

**Downward Nominal Wage Rigidity:
Evidence from the Employment Cost Index**

by David E. Lebow, Raven E. Saks, and Beth Anne Wilson
Board of Governors of the Federal Reserve System

July, 1999

Abstract:

We examine the extent of downward nominal wage rigidity using the microdata underlying the BLS's employment cost index--an extensive, establishment-based dataset with detailed information on wage and benefit costs. We find stronger evidence of downward nominal wage rigidity than did previous studies using panel data on individuals. Firms appear able to circumvent part, but not all, of this rigidity by varying benefits: Total compensation displays modestly less rigidity than do wages alone. Given our estimated amount of rigidity, a simple model predicts that the disinflation over the 1980s should have raised equilibrium unemployment notably. This prediction stands in contrast to the actual behavior of unemployment over this period: The addition of a term capturing the cost of rigidity (that rises as inflation falls) has no additional explanatory power in a standard Phillips Curve equation.

We would like to thank the Bureau of Labor Statistics for allowing the use of these data and John Ruser and Mike Lettau of the BLS for providing extensive assistance with this project. We also thank our colleagues at the Federal Reserve and participants at a BLS seminar for their comments and suggestions, and Shulamit Kahn for providing guidance in programming the econometric tests. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the staff, by the Board of Governors, or by the Federal Reserve Banks.

1. Introduction

Workers may resist pay cuts for many reasons, most obviously because cuts lead to a lower standard of living, but also because they may be perceived as unfair or demeaning. When combined with money illusion, this resistance to lower real wages translates into downward nominal wage rigidity. Such rigidity would be a concern to monetary policy if it generated a long-run negative correlation between unemployment and inflation. Specifically, if resistance to nominal pay cuts impedes the smooth adjustment of real wages to adverse shocks, then higher inflation could improve labor-market efficiency by increasing the range of real wage cuts acceptable to workers, leading to lower unemployment in equilibrium.¹

In this paper, we use a newly available data set--the microdata underlying the Bureau of Labor Statistics' employment cost index (ECI)--to shed light on two important unanswered questions that emerge from the literature on downward nominal wage rigidity. First, despite a number of recent studies, there is no consensus on the extent of such rigidity. Anecdotal evidence and interviews with selected individuals and firms strongly suggest that workers view nominal wage cuts as unfair and that firms are reluctant to impose them.² However, more quantitative studies using nationally representative data have been limited to panel data sets of households, especially the Panel Study of Income Dynamics, and have found only a modest amount of downward nominal wage rigidity.³ One commonly cited reason for the comparatively weak support for downward nominal wage rigidity could be that these data are hampered by measurement error, which could mask the extent of rigidity (Akerlof, Dickens, and Perry, 1996, Shea, 1997). As an alternative, Wilson (1999) uses payroll records over a number of years from two large firms--which, presumably, are nearly free from measurement error--and finds considerable downward nominal wage rigidity. Altonji and Devereux (1998) also find indications of substantial rigidity using data from one year at a large firm. Unfortunately, it is unclear whether these firms are representative of the overall labor market.

1. See Tobin (1972), or Akerlof, Dickens, and Perry (1996).

2. See Bewley, 1998; Kahneman, Knetsch, and Thaler, 1986; or Shafir, Diamond, and Tversky, 1997, for example.

3. See McLaughlin (1994, 1998), Lebow, Stockton, and Wascher (1995), Craig (1995), Kahn (1997), Card and Hyslop (1997).

The ECI combines the advantages of these two approaches. Because the data are derived from employers' records, they are much less prone to measurement error than are the household data. And, the sample size of the ECI is substantial, covering about 5,000 private establishments per quarter. Thus, these data have the potential to shed considerable light on the extent of downward nominal wage rigidity. This is so despite the potential drawback, discussed below, that the unit of observation in the ECI is the average hourly compensation cost in a specific *job*, rather than for an *individual*.⁴

