy in the long run. This will necessitate a reallocation of resources, and the unemployment rate will rise as part of this process of reallocation of labor across sectors. Thus, an increase in stock price dispersion will be followed by an increase in the unemployment rate.

For this paper, we updated the index used in Loungani, Rush, and Tave. The basic data consist of indexes of industry stock prices, as reported in Standard and Poor's Compustat PDE file. There are 121 industries in all, and they provide comprehensive coverage of manufacturing as well as nonmanufacturing sectors of the economy.[5] The sectoral shifts index is defined as

$$Mismatch_t = \left[\sum_{i=1}^{n} W_i (R_{it} - R_t)^2\right]^{1/2}.$$

In the equation above, R_{it} is the growth rate of industry I's stock price index, R_t is the growth rate of the S&P500 (a composite index), and W_i is a weight based on the industry's share in total employment in 1978.[6] Hence, the sectoral shifts index can be interpreted as the weighted standard deviation of industry stock returns.

An advantage of the stock price dispersion measure relative to Lilien's measure is that unlike employment changes, stock prices respond more strongly to disturbances that are perceived to be permanent (or structural in nature) than to temporary disturbances (such as those caused by business cycle fluctuations). The industry stock price represents the present value of expected profits over a long horizon. The impact of innovations in industry profits on its stock price therefore will depend on how long the shocks are expected to persist. If the shocks are purely temporary, the innovations will have little impact on the present value of expected profits and, hence, will have little impact on industries' stock prices. On the other hand, if the shocks are fairly persistent, the innovations will have a significant impact on expected future profits and will lead to large changes in industries' stock prices. Furthermore, it is these

5. The Appendix provides details on the construction of the index.

6. As a check on the robustness of these results we also reestimated some of the VARs presented below using employment shares in 1995 as weights for the dispersion index. This did not lead to a noticeable change in our results.

sorts of persistent shocks that will cause productive resources, such as capital and labor, to be displaced from the adversely affected industries. Thus, a dispersion index constructed from industries' *stock prices* automatically assigns greater weight to permanent structural changes than to temporary cyclical shocks.[7] As a consequence, a dispersion measure based on stock prices is less likely than a measure based on employment to reflect aggregate demand disturbances that result in large swings in temporary layoffs.

It is not difficult to demonstrate this difference between the two measures. In Table 1, we present the results from two three-variable VARs. The first contains a dispersion measure based on the growth rate of employment across sectors, and the second contains a dispersion measure based on stock prices; both also contain the unemployment rate and the growth rate of real GDP. Eight lags of each variable are included in both systems. Note from Panel A that the employment-based dispersion measure is significant only at the 20% level in explaining unemployment and does not help predict output at all. Instead, output growth predicts employment dispersion. By contrast, the stock market-based dispersion measure helps predict unemployment and output (the latter at a 6% level of significance), but is not explained by either of these variables.

A comparison of the variance decompositions from these two systems, reported in Panel B, also sheds light on the properties of the two indexes. When ordered first, the employment-based dispersion measure explains 20% of the variance of unemployment at the 20-quarter horizon; this falls to 3% when it is ordered last. On the other hand, even when it is placed last, the stock market-based dispersion measure still explains 30% of the variance of the error in predicting the unemployment rate at a 20-quarter horizon. Thus, the stock market index does not appear to be subject to the Abraham and Katz criticism of Lilien's

7. Presumably, similar reasoning lies behind Toledo and Marquis's (1993) use of the dispersion in capital stock changes across industries as a proxy for allocative disturbances.

TABLE 1

A COMPARISON OF DISPERSION INDEXES

A. MARGINAL SIGNIFICANCE LEVELS

Based on employment	Y	U	ED	Based on stock prices	Y	U	SD
Y	.49	.01	.04	Y	.60	.01	.30
U	.02	.01	.73	U	.01	.01	.34
ED	.87	.20	.03	SD	.06	.01	.01
Adj. R^2	.13	.98	.30	Adj. R^2	.21	.98	.31

B. VARIANCE DECOMPOSITIONS: UNEMPLOYMENT RATE

Based on employment:

Ordering:	ED, Y, U			Y, U, ED		
Qtrs	ED	Y	U	ED	Y	U
0	13	36	51	0	44	56
4	17	57	26	1	68	32
8	14	65	21	0	75	25
12	17	62	22	1	72	27
20	20	58	22	3	69	29
40	21	57	22	4	68	29

Based on stock prices:

Ordering:	SD, Y, U			Y, U, SD		
Qtrs	SD	Y	U	SD	Y	U
0	14	27	59	0	34	66
4	6	49	45	2	53	45
8	23	43	34	8	52	40
12	42	31	27	22	39	39
20	53	25	22	30	33	36
40	54	24	22	34	30	36

NOTE: Y denotes output, U denotes unemployment, ED denotes employment dispersion, and SD is stock market dispersion. The variance decompositions may not add to 100 due to rounding errors.

measure. Accordingly, we now turn to a detailed analysis of the performance of the stock market index in a larger VAR.

II. Sectoral Shifts and the Aggregate Unemployment Rate

The Basic Model

The basic model we use will be a five-variable VAR. In addition to the stock market price dispersion index and unemployment, we include three other variables—real GDP, the federal funds rate, and the S&P500 index. As mentioned above, our intent is to look at the effect of changes in the dispersion index on unemployment after we control for variables that are commonly thought to affect unemployment. Thus, the funds rate is included as a measure of monetary policy (as in Bernanke and Blinder 1992, for instance). The inclusion of real GDP controls for the stage of the business cycle; it also means that our model allows for a version of "Okun's Law." The S&P500 index is included to rule out the possibility that the dispersion index explains unemployment because it is mimicking the behavior of the stock market.[8] Both the unemployment rate and the federal funds rate are entered in levels (the latter following Bernanke and Blinder), while GDP and the S&P500 index are entered in growth rates. In addition to the basic system, we will also discuss some results from VARs that contain a somewhat different set of variables; we have refrained from including those variables in our basic system in order to keep it to a reasonable size.

Panel A of Table 2 presents marginal significance levels for our estimated equations. It shows that the dispersion index helps predict unemployment even after we account

8. Brainard and Cutler (1993) present results from different systems that contain different combinations of money growth, the price of oil, and the stock market return, in addition to a stock market-based measure of dispersion. However, their measure of dispersion is not significant at even the 20% level, once the market return variable and lagged unemployment are included in the unemployment equation.

TABLE 2

RESULTS FROM A 5-VARIABLE VAR

A. MARGINAL SIGNIFICANCE LEVELS

	UNEMPLOYMENT	OUTPUT	FUNDS RATE	S&P500	DISPERSION
UNEMPLOYMENT	.01	.16	.01	.65	.42
OUTPUT	.01	.38	.30	.92	.49
FUNDS RATE	.01	.06	.01	.45	.56
S&P500	.01	.54	.30	.84	.28
DISPERSION	.01	.38	.02	.57	.06
ADJ. R^2	.99	.30	.92	–.01	.35

B. VARIANCE DECOMPOSITIONS[a]

	UNEMPLOYMENT					OUTPUT				
QRTRS	S&P500	OUTPUT	FUNDS	UNEMP.	DISPERSION	S&P500	OUTPUT	FUNDS	UNEMP.	DISPERSION
0	3	27	9	61	0	0	100	0	0	0
4	14	33	4	46	4	5	75	14	4	2
8	12	16	24	29	19	7	57	15	4	18
12	5	8	38	18	31	8	56	15	4	18
20	3	5	52	11	28	8	54	15	5	18
40	4	6	53	11	27	8	53	16	5	19

[a] Ordering is: S&P500, output, funds, unemployment, and dispersion.

for the stage of the business cycle, as measured by real GDP growth, and the stance of monetary policy, as measured by the federal funds rate. However, the dispersion index does not help predict output.

Figure 1 shows the responses of unemployment and output to shocks to the dispersion index, along with the associated standard error bands. For comparison purposes we also show the effect of shocks to the funds rate. To avoid

FIGURE 1

Dynamic Responses of Unemployment and Output

Dispersion Shocks	Federal Funds Rate Shocks

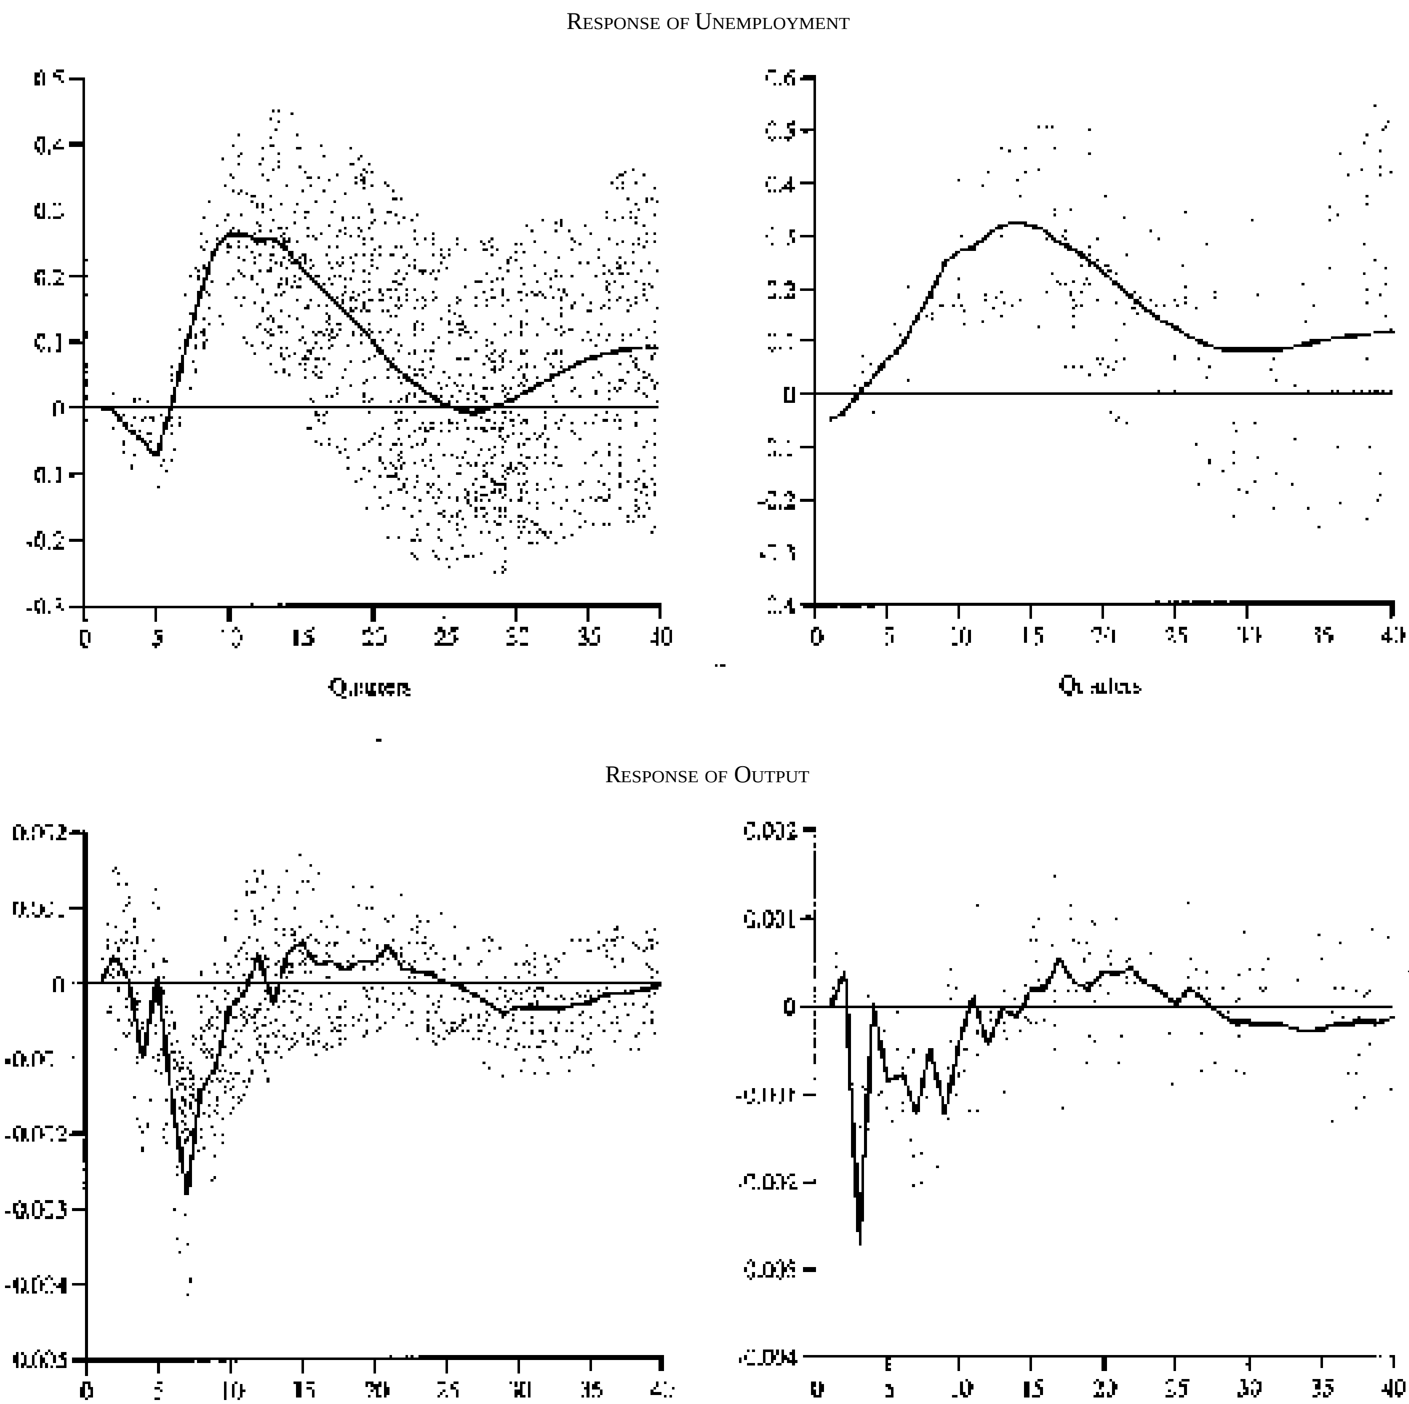

exaggerating the role of the dispersion index, we placed it last in the ordering. Specifically, the S&P500 index is placed first, followed by output, the funds rate and unemployment rate, and the dispersion index is placed last. The figure shows that the unemployment rate begins to increase about four to five quarters after a shock to the dispersion index and continues to go up for about two more years before beginning a gradual decline. This response resembles the response of the unemployment rate to funds rate shocks. Output responds to a shock to the dispersion index with a lag as well, but the response is relatively short-lived.

The associated variance decompositions are shown in the lower panel of Table 2. They show that dispersion accounts for roughly 25% to 30% of the variance of unemployment beginning about three years after the shock. The funds rate accounts for about half. In the case of output, both dispersion and the funds rate account for about 15% to 20% of the variance after the first two years.

Alternative Models

We also estimated some alternative versions of our basic VAR. We began by estimating a model that included the relative price of oil instead of the total stock market return. Our motivation here is twofold. First, we intended this to be a check on the robustness of our specification, following Loungani (1986) who showed that including this variable in the unemployment equation led to Lilien's measure of dispersion becoming insignificant. Second, even if inclusion of the oil price variable does not cause our dispersion index to become insignificant, we would like to see how much our dispersion index explains after an explicit source of sectoral reallocation is taken into account. It turns out that the dispersion variable is still significant at the 1% level in the unemployment equation, while the oil price variable has a marginal significance level of about 90%. However, including the oil price variable does lead to a reduction in the proportion of the forecast error variance of unemployment explained by the dispersion index; it falls from 31% to 22% at the 12-quarter horizon and from 27% to 20% at the 40-quarter horizon. The oil price variable accounts for roughly 5% to 6% of the error decomposition. (The dispersion index is placed last in all cases.)

The second system we estimated substituted federal defense expenditures instead of the stock market return in the original VAR. Once again, the idea was to include a variable that has been associated with a change in the sectoral allocation of labor over our sample period. The defense expenditure variable is significant at the 11% level in the unemployment equation, while the dispersion index remains significant at 1%. There is a slightly larger decline in the proportion of forecast error variance explained by the dispersion index, which now explains 21% of the variance at a 12-quarter horizon and 17% 40 quarters ahead. The defense expenditures variable explains about 15%.

Overall, we believe these results are consistent with Davis's observation (cited above) that allocative disturbances are unlikely to be associated with one particular variable.

III. SECTORAL SHIFTS AND THE DURATION OF UNEMPLOYMENT

Intuitively, it seems that sectoral shocks should lead to permanent reallocations of labor, and thus imply longer spells of unemployment than those caused by aggregate shocks. For instance, an increase in interest rates is likely to cause automobile manufacturers to respond to the temporary reduction in demand by laying off workers, who will then be hired back when demand rebounds. By contrast, a shock to the automobile sector, such as an increase in the supply of Japanese cars, is likely to lead to permanent changes in employment in the sector. As a consequence, displaced workers will have to move to other sectors. Workers who have to find jobs in other sectors will tend to stay unemployed for longer periods than those who can stay within the same sector (or even be rehired by the same firm).

Some evidence from micro data supports these intuitive ideas. Using the Michigan Panel Study of Income Dynamics, Loungani, Rogerson, and Sonn (1989) find that workers who moved across industries have longer unemployment spells than those who stayed within the same industry. Based on data from the Canadian Labor Market Activity Survey, Thomas (1996) concludes that industry movers have longer spells of unemployment than stayers, though the difference is significant only for workers who do not receive unemployment insurance.

Further evidence is provided by Brainard and Cutler (1993), who showed that (in a system that contained lagged unemployment and a measure of stock market dispersion, as well as labor market dispersion) the stock market dispersion variable entered significantly into equations that explained unemployment spells exceeding five weeks but was not significant in explaining spells up to five weeks. We extend their work by looking at variance decompositions in a five-variable VAR; it seems to us that the variance decompositions provide a more useful way of trying to judge the relative importance of sectoral shocks than *F* tests do.[9] Since our system also contains the funds rate, we are in a position to compare the effects of policy shocks and sectoral shocks as well.

9. In any case, the dispersion index is significant at 5% in all the unemployment equations we estimated, so that the *F* test cannot be used to distinguish between equations.

We have data on four different durations of unemployment: 0 to 4 weeks, 5 to 14 weeks, 15 to 26 weeks, and spells that are 27 weeks or longer. We present detailed results for spells lasting 27 weeks or more, and abbreviated results for the other three categories.

Results for a system where we have substituted the long-duration unemployment rate (that is, the rate based on unemployment that exceeds 26 weeks) for the aggregate unemployment rate are shown in Table 3.[10] The important result in Panel (A) of the Table is that lagged values of the dispersion index play a very significant role in the determination of long-duration unemployment. Furthermore, note that lags of long-duration unemployment do not influence the level of dispersion. Figure 2 shows that long-duration unemployment responds to changes in dispersion with a lag as well, although its response is somewhat more drawn out than that of overall unemployment shown in Figure 1. The variance decompositions are in Panel B. Note that dispersion accounts for a very high proportion of unemployment variation at the longer horizons: at the 20-quarter horizon, for instance, the proportion accounted for by dispersion is close to 45%.

Table 4 compares the role of the dispersion variable (Panel A) and the funds rate (Panel B) in explaining the forecast error variance of different durations of unemployment. Each column in Panel A comes from a VAR that contains the unemployment rate of the relevant duration plus the four variables in our basic system (output, the funds rate, dispersion and the stock market return). The ordering is the same as before, as well. The table shows that beyond the first two years the contribution of sectoral shifts to unemployment fluctuations rises fairly steadily with duration. For instance, comparing the 20-quarter ahead decomposition, the contribution of dispersion rises from 9% for the shortest duration to 43% for the longest duration.

10. The rate is obtained by dividing the number of unemployed workers at each duration by the total labor force.

TABLE 3

EXPLAINING LONG-DURATION UNEMPLOYMENT

A. MARGINAL SIGNIFICANCE LEVELS

	LR-UNEMP.	OUTPUT	FUNDS RATE	S&P500	DISPERSION
LR-UNEMP.	.01	.30	.85	.30	.44
OUTPUT	.10	.35	.09	.61	.39
FUNDS RATE	.21	.03	.01	.50	.50
S&P500	.29	.35	.65	.88	.44
DISPERSION	.01	.36	.06	.40	.03
ADJ. R^2	.99	.29	.91	.03	.30

B. VARIANCE DECOMPOSITIONS[a]

	LONG-DURATION UNEMPLOYMENT					OUTPUT				
QRTRS	S&P500	OUTPUT	FUNDS	LR-UNEMP	DISPERSION	S&P500	OUTPUT	FUNDS	LR-UNEMP	DISPERSION
0	1	3	3	93	0	0	100	0	0	0
4	3	29	5	63	1	8	75	14	1	2
8	7	33	8	45	7	10	59	15	1	15
12	5	17	20	25	33	11	58	15	2	15
20	5	11	28	14	43	11	56	15	2	16
40	5	11	29	14	41	11	55	15	2	17

[a]See note to Table 1.