The second question that emerges from the literature is why, despite the microeconomic evidence of at least some downward nominal wage rigidity, there is almost no evidence that this rigidity has had the predicted adverse macroeconomic consequences, at least in the post-war United States.⁵ The central prediction of models based on downward nominal wage rigidity is a negative correlation between inflation and equilibrium unemployment rates. Yet Gordon's (1998) estimates of the NAIRU (or non-accelerating inflation rate of unemployment), which he allows to vary over time, are *positively* correlated with inflation: They rise from the early 1960s through the early 1980s, and decline since then. We present a similar finding in section 5 of this paper. Similarly, when Akerlof, Dickens, and Perry (1996) embed downward nominal wage rigidity into a Phillips curve, none of the relevant parameter estimates is statistically significant, although they find that their equation does help explain the behavior of wages during the 1930s.

One factor that may help resolve this "micro-macro" puzzle is that employers may be able to use benefits--which are more difficult for workers to value and compare--to achieve the necessary flexibility that is prevented by the nominal rigidity of wages. Because our ECI data include information on a wide range of benefit costs as well as wage costs, we can use

4. Similarly, Groshen and Schweitzer (1996) analyze firm-based data that measure average wages within a job.

5. For Japan, Kimura and Ueda (1997) find that industry-level wages do not exhibit much downward rigidity. For Canada, Fortin (1996) argues that downward rigidity is substantial and has led to persistently high unemployment rates in the 1990s; but see Freedman and Macklem (1998) for a dissenting view.

our data to examine this hypothesis--that is, we can investigate whether downward nominal rigidity of total compensation (wages plus benefits) differs from rigidity of wages alone.

To anticipate our results, we find about half as many wage cuts in our data as would have been predicted in the absence of downward nominal wage rigidity--somewhat stronger evidence of rigidity than is found in the PSID. This is so despite the fact that the extent of rigidity for jobs (as in the ECI) ought to be less than for individuals (as in the PSID). In addition, we find that benefits add some additional flexibility to compensation in our data: Total compensation also exhibits some downward nominal rigidity, but to a smaller degree than do wages alone. In all, benefits appear to explain roughly one-third of the micro-macro puzzle.

The paper proceeds as follows. We describe our methodology for measuring the extent of downward nominal wage rigidity in section 2, and discuss the ECI data in section 3. We present the main results in section 4, and in section 5 we calibrate these results in terms of their predicted consequences for unemployment, and we discuss the micro-macro puzzle in more detail. Section 6 concludes.

2. Measuring the Extent of Downward Nominal Wage Rigidity

To quantify the extent of downward nominal wage rigidity we examine the distribution of wage (or compensation) changes. In the absence of downward nominal wage rigidity, we would expect the distribution of wage changes to be essentially continuous through the point of zero wage change--although there may be reasons, such as the existence of long-term contracts, for a concentration of observations at exactly zero even without such rigidity. In contrast, the presence of downward nominal wage rigidity would imply a distinct shortage of nominal wage cuts, with a corresponding pile-up of observations at zero wage change--zero being the minimum wage change acceptable to workers.

Thus, this shortage of negative wage-change observations implies a particular type of asymmetry, or right-skewness, in the distribution. However, simply showing that the wage-change distribution is skewed to the right is not enough, for the underlying distribution could be skewed even in the absence of downward nominal wage rigidity. More persuasive would be evidence that this asymmetry becomes increasingly pronounced as inflation declines.

At least five such tests have been proposed in the literature, and some of the advantages and disadvantages of these tests are summarized in table 1 below. The first, and most straightforward, test is whether the skewness statistic becomes more positive as inflation (or the median of the wage-change distribution) declines. One serious problem with this test is that the skewness statistic is extremely sensitive to observations in the tails of the distribution, and empirical wage-change distributions are quite disperse. In response to this problem, McLaughlin (1998) suggests replacing the skewness statistic with the difference between the mean and median of the wage-change distribution. This asymmetry measure is less sensitive to outliers because extreme observations affect it only by influencing the mean. However, it is not as robust to outliers as would be a pure order statistic. The third approach, Kernel estimation (Card and Hyslop, 1997), uses the right side of the wage-change distribution to determine the shape of the left side in the absence of downward rigidity. Crucially, this approach is valid only if the underlying distribution is symmetric. Finally, note that all three of these tests have the disadvantage of picking up any type of asymmetry

Table 1
Properties of Different Tests of Downward Nominal Wage Rigidity

	Specific to DNWR	Robust to outliers	Robust to underlying skewness	Robust to $\text{corr}(\sigma, \dot{p})$	Robust to noise around zero
$\text{Corr}(\text{skewness}, \dot{p})^1$	No	No	Yes	Yes	Yes
$\text{Corr}(\text{mean-median}, \dot{p})^1$	No	Somewhat	Yes	Yes	Yes
Kernel estimation ²	No	Yes	No	Yes	Yes
$\text{Corr}(\text{LSW}, \dot{p})^3$	Yes	Yes	No	Yes	No
Kahn test ⁴	Yes	Yes	Yes	No	Somewhat
Modified Kahn test	Yes	Yes	Yes	No	Largely

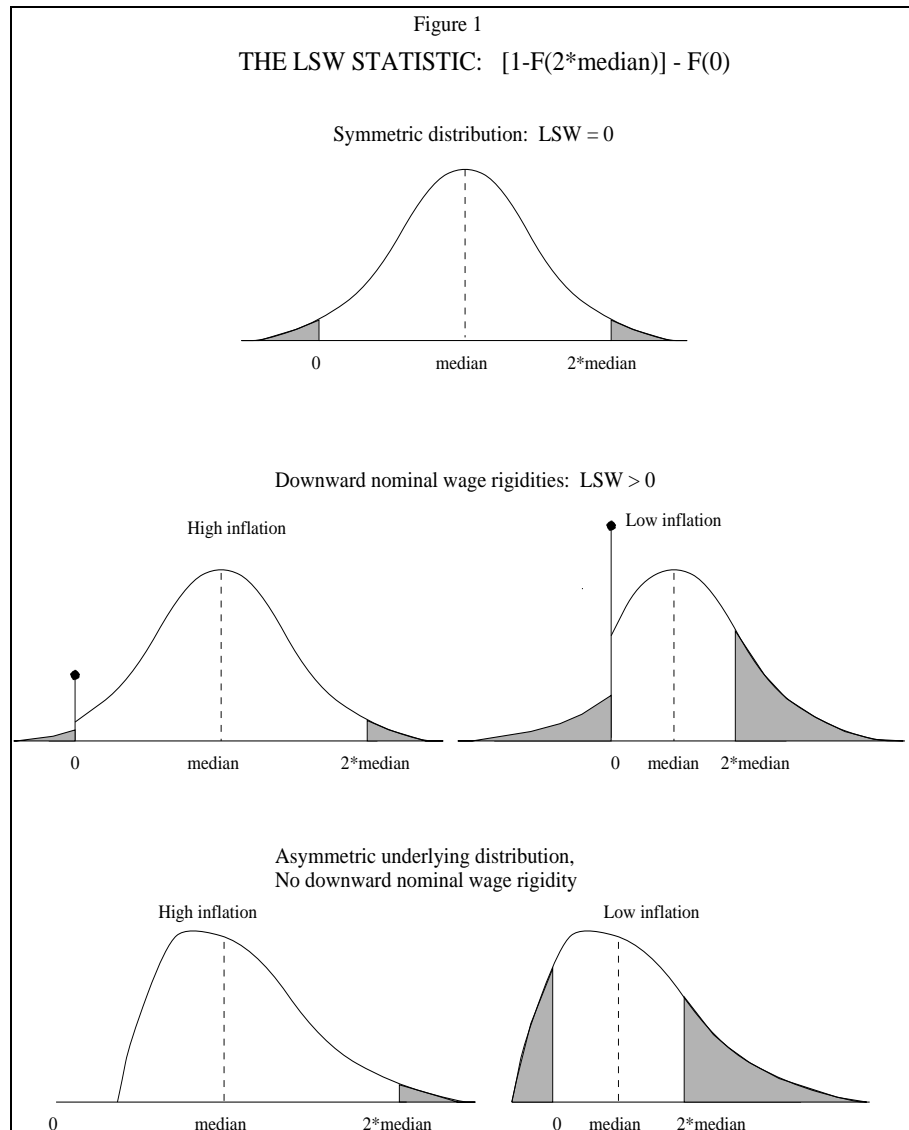
1. See McLaughlin (1998)
2. See Card and Hyslop (1997)
3. See Lebow, Stockton, and Wascher (1995)
4. See Kahn (1997)