FIGURE 2

DYNAMIC RESPONSES OF LONG-DURATION UNEMPLOYMENT AND OUTPUT

DISPERSION SHOCKS	FEDERAL FUNDS RATE SHOCKS

RESPONSE OF LONG-DURATION UNEMPLOYMENT

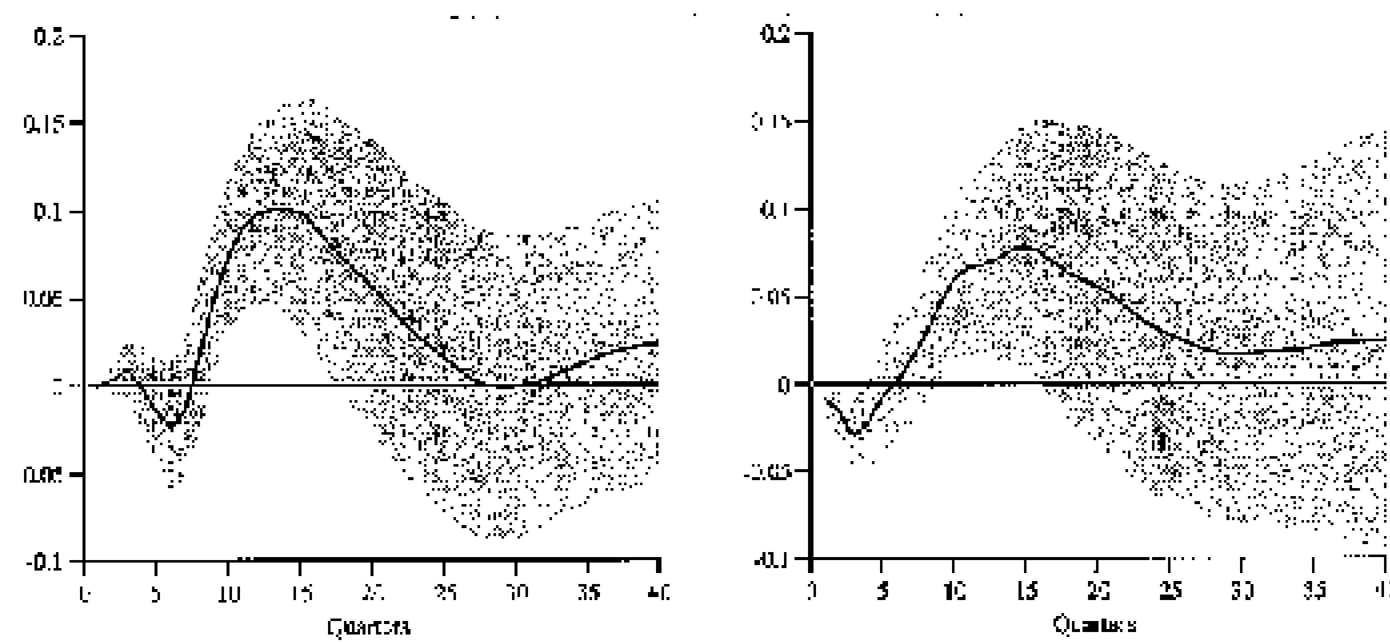

RESPONSE OF OUTPUT

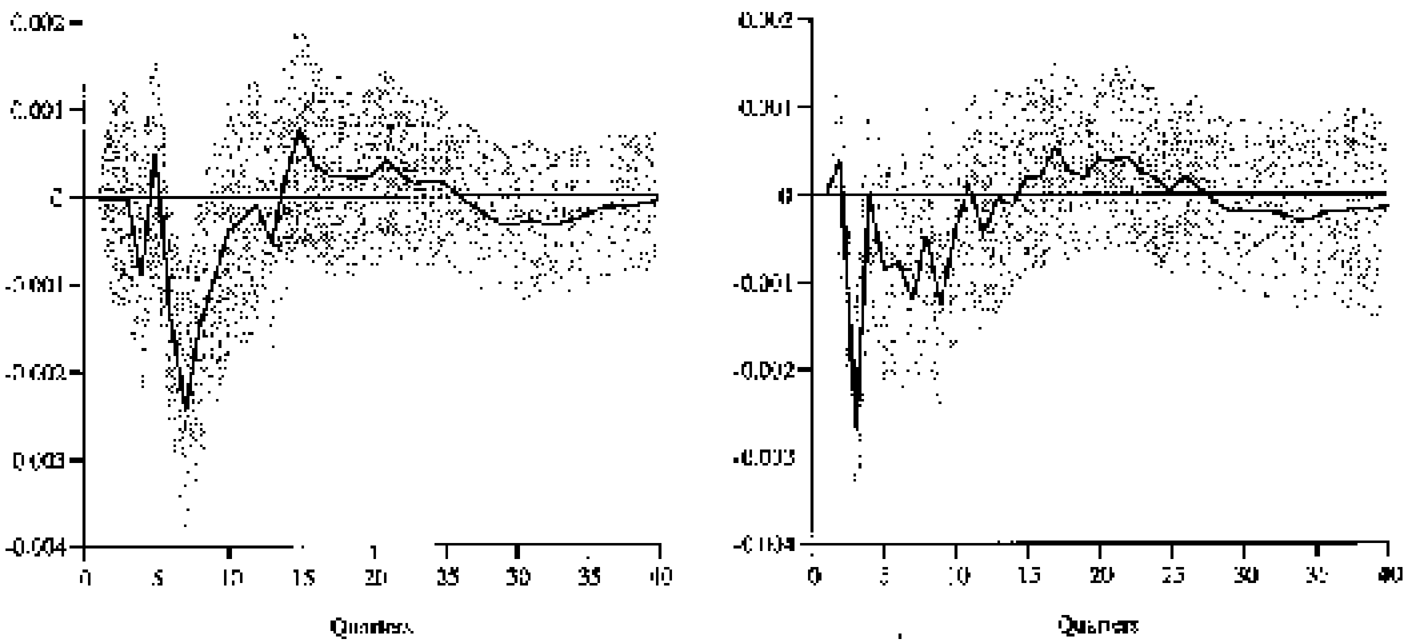

TABLE 4

EXPLAINING UNEMPLOYMENT BY DURATION

A. PROPORTION OF FORECAST ERROR VARIANCE OF THE UNEMPLOYMENT RATE EXPLAINED BY DISPERSION

	Unemployment			
Quarters ahead	Up to 5 weeks	5 to 14 weeks	14 to 26 weeks	26+ weeks
0	0	0	0	0
4	1	4	2	1
8	9	20	19	7
12	10	28	33	33
20	9	26	35	43
40	9	24	35	41

B. PROPORTION OF FORECAST ERROR VARIANCE OF THE UNEMPLOYMENT RATE EXPLAINED BY THE FUNDS RATE

	Unemployment			
Quarters ahead	Up to 5 weeks	5 to 14 weeks	14 to 26 weeks	26+ weeks
0	3	14	4	3
4	17	6	6	5
8	37	22	18	8
12	52	35	28	20
20	60	47	35	28
40	59	49	36	29

Panel B shows that the contribution of the funds rate declines as the duration of the unemployment rate rises. At a 20-quarter horizon, for instance, it falls from 60% to 28%. It is worth pointing out that the relative shares of the other variables in the system do not change as dramatically as the duration of unemployment changes.

IV. ROLE OF SECTORAL SHIFTS DURING NBER RECESSIONS

In this section we use the models we have estimated to carry out a historical decomposition of the unemployment rate. Our purpose is to examine what role, if any, sectoral shifts may have played during recessions. We also look at the role played by changes in the funds rate; this is of interest in its own right and also provides us with a benchmark for assessing the relative importance of sectoral shifts.

Figure 3 provides our results for the aggregate unemployment rate. The top panel shows the actual unemployment rate over the 1971.Q1–1995.Q4 period together with two sets of forecasts. The line labeled "Base Forecast" is the VAR's forecast for this entire period based on data up to the end of 1970 only (though the coefficients used are obtained by estimating the model over the entire period). The line labeled "Base + Dispersion" is the forecast from the VAR after it has been provided with all the innovations to the dispersion variable over this period. These innovations are the orthogonalized innovations obtained from the same ordering that was used in Table 2 and the associated Figures. The top panel shows that dispersion accounts for most of the rise in unemployment during the 1973–75 recession. Its contribution is more modest during the 1982 recession, though it does help explain part of the sharp increase in the middle of the recession. The dispersion index

FIGURE 3

HISTORICAL DECOMPOSITION OF UNEMPLOYMENT RATE

CONTRIBUTION OF DISPERSION	CONTRIBUTION OF FEDERAL FUNDS RATE

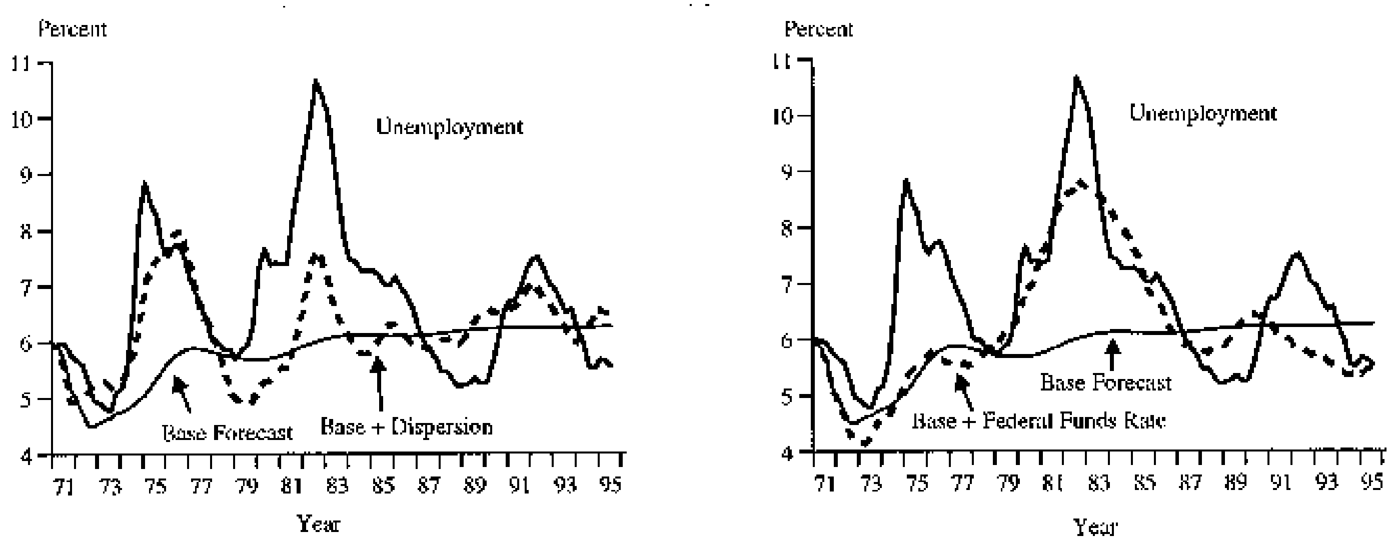

also appears to explain part of the rise in unemployment during the 1991 recession, though it does not explain the decline during the last two years or so. While we have not shown the results here, it is worth pointing out that the dispersion index accounts for somewhat less of the rise in unemployment during the 1973–75 recession in the systems where we include either defense expenditures or oil (though it still accounts for most of the increase). Its role during the 1982 recession is roughly unchanged. And in both alternative systems it helps explain some of the rise in unemployment during the 1990–91 recession, though its role is noticeably smaller than in the base system (shown in Figure 3).[11]

The lower panel of the Figure shows the contribution of the funds rate. The funds rate does not account for the rise in unemployment during the 1973–75 recession, and its contribution actually goes the wrong way during the most recent recession. However, the funds rate does an extremely good job of tracing the rise and fall of unemployment around the recessions of 1980 and 1982; this is consistent with the widespread belief that the tightening of monetary policy around this period played a big part in these recessions.

Figure 4 presents the same results for the long-duration unemployment rate. The top panel shows that dispersion accounts for the entire increase in long-duration unemployment in the 1973–75 recession (in fact, it more than accounts for the increase) and also accounts for some of the increase in unemployment during the early 1980s. However, dispersion explains only a small part of the rise in long-duration unemployment during the last recession.[12] This result is in contrast to the result in Figure 3. Since our priors are that sectoral shifts should be more closely related to long-duration unemployment, we interpret these two conflicting pieces of evidence as suggesting that sectoral shocks probably did not have a very large role to play in the 1990–91 recession.

Finally, the lower panel of the Figure shows the contribution of the funds rate to changes in unemployment over this period. Once again, the funds rate explains only what happened around the early 1980s. However, even during this period its contribution to movements in the long-duration unemployment rate is smaller than to movements in the overall unemployment rate.

11. The defense expenditure variable helps explain some of the rise in unemployment during the 1973-75 recession, but is not very important elsewhere. The oil price variable does not contribute much to movements in unemployment over this period. Again, we see these results as illustrating how difficult it is to pinpoint any particular variable as the key source of sectoral shocks.

12. This last result may appear surprising in light of the results from the variance decompositions, which suggested that dispersion plays a larger role in explaining movements in long-duration unemployment than in short-duration unemployment. However, those results pertain to the sample period as a whole and need not hold true over the course of every recession.

FIGURE 4

Historical Decomposition of Long-Duration Unemployment Rate

Contribution of Dispersion | Contribution of Federal Funds Rate

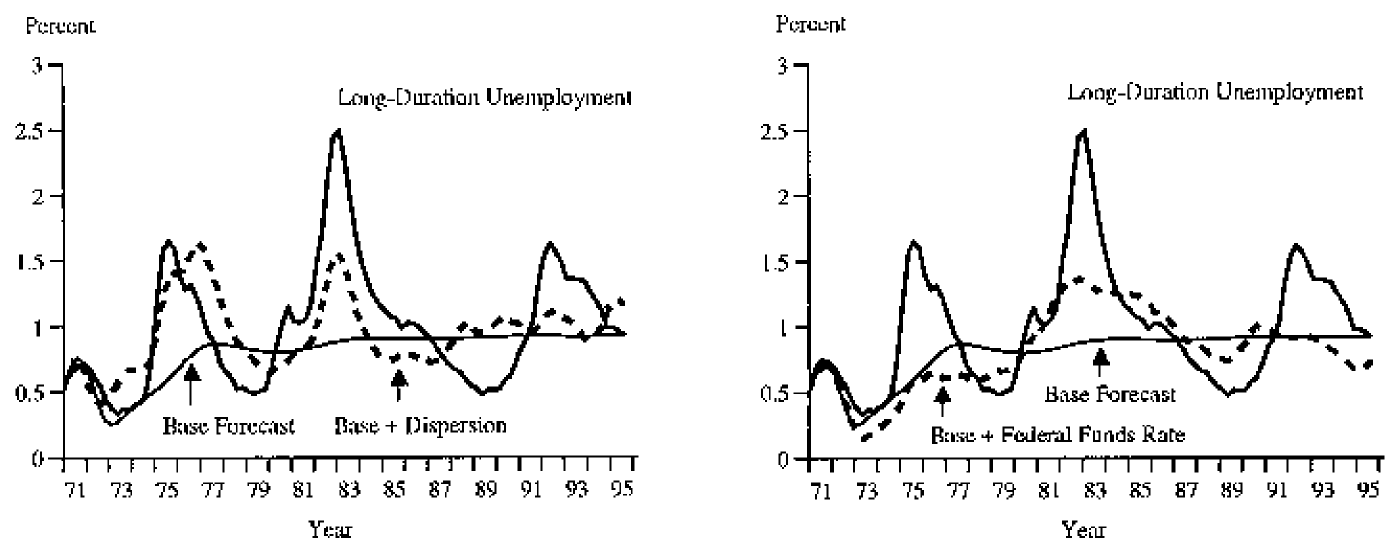

V. Summary and Conclusions

Overall, we conclude that sectoral shifts (as measured by the stock market index) explain a significant proportion of the variation in the unemployment rate. To assess the quantitative role played by sectoral shifts, it is useful to compare the contribution of the dispersion index to that of the federal funds rate, which is the leading alternate source of unemployment fluctuations considered here. Even though it is placed last in the ordering, dispersion accounts for 31% of the forecast error variance of unemployment at a 12-quarter-ahead horizon, whereas the funds rate accounts for 38%. Hence, dispersion is roughly as important as the funds rate in accounting for fluctuations in the unemployment rate over the medium term, though at longer horizons the funds rate is much more important.

The dispersion index is considerably more important when explaining movements in long-duration unemployment: except at the very short horizons, the dispersion index accounts for a larger percentage of the forecast error variance than the funds rate. At a 20-quarter horizon, for example, the respective contributions of the two variables are 43% and 28%.

It is worth emphasizing our finding that sectoral shocks play a relatively large role in explaining unemployment, even though our system includes both real GDP and the funds rate—variables that are commonly thought to have a significant effect on the unemployment rate but which have not been explicitly considered in previous analyses. In addition, we also have shown that our results are not due to the omission of other variables that could plausibly have caused sectoral shifts during particular episodes, namely, the oil price and defense expenditure variables.

The results from our exercise also provide a partial answer to an old question: Are business cycles all alike? Our historical decompositions say that recessions are not. Sectoral shifts appear to account for the 1973–75 recession, though we have not explored in detail which particular shocks may be driving the index over this period.[13] By contrast, monetary policy (as measured by the funds rate) appears to have been the key player in the 1982 recession. Neither sectoral shocks nor monetary policy appear to explain the 1990 recession, though the dispersion index does track the rise in unemployment over this period to a modest degree.

Finally, our results offer an interesting perspective on why the long-duration unemployment rate has remained high in the period since 1993. Our historical decompositions suggest that the path of the funds rate was consistent with long-duration unemployment returning to the level consistent with previous troughs in the data. However, increases in the dispersion index offset this effect, keeping long-duration unemployment higher than it would otherwise have been.[14]

13. While not the subject of our paper, the productivity slowdown that occurred around that time is consistent with the hypothesis that some kind of structural change took place over that period.

14. Valletta (1996) suggests a different explanation for the rise in long-duration unemployment. Specifically, he suggests that this increase could

APPENDIX

Construction of the Dispersion Index

Our dispersion index is constructed using the basic methodology of Loungani, Rush, and Tave (1990). Due to data constraints, our series covers 1962 to 1995. Over the lifetime of the S&P500 Composite Index, industry subgroups are added and deleted. We obtained a list of the dates of changes from S&P. The series includes only the industry indexes that were included in the composite for a given date. Three series have been omitted due to a lack of employment data: Miscellaneous, Miscellaneous (High Tech), and Conglomerates. They are not distinct industries and do not have SIC codes. Series that were deleted prior to 1973 were not in our database. There are 17 of these groups. In addition, we did not have Transport Misc. (Old). All composite indexes were dropped to avoid double counting. The index observations are the closing price of the quarter.

Weights are based on the BLS employment data by SIC industry. We determine the weight by two-digit SIC and divide that weight evenly among the component industries for that date. The weights sum to one. Two weights were constructed—one using data from 1978 (the sample midpoint) and one using data from 1995 (the sample endpoint.) The employment data for three two-digit SIC codes were available only starting in 1988. We estimated the 1978 weights for these industries using the 1988 data. For SIC 78 (Motion Pictures), we assumed that the share of the industry in the Services aggregate was the same in 1978 as in 1988. For SIC 60 and 61 (Depository and Nondepository Institutions, respectively), we found the employment for these sectors together by subtracting all other financial sectors from the Financial Sector aggregate. We assumed the share of each was the same as in 1988.

REFERENCES

Abraham, Katharine, and Lawrence Katz. 1986. "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?" *Journal of Political Economy* 94, pp. 507–522.

Barro, Robert J. 1986. "Comment on 'Do Equilibrium Real Business Theories Explain Postwar U.S. Business Cycles?'" *NBER Macroeconomics Annual* 1, pp. 135–139.

Bernanke, Ben, and Alan Blinder. 1992. "The Federal Funds Rate and the Channels of Monetary Transmission." *American Economic Review* (September) pp. 901–921.

Black, Fisher. 1995. *Exploring General Equilibrium*. Cambridge: MIT Press.

__________. 1987. *Business Cycles and Equilibrium.* New York: Basil Blackwell.

Brainard, Lael, and David Cutler. 1993. "Sectoral Shifts and Cyclical Unemployment." *Quarterly Journal of Economics,* pp. 219–243.

Campbell, Jeffrey, and Kenneth Kuttner. 1996. "Macroeconomic Effects of Employment Reallocation." *Carnegie-Rochester Conference Series on Public Policy* (June) pp. 87–117.

Cochrane, John H. 1994. "Shocks." NBER Working Paper No. 4698.

Davis, Steve. 1987. "Fluctuations in the Pace of Labor Reallocation." *Carnegie-Rochester Conference Series on Public Policy* 27 (Spring) pp. 335–402.

__________. 1985. "Allocative Disturbances and Temporal Asymmetry in Labor Market Fluctuations." Working Paper. University of Chicago Graduate School of Business.

Hall, Robert E. 1995. "Lost Jobs." *Brookings Papers on Economic Activity* No. 1, pp. 221–256.

__________. 1993. "Macro Theory and the Recession of 1990–1991." *American Economic Review Papers and Proceedings,* pp. 275–279.

Lilien, David. 1982. "Sectoral Shifts and Cyclical Unemployment." *Journal of Political Economy* 90, pp. 777–793.

Lindsey, Lawrence. 1996. "NAIRU Disrobed." *The International Economy* (March/April) pp. 8–13.

Loungani, Prakash. 1986. "Oil Price Shocks and the Dispersion Hypothesis." *Review of Economics and Statistics* (August) pp. 536–539.

__________, Mark Rush, and William Tave. 1990. "Stock Market Dispersion and Unemployment." *Journal of Monetary Economics* (June) pp. 367–388.

__________, Richard Rogerson, and Yang Hoon Sonn. 1989. "Unemployment and Sectoral Shifts: Evidence from PSID Data." Working Paper. University of Florida.

Perry, George, and Charles Schultze. 1993. "Was This Recession Different? Are They All Different?" *Brookings Papers on Economic Activity* 1, pp. 145–195.

Polivka, Anne E., and Stephen M. Miller. 1995. "The CPS after the Redesign: Refocusing the Economic Lens." Mimeo. Bureau of Labor Statistics.

Romer, Christina, and David Romer. 1989. "Does Monetary Policy Matter: A New Test in the Spirit of Friedman and Schwartz," *NBER Macroeconomics Annual,* pp. 121–170.

Thomas, Jonathan. 1996. "An Empirical Model of Sectoral Movements by Unemployed Workers," *Journal of Labor Economics* (January) pp. 126–153.

Toledo, Wilfredo, and Milton Marquis. 1993. "Capital Allocative Disturbances and Economic Fluctuations." *Review of Economics and Statistics* (May) pp. 233–240.

Valletta, Robert G. 1996. "Has Job Security in the U.S. Declined?" Federal Reserve Bank of San Francisco *Economic Letter* 96-06 (February 16).

Yellen, Janet. 1989. "Comment on 'The Beveridge Curve' by Olivier Blanchard and Peter Diamond." *Brookings Papers on Economic Activity* 1, pp. 65–71.

be related to changes in job security, since many workers who lost jobs during this period were from groups that in the past did not have to think about job search.

The Effects of Industry Employment Shifts on the U.S. Wage Structure, 1979–1995

Robert G. Valletta

Economist, Banking and Regional Section. John DiNardo provided very helpful discussions of the statistical approach used. I am grateful to Fred Furlong, Ken Kasa, and especially Mary Daly for their careful review of this paper. I also thank Nan Maxwell and session participants at the 1996 Western Economic Association Meetings for comments on an early version. Randy O'Toole provided useful research assistance.

The trend toward increasing U.S. wage inequality during the 1980s is well documented. I investigate the role of employment shifts from goods-producing to service-producing industries in contributing to increased inequality during the period 1979–1995. Earlier analyses revealed that average earnings are lower, and earnings inequality is higher, for service-producing workers than for goods-producing workers. For both reasons, an increasing share of service employment may increase earnings inequality.

I analyze the effect of broad industry employment shifts by using a recently developed statistical technique, which I term "conditionally weighted density estimation." This technique enables investigation of the effects of changing industry employment shares on the complete distribution of earnings, conditional on changes in other earnings-related characteristics. The results show at most a small effect of industry employment shifts on growing inequality in male hourly earnings.

The trend toward increasing U.S. wage inequality during the 1980s is well documented and extensively analyzed (for example, Bound and Johnson 1992; DiNardo, Fortin, and Lemieux 1996; Juhn, Murphy, and Pierce 1993; Karoly 1992; Katz and Murphy 1992). During this decade, earnings inequality increased both across and within industry sectors and worker groups, and the return to measurable skills (particularly formal education) increased substantially. Research in this area has focused largely on assessing the contribution to rising earnings inequality of factors such as the declining real minimum wage, declining unionism, changing supply and demand across worker groups, increased international trade, and skill-biased technological change. Each of these factors appears to have played a role in increased U.S. earnings inequality during the decade.

An additional factor that may have contributed to increased earnings inequality during the 1980s and earlier, however, is the substantial employment shift in the U.S. from goods-producing to service-producing industries. A common stereotype associated with service-producing jobs is that they pay less than goods-producing jobs. Consistent with this belief, studies such as Blackburn (1990) report that average earnings are lower, and earnings inequality higher, in service-producing jobs than in goods-producing jobs. For both reasons, an increasing share of service employment may increase earnings inequality. Thus, in popular and academic discussion, the shift from goods to services has been cited as a reason for increased inequality and a declining middle class (for example, see Bluestone and Harrison 1982, 1988).

In this paper, I examine the contribution of such industry employment shifts to changing earnings inequality from 1979 to 1995. As described in more detail in Section I, several papers have examined this issue. For example, Maxwell (1989, 1990) and Bluestone and Harrison (1988) both included measures of relative manufacturing employment in their analyses of changing inequality and low-wage employment over the periods 1947–85 and 1963–86, respectively. Each found that employment shifts out of manufacturing have played an important role. However, the use of aggregate time-series data may obscure the role of underlying forces such as changing skill attributes. Blackburn (1990) examined the impact of industry employment shifts on earnings inequality using individual level data from various March Current Population Survey files and found them to have a noticeable but limited influence. In contrast, Murphy and Welch (1993) and Juhn, Murphy, and Pierce (1993) found no effect of industry shifts on average wages and the variance of the wage distribution, respectively.