rather than that due specifically to downward nominal wage rigidity, namely, a shortage of observations below zero.

In contrast, the fourth alternative--one of two that we use in this paper--makes use of an asymmetry measure that is specific to downward nominal wage rigidity. We call this the LSW statistic (see Lebow, Stockton, and Wascher, 1995). This statistic is defined as the cumulative frequency of the wage-change distribution above twice the median minus the cumulative frequency of the distribution below zero:

$$\text{LSW} \equiv [1 - F(2 * \text{median})] - F(0) .$$

Figure 1 helps explain this test. Because twice the median and zero are equidistant from the median, a symmetric distribution will have an equal mass in the right and left tails, and the LSW statistic will be zero. But with downward nominal wage rigidity, there will be a shortage of nominal wage cuts, and the LSW



statistic will be positive.⁶ This measure of asymmetry will become larger as inflation (and the median of the distribution) declines. In contrast to the previous tests, the LSW statistic measures precisely the asymmetry generated by downward nominal wage rigidity--a shortage of observations less than zero. And, because it is a pure order statistic, the LSW statistic is unaffected by extreme observations. However, the LSW statistic has the drawback that it is not robust to asymmetry in the underlying wage-change distribution. That is, suppose the wage-change distribution is right-skewed independent of downward nominal wage rigidity. Then as inflation falls and the distribution shifts to the left, the LSW statistic would change even if the shape of the distribution did not (figure 1, lower panel).⁷

The last test, which we will most emphasize, was proposed by Kahn (1997). This test compares heights of histogram bars a given distance from the median of the wage-change distribution in years when those bars fall below nominal zero with their height when they fall above nominal zero. Specifically, we construct a histogram for each year's distribution (with bars 1 percentage point wide), and estimate the following system of equations:


$$\begin{aligned}
 PROP2_t &= p_2 + np_2 DNEG2_t + [z - n \sum_{j=3}^m p_j] DZERO2_t \\
 PROP3_t &= p_3 + np_3 DNEG3_t + [z - n \sum_{j=4}^m p_j] DZERO3_t \\
 &\vdots \\
 &\vdots \\
 &\vdots \\
 PROPm_t &= p_m + np_m DNEGm_t + [z] DZEROm_t ,
 \end{aligned} \tag{1}$$

6. The statistical significance of the LSW statistic can be constructed as follows. For a symmetric distribution, observations of greater than twice the median and less than zero will occur with the same probability and LSW=0. If we observe a fraction a of observations in the right tail, then the binomial formula gives confidence bounds around our right-tail probability estimate as $a \pm z(\alpha/2) \sqrt{a(1-a)/N}$, where $z(\alpha/2)$ is the cdf of the standard normal distribution and N is the number of observations. If we observe a fraction b of observations in the left tail, we can construct an equivalent confidence bound for the left-tail probability. We can then test whether the two probabilities are equal.

7. For the PSID, McLaughlin (1998) presents evidence suggesting that the underlying wage-change distribution indeed may be skewed to the right.

where $PROPr_t$ is the proportion of observations in the bar r percentage points below the median in year t , $DNEGr_t$ is a dummy variable that takes on the value 1 when $PROPr_t$ is completely below zero, and $DZEROr_t$ is a dummy variable that takes on the value 1 when $PROPr_t$ contains zero.