From an academic perspective, then, the exact contribution of industry employment shifts to rising earnings inequality remains an open question. I attempt to resolve this debate by applying a recently developed methodology that is particularly well suited to analyzing the contribution of broad economic changes to earnings inequality. The technique—which I call "conditionally weighted density estimation"—was recently developed by DiNardo, Fortin, and Lemieux (1996) and applied to the analysis of increased earnings inequality. Their technique enables estimation of the effects of broad economic changes on the entire distribution of earnings. Most studies of rising earnings inequality have focused on explaining changes in the mean or variance of the distribution, or changes in expected wage differentials across labor market groups. In contrast, the conditionally weighted density approach is far less restrictive and applies particularly well when there is no strong *a priori* knowledge about what portions of the earnings distribution are most affected by the factor being examined. For example, conditional weighted density estimation enables examination of whether wages have become more disperse due to widening of the tails or movement from the middle to the tails, a distinction that is important for distinguishing among different explanations of increased inequality (one of which is the "deindustrialization" hypothesis of Bluestone and Harrison 1982).

In general, the technique of DiNardo, et al., enables estimation of a distribution under counterfactual assumptions about the state of the world, which in turn reveals the distributional impact of the state of the world as it actually evolved. My focus is on the effect of changing industry employment shares. In particular, the technique enables me to answer the question, "How would the distribution of earnings look in 1995 if industry employment shares had remained as they were in 1979?" Furthermore, it produces two depictions of how the earnings distribution has been altered by the modeled changes: (1) a visual depiction obtained through comparison of kernel density estimates of the earnings distribution; (2) quantitative comparison based on calculation of parametric inequality measures (standard deviation, quantile dispersion measures, the Gini coefficient, etc.). Both depictions are based on a comparison of calculations that use the original data and survey sampling weights with calculations for which the sampling weights are modified by estimated conditioning weights. This procedure is described heuristically in Section II, with analytic details provided in the Appendix.

To estimate the role of changing industry employment shares, I use data from the 1979 and 1995 Current Population Surveys, as described in Section III. Much of the literature focuses on widening earnings inequality during the 1980s. However, a recent paper by Karoly (1996) finds that increasing inequality continued during the early 1990s. Despite this continued increase during the period covered by my analysis, and despite finding that the service sector exhibits lower average earnings and higher earnings variation, I find at most a small independent impact of industry employment shifts on dispersion in the lower half of the male earnings distribution. These results are described in detail in Section IV of the paper, with conclusions provided in Section V.

I. Earnings Inequality and Changing Industry Employment

A large number of papers in recent years have attempted to attribute increasing earnings inequality during the 1980s to a variety of observable factors (e.g., Bound and Johnson 1992, Katz and Murphy 1992, Blackburn, Bloom, and Freeman 1990, Juhn, Murphy, and Pierce 1993). These authors typically focused on regression-based decompositions or similar analysis based on worker groups defined by earnings-related characteristics, using either aggregated time-series data or yearly individual data.

One recent methodological advance in this literature is the application of kernel density estimation, which provides visual depiction of the entire distribution of earnings. The use of kernel density estimation as an exploratory data analysis tool has long been recognized (see Silverman 1986). In the analysis of changing earnings inequality, kernel density estimates provide a useful visual depiction of how the distribution of earnings has changed over time and where in the distribution the largest changes have been concentrated. Given the lack of strong prior knowledge on where in the distribution the largest changes have occurred, and the focus in the literature on parametric measures such as the variance in earnings, this is an important advance. For example, Levy and Murnane (1992) noted that standard scalar measures of inequality may not distinguish among alternative sources of increasing inequality that have differing economic and social implications, since these measures do not identify the portion of the earnings distribution on which changes have occurred.

Burkhauser, et al., (1996) recently applied kernel density estimation to the analysis of changing inequality. They examined changes in the distribution of family earnings in the U.S., U.K., and Germany during the 1980s. In this form, kernel density estimation serves essentially as a smoothed histogram, thereby providing visual insight into changing inequality. Burkhauser, et al., found that rising inequality in family earnings in the U.S. was characterized primarily by large but unequal income gains in the middle of the family income distribution.

Although kernel density estimation is useful for such exploratory analysis and visual characterization of distributions, its direct use as an analytical tool is limited. In contrast, conditional density estimation enables a full range of analytical applications. Conditional density estimation methods proceed by reestimating the entire distribution of earnings after accounting for various earnings determinants, or by reweighting the distribution according to conditional probabilities. For example, Juhn, Murphy, and Pierce (1993) applied a regression-based conditioning approach. They used the cumulative distribution function of residuals obtained from wage equations to decompose changes in inequality measures into portions due to changes in observable personal characteristics, changes in the returns to those characteristics, and changes in the distribution of unobservables. They found an increasing contribution of unobservables to rising earnings inequality in the 1980s.

DiNardo, Fortin, and Lemieux (1996; henceforth DFL) and DiNardo and Lemieux (1994) improved on previous methods by conditioning through the use of estimated weights. They combined the estimated conditioning weights with sample survey weights to produce an adjusted earnings distribution. This is a flexible procedure that provides semiparametric estimates of the entire distribution of earnings under various counterfactual assumptions. The adjusted distribution can be compared with the original distribution both visually, using appropriately reweighted kernel density estimates, and quantitatively, by comparing dispersion measures from the adjusted and unadjusted distributions.

DFL used their technique to estimate how much earnings inequality would have risen between 1979 and 1992 if the real minimum wage and union membership density in the U.S. had remained at their 1979 levels. Comparison to the actual amount by which earnings inequality rose revealed the impact of the declining minimum wage and declining union membership, conditional on changes in other important variables (such as individual skill attributes). Because the minimum wage affects only the lower portion of the earnings distribution, the technique's ability to reveal features of the entire distribution is particularly salutary. Both papers reported important contributions of a declining real minimum wage and declining unionism to increasing U.S. earnings inequality during the period 1979–88.

These authors, however, did not examine the role of changing industry employment patterns. During most of the postwar period, the share of service-producing jobs in the U.S. has increased substantially. These shifts will alter the distribution of earnings if either the level or dispersion in earnings is different across the goods-producing and service-producing sectors.

Previous work that analyzed the effect of industry employment shifts on earnings inequality typically used aggregated data. Using aggregate time series data, Maxwell (1989) found that the increasing share of service sector employment relative to manufacturing employment explains a substantial portion of increasing inequality over the period 1947–1985; she attributed much of this to declining unionization (Maxwell 1990). Also using aggregate data, Bluestone and Harrison (1988) found a corresponding effect on low-wage employment for the period 1963–86.

In contrast, Blackburn (1990) examined the influence of changing industry structure and other factors on earnings inequality using individual level data from various March Current Population Survey files and found only a limited impact of industry employment shifts. Similarly, Murphy and Welch (1993) and Juhn, Murphy, and Pierce (1993) found no effect of industry shifts on average wages and the variance of the wage distribution, respectively. Furthermore, Schweitzer and Dupuy (1995) used kernel density techniques and found substantial convergence in the goods-producing sector and service-producing sector wage distributions through 1993, which suggests a limited impact of industry employment shifts on inequality. Thus, evidence on the role of industry employment shifts in increased earnings inequality is mixed.

I build on previous work by using weighted density estimation to assess the contribution of changing industry employment shares to increasing earnings inequality. As noted, this enables more flexible and detailed assessment of the impact of industry shifts on the structure of earnings than do other approaches.

II. METHODS

Kernel Density Estimation

Kernel density estimation is a flexible, largely nonparametric means of estimating the underlying distribution from which an empirical distribution is sampled.[1] The estimated densities essentially serve as "smoothed histogram" representations of a distribution, and as such are useful for exploratory data analysis. This subsection describes the basics of kernel density estimation, and the next

---

1. Silverman (1986) discusses non-parametric density estimation in detail, and Delgado and Robinson (1992) provide a useful summary of econometric applications.

two subsections describe the estimation and incorporation of conditioning weights into density estimation.

The kernel density estimate of a univariate distribution based on a random sample $(W_1, \ldots, W_n)$ of size n with sampling weights $\theta_1, \ldots, \theta_n$ (normalized so that $\Sigma_i \theta_i = n$) is:

$$(1) \quad f_h(w_j) = \frac{1}{n} \sum_{i=1}^{n} \frac{\theta_i}{h} K\left(\frac{w_j - W_i}{h}\right) \quad \text{for } j = 1,2,\ldots, m.$$

In this expression, K is the kernel function, h is the bandwidth, and m is the number of points at which the density function is evaluated.[2] Several alternatives are available for the function K, although they typically are probability density functions (and therefore are symmetric and integrate to 1 over the range of W). For each evaluation point w_j, these functions assign to the W's estimation weights that decline (smoothly or abruptly) as the W's move farther from w_j. The subscript j denotes evenly spaced values of w, with the choice of m depending largely on computing resources and the data. The full estimation essentially involves sliding a window (of width $2h$) across the range of W_i, with m density estimates computed at equal intervals.

The choice of h has been subject to substantial discussion in the literature and is generally acknowledged to be more important than the choice of kernel function. Various "optimal bandwidth selection" rules are available. Rather than investigating this issue in detail, I follow DiNardo and Lemieux (1994) in setting the bandwidth equal to .075 for all ln(hourly earnings) estimates provided below. This falls within the range of bandwidths selected by the optimal method of Sheather and Jones (1991) for similar data in DFL. This bandwidth also does a good job of capturing important visual features of the distribution of hourly earnings, such as the spike at the minimum wage. I use the Epanechnikov kernel function, which yielded results identical to a Gaussian kernel in comparison tests.

Conditional Weighted Density Estimation

In this section, I describe how simple estimated reweighting functions can be obtained and applied to the estimation of earnings distributions that embody counterfactual assumptions. In the text I describe these procedures heuristically; the exact derivation—which is conceptually simple but notationally complex—is provided in the Appendix.

Consider the distribution of wages w in year t, conditional on individual characteristics X and a measure of industry employment patterns E:

$$(2) \quad f_t(w) \equiv f(w; t_w = t, t_{E|X} = t, t_X = t).$$

This identity is notational; it shows that the distribution of w is defined in year t, conditional on the distribution of X and E (conditional on X) in the same year. In the empirical work, I focus on $t_w = 1995$, and I measure industry employment patterns by a dummy variable indicating whether each worker is in the broad goods-producing or service-producing sector.

The essence of the test is to investigate the effect of holding $t_{E|X}$ at earlier year (1979) levels—i.e., to estimate what the distribution of earnings would be if the distribution of goods-producing versus service-producing jobs had remained the same as in 1979. The simplest way to do this is to upweight individuals in the goods sector by a factor that is proportional to the decrease in the share of goods-producing jobs in the economy (and similarly downweight service sector workers). However, this simple test ignores any changes in the relationship between earnings-related characteristics and the probability of being in different broad industry sectors. For example, if the movement to services was exclusively by low-skilled workers, then the shift toward services did not have a substantial independent effect on the earnings distribution. Thus, we need to estimate the 1995 distribution of earnings with the industry employment distribution, *and its relationship to X,* held to its 1979 level.

In terms of the notation in (2), we are interested in:

$$(3) \quad f(w; t_w = 95, t_{E|X} = 79, t_X = 95).$$

This expression represents the density that would be observed if the probability (conditional on individual characteristics X) of being employed in goods-producing industries retained its 1979 level and structure, but workers were otherwise paid according to the earnings schedule prevailing in 1995. As shown in the Appendix, this distribution can be expressed as the original unconditional distribution of earnings in 1995, with observations reweighted by a function $\psi_{E|X}$. This function represents the change in the probability between 1979 and 1995 that an observation defined by characteristics X is observed in the goods-producing or service-producing sector.

Intuitively, to obtain the density of earnings that would prevail if the structure of conditional industry affiliation remained as it was in 1979, we downweight individuals in the 1995 sample whose characteristics would have made them less likely to work in the same sector in 1979 as they worked in 1995. These conditional probabilities can be estimated as the fitted values obtained from standard binary variable models; in the empirical work below, I use the probit model to estimate these conditional probabilities.

2. See Silverman (1986) for a detailed discussion of kernel density techniques.

As long as the unconditional probability of being in the goods sector or the relationship of the *X*'s to that probability have changed, then the estimated weighting function $\hat{\psi}_{E|X}$ will differ from one, and the counterfactual density will differ from the observed density. In general, because the probability of working in the goods-producing sector declined between 1979 and 1995, compared to the unconditional density the reestimated density will attach more weight to individuals currently working in the goods-producing sector and less weight to those in the service-producing sector.

In addition to accounting for the impact of changes in industry employment shares, the technique enables us to account explicitly for the impact of changes in the *X* vector of earnings-related characteristics. This serves as a useful basis for comparison and also enables us to account for interactions between the *X*'s and industry structure (as described in the next subsection).

The distributional effect of changes in the *X*'s can be modeled by again estimating weights and applying them to the 1995 earnings distribution (see DFL for details). This estimated weighting function—$\psi_X(X)$—is equal to the relative probability of observing an individual with characteristics *X* in the 1979 versus the 1995 sample, normalized by the unconditional probabilities of being in either sample. As long as the distribution of *X*'s changed between the two years (for example, through higher average educational attainment), the weights ψ_X will alter the estimated distribution. In the empirical work, the function ψ_X is calculated based on fitted values from probit equations that estimate the probability of observing an individual with characteristics *X* in the 1979 versus the 1995 data set. Workers in the 1995 sample with characteristics that make them relatively more likely to be observed in the 1979 sample will receive more weight in the conditional density estimation than they do in the unconditional density estimation.[3]

Reweighted Estimates and Comparisons

I now describe how the conditioning weights are used in the estimation. Briefly, the estimated conditioning weights are used to modify the sampling weights; I term this process "conditional weighted kernel density estimation." Comparison of the original and adjusted distributions reveals the effects of interest.

For each of the estimated weighting functions ($\psi_{E|X}$ and ψ_X), the conditional weighted kernel density estimates are obtained by multiplying the sampling weights for each observation (θ) by the estimated conditioning weights ($\hat{\psi}$). The combination of sampling weights and conditioning weights produces three distributions of earnings:

(1) Population weighted distribution: $f(w,\theta)$
(2) Distribution adjusted for industry employment structure: $f_e(w,\theta\cdot\hat{\psi}_{E|X})$
(3) Distribution adjusted for individual characteristics: $f_x(w,\theta\cdot\hat{\psi}_{E|X}\cdot\hat{\psi}_X)$

where

θ = survey sampling weight
$\hat{\psi}_{E|X}$ = estimated conditioning weight for industry employment structure
$\hat{\psi}_X$ = estimated conditioning weight for individual characteristics.

These new weights can be incorporated directly into the estimation of the kernel densities, which requires only slight modification of equation 1:

$$(4)\quad f_h(w_j) = \frac{1}{n}\sum_{i=1}^{n}\frac{\theta_i\psi_i}{h}K\left(\frac{w_j - W_i}{h}\right)\quad for\ j = 1,2,\ldots,m.$$

The result is a different kernel density estimate for each weighting scheme. Graphical depiction and comparison of the sample weighted and conditionally weighted kernel estimates provide a visual representation of the impact of changing industry distribution and individual characteristics.

Furthermore, the reweighting procedure enables calculation of the effect of the modeled change on any distributional statistic: moments (such as the mean and variance), quantile differences (the difference in earnings measured at specific cumulative points on the distribution), and parametric inequality indices (for example, the Gini and Theil indices). This procedure is particularly simple. Distributional statistics for the adjusted distribution are obtained by replacing the population weights by their product with the estimated conditioning weights when calculating the distributional statistics, a procedure easily handled by software that allows weighted tabulations.

One objection to the procedure outlined above is that it gives industry employment shifts precedence over changing individual characteristics in assessing the contribution of each factor. Given this ordering, any interactions between the two factors—for example, due to increasing concentration of unskilled workers in service industries—will be attributed to industry employment shifts. A useful check on the results, then, is to reverse the ordering of the estimation, which entails accounting for the impact of the *X*'s

3. In attempting to control for changes in educational attainment over their sample frame, Schweitzer and Dupuy (1995, p. 20) apply a restricted version of conditional weighted kernel density estimation: for each observation, they scale the sampling weight up or down to reflect a larger or smaller number of individuals with similar educational attainment in the base year.

first and then assessing the impact of changing industry employment shares. I report results from this procedure below; it requires reformulation of the conditioning weights, as described in DFL.

In terms of the treatment of the conditioning variables (X), the conditional weighting procedure is closely related to standard regression-based decompositions of variance. Regressions typically are used, however, to estimate the mean of a distribution. The advantage of the weighted kernel density procedure is that it estimates the entire conditional distribution, as opposed to analyzing distributional characteristics one-by-one, and therefore provides a more flexible method than regression techniques for investigating distributional changes. Regression techniques would require a potentially lengthy search for the exact effect of industry employment shifts on the wage structure; conditionally weighted density estimation provides an immediate visual representation, and it enables estimation of any desired dispersion measure.

III. DATA

The data used in this study are the merged outgoing rotation group files, or Annual Earnings Files, from the Current Population Survey for the years 1979 and 1995 (CPS–AEF). Each month, members of the outgoing rotation group of CPS sample households (about one quarter of the sample) are asked questions concerning earnings on their current job. Pooled over the 12 months in a year, these files provided me with approximately 150,000 observations per year, after sample restrictions. I dropped observations with allocated values for earnings or hours and limited the analysis to individuals aged 16–64. To focus clearly on the goods/services distinction, I eliminated agricultural workers from the sample. I inflated 1979 earnings to 1995 levels using the GDP deflator for personal consumption expenditures,[4] and I dropped earnings observations with values below $1/hour and above $200/hour (in 1995 dollars).[5]

I focus on hourly earnings data from the CPS–AEF for several reasons. First, this provides a large, representative data set for a period characterized by substantial changes in earnings inequality. An alternative is to use data from the March CPS Annual Demographic Surveys, which collect information on labor market experience and earnings in the entire previous year. Although these data are extensively used in the study of inequality (for example, in Blackburn 1990, Burtless 1990, Juhn, Murphy and Pierce 1993, and Katz and Murphy 1992), the yearly earnings data are affected by job changes and labor supply factors. Also, formation of point-in-time earnings measures in the March CPS requires dividing by weeks worked and hours worked, which may introduce measurement error. The primary alternative—the use of yearly earnings for full-time, full-year workers—would narrow the sample undesirably for my experiment.

I begin my analysis in 1979 rather than earlier in the 1970s because the rate at which inequality increased was faster in the 1980s than in the 1970s, particularly for low-skilled workers (Bound and Johnson 1992). The 1995 data are the most recent available, and they produce the added benefit of enabling comparison across similar points in the business cycle: the unemployment rate was 5.8% in 1979 and 5.6% in 1995. Furthermore, Burtless (1990) found that cyclical effects on inequality were small to nonexistent in the 1980s. Given that he focused on yearly wage and salary earnings, this concern is mitigated further by my use of hourly earnings data, which are relatively insensitive to variation in hours and weeks worked over the year.

I use the CPS sample earnings weights for all estimates reported in this paper. Unlike DFL, however, I do not weight by hours worked. This enables greater flexibility in isolating job composition shifts associated with the shift from goods-producing to service-producing industries. For example, if service sector jobs are more likely to be part-time, and if part-time jobs pay less than full-time jobs, weighting by hours worked would undesirably downweight the wage inequality created by such shifts. Thus, my focus is on the distribution of earnings by job, rather than by hour.

IV. RESULTS

Summary Statistics and Densities

Table 1 shows summary statistics for ln(hourly earnings) for the 1979 and 1995 samples, with separate panels stratified by sex. I list mean ln(earnings), and the standard deviation as a simple dispersion measure. I provide a major industry breakdown in addition to the overall goods/services distinction, and I show employment shares by industry. These figures show substantially higher hourly earnings dispersion in 1995 than in 1979, a large reduction in mean real earnings for men, and a small increase in mean real earnings for women. There was a substantial decline between those years in the share of goods-producing jobs in

4. The deflator does not affect the dispersion measures, but it is useful for comparisons of means over time.

5. I did not directly account for top-coding of weekly earnings. Although this may affect the dispersion measures, the top-code is roughly at the same level in real terms in 1979 and 1995. To the extent that the share of very high wage workers increased over the period, the estimates in this paper may understate increasing dispersion due to increased mass in the upper tail.

TABLE 1

Mean and Standard Deviation of ln(Hourly Earnings), Industry Employment Shares, by Sex, Year, and Industry

A. Men

	1979			1995		
Industry	Mean	S.D.	Share	Mean	S.D.	Share
Total	2.55	.497	1.0	2.47	.641	1.0
Goods-Producing	2.62	.434	.430	2.52	.550	.339
Mining	2.78	.398	.016	2.65	.576	.009
Construction	2.62	.471	.097	2.47	.537	.091
Durable Manufacturing	2.64	.406	.208	2.56	.540	.152
Nondurable Manufacturing	2.57	.448	.108	2.49	.568	.087
Service-Producing	2.50	.534	.570	2.45	.681	.661
Trans., Comm., & Public Utilities	2.73	.446	.095	2.63	.602	.102
Wholesale Trade	2.58	.477	.048	2.49	.619	.053
Retail Trade	2.20	.459	.140	2.09	.621	.160
Finance, Insurance, & Real Estate	2.71	.559	.038	2.71	.679	.046
Services	2.46	.555	.178	2.49	.693	.240
Government	2.72	.433	.071	2.73	.558	.059
Total Observations		74,671			83,931	

B. Women

	1979			1995		
Industry	Mean	S.D.	Share	Mean	S.D.	Share
Total	2.16	0.44	1.0	2.23	0.61	1.0
Goods-Producing	2.21	.351	.197	2.25	.519	.132
Mining	2.49	.413	.003	2.52	.542	.002
Construction	2.28	.389	.010	2.34	.546	.012
Durable Manufacturing	2.27	.342	.091	2.30	.497	.057
Nondurable Manufacturing	2.14	.337	.094	2.19	.524	.063
Service-Producing	2.14	.459	.803	2.18	.540	.868
Trans., Comm., & Public Utilities	2.43	.408	.041	2.45	.543	.046
Wholesale Trade	2.22	.363	.024	2.28	.513	.024
Retail Trade	1.89	.349	.197	1.84	.504	.191
Finance, Insurance, & Real Estate	2.22	.354	.081	2.39	.538	.083
Services	2.18	.436	.414	2.31	.625	.474
Government	2.37	.397	.047	2.47	.502	.051
Total Observations		62,681			82,153	

Note: All tabulations are weighted by the CPS earnings weight, and 1979 earnings were inflated to 1995 levels using the GDP deflator for personal consumption expenditures.

the economy, particularly for men. Most of this change arose from a decline in manufacturing jobs (particularly durable manufacturing) and a rise in jobs in the services (narrowly defined) sector.

A key revelation from Table 1 is little or no convergence in the goods-producing and service-producing earnings distributions between 1979 and 1995. For men, there was a substantial decrease in the mean and substantial increase in the standard deviation in both broad sectors, and mean earnings are much higher in the goods-producing sector than in the service-producing sector. In contrast, for women mean earnings are very similar across the two broad sectors, and became more so over the period. Like men's jobs, however, for women earnings dispersion increased for all sectors, and the service sector as a whole (and for most subcategories) exhibits higher dispersion than does the goods sector. In general, this table is consistent with the view that the shift from goods-producing to service-producing jobs has increased inequality in hourly earnings, although this effect is likely to be much more pronounced for men.

Figures 1 and 2 show kernel density estimates of several unadjusted earnings distributions. Figure 1 shows the 1979 and 1995 distributions of hourly earnings, for men in Panel A and women in Panel B. These figures confirm the pattern identified in Table 1 of increasing dispersion for both men and women, and also the declining mean for men, between 1979 and 1995. For men, Figure 1 reveals that much of the increased dispersion is due to a shift from the middle of the distribution to the lower part, although there also is some added mass in the right tail of the 1995 distribution. For women, there appears to be a more uniform increase in dispersion across the upper and lower portions of the distribution. Furthermore, although not explicitly labeled in these figures, each distribution exhibits a pronounced spike at the real minimum wage, which declined substantially between 1979 and 1995. In their formal analysis, DFL attribute much of the increased inequality during 1979–1992 to the declining real minimum wage; Figure 1 also illustrates this effect, extended out by three years to 1995.

Figure 2 shows the earnings distributions for the goods-producing and service-producing sectors, by year and sex. For both men and women, the distributions in the two sectors have become more alike over time. However, each of these distributions became more disperse between 1979 and 1995, with the degree of dispersion remaining higher in services than in goods in all cases.[6]

FIGURE 1

EARNINGS DISTRIBUTIONS, 1979 AND 1995

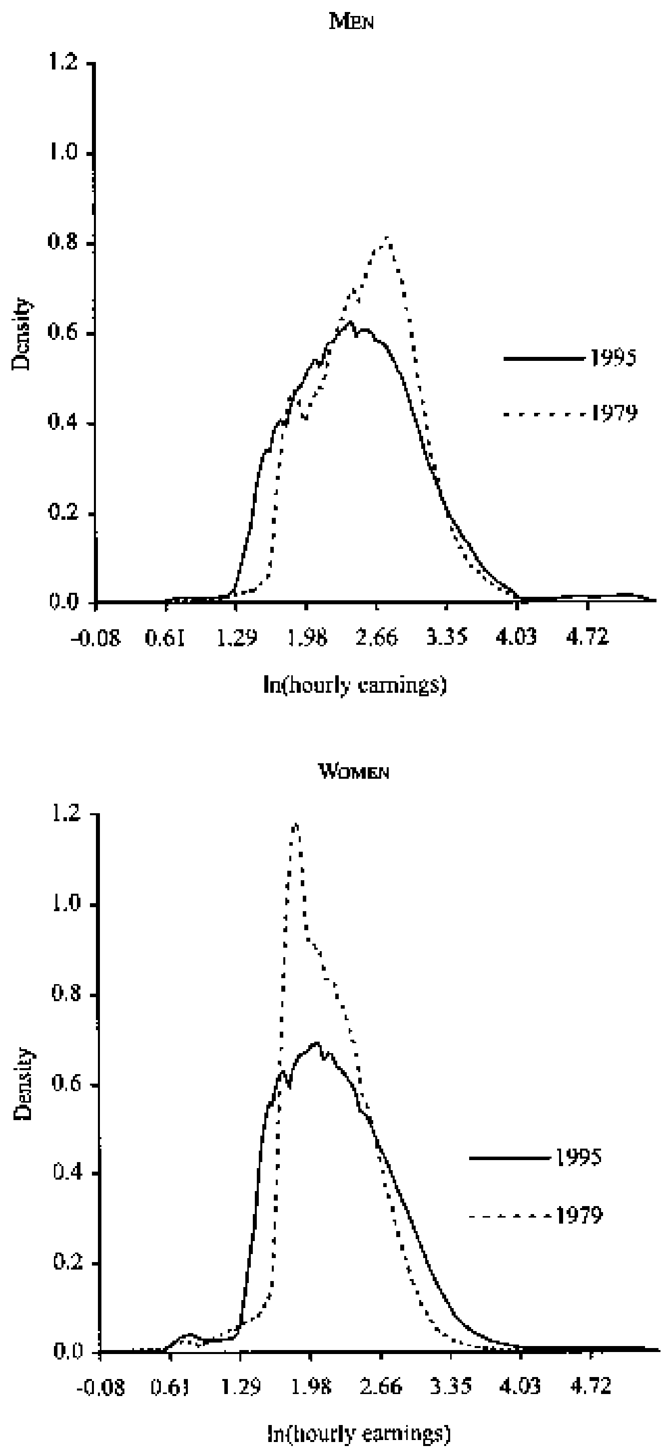

6. These tabulations and figures conflict somewhat with the results of Schweitzer and Dupuy (1995), who reported a substantial increase in overlap of the goods and service sector earnings distributions from 1979 to 1993. To the extent that my results differ from theirs, it is probably due to different sample definition: they used March CPS data on full-time workers who worked at least 39 weeks in the previous year, and they pooled men and women in their sample.

FIGURE 2

Earnings Distributions, 1979 and 1995, by Sex

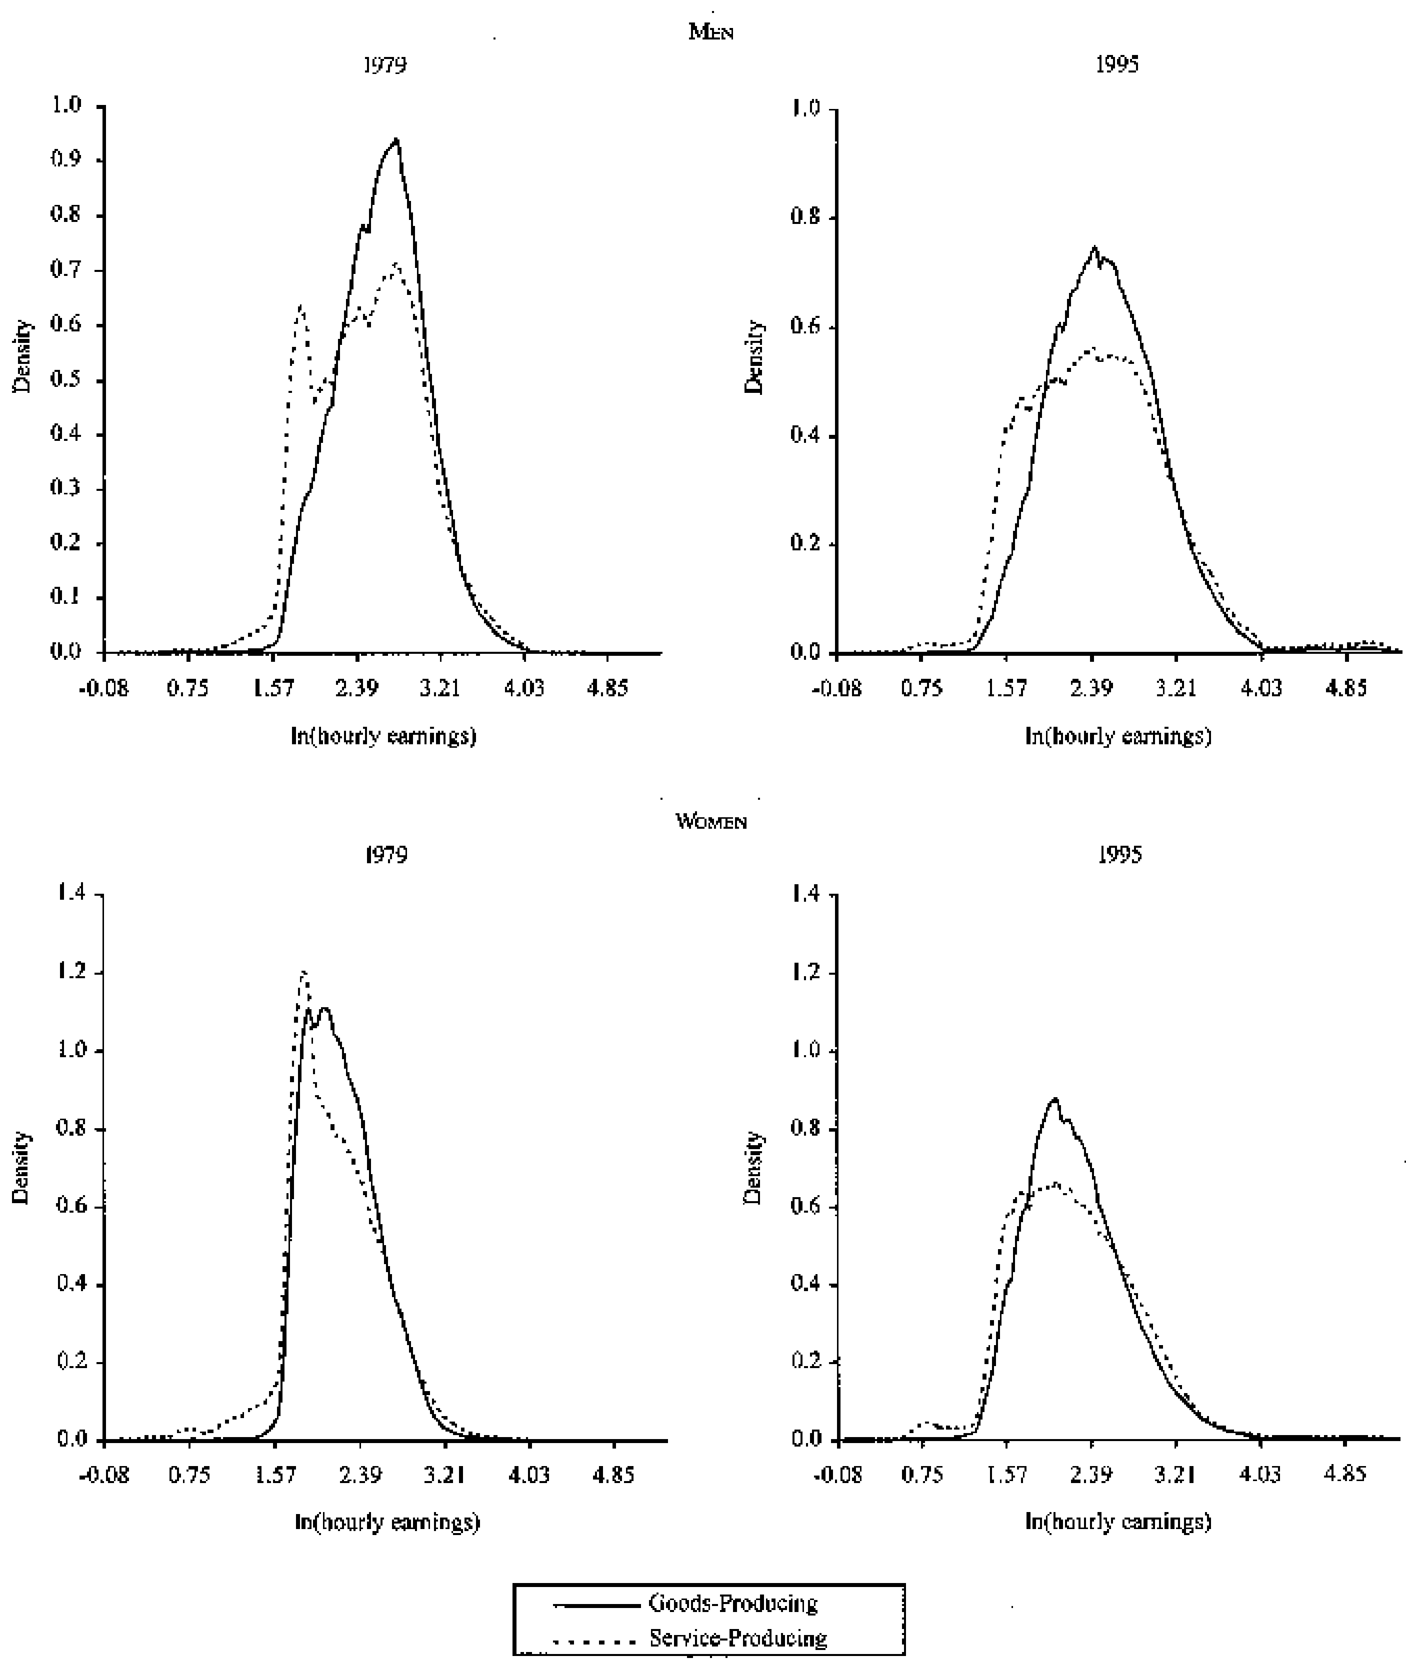

One difference between the two broad sectors that may help to explain the different earnings distributions is in the share of part-time jobs. Table 2 lists the mean and variance of earnings by broad sector and part-time status (and by sex). For men and women, mean earnings are lower, and the variance in earnings is higher, in part-time than in full-time jobs. Several changes occurred between 1979 and 1995 for both men and women. The most noticeable change is a large increase in the variance of earnings within all part-time categories listed; this increase dwarfs the increased

TABLE 2

MEAN AND STANDARD DEVIATION OF LN(HOURLY EARNINGS), BY PART TIME AND GOODS/SERVICES STATUS

A. MEN

	1979			1995		
INDUSTRY	MEAN	S.D.	SHARE[a]	MEAN	S.D.	SHARE[a]
TOTAL	2.55	.497	1.0	2.47	.641	1.0
Full Time	2.61	.464	.912	2.52	.547	.851
Part Time	2.01	.498	.088	2.23	.990	.149
GOODS-PRODUCING	2.62	.434	.430	2.52	.550	.339
Full Time	2.64	.421	.969	2.53	.498	.917
Part Time	2.12	.542	.031	2.46	.954	.083
SERVICE-PRODUCING	2.50	.534	.570	2.45	.681	.661
Full Time	2.58	.496	.870	2.51	.573	.818
Part Time	1.99	.487	.130	2.18	.990	.182
TOTAL OBSERVATIONS		74,671			83,931	

B. WOMEN

	1979			1995		
INDUSTRY	MEAN	S.D.	SHARE[a]	MEAN	S.D.	SHARE[a]
TOTAL	2.16	.441	1.0	2.23	.605	1.0
Full Time	2.23	.411	.724	2.31	.520	.700
Part Time	1.96	.457	.276	2.05	.738	.300
GOODS-PRODUCING	2.21	.351	.197	2.25	.519	.132
Full Time	2.22	.343	.915	2.26	.472	.877
Part Time	2.08	.415	.085	2.19	.780	.123
SERVICE-PRODUCING	2.14	.459	.803	2.18	.540	.868
Full Time	2.23	.431	.679	2.32	.528	.673
Part Time	1.95	.458	.321	2.05	.735	.327
TOTAL OBSERVATIONS		62,681			82,153	

Note: All tabulations are weighted by the CPS earnings weight, and 1979 earnings were inflated to 1995 levels using the GDP deflator for personal consumption expenditures.

[a] The full-time/part-time employment shares sum to 1 within each industry category (total, goods, services).

variance for full-time jobs. Mean earnings in part-time jobs also increased, particularly for men. Similarly, the part-time job share increased by 4 to 6 percentage points in most categories; the exception is service-producing women, for whom the share of part-time jobs remained essentially constant between 1979 and 1995. This latter fact suggests that any industry shift effects on inequality associated with increased part-time work will be greater for men than for women.

The following section presents conditionally weighted results. However, it is illustrative first to investigate the unconditional effect of the goods/services shift. This unconditional effect is obtained by upweighting the 1995 goods-producing sector observations by the relative goods share in 1979 versus 1995 (and downweighting the service sector observations by a similarly formed ratio for that sector). Figure 3 shows the impact on the male earnings distribution of this reweighting scheme, which does not account for any changes in the distribution of or returns to other earnings related characteristics. Relative to the actual distribution, the adjusted distribution has more weight around the median and less in the lower portion. This first pass at depicting the impact of the goods/services shift is consistent with the stereotypical view that the growing services share (as embodied in the solid "actual" line) is partially responsible for the erosion of the middle-class job base.

FIGURE 3

UNCONDITIONAL EFFECT
OF GOODS/SERVICES SHIFT, MEN

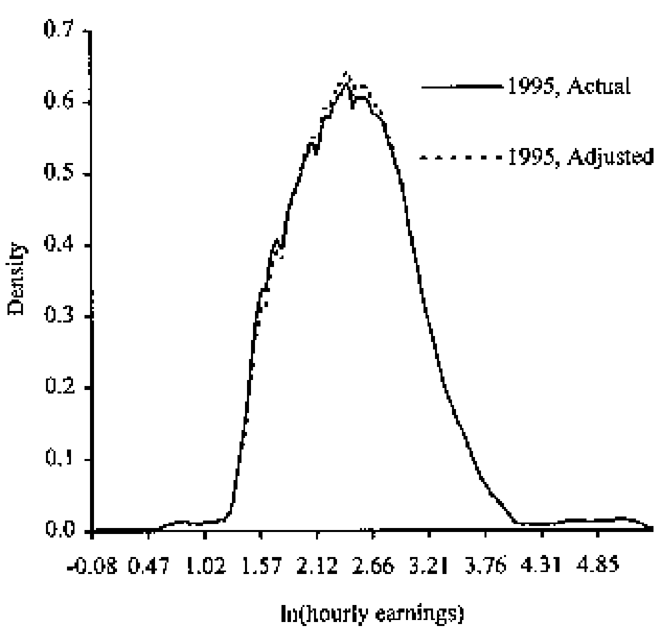

However, the corresponding figure for women (not shown) exhibits a much smaller unconditional impact of the goods/services shift.

Conditionally Weighted Density Estimates

The tabulations and densities in the previous section show the unconditional difference in the earnings distribution over time and across the goods-producing and service-producing sectors. It is likely, however, that the distribution of earnings-related characteristics differs across the two sectors, and that the earnings distributions conditional on these characteristics will differ less than the unconditional distributions. It is therefore important to condition on observables. To this end, I use a basic vector of X variables that includes a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence and marital status. Also, because the results in Table 2 suggest the potential importance of shifts between full-time and part-time jobs, I add a dummy for part-time work in additional analyses.

Figures 4 (men) and 5 (women) present the key results. For each panel, comparison of the solid line to the counterfactual dotted line shows the impact of the modeled change (industry structure or individual characteristics) as it actually evolved. Panel A in both figures shows the effect of accounting for the net shift from goods-producing to service-producing jobs between 1979 and 1995. This effect is estimated by reweighting the distribution through use of the conditioning weight $\psi_{E \div X}$. Panel B for each of these figures shows the effect of changing individual characteristics, which is estimated by use of the conditioning weight ψ_X.

For men, the adjusted distribution in Figure 4A reveals a small but discernible impact of industry employment shifts on the distribution of earnings. The adjusted distribution has slightly less mass in the lower portion and slightly more at or just above the middle; other portions of the adjusted and unadjusted distributions are nearly identical. This mass shift in the lower and middle portions is consistent with but smaller than the unconditional effect of the goods/services shift depicted in Figure 3. For women (Figure 5A), the conditional effect is barely discernible.

Figures 4B and 5B illustrate the impact of changing individual characteristics on the male and female earnings distributions. Their main impact for both men and women, as revealed by the comparison of the solid (actual) line to the dotted (adjusted) line, was to shift the distribution to the right. Also, the change in female characteristics and returns to them stretched the distribution somewhat from the median to the right.

FIGURE 4A

EFFECT OF GOODS/SERVICES SHIFT, MEN

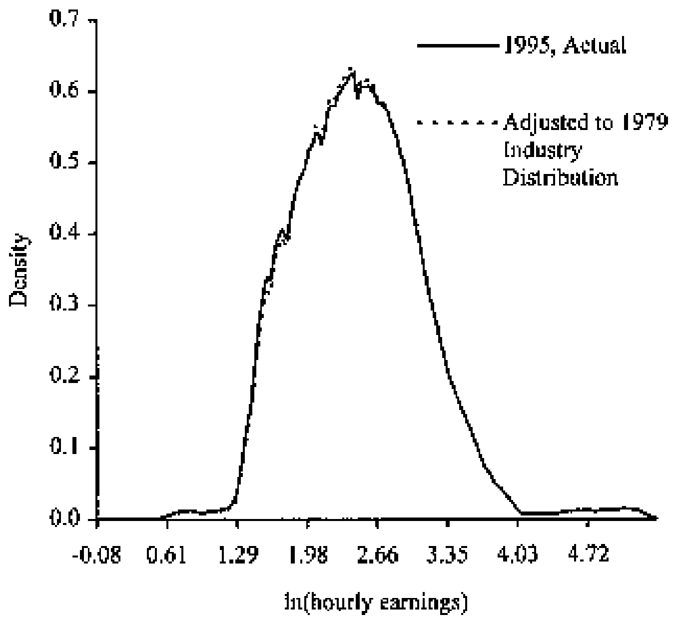

FIGURE 4B

EFFECT OF CHANGING *X*'S, MEN

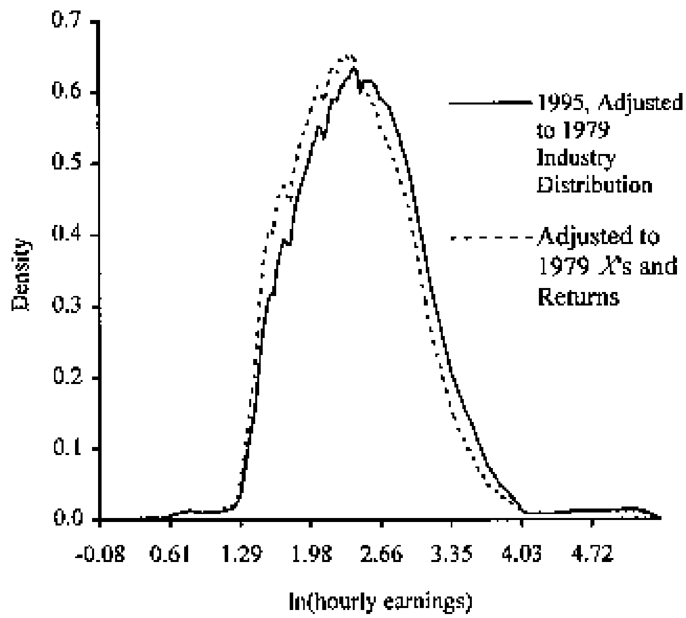

FIGURE 5A

EFFECT OF GOODS/SERVICES SHIFT, WOMEN

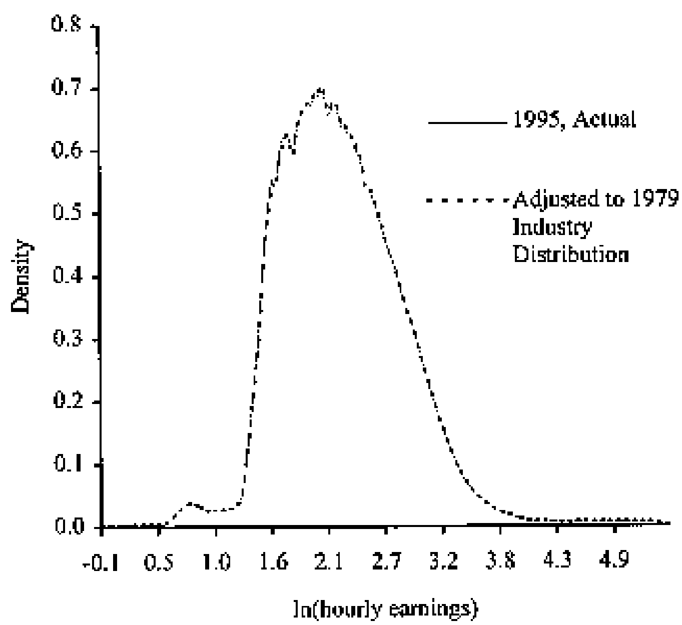

FIGURE 5B

EFFECT OF CHANGING *X*'S, WOMEN

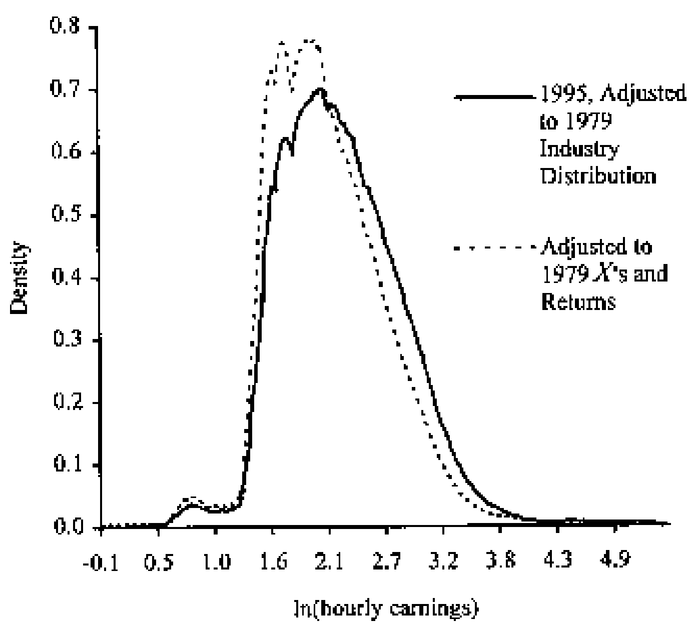

These figures provide a useful visual representation of the measured effects. As noted in Section I, however, the conditional weighted densities also can be used to assess the quantitative contribution of changing industry structure to changes in mean earnings and various dispersion measures. These results are reported in Table 3, for men and women separately. I analyze changes in the mean, standard deviation, and several quantile differences. I also analyze changes in two commonly used parametric inequality measures, the Gini coefficient and Theil's entropy measure.[7] For each measure, I list the total change in the measure between 1979 and 1995 and the amount explained by changing industry employment shares and changing individual characteristics.[8]

The results reported under the column "Goods vs. Services" in Table 3 show that broad industry employment shifts explain a small to moderate amount of the change in several earnings distribution measures for men. Broad industry shifts explain about 10% of the declining mean and 5% of the rising standard deviation. The largest impact on the changing quantile differences is for the 10–50 differential: the goods/services shift explains 43% of the increased dispersion in that range of the distribution. However, the lower portion of the male distribution changed far less than the upper portion; for example, the widening in the 10–50 differential is less than a quarter of the widening in the 50–90 differential. Among other measures, the goods/services shift also explains approximately 14% of the increase in the 10–90 differential. For women, the goods/services shift offset the increase in mean earnings somewhat, but had very little effect on the dispersion measures.