We explain these equations with the help of figure 2. Consider the equation for the height of the bar 3 percentage points below the median, $PROP3_t$. In years when the median is greater than 3, all observations in this bar are positive, $DNEG3_t$ and $DZERO3_t$ are both zero, and we estimate the height of the bar to be p_3 (top panel). But when the median is less than 3 and all observations in the bar are negative (lower panel), then $DNEG3_t$ equals one and we estimate the height of the bar to be $(1+n)p_3$. The parameter n --the primary parameter of interest--captures the extent to which the histogram bar is altered when it falls below nominal zero. If $n=0$, then the bar is the same height in either case, and there is no downward



nominal rigidity. If $n=-1$, there are no negative wage changes--an extreme case of downward nominal wage rigidity. We have drawn figure 2 for the intermediate case of $n=-1/2$. Note that we constrain n to be the same in each equation. The Kahn test has one equation for each bar that falls above zero in some years and below zero in others (bars in the dotted box in figure 2), for these are the bars that help identify n .

In years when the median wage change is 3 percent, $DNEG3_t$ is zero but $DZERO3_t$ equals one, and the bar is larger than p_3 (middle panel). The term $n \sum_{f=-\infty}^j p_f$ reflects the assumption that any observations that are prevented from showing negative wage change by downward rigidity will appear as zero wage change instead.⁸ This parameterization assumes that the histogram bar at zero wage change will become larger as inflation declines because more of the distribution will fall below nominal zero, and the rigidity will therefore affect a larger portion of the distribution. We also allow the bar that includes zero to be boosted by a constant amount in each year, z . This parameter accounts for long-term contracts or other reasons besides downward rigidity for a concentration of observations at zero nominal wage change.

The Kahn test has many advantages. It is structured to measure precisely the hypothesized effects of downward nominal wage rigidity on the wage-change distribution, is robust to outliers, and does not assume that the underlying distribution is symmetric. The only distributional assumption made is that bars a given distance from the median would have the same height in all years if not for downward rigidity.⁹

One potential problem with the Kahn test, though, is that the test is not completely robust to noise around zero wage change. Specifically, the Kahn test assumes that observations removed from the left tail of the distribution owing to downward rigidity instead

8. With downward nominal wage rigidity, some workers who would ordinarily receive a nominal wage cut might not receive zero wage change, but might instead lose or leave their jobs. To the extent that this occurs, the shortage of individuals' nominal wage cuts will not be matched by a pile-up at zero. But, given that our data are for jobs and not individuals, this is unlikely to affect our results.

9. One reason this assumption may be questionable is if the distribution becomes less disperse as inflation falls and the distribution shifts left (see Card and Hyslop, 1997). In fact, our distributions are about equally disperse in high- and low-inflation years, so our tests are unlikely to be seriously affected by this issue.

appear in the histogram bar that includes zero wage change. True, the histogram bars are 1 percentage point wide, so the bar containing zero also includes observations close to zero but not precisely equal to zero. (This is the sense in which we say in table 1 that the Kahn test is “somewhat” robust to noise around zero wage change.) But if observations that would have been in the left tail in the absence of downward nominal wage rigidity are instead *near* zero, but not precisely in the histogram bar that includes zero, then the Kahn test will understate the amount of downward rigidity.

To deal with this problem, we also present results from a modification of the Kahn test in which we relax the assumption that all observations truncated due to downward nominal wage rigidity are piled up in the zero bar. In this more general setup, a fraction $\theta \leq 1$ of those observations are included in the bar that contains zero. Another fraction δ are included in the bar that contains negative one, and the remainder $(1-\theta-\delta)$ is put in the bar containing positive one:

$$\begin{aligned}
PROP2_t &= p_2 + np_2 DNEG2_t + [z - \theta n \sum_{j=3}^m p_j] DZERO2_t \\
&\quad - \delta n \sum_{j=2}^m p_j DNEG12_t - (1 - \theta - \delta) n \sum_{j=4}^m p_j DPOS12_t \\
PROP3_t &= p_3 + np_3 DNEG3_t + [z - \theta n \sum_{j=4}^m p_j] DZERO3_t \\
&\quad - \delta n \sum_{j=3}^m p_j DNEG13_t - (1 - \theta - \delta) n \sum_{j=5}^m p_j DPOS13_t \\
&\quad \cdot \\
&\quad \cdot \\
&\quad \cdot \\
PROPM_t &= p_m + np_m DNEGM_t + [z] DZEROM_t \\
&\quad - \delta n \sum_{j=m}^m p_j DNEG1m_t,
\end{aligned} \tag{2}$$

where $DNEG1r_t$ is a dummy variable that is 1 when $PROPr_t$ contains negative one, and $DPOS1r_t$ is a dummy variable that is 1 when $PROPr_t$ contains positive one. Obviously, when $\theta=1$ and $\delta=0$, equations (2) reduce to equations (1). We will use this more general, “modified” Kahn test in section 4.2 below.

3. The ECI data

The BLS’s ECI program collects data quarterly on the cost of wages and benefits in specific job categories in an establishment. The sample is intended to be representative of the private nonfarm sector of the economy (excluding households and the self employed).¹⁰ Each establishment remains in the sample for around three to four years before it is replaced by another establishment within the same industry. Because the compensation cost information are collected from the establishment, they should be less affected by measurement error than are compensation data from surveys of individuals (like the PSID).

The dataset has detail on wages and salaries and a number of benefit categories including cash bonuses, insurance, retirement plans, and legally required benefits including social security and federal and state unemployment. Not all benefits are equal, however, when it comes to thinking about circumventing wage rigidity. The growing emphasis on “Pay for Performance” systems suggests that manipulating bonuses in order to target pay to individuals is fairly routine for firms. Firms may also target larger groups of employees by varying benefits such as insurance, and retirement plans. However, the firm has no discretion in changing legally required benefits and certainly cannot use them to single out individuals or even groups of workers. For these reasons, we have chosen to separate out the legally required benefits and to consider three groups of compensation measures: *Wages and salaries* includes straight-time hourly wage and salary costs, including commissions.¹¹ *Total compensation excluding legally required benefits* adds most benefit costs over which the firm

10. The ECI also includes data, which we do not use in this paper, on compensation costs of state and local governments.

11. The distinction between workers who are paid an hourly wage and those who are paid a salary is not recorded in the ECI dataset, so we were unable to make this comparison in our analysis. Nor are commissions recorded separately in the ECI.

has some discretion, including paid leave (vacation, holiday, and sick leave), supplemental pay (nonproduction bonuses, overtime, and shift differentials), health and life insurance, retirement and savings plans, and severance pay.¹² *Total compensation* also includes the cost of several benefits that are required by law: social security and Medicare, state and federal unemployment insurance, and workers' compensation.¹³

The ECI sample collects data on an average of about four jobs per establishment in about 5,000 establishments each quarter, yielding roughly 18,000 total observations per quarter. We examine twelve-month (log) changes in compensation costs, and we picked June-to-June observations to make use of the maximum number of years available. We do not use observations that have been excluded from the calculation of the published ECI, nor do we use observations with values imputed by the BLS (as opposed to values reported by the firm).¹⁴ When we restrict our sample to contain only those observations with consecutive June-to-June, non-imputed values for all benefits, we end up with roughly 5,000 observations per year for 18 years.