The final column of Table 3 shows the contribution of changing individual characteristics to the mean and dispersion measures. The impact of changing characteristics on mean male earnings was counterfactual. Also, although changing characteristics explain a substantial amount of the change in the 10–50 and 25–75 differentials, they explain very little of the increase in the other dispersion measures.

In contrast, for women the increase in mean earnings is more than fully explained by changing individual characteristics. These characteristics also explain a substantial portion of the change in various dispersion measures, including nearly 20% of the change in the standard deviation and the 10–50 differential and approximately 30% of the changes in the 5–95 and 25–75 differentials.

7. The 10–90 differential, for example, is defined as (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)); the other quantile measures are defined similarly. See DFL for a definition of the Gini and Theil indices.

8. Unlike DFL, I do not provide a full decomposition of the change in each measure.

TABLE 3

CONTRIBUTION OF CHANGING INDUSTRY SHARES AND INDIVIDUAL ATTRIBUTES TO THE CHANGING EARNINGS DISTRIBUTION, 1979–1995

		Portion explained by:	
Statistic	Total Change[a]	Goods vs. Services	Individual Characteristics[b]
A. Men			
Mean	–.079	–.008 (.102)	.090 (–1.14)
Standard Deviation	.143	.008 (.052)	.014 (.095)
10–90[c]	.259	.037 (.143)	.010 (.038)
10–50	.047	.020 (.426)	.017 (.368)
50–90	.213	.017 (.081)	–.007 (–.034)
25–75	.178	–.001 (–.005)	.037 (.209)
5–95	.379	.026 (.069)	.032 (.083)
Gini	.109	.004 (.042)	.002 (.021)
Theil	.178	.007 (.041)	–.004 (–.024)
B. Women			
Mean	.076	–.005 (–.068)	.128 (1.68)
Standard Deviation	.164	.004 (.024)	.030 (.184)
10–90	.466	.030 (.065)	.057 (.123)
10–50	.288	.021 (.075)	.053 (.185)
50–90	.179	.009 (.050)	.004 (.022)
25–75	.212	–.002 (–.012)	.072 (.339)
5–95	.515	0 (0)	.151 (.292)
Gini	.121	.002 (.016)	.016 (.129)
Theil	.180	.003 (.015)	.017 (.094)

NOTE: Percentage of total change explained is shown in parentheses.

[a] Difference between statistic in 1979 and 1995.

[b] The individual attributes include a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence and marital status.

[c] This is defined as the change between 1979 and 1995 in (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)). The other quantile measures are defined similarly.

The conditional results presented thus far are based on specifications that exclude any control for hours worked. However, as suggested by the tabulations presented in Table 2, there may be potentially important interactions between industry structure, earnings, and the share of part-time jobs. I therefore estimated additional models with a dummy variable for part-time work added to the list of individual characteristics. The results from this model are presented in Table 4.[9] The results in the second column reveal that inclusion of the part-time dummy in the conditioning equation reduces the estimated impact of the goods/services shift for men. Although the goods/services shift still explains about 26% of the increase in the 10–50 differential, the other estimated impacts are close to zero in percentage terms.

In contrast, inclusion of the part-time dummy substantially increases the share of the change in dispersion accounted for by individual characteristics. For men in Panel A, individual characteristics explain approximately 45% of the increase in the standard deviation and the Gini and Theil indices, and from 15% to more than 100% of the increase in the quantile dispersion measures. For women in Panel B, individual characteristics explain approximately 30% of the increase in the standard deviation and the Gini and Theil indices, and from 10 to 55% of the increase in the quantile dispersion measures.

I performed two primary checks of the robustness of these results. First, the results may be sensitive to the ordering of attribution imposed. Above, I assessed the contribution of the goods/services shift first, and the contribution of the *X*'s second; with this ordering, any interaction effects between the two are attributed to the goods/services shift. Therefore, I also performed the analysis in reverse order, letting the *X*'s affect the distribution first. This did not noticeably change the results for women. For men, however, this order reversal largely eliminated the impact of the goods/services shift on the 10–50 differential. Also, although the order reversal substantially increased the effect on the 25–75 differential in the model that excludes the part-time dummy, it did not do so in the model that includes the part-time dummy. Thus, it appears that in regard to their impact on male earnings inequality, there are important interaction effects between individual characteristics—particularly working part time—and the probability of working in goods versus services.

Another objection to the basic framework is that it does not account for changes in the general structure of the economy between 1979 and 1995. One way to assess how im-

9. I do not report corresponding kernel density estimates in additional figures, because in this model the actual density and density adjusted for industry shifts are indistinguishable.

TABLE 4

CONTRIBUTION OF CHANGING INDUSTRY SHARES AND INDIVIDUAL ATTRIBUTES TO THE CHANGING EARNINGS DISTRIBUTION, 1979–1995, PART TIME DUMMY ADDED

Statistic	Total Change[a]	Portion explained by: Goods vs. Services	Portion explained by: Individual Characteristics[b]
A. Men			
Mean	–.079	–.002 (.031)	.094 (–1.18)
Standard Deviation	.143	.002 (.015)	.066 (.460)
10–90[c]	.259	.019 (.072)	.086 (.331)
10–50	.047	.012 (.257)	.055 (1.17)
50–90	.213	.007 (.032)	.031 (.147)
25–75	.178	–.004 (–.021)	.040 (.224)
5–95	.379	.002 (.005)	.177 (.466)
Gini	.109	.001 (.011)	.046 (.423)
Theil	.178	.002 (.012)	.083 (.467)
B. Women			
Mean	.076	–.001 (–.020)	.118 (1.55)
Standard Deviation	.164	.002 (.013)	.053 (.322)
10–90	.466	.010 (.022)	.099 (.212)
10–50	.288	.007 (.024)	.079 (.274)
50–90	.179	.003 (.019)	.020 (.111)
25–75	.212	.000 (.002)	.115 (.540)
5–95	.515	0 (0)	.174 (.338)
Gini	.121	.001 (.009)	.034 (.285)
Theil	.180	.002 (.010)	.051 (.284)

Note: Percentage of total change explained is shown in parentheses.

[a] Difference between statistic in 1979 and 1995.

[b] The individual attributes include a linear measure of educational attainment, potential experience and its square, two race dummies, three region dummies, and dummies for SMSA residence, marital status, and whether worked part-time.

[c] This is defined as the change between 1979 and 1995 in (ln(earnings at the 90th percentile) – ln(earnings at the 10th percentile)). The other quantile measures are defined similarly.

portant such changes might be is to reverse the temporal ordering of the analysis: i.e., rather than imposing the 1979 industry and characteristics structure on the 1995 distribution of earnings, impose the 1995 structure on the 1979 distribution of earnings. Again, the results differ across the models that include or exclude the part-time dummy. In the model that excludes it, the results are very similar to those using the original temporal ordering. In the model that includes the part-time dummy, however, the estimated goods/services shift impact on the 10–50 differential is largely eliminated but replaced by a comparable impact on the 25–75 differential.

Overall, the estimated small effect of the goods/services shift on earnings dispersion in the lower half of the male distribution seems sensitive to the treatment of part-time work in the model. This finding, combined with the absence of an effect for women, suggests that the increase in part-time work by men in the services industry, as identified in Table 2, played a key role in any existing industry shift effects on earnings inequality. Furthermore, the most important measured characteristic in these models is part-time work. Inclusion of the part-time dummy in the model increases the share of increased dispersion explained by individual characteristics substantially, to nearly one-half for men and nearly one-third for women.

V. Conclusions

In this paper, I investigated the extent to which a substantial net shift from goods-producing to service-producing jobs altered the U.S. distribution of individual earnings between 1979 and 1995. Relative to previous work in this area, my paper's primary contribution is to apply recently developed semi-parametric estimation techniques to the problem. The analyses revealed four key empirical findings:

(1) Average earnings are lower and the dispersion of earnings is higher in service-producing than in goods-producing jobs.

(2) Consistent with (1), the unconditional effect of the shift from goods-producing to service-producing jobs has been to increase dispersion in the lower half of the earnings distribution.

(3) When individual characteristics are introduced into the model, a smaller but detectable impact on the lower half of the male earnings distribution remains. The quantitative impact was to increase the 10–50 earnings differential by several percentage points, nearly half the total change. However, this change is small relative to the large changes that occurred in upper half of the male earnings distribution. No similar effect was found for women.

(4) Result (3) is sensitive to controlling for part-time work. Although the estimated impact of the goods/services shift on the 10–50 differential largely withstands inclusion of a part-time dummy, additional checks revealed that this result is not fully robust to reversing the ordering of attribution or temporal ordering in the model.

The results from this analysis provide at best weak evidence in support of the view that the shift from goods-producing to service-producing jobs made an independent contribution to the erosion of middle-class earnings in the U.S. To the extent that an effect was isolated, its largest contribution was in the lower portion of the male distribution, which is consistent with the stereotype that shrinkage of the manufacturing sector has helped to erode the middle-class job base. However, this effect appears largely due to increased incidence of part-time work by men, particularly in the service-producing sector, which exhibited a sharp increase in earnings variance in part-time jobs. To the extent that increased part-time work by men was voluntary, this trend has limited adverse implications. However, if this trend reflects demand-side constraints, it may bode poorly for men stuck in part-time jobs. Furthermore, the increased incidence of part-time work appears to have made a large contribution to growing inequality for both men and women. The exact contribution of part-time work to growing inequality merits further investigation.

Overall, my results are much closer to those of authors (such as Juhn, Murphy, and Pierce 1990) who find no industry shift effects on earnings inequality than to those of authors (such as Maxwell 1989, 1990) who find large industry shift effects on earnings inequality. However, one key drawback of my approach is its broad measure of industry sectors (goods-producing vs. service-producing). It might be interesting to incorporate a finer industry breakdown into the analysis, although this extension may be problematic for the conditional weighted density estimation framework. In the meantime, additional applications of the conditionally weighted approach, as developed in DFL, appear warranted. This procedure is relatively easy to implement, and it is very powerful in regard to uncovering distributional changes and in its ability to perform additional tests on the altered distributions.

APPENDIX

Derivation of the Conditioning Weights

This appendix provides the derivation of the conditioning weights, $\psi_{E|X}$ and ψ_X, described heuristically in Section II. This discussion largely follows that in DiNardo, Fortin, and Lemieux (1996; DFL), although they provide a more complete and therefore more complex decomposition of changing earnings inequality.

Using the notation in the text, consider the distribution of wages *w* in year *t*, conditional on a vector of individual characteristics *X* and a dummy variable (*E*) indicating whether the worker is in a goods-producing sector job:

$$\text{(A1)} \qquad f_t(w) \equiv f(w; t_w = t, t_{E|X} = t, t_X = t).$$

DFL show that a distribution such as (A1) can be expressed as:

$$\text{(A2)} \qquad f_t(w) = \iint f(w|E, X, t_w = t)\, dF(E|X, t_{E|X} = t)\, dF(X|t_X = t).$$

In this equation, $f(.)$ is the conditional distribution of *w* and $F(.)$ is the joint distribution of *w*, *E*, and *X*. The right-hand side of (A2) indicates that the distribution of earnings in a given year can be expressed as the conditional distribution multiplied by the marginals (the first of which also is conditional) and integrated over *E* and *X*.

We are interested in the distribution of w if the distribution of *E* conditional on *X* is held to its 1979 structure:

$$\text{(A3)} \qquad f(w; t_w = 95, t_{E|X} = 79, t_X = 95).$$

Using (A2), this distribution can be expressed as:

$$\text{(A4)} \qquad f_t(w; t_w = 95, t_{E|X} = 79, t_X = 95)$$

$$= \iint f(w|E, X, t_w = 95)\, dF(E|X, t_{E|X} = 79)\, dF(X|t_X = 95)$$

$$= \iint f(w|E, X, t_w = 95)\, \psi_{E|X}(E, X)\, dF(E|X, t_{E|X} = 95)\, dF(X|t_X = 95),$$

where $\psi_{E|X}(E,X)$ is a reweighting function to be defined momentarily. It is very important to note that except for $\psi_{E|X}$, (A4) is identical to (A2) with $t = 95$— i.e., the distribution that we are interested in is equal to the unconditional distribution of earnings in 1995, with observations reweighted by the function $\psi_{E|X}$. If we can estimate $\psi_{E|X}$, it is straightforward to incorporate it and obtain the counterfactual distribution expressed in (A4) by using the observed univariate, unconditional distribution of wages in 1995.

The reweighting function is defined (identically) as:

$$\text{(A5)} \qquad \psi_{E|X}(E, X) \equiv \frac{dF(E|X, t_{E|X} = 79)}{dF(E|X, t_{E|X} = 95)}$$

$$= E \left(\frac{\Pr(E = 1|X, t_{E|X} = 79)}{\Pr(E = 1|X, t_{E|X} = 95)} \right)$$

$$+ (1 - E) \left(\frac{\Pr(E = 0|X, t_{E|X} = 79)}{\Pr(E = 0|X, t_{E|X} = 95)} \right).$$

The first line identity in (A5) is obtained by substituting the expression on the right side into (A4) and canceling out the denominator. The second line is derived by noting that *E* only takes the values 0 or 1, so that:

$$\text{(A6)} \qquad dF(E|X, t_{E|X} = t) \equiv E \Pr(E = 1|X, t_{E|X} = t) + (1 - E) \Pr(E = 0|X, t_{E|X} = t).$$

The second equality in (A5) follows from the recognition that one term in this expression will always equal zero.

This weight represents the change in the probability between 1979 and 1995 that an observation defined by characteristics *X* is in the goods-producing or service-producing sector. The probabilities in (A6) are easily recognized as expressions from standard binary dependent variable models. These conditional probabilities can be obtained by estimating a model such as a probit or logit and then using the fitted values. I use the probit equation

$$\text{(A7)} \qquad \Pr(E = 1|X, t_{E|X} = t) = \Pr(\varepsilon > -H(X)\beta)$$

$$= 1 - \Phi(-H(X)\beta)$$

to obtain the structure of $E|X$ at time *t*, where ε is a normally distributed random variable. In (A7), $H(X)$ is a vector function of *X* designed to capture the conditional relationship being modeled, and β is a vector of estimated coefficients. This equation is estimated for both the 1979 and 1995 samples, and the coefficients are retained. We use these results to fit the probabilities in (A5) using the 1995 sample *X*'s, combined with the 1979 coefficients for the numerator and the 1995 coefficients for the denominator. The resulting estimated weights, $\hat{\psi}_{E|X}$, are incorporated into the kernel density estimation or into the tabulation of distributional statistics, as described in Section II.

A modification of this procedure enables us to account for the impact of changes in the *X* vector of earnings-related characteristics. In this case, the weighting function is obtained through a simple application of Bayes' Law:

$$\text{(A8)} \qquad \psi_X(X) = \frac{\Pr(t_X = 95)}{\Pr(t_X = 79)} \, \frac{\Pr(t_X = 79|X)}{\Pr(t_X = 95|X)} .$$

This function represents the relative probability of observing an individual with characteristics X in the 1979 versus the 1995 sample, normalized by the unconditional probabilities of being in either sample. As long as the distribution of X's changed between the two years (for example, through higher average educational attainment), the weights ψ_X will alter the estimated distribution.

The function ψ_x is estimated by pooling the 1979 and 1995 data sets, and then estimating a binary dependent variable model for a dummy variable indicating the sample from which the observation is obtained. The conditional probabilities $\Pr(t_X = t|X)$ are obtained by forming fitted probabilities for workers in the 1995 sample, based on their X values. The unconditional probabilities, $\Pr(t_X = t)$, are simply the weighted shares of the 1979 and 1995 samples in the pooled sample. Estimation is then performed on the 1995 sample, with the estimated weights $\hat{\psi}_X$ modifying the sampling weights (as described in Section II).

Two points should be noted. First, the conditional probability estimating equation (A7) has no behavioral interpretation; it simply permits conditioning out the effect of covariates (X) which may be related to the factor (industry employment shifts) whose effect we are attempting to estimate. Second, due to potentially important interactions between the effects of industry shifts and changing individual characteristics, I also estimated models that reverse the order of attribution, by first assessing the contribution of the X's, and then assessing the contribution of E. The exact procedure is described in DFL.

REFERENCES

Blackburn, McKinley L. 1990. "What Can Explain the Increase in Earnings Inequality Among Males?" *Industrial Relations* 29 (3) pp. 441–456.

__________, David E. Bloom, and Richard B. Freeman. 1990. "The Declining Economic Position of Less Skilled American Men." In *A Future of Lousy Jobs,* ed. Gary Burtless, pp. 31–67. Washington, D.C.: Brookings.

Bluestone, Barry, and Bennett Harrison. 1988. "The Growth of Low-Wage Employment: 1963–86." *American Economic Review* 78 (2) pp. 124–128.

__________, and __________. 1982. *The Deindustrialization of America.* New York: Basic Books.

Bound, John, and George Johnson. 1992. "Changes in the Structure of Wages in the 1980s: An Evaluation of Alternative Explanations." *American Economic Review* 82 (3) pp. 371–392.

Burkhauser, Richard V., Amy D. Crews, Mary C. Daly, and Stephen P. Jenkins. 1996. *Income Mobility and the Middle Class.* Washington, D.C.: American Enterprise Institute.

Burtless, Gary. 1990. "Earnings Inequality over the Business and Demographic Cycles." In *A Future of Lousy Jobs,* ed. Gary Burtless, pp. 77–116. Washington, D.C.: Brookings.

Delgado, Miguel A., and Peter M. Robinson. 1992. "Nonparametric and Semiparametric Methods for Economic Research." *Journal of Economic Surveys* 6 (3) pp. 201–250.

DiNardo, John, Nicole M. Fortin, and Thomas Lemieux. 1996. "Labor Market Institutions and the Distribution of Wages, 1973–1992: A Semiparametric Approach." *Econometrica* 64 (5) pp. 1001–1044.

__________, and Thomas Lemieux. 1994. "Diverging Male Wage Inequality in the United States and Canada, 1981–88: Do Unions Explain the Difference?" Irvine Economic Paper No. 93–94-16 (June). University of California, Irvine.

Juhn, Chinhui, Kevin M. Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy* 101 (3) pp. 410–442.

Karoly, Lynn. 1996. "Anatomy of the U.S. Income Distribution: Two Decades of Change." *Oxford Review of Economic Policy* 12 (1) pp. 77–96.

__________. 1992. "Changes in the Distribution of Individual Earnings in the United States: 1967–86." *Review of Economics and Statistics* 74 (1) pp. 107–115.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–87: Supply and Demand Factors." *Quarterly Journal of Economics* 107 (1) pp. 35–78.

Levy, Frank, and Richard J. Murnane. 1992. "U.S. Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations." *Journal of Economic Literature* 30 (September) pp. 1333–1381.

Maxwell, Nan. 1990. *Income Inequality in the United States,1947–85.* Westport, CT: Greenwood Press.

__________. 1989. "Demographic and Economic Determinants of United States Income Inequality." *Social Science Quarterly* 70 (2) pp. 245–264.

Murphy, Kevin M., and Finis Welch. 1993. "Industrial Change and the Rising Importance of Skill." In *Uneven Tides:Rising Inequality in America,* eds. Sheldon Danziger and Peter Gottschalk. New York: Russell Sage Foundation.

Schweitzer, Mark E., and Max Dupuy. 1995. "Sectoral Wage Convergence: A Nonparametric Distributional Analysis." Working Paper 95-20 (December). Federal Reserve Bank of Cleveland.

Sheather, S.J., and M.C. Jones. 1991. "A Reliable Data-based Bandwidth Selection Method for Kernel Density Estimation." *Journal of the Royal Statistical Society* B 53 (3) pp. 683–690.

Silverman, B.W. 1986. *Density Estimation for Statistics and Data Analysis.* London: Chapman and Hall.

Interpreting Procyclical Productivity: Evidence from a Cross-Nation Cross-Industry Panel

J. Bradford De Long
and Robert J. Waldmann

De Long is at U.C. Berkeley, the National Bureau of Economic Research, and the Federal Reserve Bank of San Francisco; Waldmann is at the European University Institute.

The authors would like to thank Hoang Quan Vu for research assistance, the European University Institute and the National Science Foundation for research support, and Olivier Blanchard, Susanto Basu, Ricardo Caballero, Tim Cogley, John Fernald, Richard Freeman, Lars Jonung, Richard Lyons, Glenn Rudebusch, Alan Stockman, Larry Summers, and especially Larry Katz for helpful and interesting discussions.

We use an international panel data set of value added by industry to see if labor productivity is procyclical in response to demand shocks. It is:holding fixed our proxy for supply-side factors—the value added levels of an industry in other nations—industry-level productivity rises when value added in the rest of manufacturing rises.

Moreover, increases in unemployment are associated with a lowered degree of procyclicality in the U.S. and with heightened procyclicality in Europe. This suggests that procyclical productivity arises primarily from "labor hoarding" by firms in the U.S. that wish to avoid future training costs and primarily from "job hoarding" by workers in Europe who wish to avoid unemployment.

Labor productivity is procyclical, rising in business expansions and falling in recessions.[1] Some believe that procyclical productivity is ingrained in the technology of production. But a standard view of procyclical productivity sees it as a consequence, not a cause, of changes in activity. Labor productivity falls when output falls because firms retain more workers than required to produce low current output. They do this to avoid the costs of laying workers off now and hiring replacements in the future when activity recovers.[2] Procyclical productivity does not cause but results from business cycles because firms "value the match" that they have made with their employees.[3]

This account has been challenged by real business cycle theories which speculate that the shocks driving the business cycle are not shocks to demand, but are instead technology-driven shocks to productivity in particular industries (for example, Kydland and Prescott 1982; Long and Plosser 1983). Such industry-specific technology shocks directly cause an increase in production in the affected industry. They cause increased production in other industries by (i) increasing the wealth of consumers, (ii) increasing demand for intermediate inputs used in the directly affected industry, and (iii) increasing demand for (gross) complements of the output of the directly affected industry.[4]

Such theories have been criticized on the grounds that they cannot account for correlations in *productivity* (Summers 1986) though they might account for correlations in *output* across industries. Demand spillovers from positive technology shocks in one industry should lead to reduced labor productivity in other industries.[5] But production and

1. See, for example, Hultgren (1960), Okun (1962), Shapiro (1987, 1993), Caballero and Lyons (1990, 1992), Bils and Cho (1992), Solon, Barsky, and Parker (1994), Basu and Fernald (1995 and forthcoming).

2. See Holt, et al. (1960), Oi (1962), and Okun (1962).

3. This literature is reviewed by Fair (1969), Hamermesh (1976), and Nickell (1986). See also Aizcorbe (1992), Rotemberg and Summers (1988), Summers and Wadhwani (1988), and Medoff (1979).

4. There are also views in which procyclical productivity is both cause and effect, and multiple equilibria are possible. See Benhabib and Farmer (1996), or Murphy, Shleifer, and Vishny (1989).

5. Under the assumption that the short-run marginal product of labor is decreasing.

productivity are positively correlated across industries. Some have made the assumption that supply-side shocks are industry-specific, while demand shocks are aggregate. They have regressed productivity in an industry on total productivity or production to show that demand shifts—not shifts in supply—underlie procyclical productivity.[6]

This, however, is not a convincing refutation of supply-side theories. There are supply-side shocks that affect labor productivity in many industries at once: the oil price shocks of 1973 and 1979, for example. Such shocks generate procyclical productivity in many industries: producers shift away from the now-expensive factor of energy and use labor more intensively. A real business cycle-driven response of economies to oil shocks creates aggregate movements in productivity and output. Economists assuming that "supply" is industry-specific and "demand" aggregate would falsely interpret such movements as evidence that procyclical productivity was demand-driven.

This paper attempts to take some steps toward disentangling demand- and supply-driven components of procyclical productivity without making the possibly dangerous assumption that everything aggregate is demand. It uses Alan Stockman's insight that there are a great many technology and cost shocks—like the oil shocks of 1973 and 1979—that directly affect productivity in many *industries* and also affect productivity in many *nations*. A cross-industry cross-nation panel of data on value added by industry can be used to separate the effects of demand and supply shocks if supply shocks are truly "technological"—that is, they affect the production process no matter where in the world it happens to be located.[7]

Industrial value added shifts correlated with value added shifts in other industries in the same economy, and yet not correlated with industrial value added shifts in other nations, are candidates for the label "demand." How could a change in an industry's *technology* of production affect other industries in the same country but not the same industry in other countries? Industrial value added shifts correlated with value added shifts in the same industry in other nations, but not correlated with value added shifts in other industries in the same country, are candidates for the label "supply."

The effects of idiosyncratic national aggregate demand shocks can be determined because such shocks are both intersectoral and nation-specific.

Needless to say, we do not believe that "supply" shifts caused by technology or even changes in prices diffuse instantly across the nations of our sample. However, we do believe that such shifts ought to spread over the countries in our sample within a few years. And it is certainly important to control for such supply shifts before concluding that procyclical productivity is demand-driven.

We are concerned that much of the evidence on procyclical productivity is driven by relatively low-frequency changes in productivity and output growth—which was high in the 1960s, low during a period from the early 1970s to the early 1980s, and moderate through the later 1980s. Such low-frequency changes might well be supply-driven, rather than demand-driven. We do not believe that supply-side effects will be adequately removed by using instruments, for example U.S. military spending, that while not causally related to supply factors nevertheless have much of their own variance produced by low-frequency movements.

The identifying assumptions we require are relatively minor. One need not assume that technological progress is uniform across countries. One need only assume that there are no technological or other supply-side shocks that are (a) specific to a single country, yet (b) affect a broad range of industries within manufacturing.[8]

We find that even after controlling for industry-specific cross-nation shocks, sectoral productivity growth remains positively correlated with aggregate manufacturing output. This suggests that increased aggregate demand does lead to increased labor productivity, and that there is a component of procyclical productivity that could be accounted for by an old-fashioned Keynesian "labor hoarding" story, or by some other model in which firms and workers value their match.

We go on to investigate the cross-nation pattern of procyclical productivity. If firm-side labor hoarding—due to workforce finding and training costs—is important, productivity should be more procyclical when unemployment is low.[9] When unemployment is high, laid off workers are less likely to find new jobs and are more likely to be avail-

6. See Hall (1986); Domowitz, Hubbard, and Petersen (1988); Caballero and Lyons (1990 and 1992); and Shapiro (1987). These studies place the measured Solow residual on the left hand side of their equations, while we focus on labor productivity.

7. See Stockman (1988).

8. We use labor productivity and not total factor productivity as a dependent variable, and so our results on procyclical productivity cannot be attributed to market power. Our calculations are not affected by deviations of prices from marginal products.

9. In the United States, labor productivity appears less procyclical in highly unionized industries (see Medoff, 1979; Freeman and Medoff, 1984). This might arise because unionized workers share rents, would be likely to suffer a cut in wages if they took jobs in non-union establishments, and so wait to be around. Thus the firm is free to lay them off temporarily when demand is momentarily slack without risking the loss of the value of the match.

able when the firm wishes to recall them, so the incentives for the firm to engage in labor hoarding are diminished.

If worker-side job hoarding—firing costs that firms bear when workers are laid off but avoid when workers voluntarily quit—is important, then productivity should be more procyclical when unemployment is high: workers will then resist dismissals with more determination, quits will be rare, and restrictions on layoffs may bind more when unemployment is high.[10]

If procyclical labor productivity is simply a consequence of increasing returns to scale, as in Caballero and Lyons (1990 and 1992) or many others, the procyclicality of production should be unaffected by the level of the unemployment rate.

We find that in the United States labor productivity is less procyclical when unemployment is high. In Germany and—less strongly—in Britain and in Europe as a whole, however, there is some weak evidence that productivity is more procyclical when unemployment is high. There are differences between the U.S. and Europe in sign and strength of the relation between the degree of procyclicality in productivity and unemployment.

This difference suggests that demand-driven procyclical productivity may spring more from labor hoarding in the United States and more from job hoarding in Europe.[11] The dependence of the cyclical behavior of productivity in these nations on labor market conditions raises the possibility that procyclical productivity arises from national institutions that mold the dynamic relationships between workers and firms and is not simply the result of an increasing returns to scale technology.

After this introductory section, Section I describes the data used in this paper. Section II presents the evidence on the existence of procyclical productivity in response to demand shocks. Section III correlates the degree of procyclical productivity with the unemployment rate. It leads to the tentative conclusion that "worker hoarding" by firms is relatively more important as a cause of procyclical productivity in response to demand shocks in the United States, while "job hoarding" by workers is relatively more important as a cause of procyclical productivity in response to demand shocks in Germany and perhaps in Britain. Section IV concludes.

I. An International Intersectoral Panel of Value Added by Industry

We use the OECD International Sectoral Data Bank as our primary data source (Meyer-zu-Schlochtern, 1988). Our data set contains annual data on real value added, employment, and capital by industry for fourteen OECD nations. Since our approach requires a balanced panel and we strongly desire sample of long length, we are forced to focus on seven nations for which data on real value added are available from the 1960s onward—Belgium, Finland, France, Germany, Norway, the United Kingdom, and the United States.[12] Data are available for seven ISIC industries within the manufacturing sector: food, textiles, paper, chemicals, non-metallic minerals (i.e., stone, clay, and glass), basic metal production, and mechanical equipment. Observations are available from 1966 to 1987.

Unfortunately, changes in data collection and definition keep us from extending our sample beyond 1987 while still retaining data from the late 1960s and early 1970s: we have chosen to maximize our sample length.[13]

The OECD international sectoral database includes employment by industry, but it does not include average hours worked by industry. We augment the data by multiplying employment by average hours worked in manufacturing.[14] This procedure assumes a perfect correlation between average hours worked in different industries. Thus it induces positively correlated measurement error between total hours worked in different industries.

Since hours are correlated with value added, this measurement error induces a negative correlation between value added per man-hour in one industry and in another. The

10. See Blanchard and Summers (1986), Bentolila and Bertola (1990), and Krugman (1988).

11. Abraham and Houseman (1989) use a panel of ten matched U.S. and German manufacturing industries, and find that the immediate effect of a reduction in shipments on employment is much smaller than Germany. They interpret their findings as implying that German firms, because of worker job hoarding, are less free in the short run to use layoffs to adjust unemployment. They also find that the workforce adjustment process was slower in Germany after 1972. In 1972 German legal restrictions on layoffs were significantly strengthened by the Works Constitution Act, and the post-1972 period has seen higher unemployment. Abraham and Houseman, however, are unable to control for changes in technology and costs—particularly the cost of oil—and are forced to assume that production is exogenous. Our broader panel of countries should make it possible to control for such factors to some degree.

12. Netherlands data are also available for the 1960s. Unfortunately a change in definitions in 1970 makes Dutch data from the 1960s incomparable to data from 1970 on.

13. We have experimented with alternative data definitions that omit the first years of our sample and contain more recent observations. The statistical results we obtain are very similar with one notable exception: the United Kingdom's pattern of procyclical productivity is closer to that of the United States and further from that of the rest of Europe the more recent the data. We speculate that this reflects changes in the British economy as a result of Prime Minister Margaret Thatcher's attempts in the 1980s to curb the power of British unions and make the British labor market more "competitive."

14. Unpublished data were kindly provided by Robert Gordon.

use of average hours in manufacturing, instead of average hours in each industry, biases the data against revealing procyclical labor productivity.[15]

We do not possess data on average hours worked in Finland. Reported average hours worked can be found in the I.L.O.'s *Labor Statistics Yearbook,* but reported average hours from this source show a large increase, from 38.5 hours per week in 1978 to 41 hours per week in 1979. This shift is large relative to other variations, and we believe it reflects a change in coverage. Finland, therefore, was excluded from all regressions that required average hours worked.

In addition, data on employment in basic metals and in equipment are not available for France or Belgium in the 1960s. Regressions using these industries as dependent variables therefore use data since 1970 only.

II. Procyclical Productivity and Aggregate Demand

Nation-Specific Aggregate Demand Movements

It is fruitless to try to use nation-specific movements in manufacturing value added to identify demand-driven movements in productivity unless such movements exist. Stockman (1988) has already used an international intersectoral panel to identify demand-specific and supply-specific movements in output. Stockman assumed that all industries have the same cyclical responsiveness to shifts in aggregate demand—that, in the language of finance, all have the same β_i with respect to aggregate output—and that all countries have the same γ_i responsiveness to international supply shocks.[16] In spite of these restrictive assumptions on the form of his nation- and industry-specific components, Stockman found that 12.2% of variance of industry value added is accounted for by nation-specific components that are orthogonal to industry-specific value added movements, and that 14% of variance is accounted for by industry-specific components orthogonal to nation-specific output movements.

Waldmann (1991), using the OECD database, estimated nation and industry effects without imposing the assumption that β_i and γ_i coefficients were constant across industries. He found that orthogonal nation effects account for 17% of the variance in real value added, while orthogonal industry effects account for only 9.5% of the variance. He also found that orthogonal nation effects accounted for a very small fraction of the variance in real value added in small open economies such as Belgium and Finland; this is reassuring, because standard open-economy models suggest that countries like Belgium and Finland should not have a significant nation-specific business cycle. Results from Waldmann (1991) are reproduced as Table 1.

The divergence of the strength of nation-specific movements in manufacturing value added leads us to anticipate that our attempts to identify demand-driven procyclical productivity will have almost no power in small open economies like Belgium, Finland, and Norway. The existence of large nation-specific components in value added for larger countries leads us to anticipate that our procedures will have considerable power for large countries—the polar case of the United States, and also France, Germany, and the United Kingdom—where spillovers of demand shocks are smaller, and where there appears to be more of a nation-specific business cycle.[17]

Initial Regressions

The growth of value added per man-hour for industry i in nation n was regressed on the growth of value added of the rest of manufacturing in nation n, and on the average growth of production per man-hour in industry i in other countries, as described in equation 1:

$$(1)\quad \Delta\{\log(Y/N_{int})\} = c_{ni} + \beta_{ni}[\Delta\{\log(Y_{nt} - Y_{int})\}] + \gamma_{ni}[\Delta\{\log(Y/N_{i(-n)t})\}] + \varepsilon_{int}\,,$$

15. An additional data problem is posed by the fact that labor productivity data are not available whenever value added data are available. Oddly, the OECD does not report total manufacturing employment in the U.S. before 1968. The ISIC classifications used by the OECD do not correspond exactly to SIC classifications, and so comparable data cannot be added from BLS sources. The OECD does, however, provide data on wage and salary employment in U.S. industries for the 1960s. Fitted values from a regression of total employment on wage and salary employment were therefore used as a proxy for total employment. The R^2 of these regressions ranges from 99.5% to 99.9%. We conclude that it is unlikely that this neglect of the self-employed induces significant biases.

16. Another possibly dubious assumption. For example, the U.S. imposed oil price controls after the 1973 oil shock, and so the real price of oil in the U.S. did not rise as much as in other nations. We would be surprised if substitution away from intensive use of energy proceeded as fast in the U.S. as in Europe after 1973.

17. It would not be appropriate to draw the conclusion that aggregate demand shocks account for twice as much of the variance in the typical industry's value added growth rate as supply shocks. Undoubtedly, most of both supply- and demand-side shocks are left unidentified by our procedures. We wish only to maintain that the 17% of industry value added growth rate variance that is (a) correlated with changes in the rest of manufacturing production in the same country while (b) orthogonal to changes in value added in the same industry in other countries is not supply. (Conversely, the 9.5% of industry value added growth rate variance that is (c) correlated with changes in value added in the same industry in other countries but (d) orthogonal to changes in the rest of manufacturing production in the same country is not demand.)

TABLE 1

Share of Industry Value Added Growth Variance Accounted for by Orthogonal Country and Industry Effects

		USA	DEU	FRA	BEL	FIN	NOR	UK	Average	Ratio
Food	Country	0.030	0.198	0.183	0.072	0.152	0.048	0.338	0.101	2.267
	Industry	0.062	0.032	0.031	0.027	0.040	0.029	0.219	0.044	
Textiles	Country	0.389	0.207	0.093	0.071	0.291	0.034	0.693	0.248	5.075
	Industry	0.007	0.002	0.057	0.129	0.029	0.053	0.003	0.050	
Paper	Country	0.240	0.213	0.161	0.003	0.186	0.042	0.124	0.147	1.157
	Industry	0.021	0.049	0.011	0.177	0.288	0.137	0.108	0.127	
Chemicals	Country	0.169	0.022	0.025	0.137	0.062	0.023	0.182	0.091	0.521
	Industry	0.096	0.376	0.210	0.002	0.320	0.168	0.082	0.174	
Stone, Clay,	Country	0.589	0.112	0.038	0.048	0.312	0.052	0.257	0.214	3.456
and Glass	Industry	0.035	0.097	0.150	0.056	0.002	0.063	0.040	0.062	
Basic Metals	Country	0.258	0.095	0.099	0.076	0.021	0.004	0.279	0.126	1.266
	Industry	0.076	0.017	0.109	0.225	0.094	0.147	0.014	0.099	
Mechanical	Country	0.730	0.276	0.303	0.109	0.025	0.009	0.218	0.295	5.698
Equipment	Industry	0.038	0.039	0.020	0.106	0.002	0.121	0.063	0.095	
Average	Country	0.375	0.137	0.109	0.083	0.121	0.022	0.292	0.170	1.789
	Industry	0.056	0.118	0.097	0.099	0.130	0.118	0.049	0.095	

where Y/N denotes value added per man-hour; subscripts i, n, and t run over industries, nations, and years, respectively; Y_{nt} refers to value added in all of manufacturing in country n in year t; $Y_{nt} - Y_{int}$ denotes value added in manufacturing in country n in year t in all industries except industry i; and a subscript ($-n$) denote averages over the other countries in the sample (i.e., excluding country n) for an industry i. Results from estimating equation 1 are reported in Tables 2 through 4.

Table 2 reports the β coefficients, which measure the sensitivity of industry-level value added per man hour to movements in value added in the rest of manufacturing in the same nation (holding constant value added in that particular industry in other nations). The β coefficients on the growth of manufacturing are generally positive. The precision-weighted average is positive for all nations. The precision-weighted average across countries of β coefficients for a given industry vary strikingly.

Under the assumption that the disturbance terms for different industries are independent, standard errors for the precision-weighted averages within industries and within nations of the β coefficients were calculated and are reported in Table 2. However, this assumption is not valid. Instead, seemingly-unrelated-regressions procedures were used to test the null hypothesis that labor productivity is not procyclical when controlling for value added growth in the same industry in other countries. Data on different industries were stacked, and equation 1 was reestimated with the β coefficient restricted to be the same in different industries.[18]

Table 2 also reports seemingly-unrelated-regressions estimated coefficients on the growth of the rest of manufacturing in the same country. The coefficients are all positive and are similar to the precision-weighted national average OLS estimated β coefficients. Their rank order is almost unchanged. All coefficients, save that of France, are within one standard error of the precision-weighted national average coefficients. The reported standard errors are somewhat larger for the SUR estimates.[19]

18. Unfortunately, when seemingly-unrelated-regressions procedures are used and international averages are included, the sample contains only those years in which all industries in all countries report data. The absence of data from the 1960s on employment in the basic metals and metal equipment industries in France and Belgium leaves only five industries in six countries.

19. Since the true standard errors must be smaller, this implies that the reported standard errors of the weighted averages are understated.

TABLE 2

β COEFFICIENTS OF VALUE ADDED PER MAN-HOUR REGRESSED ON THE GROWTH OF MANUFACTURING IN THE SAME COUNTRY

	OBSERVATIONS	USA	DEU	FRA	BEL	NOR	UK	AVERAGE
FOOD	22	–0.100 (0.127)	–0.042 (0.088)	0.172 (0.226)	–0.033 (0.123)	0.363 (0.333)	0.024 (0.076)	–0.006 (0.047)
TEXTILES	22	0.011 (0.123)	–0.093 (0.163)	0.338 (0.213)	0.249 (0.234)	–0.025 (0.262)	0.292 (0.185)	0.093 (0.073)
PAPER	22	0.048 (0.166)	0.224 (0.131)	0.186 (0.294)	–0.097 (0.195)	0.117 (0.244)	0.495 (0.172)	0.181 (0.074)
CHEMICALS	22	0.017 (0.151)	0.078 (0.238)	–0.077 (0.229)	0.860 (0.411)	0.465 (0.358)	0.247 (0.214)	0.127 (0.093)
STONE, CLAY AND GLASS	22	0.201 (0.103)	0.290 (0.122)	0.479 (0.242)	0.393 (0.267)	–0.292 (0.357)	0.345 (0.206)	0.259 (0.067)
BASIC METALS	14	0.526 (0 445)	–0.228 (0.440)	–0.346 (0.422)	0.091 (0.345)	–0.216 (0.877)	0.288 (0.651)	0.025 (0.189)
MECHANICAL EQUIPMENT	14	0.330 (0.213)	0.137 (0.113)	–0.130 (0.283)	0.169 (0.273)	0.113 (0.191)	0.147 (0.191)	0.142 (0.074)
AVERAGE		0.079 (0.055)	0.088 (0.050)	0.150 (0.096)	0.088 (0.081)	0.107 (0.109)	0.153 (0.057)	0.106 (0.027)
SUR ESTIMATE[a]		0.037 (0.079)	0.085 (0.069)	0.332 (0.084)	0.077 (0.089)	0.203 (0.153)	0.117 (0.066)	0.127 (0.033)

NOTES: Standard errors in parentheses.
Regressions control for average productivity growth in the same industry in other countries.
[a]Equation for five industries estimated by SUR, restricting β to be the same in each industry. Regression does not use data from basic metals or mechanical equipment.

The disturbances in different countries are nearly orthogonal by construction, because international averages of productivity growth in the same industry in the other countries are included in the regressions. Therefore, a standard error can be calculated for the grand precision-weighted average of the SUR β coefficients estimated for each nation. The grand precision-weighted average is 0.127, with a standard error of 0.033. Labor productivity thus remains procyclical after controlling for the average rates of industry productivity growth in different countries.

Table 3 reports the estimated γ coefficients, which capture the responsiveness of productivity growth in an industry to value added growth in the same industry in other countries (holding constant value added in the rest of the manufacturing sector of that particular country). The industries that appear most sensitive to "supply" conditions, as captured by the growth of value added in the same industry in other countries, are the chemicals and the nonmetallic minerals industries. The industries that appear least sensitive are the food products and textiles industries.

Table 4 shows that the fraction of the variance in productivity accounted for by orthogonal nation-specific effects is much smaller than the fraction of the variance in value added explained by orthogonal nation-specific effects in Table 1. Orthogonal nation-specific effects account for 4.98% of the variance, including the variance "explained"

TABLE 3

Value Added per Hour on the Growth of the Industry in Other Countries

	USA	DEU	FRA	BEL	NOR	UK	Precision-Weighted Average
Food	0.278	0.282	0.694	0.647	–0.631	0.097	0 .244
	(0.618)	(0.282)	(0.688)	(0.471)	(0.832)	(0.268)	(0.164)
Textiles	0.292	0.325	0.728	0.423	0.347	0.619	0.445
	(0.424)	(0.398)	(0.457)	(0.635)	(0.536)	(0.513)	(0.195)
Paper	0.744	0.667	0.039	0.331	0.946	1.086	0.637
	(0.511)	(0.275)	(0.466)	(0.406)	(0.441)	(0.449)	(0.163)
Chemicals	0.391	0.999	0.761	0.023	0.794	0.938	0.720
	(0.303)	(0.345)	(0.258)	(0.648)	(0.426)	(0.343)	(0.140)
Stone, Clay,	0.545	0.586	1.000	0.500	1.236	0.667	0.670
and Glass	(0.220)	(0.186)	(0.316)	(0.433)	(0.465)	(0.349)	(0.114)
Basic Metals	0.361	0.251	0.412	1.101	1.157	1.070	0.589
	(0.763)	(0.358)	(0.299)	(0.384)	(0.813)	(0.833)	(0.181)
Mechanical	0.163	0.466	0.426	1.323	0.154	1.208	0.470
Equipment	(0.704)	(0.231)	(0.447)	(0.828)	(0.405)	(0.678)	(0.168)
Precision-Weighted	0.458	0.526	0.632	0.623	0.645	0.621	0.569
Average	(0.146)	(0.102)	(0.138)	(0.186)	(0.190)	(0.152)	(0.058)

Note: Standard errors in parentheses.

for industries which have negative estimated coefficients on value added in other industries. In each country, it is far lower than the partial R^2 for the regressions in Waldmann (1991), with value added as the dependent variable.

Table 4 also shows that the fraction of total variance in value added per man-hour explained by orthogonal industry effects is on the order of one-tenth, somewhat higher than the fraction of the variance of value added accounted for by orthogonal industry effects in Waldmann (1991). Since employment in different industries is highly correlated, this finding that the nation effects on labor productivity are smaller than nation effects on value added is not unexpected.

Omitted Variable Bias

The presence of significant nation effects on labor productivity would appear to be evidence in favor of labor hoarding-based, job hoarding-based, or increasing returns to scale-based interpretations of procyclical labor productivity. Increased aggregate demand causes increased labor productivity, even controlling for cost and supply shocks.

But before the orthogonal nation-specific effects can be interpreted as effects of aggregate demand, at least some attempts to control further for supply shocks would be desirable. We examined three possible sets of omitted supply-side variables.

The first avenue of approach was that perhaps the rate of productivity growth in the nations of the sample shifts over time. A linear trend was added to the regressions which—since the dependent variable is a growth rate—corresponds to allowing for quadratic trend in the level of productivity. Such a trend has almost no effect on the estimated β coefficients on national value added in the rest of manufacturing. For example, the seemingly-unrelated-regressions estimated β coefficient for the United States is 0.032 instead of 0.037, and the summary precision-weighted average of the coefficients on national manufacturing remains 0.127 (results not shown).

The second avenue was to try to control explicitly for the oil shocks of the 1970s. Controlling for average productivity growth in the same industry in different countries is to some degree a control for the effects of the oil shocks of the 1970s. But oil shocks may well have had different effects on different nations—on oil exporters such as England and Norway, for example.

Since the complex pattern of effects of the oil shocks on an industry may not have been captured by a single γ coefficient, we reestimated the productivity regressions including as additional explanatory variables the change and the lagged change in the price of oil. Once again, the additional regressors had little effect on the estimated β coefficients. The summary precision-weighted average of the seemingly-unrelated-regressions coefficients on national value added in the rest of manufacturing increases from 0.127 to 0.128 (results not shown).

The third avenue considered was to include estimates of the capital stock in order to examine the procyclicality not of labor productivity but of total factor productivity—the Solow residual. Up to this point the Solow residual has been neglected for three reasons. First, the data set does not contain adequate data on shares of labor and capital in pre-tax value added; second, OECD studies based on the data warn that reported factor shares are unreliable (see Meyer-zu-Schlochtern, 1988).[20] Third, Solow residuals exhibit spurious cyclicality if firms possess market power (Hall, 1986 and 1988).[21]

To investigate whether the omission of capital stock variables was biasing our results, we assumed that the elasticity of value added with respect to labor and capital was constant and imposed constant returns to scale to arrive at

20. In fact, such studies throw away the reported factor share data and instead arbitrarily assume that the share of labor is 75%.

21. A corrected Solow residual could be constructed under the assumption that the ratio of price to marginal cost is constant, but there is little reason to believe this assumption (see Domowitz, Hubbard, and Petersen, 1988).

TABLE 4

PARTIAL R^2'S OF ORTHOGONAL COUNTRY AND INDUSTRY EFFECTS, AS DETERMINANTS OF VALUE ADDED PER HOUR GROWTH

		USA	DEU	FRA	BEL	NOR	UK	AVERAGE	RATIO
FOOD	Country	0.031	0.011	0.027	0.003	0.058	0.005	0.035	0.981
	Industry	0.010	0.050	0.098	0.090	0.028	0.007	0.036	
TEXTILES	Country	0.001	0.016	0.104	0.055	0.001	0.108	0.052	1.191
	Industry	0.024	0.034	0.104	0.022	0.022	0.063	0.043	
PAPER	Country	0.003	0.087	0.021	0.013	0.009	0.139	0.049	0.536
	Industry	0.085	0.174	0.001	0.034	0.184	0.098	0.092	
CHEMICALS	Country	0.001	0.003	0.004	0.165	0.067	0.034	0.074	0.580
	Industry	0.077	0.271	0.295	0.000	0.137	0.190	0.127	
STONE, CLAY	Country	0.113	0.121	0.082	0.080	0.025	0.085	0.075	0.438
AND GLASS	Industry	0.182	0.212	0.208	0.049	0.267	0.112	0.170	
BASIC	Country	0.086	0.023	0.051	0.003	0.005	0.013	0.026	0.194
METALS	Industry	0.014	0.043	0.146	0.373	0.155	0.108	0.135	
MECHANICAL	Country	0.175	0.084	0.018	0.027	0.030	0.031	0.070	0.652
EQUIPMENT	Industry	0.004	0.234	0.076	0.179	0.012	0.166	0.107	
AVERAGE	Country	0.073	0.039	0.048	0.071	0.040	0.049	0.050	0.435
	Industry	0.039	0.160	0.130	0.114	0.141	0.116	0.115	
RATIO		1.872	0.244	0.369	0.623	0.284	0.422		

a Cobb-Douglas production function in which value added per man-hour is a function also of the capital/labor ratio.[22] This led to equation 2:

$$(2)\quad \Delta\{\log(Y/N_{int})\} = c_{ni} + \beta_{ni}[\Delta\{\log(Y_{nt} - Y_{int})\}] + \gamma_{ni}[\Delta\{\log(Y/N_{i(-n)t})\}] + \delta_{in}[\Delta\{\log(K_{int}/E_{int})\}] + \varepsilon_{int}\,,$$

where K stands for the real capital stock and E for the level of employment.[23]

The results of regressions including the capital/labor ratio were disappointing. Summary β, γ, and δ coefficients for nations and industries are reported in Table 5. The estimated coefficients on the capital/labor ratio are often implausible: for 13 of 35 underlying regressions, the coefficient on the capital/labor ratio is negative. The precision-weighted average coefficient is negative for France; for England the average coefficient is enormous and implausible.

We ascribe these disappointing results to the fact that the variance of changes in capital stocks is low, and so changes in the capital/labor ratio are nearly the negative

22. Data on capital stocks are not available for Finland. However, the absence of reliable average hours data for Finland makes its inclusion impossible in any event. Data on capital stocks in Norway in the 1960s also do not exist in the data set, reducing the number of countries in the sample to five.

23. The ratio of capital per worker, rather than the ratio of capital per man-hour, is used on the assumption that the work week of capital is the same as work week of workers. Under the alternative assumption that the work week of capital is fixed, and thus that the appropriate capital/labor ratio is capital divided by hours worked, the estimated elasticity of value added with respect to capital is negative for most industries and most countries.

TABLE 5

VALUE ADDED PER HOUR REGRESSED ON GROWTH OF THE REST OF MANUFACTURING, ON INDUSTRY GROWTH IN OTHER COUNTRIES, AND ON THE CAPITAL/LABOR RATIO

	USA	DEU	FRA	BEL	UK	AVERAGE
β	0.247 (0.080)	0.094 (0.053)	0.193 (0.094)	0.062 (0.072)	0.174 (0.045)	
γ	0.518 (0.144)	0.428 (0.105)	0.581 (0.142)	0.781 (0.181)	0.284 (0.119)	
δ	0.245 (0.143)	0.272 (0.087)	–0.025 (0.182)	0.325 (0.144)	0.795 (0.091)	
γ estimated by SUR	0.089 (0.089)	0.045 (0.074)	0.286 (0.086)	0.000 (0.075)	0.242 (0.066)	0.133 (0.034)

	FOOD	TEXTILES	PAPER	CHEMICALS	STONE CLAY, & GLASS	BASIC METALS	MECHANICAL EQUIPMENT
β	–0.010 (0.038)	0.270 (0.088)	0.261 (0.083)	0.448 (0.107)	0.373 (0.085)	0.390 (0.161)	0.216 (0.103)
γ	0.097 (0.135)	0.421 (0.185)	0.527 (0.189)	0.665 (0.137)	0.645 (0.126)	0.608 (0.176)	0.455 (0.186)
δ	0.777 (0.110)	0.371 (0.130)	0.098 (0.171)	0.484 (0.150)	0.206 (0.109)	0.478 (0.168)	0.330 (0.162)

NOTES: Standard errors in parentheses.
Regressions control for average productivity growth in the same industry in other countries.

of changes in employment.[24] This interpretation is supported by the finding of similar results when the capital/labor ratio is replaced by 1/employment (results not shown): the coefficients on 1/employment are in fact greater than the coefficients on the capital labor ratio. We conclude that the OECD estimates of capital are not useful in attempting to analyze labor productivity over the 1960s, 1970s, and 1980s.[25]

Each of the different sets of regressions that we have run is vulnerable to criticism based on the omission of a potential supply side effect. The similarity of results that devote different degrees of effort to controlling for such effects in the specifications we have tried, together with the finding that all regressions show significant nation effects, suggest that these criticisms may not be crippling.

The inclusion of valid supply-side variables should reduce estimated nation effects even if the supply-side variables are relatively poor proxies for the true underlying determinants of procyclical productivity. This does not take place. Thus we are more confident that the estimated association of value added and productivity shows that increased demand leads to increased labor productivity.

Instrumental Variables Estimates of Procyclical Labor Productivity

An appropriate measure of the magnitude of demand-driven productivity changes in a given industry is the elasticity of hours worked with respect to value added. In Table 6, we estimate this elasticity by instrumenting the demand for each industry's value added by the growth of manufacturing value added in the same country outside that industry. The estimates reflect not only long run elasticities of hours with respect to value added, but also the effects of labor or job hoarding, which should reduce the elasticity of hours worked with respect to value added.

In Table 6 no additional regressors are included to control for industry-specific shocks. The precision-weighted average of the estimated elasticities are all near one-half. Even for Germany, the nation with the highest value, the precision-weighted average coefficient is significantly less than one—implying that labor productivity is procyclical. The U.S. has the second highest average elasticity, noticeably greater than the elasticity estimated for any European country save Germany. The fact that total hours adjust the same amount in the United States and Germany confirms the results of Abraham and Houseman (1989).[26]

Table 7 adds the average change in value added and in hours worked in the same industry in other countries as additional regressors to control for supply shocks. With these variables included, estimated elasticities differ more across countries: removing averages highlights national differences. All precision-weighted national averages of elasticities are significantly less than one. Germany continues to exhibit the highest average estimated elasticity, with the lowest for Belgium. Such elasticities also suggest that labor productivity is procyclical after controlling for cross-national supply shocks.

Table 8 reports national average coefficients from instrumental variables regressions of the elasticity not of man-hours but of employment with respect to value added, controlling for average growth of employment and value added in the same industry in other countries.[27] The U.S. has a markedly greater precision-weighted average estimated elasticity than most European countries. For five of the six European countries the estimate is on the order of 0.3 or smaller: Finland is the exception.[28]

This is of interest: labor and job hoarding work to prevent layoffs, not necessarily to keep hours unchanged. The differing elasticities suggest that there may be some returns to pursuing institution-based explanations of procyclical productivity. The next section correlates the degree

24. In this case, the large coefficient estimated for England perhaps reflects the fact that the capital/labor ratio is picking up the large negative disturbance to employment following the accession of Margaret Thatcher (Layard and Nickell, 1989). There was a large drop in English manufacturing employment in 1982. This huge drop in employment corresponds to a huge increase in the measured capital/labor ratio and to a huge increase in production per man-hour. Such an increase can be interpreted as showing that employment follows value added with a lag due to job hoarding, or that by 1983 Thatcher had finally terrorized the unions enough that she and private firms could fire workers in droves (see Bertola and Bentolila, 1987). Neither explanation has anything to do with the effect the capital/labor ratio is supposed to capture—that workers can produce more value added working with machines than without them.

25. It is reassuring to note that the inclusion of the capital/labor ratio does not change the measured cyclicality of labor productivity enormously. Labor productivity remains procyclical controlling for growth in the same industry in other countries when the capital/labor or the 1/employment ratios are included in the regressions.

26. The much lower elasticities estimated for European countries other than Germany suggest that, as Abraham and Houseman note, their results may have been caused in part by the fact that Germany reports hours actually worked, while other countries instead report hours paid.

27. Results are similar without the controls.

28. One possible problem with this result is that changes in average hours worked reflect not only the adjustment of hours worked by normally full-time workers, but also changes in the proportion of full- and part-time workers. The result may simply show that in the U.S. part-time work is more cyclical than in Europe, so average hours worked are less cyclical.

TABLE 6

Elasticity of Hours Worked with Respect to Value Added, Not Controlling for Industry Effects

	USA	DEU	FRA	BEL	NOR	UK	Precision-Weighted Average
Food	1.747	1.089	0.265	1.016	0.433	0.907	0.722
	(1.562)	(0.313)	(0.346)	(0.621)	(0.289)	(0.271)	(0.133)
Textiles	0.955	1.076	0.697	0.725	1.040	0.719	0.810
	(0.159)	(0.239)	(0.158)	(0.201)	(0.514)	(0.153)	(0.075)
Paper	0.661	0.646	0.731	1.048	0.514	0.270	0.480
	(0.135)	(0.103)	(0.282)	(0.372)	(0.330)	(0.093)	(0.055)
Chemicals	0.826	0.648	0.768	0.343	0.343	0.501	0.536
	(0.172)	(0.149)	(0.161)	(0.102)	(0.189)	(0.116)	(0.055)
Stone, Clay	0.729	0.615	0.451	0.492	0.744	0.551	0.579
and Glass	(0.082)	(0.081)	(0.085)	(0.120)	(0.372)	(0.111)	(0.039)
Basic	0.561	0.702	1.128	0.553	0.982	0.519	0.574
Metals	(0.068)	(0.166)	(0.326)	(0.231)	(1.092)	(0.183)	(0.056)
Mechanical	0.729	0.704	0.109	0.678	0.394	0.529	0.692
Equipment	(0.106)	(0.073)	(0.401)	(0.311)	(0.340)	(0.167)	(0.055)
Precision-Weighted	0.682	0.683	0.573	0.492	0.480	0.485	
Average	(0.041)	(0.043)	(0.063)	(0.066)	(0.119)	(0.050)	

Notes: Standard errors in parentheses.
Regressions do not include the average growth of value added or of hours worked in the same industry in different countries.

of procyclical productivity with unemployment. It argues that procyclical productivity is driven by institutional interactions of workers and firms, and not by technological interactions of workers and machines.

III. Procyclical Productivity and the Unemployment Rate

The previous section establishes a presumption that a component of procyclical productivity is independent of supply-side shocks and is, instead, a consequence of shifts in demand. There are at least three interpretations of how such demand-driven procyclical productivity comes about. First, there may be increasing returns. Second, firms may hoard labor. Third, workers may hoard jobs.

Each interpretation leads to its own predictions of the likely cross-country pattern of procyclical productivity, and of the shifts over time in the cyclicality of productivity. “Job hoarding” by workers is likely to show itself most clearly in European countries, which have stronger labor movements and job protection legislation than the United States (see Cross, 1985; Bentolila and Bertola, 1990; Clarke, 1988; and Lazear, 1990).

Section II noted that the U.S. shows more adjustment of employment to shifts in demand than does Europe. It is difficult to see how increasing returns could produce such a pattern: European industry would have to have more sharply increasing returns than U.S. industry. It seems more straightforward to conclude that the European labor market has institutions that cause more labor hoarding, or job hoarding, than those of the United States.

The cross-country pattern alone does not tell us whether procyclical productivity arises because of hiring costs—firms hoarding workers because they fear they will not find personnel when the economy recovers—or because of firing costs—workers hoarding jobs because their positions in the labor market are valuable assets in which they have quasi-property rights.

TABLE 7

Elasticity of Hours Worked with Respect to Value Added, Controlling for Industry Effects

	USA	DEU	FRA	BEL	NOR	UK	Precision-Weighted Average
Food	2.155	1.057	–0.016	0.365	0.534	0.607	0.181
	(3.338)	(0.518)	(0.169)	(0.591)	(0.678)	(0.440)	(0.143)
Textiles	1.036	0.892	0.406	0.674	4.334	0.552	0.634
	(0.246)	(0.306)	(0.197)	(0.364)	(13.114)	(0.145)	(0.096)
Paper	0.843	0.688	0.791	0.771	0.943	0.524	0.672
	(0.396)	(0.143)	(0.538)	(1.752)	(1.289)	(0.242)	(0.114)
Chemicals	0.859	2.481	2.145	0.199	0.344	0.428	0.360
	(0.277)	(1.748)	(1.205)	(0.130)	(0.413)	(0.244)	(0.102)
Stone, Clay, and Glass	0.705	0.764	0.966	0.312	1.440	0.420	0.653
	(0.083)	(0.162)	(0.302)	(0.225)	(1.278)	(0.162)	(0.063)
Basic Metals	0.637	–1.221	–6.438	1.098	1.873	0.711	0.672
	(0.260)	(1.417)	(43.197)	(0.615)	(5.183)	(0.367)	(0.199)
Mechanical Equipment	0.646	0.809	1.393	0.656	0.305	0.390	0.650
	(0.129)	(0.290)	(1.025)	(0.502)	(0.527)	(0.481)	(0.109)
Precision-Weighted Average	0.720	0.754	0.338	0.310	0.451	0.502	0.578
	(0.063)	(0.094)	(0.114)	(0.102)	(0.278)	(0.086)	(0.038)

Notes: Standard errors in parentheses.
Regressions include the average growth of value added or of hours worked in the same industry in different countries.

Distinguishing between Labor Hoarding and Job Hoarding

More information on the relative importance of labor hoarding, job hoarding, and increasing returns can be gained by looking at shifts in the cyclicality of labor productivity within each country. Increasing returns suggests no link between procyclical productivity and macroeconomic variables. But if labor hoarding is the cause of procyclical productivity, then productivity will be less cyclical in periods of high unemployment. In a time of high unemployment firms need not fear that workers will find new jobs and be unavailable when business picks up. Firms are therefore more likely to use temporary layoffs to manage their costs when the unemployment rate is chronically high.

By contrast, if workers resist layoffs—and "hoard" their jobs—because they are well organized or because of employment protection legislation, labor productivity will be more procyclical when unemployment is high. At a low unemployment rate quits will be sufficient for firms wishing to reduce work forces to do so by attrition. Unions are unlikely to spend political capital resisting layoffs when members can easily find other good jobs.

In the United States, labor productivity is less cyclical in unionized industries (Medoff, 1979; Freeman and Medoff, 1984). This suggests that labor hoarding is more important than job hoarding: if workers resisted layoffs, they would be more able to do so in highly unionized industries. If job hoarding were an important cause of U.S. procyclical productivity, labor productivity would be more cyclical in highly unionized industries. This cross-sectional pattern leads to the prediction that labor productivity will be less procyclical in the U.S. when the unemployment rate is chronically high.

By contrast, high unemployment should increase the procyclicality of productivity in Europe. In Europe, powerful union movements and legal restrictions on layoffs are likely to make job hoarding important. When unemploy-

TABLE 8

ELASTICITY OF EMPLOYMENT WITH RESPECT TO VALUE ADDED, CONTROLLING FOR INDUSTRY EFFECTS

	USA	DEU	FRA	BEL	NOR	UK	FIN	PRECISION-WEIGHTED AVERAGE
FOOD	0.851	0.029	–0.286	0.448	–0.126	–0.047	0.522	–0.035
	(1.286)	(0.247)	(0.182)	(0.777)	(0.725)	(0.326)	(0.321)	(0.120)
TEXTILES	0.756	0.459	0.413	0.508	–2.247	0.398	0.991	0.535
	(0.181)	(0.219)	(0.324)	(0.342)	(11.236)	(0.126)	(0.302)	(0.083)
PAPER	0.338	0.307	0.336	–4.276	0.171	0.155	0.351	0.289
	(0.308)	(0.135)	(0.280)	(55.708)	(1.044)	(0.201)	(0.264)	(0.092)
CHEMICALS	0.685	1.861	3.062	0.185	–0.037	0.239	1.062	0.278
	(0.279)	(1.685)	(4.559)	(0.163)	(0.308)	(0.199)	(0.639)	(0.106)
STONE, CLAY, AND GLASS	0.598	0.338	1.015	–0.223	0.649	0.244	0.577	0.511
	(0.065)	(0.165)	(0.612)	(0.385)	(0.934)	(0.142)	(0.157)	(0.052)
BASIC METALS	0.391	–0.422	–8.933	0.490	0.775	0.473	–0.285	0.340
	(0.181)	(0.895)	(77.390)	(0.418)	(2.610)	(0.403)	(0.533)	(0.150)
MECHANICAL EQUIPMENT	0.597	0.332	0.664	0.063	4.957	0.287	0.177	0.443
	(0.145)	(0.177)	(0.404)	(0.455)	(39.745)	(0.224)	(0.877)	(0.095)
PRECISION-WEIGHTED AVERAGE	0.591	0.308	0.114	0.209	0.025	0.279	0.558	0.409
	(0.052)	(0.079)	(0.128)	(0.124)	(0.261)	(0.071)	(0.110)	(0.032)

NOTES: Standard errors in parentheses.
Regressions include the average growth of value added or of hours worked in the same industry in different countries.
The sample is 1970–84 for basic metals and mechanical equipment industries. The sample is 1963–84 for other industries.

ment is high workers are less likely to quit. And workers are more likely to resist layoffs when unemployment is high and makes their jobs valuable property.[29]

In either case, a significant effect of unemployment on the cyclicality of labor productivity is evidence that hiring and firing costs are among the causes of procyclical productivity. The absence of an effect would be evidence that the cause may be technological change and increasing returns.

Estimating the Effect of Unemployment on the Procyclicality of Productivity

The interaction of value added growth and the unemployment rate was added to the independent variables of equation 1, giving equation 3:

$$(3)\quad \Delta\{\log(Y/N_{int})\} = c_{ni} + \beta_{ni}[\Delta\{\log(Y_{nt} - Y_{int})\}] + \gamma_{ni}[\Delta\{\log(Y/N_{i(-n)t})\}] + \mu[U_{n[t-\lambda]} - \mathrm{Avg}(U_{nt})][\Delta\{\log(Y_{nt} - Y_{int})\} - \mathrm{Avg}(\Delta\{\log(Y_{nt} - Y_{int})\})] + \varepsilon_{int} ,$$

where $U_{n[t-\lambda]}$ is the unemployment rate in nation n lagged λ years. We estimate equation 3 for λ equal to 1 and 2—with value added growth in the rest of manufacturing interacted with unemployment lagged one and two years. We lag unemployment one year to reduce correlations between

29. Just as anticipated future hiring costs may prevent layoffs during recessions, anticipated firing costs may reduce hiring during expansions. An extensive literature discusses the possibility that increased unemployment may have caused the constraints on layoffs in Europe to become binding (see Blanchard and Summers, 1986 and 1988; Bertola and Bentolila, 1990; Krugman, 1988; Freeman, 1988). It has been noted that aggregate employment and unemployment fluctuations have become more persistent in Europe in the 1980s (Blanchard and Summers). It is important to learn if this reflects greater persistence in demand fluctuations, or instead a reduced response of employment to demand fluctuations.

this period's disturbance and this period's unemployment rate. We lag unemployment two years for two reasons. First, since all data are averages over a year of continuous-time processes, a two-year lag is needed to purge the correlation of current disturbance terms.

Second, use of unemployment lagged two years serves as a specification check: we believe that the degree of procyclical productivity changes relatively slowly, as workers' and firms' perceptions of the ease of finding new jobs or new workers shifts. If results differed depending on the exact lag of unemployment, we would no longer believe our specification.

In equation 3, averages are taken over 1963–84 for each individual nation. Average unemployment and rates of value added growth in the rest of manufacturing were subtracted in the second line of equation 3 to make the estimates of β and γ comparable to those estimated for equation 1. For each country, the system of equation 3 for the five industries (food, textiles, paper, chemicals, and non-metallic minerals) was estimated by seemingly-unrelated-regressions procedures, restricting μ to be the same across industries.[30]

30. Similar results were obtained by estimating equation 3 by OLS for these industries, and for basic metals and mechanical equipment, and calculating the precision-weighted national average of the estimates of μ.

For those four nations with data available on employment in metals and equipment in the 1960s, the system was estimated for all seven industries as well, restricting μ to be the same across industries.

The first set of estimated μ interaction coefficients are presented in Table 9. For each country-industry pair, it presents the values of the interaction coefficients from regressions of the growth of value added per hour worked on value added growth in other industries in that nation, the average of value added growth in the same industry in other nations, and the interaction of the unemployment rate level with national value added growth. For the United States the interaction term is negative and significant. This provides some evidence suggesting that labor hoarding is a dominant cause of U.S. procyclical productivity.

For Germany, the interaction is *positive* and significant, suggesting that job hoarding is a predominant cause of procyclical productivity and is more prevalent during periods of chronically high unemployment. For Britain the interaction is positive, but its significance is borderline and changes from specification to specification.

For France and Belgium, the coefficient is negative and insignificant. For Norway, it is far from significant with a huge standard error, and its sign depends on the specification. The failure of a pattern to emerge for the small open economies of Norway and Belgium is not unexpected. The

TABLE 9

INTERACTION COEFFICIENTS OF VALUE ADDED GROWTH AND UNEMPLOYMENT, WITH INDUSTRY VALUE ADDED PER MAN-HOUR AS THE DEPENDENT VARIABLE

	USA	DEU	FRA	BEL	NOR	UK	EUROPE POOLED	DIFFERENCE BETWEEN USA AND EUROPE
				FIVE INDUSTRIES—SIX COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.104 (0.042)	0.173 (0.049)	–0.030 (0.041)	–0.055 (0.047)	0.092 (0.347)	0.024 (0.028)	0.022 (0.019)	0.127 (0.047)
UNEMPLOYMENT LAGGED 2 YEARS	–0.088 (0.048)	0.206 (0.063)	–0.039 (0.044)	–0.039 (0.053)	–0.025 (0.484)	0.048 (0.027)	0.033 (0.020)	0.121 (0.052)
				SEVEN INDUSTRIES—FOUR COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.071 (0.036)	0.145 (0.038)			–0.270 (0.230)	0.049 (0.027)	0.077 (0.022)	0.148 (0.042)
UNEMPLOYMENT LAGGED 2 YEARS	–0.056 (0.040)	0.190 (0.045)			0.036 (0.325)	0.070 (0.022)	0.092 (0.020)	0.148 (0.044)

interaction term is only identified by the orthogonal nation-specific shock to aggregate demand, and these small open economies possess only small nation-specific movements in total manufacturing value added.[31] The failure of a pattern to emerge for France is disappointing, for France is large and has pursued independent macroeconomic policies over the past third of a century. We expected to see stronger results.

However, the difference between the interaction coefficients estimated for the United States and those estimated for a pooled sample of European countries is large and highly significant. Procyclical productivity is weaker in the United States when unemployment is high, but it is not weaker in Europe. It is difficult to argue that the same "labor hoarding" that appears to generate procyclical productivity in the U.S. generates it in Europe as well.

The dependent variable in equation 3 is production per man-hour. Since hiring and firing costs are likely to depend on the change not in man-hours but in employment, it is interesting to compare the behavior of production per worker with the behavior of production per man-hour.[32] Similarity in coefficients would suggest that the results in Table 9 are not simply due to changes in the labor force or differences in the reporting of hours worked.

Equation 3 was thus reestimated, replacing value added per man-hour by value added per worker. Table 10 reports the results, which are indeed similar to those reported in Table 9. The procyclicality of value added per worker undergoes the same shifts with changing unemployment as does the procyclicality of value added per man-hour. The interaction term is significantly negative only for the U.S., for

31. The magnitude, however, of the Norwegian interaction coefficient is deserving of explanation. We tentatively ascribe the high magnitude to the fact that the Norwegian unemployment rate exhibits a very small rise in the 1970s, and is therefore highly collinear with a post-North Sea oil discovery dummy variable. Under this interpretation, the coefficient is capturing the fact that Norwegian productivity became much more sensitive to the level of production in Norway after the discovery of North Sea oil. If this interpretation is correct, the coefficient carries little information about the magnitude of job hoarding in Norway.

32. Ideally, one would want to examine labor hoarding by examining hours worked by workers who normally work full time—thus obtaining a more direct measure of overtime and slack time hours. As noted above, differences between countries in the cyclicality of production per worker can reflect differences in the cyclicality of part time work. Changes over time within a country also reflect, among other things, the entry of women into the labor force.

TABLE 10

Interaction Coefficients of Value Added Growth and Unemployment, with Industry Value Added per Worker as the Dependent Variable

	USA	DEU	FRA	BEL	NOR	UK	FIN	Europe Pooled	Difference between USA and Europe
					Five Industries—Seven Countries				
Unemployment Lagged 1 Year	–0.105 (0.048)	0.158 (0.054)	0.002 (0.037)	0.020 (0.043)	0.035 (0.312)	0.017 (0.030)	0.045 (0.052)	0.033 (0.018)	0.138 (0.051)
Unemployment Lagged 2 Years	–0.099 (0.054)	0.093 (0.072)	–0.020 (0.040)	0.026 (0.048)	0.327 (0.438)	0.041 (0.031)	0.026 (0.065)	0.026 (0.020)	0.125 (0.057)
					Seven Industries—Five Countries				
Unemployment Lagged 1 Year	–0.075 (0.037)	0.144 (0.045)			–0.206 (0.225)	0.051 (0.026)	0.022 (0.041)	0.068 (0.026)	0.143 (0.045)
Unemployment Lagged 2 Years	–0.069 (0.041)	0.083 (0.059)			0.365 (0.308)	0.067 (0.023)	0.017 (0.051)	0.073 (0.020)	0.142 (0.045)

which it is virtually unchanged. For the European countries the coefficients are somewhat smaller, but still positive.[33]

Errors in and Omissions of Variables

Controlling for average productivity growth in the same industry in different countries has the important advantage of controlling for supply shocks. But from a Keynesian standpoint, it would be disturbing if results were substantially changed if the international average growth rates of value added in individual industries were excluded from the list of independent variables. It is also possible that the average of growth in the same industry in other countries is not an appropriate measure of supply and cost shocks: perhaps nation-specific shocks—like the discovery of North Sea oil—contaminate the results for other countries. To the extent that nation-specific industry value added movements reflect the discovery of a nation-specific shock, like the discovery of North Sea oil for Norway, the average across nations of value added growth in an industry is a poor measure of true supply shocks.

These considerations led us to repeat the interaction regressions without controlling for average growth in the same industry in other countries, as shown in equation 4:

$$(4) \qquad \Delta\{\log(Y/N_{int})\} = c_{ni} + \beta_{ni}[\Delta\{\log(Y_{nt} - Y_{int})\}] + \mu[U_{n[t-\lambda]} - \text{Avg}(U_{nt})] [\Delta\{\log(Y_{nt} - Y_{int})\} - \text{Avg}(\Delta\{\log(Y_{nt} - Y_{int})\})] + \varepsilon_{int}\,.$$

Except for Norway itself, the interaction terms were virtually unchanged, as Table 11 shows.

Omitted variables might corrupt our results. Productivity might be more cyclical in Germany and Europe when unemployment is high simply because both the unemployment rate and the cyclicality of labor productivity have increased for other reasons. It is easy to see how the cyclicality of labor productivity could have a positive trend if, say, the ratio of administrative to production workers increases over time. Given the time pattern of European unemployment, the interaction terms in the regressions come close to comparing the cyclicality of labor productivity in the later half of the sample to the cyclicality in the earlier half of the sample.

While the use of a disaggregated dependent variable—of sector-specific value added—gives greater precision, it does not increase the ability to discriminate between increased unemployment and the effect of time. In the case of the United States, the time pattern of unemployment makes it correlated with lagged oil shocks; lagged oil shocks might have reduced the cyclicality of labor productivity.

In each case it is possible in principle to control for omitted variable bias by including an additional independent variable: the omitted variable interacted with the growth in manufacturing value added. But such regressions are likely to lack power.

We use an alternative procedure. If German unemployment is standing in for an omitted variable, this omitted variable should also have been in operation in other countries. If a secular increase in the amount of overhead labor is making productivity more procyclical, and if the German unemployment rate is correlated with this omitted variable, then a regression of productivity growth in an *American* industry on the growth of value added in the rest of *American* manufacturing and the interaction with *German* unemployment should produce the same, positive, interaction coefficient.

But Table 12 shows that interacting the growth of manufacturing value added in a country with the *German* unemployment rate rather than the national unemployment rate does not cause the interaction terms to mimic the German pattern. The coefficient drops for England, remains negative for Norway, and for Belgium and France switches from negative to positive but remains insignificant.

These results do not suggest that the positive effect of German unemployment on the cyclicality of German labor productivity is due to the correlation of German unemployment and another factor causing increased cyclicality of labor productivity.[34]

33. Finland can be included in regressions using value added per worker because average hours are not needed. The large standard error for Finland presumably reflects the difficulty of identifying national demand in a small open economy. The use of production per worker instead of production per man-hour reduces the spread of the European coefficients, making the contrast with the United States more striking.

34. The analogous question can be asked about the negative coefficients on the interaction term found for the United States: perhaps they reflect the fact that U.S. unemployment is highly correlated with lagged oil shocks. If so, regressions of other countries' productivity growth on the interaction of their growth of the rest of manufacturing and the *United States* unemployment rate should be negative.

However, when such regressions are estimated the interaction coefficient for Germany remains positive and significant (results not shown). The coefficient for England falls and is not significant, but remains positive. For other countries, coefficients remain insignificant and negative. The summary precision-weighted average coefficient on the interaction of *United States* unemployment and national rates of growth in the rest of manufacturing is positive, the opposite of what one would have expected according to the omitted variable-bias story.

TABLE 11

INTERACTION COEFFICIENTS OF VALUE ADDED GROWTH AND UNEMPLOYMENT, WITH INDUSTRY VALUE ADDED PER MAN-HOUR AS THE DEPENDENT VARIABLE, NO INDUSTRY CONTROLS

	USA	DEU	FRA	BEL	NOR	UK	EUROPE POOLED	DIFFERENCE BETWEEN USA AND EUROPE
				FIVE INDUSTRIES—SIX COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.122	0.139	–0.028	–0.009	–0.378	0.039	0.030	0.151
	(0.042)	(0.057)	(0.054)	(0.041)	(0.376)	(0.029)	(0.020)	(0.047)
UNEMPLOYMENT LAGGED 2 YEARS	–0.115	0.168	–0.041	–0.002	–0.632	0.057	0.042	0.157
	(0.047)	(0.072)	(0.058)	(0.047)	(0.530)	(0.026)	(0.020)	(0.051)
				SEVEN INDUSTRIES—FOUR COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.090	0.087			–0.246	0.047	0.053	0.143
	(0.035)	(0.049)			(0.238)	(0.027)	(0.024)	(0.042)
UNEMPLOYMENT LAGGED 2 YEARS	–0.083	0.124			–0.169	0.064	0.071	0.154
	(0.038)	(0.058)			(0.322)	(0.024)	(0.022)	(0.044)

TABLE 12

INTERACTION COEFFICIENTS OF VALUE ADDED GROWTH AND *GERMAN* UNEMPLOYMENT, WITH INDUSTRY VALUE ADDED PER MAN-HOUR AS THE DEPENDENT VARIABLE

	USA	DEU	FRA	BEL	NOR	UK	EUROPE POOLED	DIFFERENCE BETWEEN USA AND EUROPE
				FIVE INDUSTRIES—SIX COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.105	0.139	0.002	0.008	–0.058	0.064	0.052	0.157
	(0.029)	(0.057)	(0.062)	(0.073)	(0.092)	(0.038)	(0.025)	(0.039)
UNEMPLOYMENT LAGGED 2 YEARS	–0.122	0.168	0.047	0.033	–0.059	0.052	0.061	0.183
	(0.041)	(0.072)	(0.079)	(0.086)	(0.114)	(0.038)	(0.028)	(0.049)
				SEVEN INDUSTRIES—FOUR COUNTRIES				
UNEMPLOYMENT LAGGED 1 YEAR	–0.086	0.087			–0.042	0.067	0.049	0.134
	(0.024)	(0.049)			(0.055)	(0.035)	(0.025)	(0.035)
UNEMPLOYMENT LAGGED 2 YEARS	–0.113	0.124			0.003	0.057	0.063	0.176
	(0.032)	(0.058)			(0.066)	(0.035)	(0.028)	(0.042)

NOTE: Regressions do not control for industry effects.

One final errors-in-variables problem is somewhat subtle, but easy to evaluate. The decomposition of productivity growth into trend and cycle would be difficult even if many more years of data were available. The regressions reported above do not reveal whether the interaction term reflects a change in the cyclicality of productivity, or simply reflects changes in the trend in productivity growth which happen to be correlated with changes in the average decade-to-decade level of the unemployment rate. A confident interpretation of the coefficient on the interaction term would require many more years of data with high and with low unemployment.

It is possible with available data to control for some obvious factors which could have changed both the trend of value added and of productivity. Inclusion of a time trend had very little effect; inclusion of the capital/labor ratio had little effect also (results not shown).

It is possible to control directly for changes in the trend of productivity and value added which happen to be correlated with average unemployment rates by including unemployment itself in the regressions.[35] Including the unemployment rate does not affect the interaction coefficients. As reported in Table 13, the coefficient remains significantly negative in the U.S. and positive in Germany. The difference between pooled Europe and the U.S. remains large and significant.

Assessment

None of the explorations and alternatives considered in the second half of this section shake the finding that the effect of unemployment on the cyclicality of productivity is different in the U.S. and in Europe. In the U.S., high unemployment is correlated with low cyclicality in productivity. This reinforces the cross-sectional evidence that labor hoarding by firms is an important component of procyclical productivity in the U.S. In Europe by contrast, the correlation between high unemployment and the cyclicality of labor productivity is positive or statistically insignificant. This suggests that the importance of job hoarding by workers is greater in Europe. This is as one would have expected from the literature on labor market institutions.

35. The failure of the inclusion of the unemployment rate to materially affect the interaction coefficients should not come as a surprise. Earlier regressions did not directly control for the level of unemployment, but they did remove sample means from growth rates. The effect on the interaction coefficient of the inclusion of the unemployment level is thus proportional to the covariance of output growth and the squared deviation of unemployment from its sample average, which is close to zero.

TABLE 13

Interaction Coefficients with the Level of Unemployment Added to the List of Independent Variables

	USA	DEU	FRA	BEL	NOR	UK	Europe Pooled	Difference between USA and Europe
				Five Industries—Six Countries				
Unemployment Lagged 1 Year	–0.106 (0.041)	0.170 (0.051)	0.002 (0.057)	–0.061 (0.060)	0.127 (0.358)	0.021 (0.031)	0.035 (0.022)	0.142 (0.046)
Unemployment Lagged 2 Years	–0.082 (0.049)	0.186 (0.067)	–0.033 (0.066)	–0.018 (0.077)	0.044 (0.497)	0.042 (0.033)	0.046 (0.026)	0.128 (0.055)
				Seven Industries—Four Countries				
Unemployment Lagged 1 Year	–0.065 (0.035)	0.138 (0.040)			–0.389 (0.246)	0.029 (0.030)	0.064 (0.024)	0.129 (0.043)
Unemployment Lagged 2 Years	–0.052 (0.040)	0.162 (0.047)			0.164 (0.335)	0.064 (0.024)	0.085 (0.022)	0.136 (0.046)

IV. Conclusion

This paper has reported evidence that procyclical productivity is more than the consequence of supply-side shocks propagating through a standard real business cycle model. Such theories can account for a correlation of sectoral productivity growth with aggregate value added, and can—if cost shocks affect an intermediate input, like oil, necessary for production in many sectors—account for a correlation of sectoral productivity growth with aggregate productivity.

One explanation for procyclical productivity in response to shifts in demand uncorrelated with shifts in industry supply is that a firm receives surplus from keeping a stock of workers—and that a worker receives surplus from keeping an existing job. Thus labor "hoarding" by firms and job "hoarding" by workers underlies procyclical productivity. We have not built a model of the labor market. Nevertheless, the correlations make us optimistic about the utility of such models.

The differences across countries in the elasticities of labor input with respect to value added lend some support to the view that procyclical productivity reflects the strength of attachment of workers to jobs. In the United States, the response of employment to changes in value added appears much greater than in European countries. This difference might be caused by stronger union movements and employment protection legislation in Europe making "job hoarding" a more important factor in Europe. Real business cycle theories are silent on the causes of such cross-national differences.

Moreover, the *level* of the unemployment rate appears to have an effect on the degree to which productivity is procyclical. In the United States, higher unemployment levels correspond to significantly lower procyclicality. This might be explained in a model in which firms do not have to worry about permanently losing the ability to reemploy laid-off workers when unemployment is high. In Europe, however, increased unemployment does not seem to correspond to less procyclical labor productivity. British and German labor productivity appears more, not less, procyclical under high unemployment.

This difference between the effect of unemployment on the cyclicality of productivity might be accounted for by the greater ability of European workers to resist layoffs, and their determination to do so in times of high unemployment, in a model in which labor market institutions had effects on the organization and level of real production. By contrast, it is difficult to think how to begin to construct an explanation of this cross-Atlantic pattern based on supply shocks or on increasing returns to scale. The pattern suggests that it is worth investigating whether procyclical productivity arises from institutionally influenced hiring and firing costs, and reflects the relationship between workers and firms—and not the relationship between workers and machines.

References

Abraham, K., and S. Houseman. 1989. "Employment and Hours Adjustment: a U.S./German Comparison." Mimeo. University of Maryland.

Aizcorbe, Ana. 1992. "Procyclical Labor Productivity, Increasing Returns to Labor, and Labor Hoarding in U.S. Automobile Assembly Plant Employment." *Economic Journal* 102, pp. 860–873.

Basu, Susanto, and John Fernald. Forthcoming. "Returns to Scale in U.S. Production: Estimates and Implications." *Journal of Political Economy*.

__________, and __________. 1995. "Aggregate Productivity and the Productivity of Aggregates." Cambridge: NBER Working Paper #5382.

__________, and Miles Kimball. 1994. "Cyclical Productivity with Unobserved Input Variation." Mimeo. Palo Alto: Hoover Institution.

Benhabib, Jess, and Roger Farmer. 1996. "Indeterminacy and Sector-Specific Externalities." *Journal of Monetary Economics*.

Bentolila, S., and G. Bertola. 1990. "Firing Costs and Labor Demand in Europe: How Bad Is Eurosclerosis?" *Review of Economic Studies* 57 (July) pp. 381–402.

Bils, Mark, and Jang-Ok Cho. 1992. "Cyclical Factor Utilization." *Journal of Monetary Economics* 33, pp. 319–354.

Blanchard, O. J., and L. H. Summers. 1988. "Why Is Unemployment so High in Europe?: Beyond the Natural Rate Hypothesis." *American Economic Association Papers and Proceedings* 78, pp. 182–187.

__________, and __________. 1986. "Hysteresis and the European Unemployment Problem." *NBER Macroeconomics Annual 1986*. Cambridge: MIT Press.

Caballero, R. J., and R. K. Lyons. 1992. "External Effects in U.S. Procyclical Productivity." *Journal of Monetary Economics* 29, pp. 209–226.

__________, and __________. 1990. "Internal versus External Economies in European Industry." *European Economic Review* 34 (June) pp. 885–906.

Clarke, O. 1988. "Employment Adjustment: An International Perspective." *Labour* 2, pp. 3–29.

Cross, M. 1985. *Managing Work Force Reduction: an International Survey*. New York: Praeger.

Domowitz, I., R. G. Hubbard, and B. Petersen. 1988. "Market Structure and Cyclical Fluctuations in U.S. Manufacturing." *Review of Economics and Statistics* 70, pp. 55–66.

Fair, R.1969. *The Short-Run Demand for Workers and Hours*. Amsterdam: North-Holland.

Freeman, R. B. 1988. "Labour Markets." *Economic Policy* 7, pp. 64–80.

__________, and J. L. Medoff. 1984. *What Do Unions Do?* New York: Basic Books.

Hall, R. 1988. "The Relation between Price and Marginal Cost in U.S. Industry." *Journal of Political Economy* 96, pp. 921–947.

__________. 1986. "Market Structure and Macro Fluctuations." *Brookings Papers on Economic Activity* 1, pp. 285–322.

Hamermesh, D. 1976. "Econometric Studies of Labor Demand and Their Application to Policy Analysis." *Journal of Human Resources* 11, pp. 507–525.

Holt, C. F., F. Modigliani, J. Muth, and H. Simon. 1960. *Planning, Production Inventories, and Work Force.* New York: Prentice-Hall.

Hultgren, T. 1960. *Changes in Labor Cost during Cycles in Production and Business.* Cambridge: NBER Occasional Paper #72.

International Labour Office. 1986. *Annotated Bibliography on Working Time.* Geneva.

Krugman, P. 1988. "Slow Growth in Europe: Conceptual Issues." In *Barriers to European Growth: A Transatlantic View,* eds. C. Schultze and R. Lawrence. Washington, D.C.: Brookings Institution.

Kydland, F. E., and E. C. Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica* 50, pp. 1345–1370.

Layard, R., and S. Nickell. 1989. "The Thatcher Miracle?" *American Economic Review* 79 (May), pp. 215–219.

Lazear, E. 1990. "Job Security Provisions and Employment." *Quarterly Journal of Economics* 105 (August) pp. 699–726.

Long, J. B., and C. I. Plosser. 1983. "Real Business Cycles." *Journal of Political Economy* 91, pp. 39–59.

Medoff, J. L. 1979. "Layoffs and Alternatives under Trade Unions in the United States." *American Economic Review* 69, pp. 380–395.

Meyer-zu-Schlochtern, F.J.M. 1988. *An International Sectoral Data Base for Thirteen OECD Countries.* OECD Department of Economics and Statistics Working Paper #57.

Murphy, Kevin, Andrei Shleifer, and Robert Vishny. 1989. "Building Blocks of Market Clearing Business Cycle Models." *NBER Macroeconomics Annual.*

Nickell, S. 1986. "Dynamic Models of Labour Demand." In *The Handbook of Labor Economics,* eds. O. Ashenfelter and R. Layard, vol. 1. Amsterdam: North Holland.

Oi, W. 1962. "Labor as a Quasi Fixed Factor." *Journal of Political Economy* 70, pp. 538–555.

Okun, A. 1962. "Potential GNP: Its Measurement and Significance." Reprinted in Arthur Okun. 1970. *The Political Economy of Prosperity.* Washington: Brookings Institution.

Rotemberg, J., and L. Summers. 1988. "Labor Hoarding, Inflexible Prices, and Procyclical Productivity." NBER Research Working Paper #2591.

Shapiro, Matthew. 1993. "Cyclical Productivity and the Workweek of Capital." *American Economic Review* 83, pp. 229–233.

__________. 1987. "Are Cyclical Fluctuations Due More to Supply Shocks or Demand Shocks?" *American Economic Review Papers and Proceedings* 77 (May) pp. 118–124.

Soligo, R. 1966. "The Short Run Relationship between Employment and Output." *Yale Economic Essays* 6, pp. 161–215.

Solon, Gary, Robert Barsky, and Jonathan Parker. 1994. "Measuring the Cyclicality of Real Wages: How Important Is Composition Bias?" *Quarterly Journal of Economics* 109, pp. 1–26.

Stockman, A. C. 1988. "Sectoral and National Aggregate Disturbances to Industrial Output in Seven European Countries." *Journal of Monetary Economics* 21, pp. 387–409.

Summers, L.H. 1986. "Some Skeptical Observations on Real Business Cycle Theory." *Federal Reserve Bank of Minneapolis Quarterly* 10, pp. 23–27.

__________, and S. Wadhwani. 1988. *Some International Evidence on Labour Cost Flexibility and Output Variability.* London School of Economics: Centre for Labour Economics Working Paper #981.

Waldmann, R.J. 1991. "Assessing the Relative Sizes of Industry and Nation-Specific Shocks to Output." European University Institute Working Paper ECO #91/41, 1991.