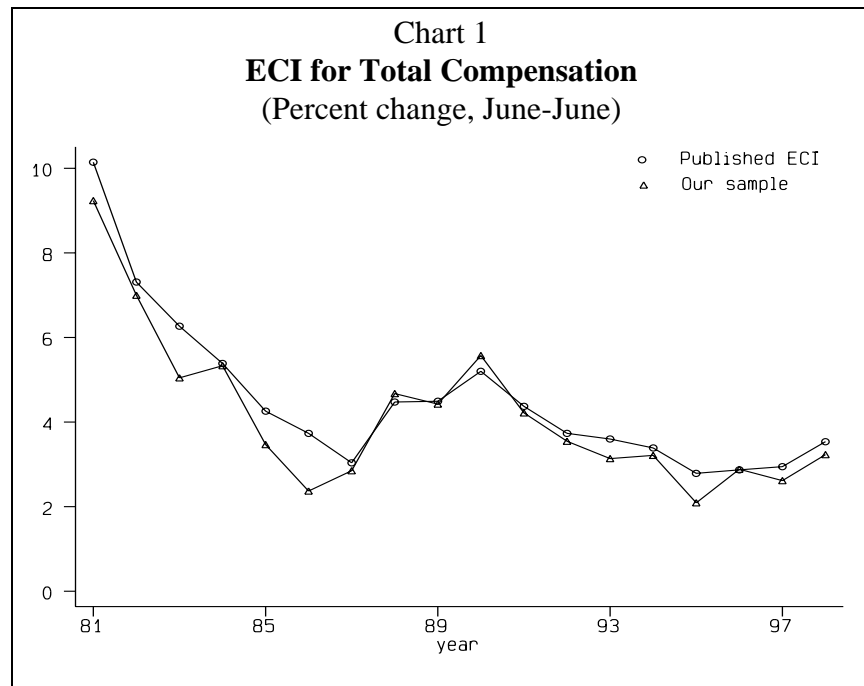
Chart 1 shows the mean June-to-June change in the ECI for total compensation from our sample, along with changes in the published ECI. As can be seen, our micro data

12. The ECI is constructed with an eye toward capturing changes in compensation *policy*. Thus, many of the benefits costs do not fluctuate as much as one might expect, and the ECI may not capture some reductions in benefits costs. For example, the ECI measures overtime costs as the overtime premium; if a firm reduces the number of overtime hours that their employees work, this is not recorded as a cost reduction in the ECI. However, a decline in base wages would lead as well to a drop in the cost of overtime, because a deliberate policy change has reduced the cost of those benefits. In addition, if a change in benefits usage is believed to be *induced* by a change in benefits policy, then the change in usage would be included in the ECI. For example, the cost of health insurance, as with paid leave, often is measured for a fixed structure of usage. But if a relatively low-cost HMO begins to be included among the choice of health plans, and many workers switch to the HMO, then the ECI would include the switch to the HMO as a reduction in health insurance costs.

13. Our compensation measures exclude a few small benefit categories that have too many values imputed. These benefits are “other” paid leave, railroad retirement and insurance, “other” legally required benefits, and long-term disability.

14. Taking out the imputed data does not change the industry or occupation mix of our sample, nor does it qualitatively change any of the results in this paper.

aggregate to be very close to the published index in most years. To calculate these averages, as with all results in this paper, we weight each job's wage change by the number of employees in the overall population represented by that job.




3.1 Jobs versus individuals

The unit of observation in the ECI is the average cost per hour of a given component of compensation for a specific, narrowly defined job within that establishment. Examples of jobs might be accountant grade 1, junior electrician, or administrative secretary. The median number of employees in a job is 7 in establishments of 100 or more employees, and is 2 in smaller establishments.

Because the ECI measures costs as averages across individuals within a job, the ECI data cannot distinguish between average pay changes that reflect variations in the composition of workers within a job, from those that reflect actual pay changes to each employee. For example, if the employee with the most tenure and highest pay in a given job category were to retire, the average pay in the job could fall and be recorded in our data as a “nominal wage cut,” even though no individual actually received a cut in pay. Such job turnover would tend to obscure evidence of downward nominal wage rigidity in these data.

Of course, one could also give examples in which averaging individual pay changes within jobs generates fewer, not more, nominal wage cuts for jobs than for individuals. One bit of empirical evidence that downward nominal wage rigidity is less pronounced at the job

level, as in this paper, than for individuals, comes from Wilson (1999). Using a dataset that tracks individuals and jobs within a large service-sector firm, she constructed job averages from data on individuals. The top panel of chart 2 shows the distribution of wage changes for approximately 6,000 individuals in her firm from 1982-1994. These data exhibit the signs of downward nominal wage rigidity--a scarcity of



observations below zero and a pronounced spike at zero. The lower panel of the chart shows the distribution of wage changes constructed by averaging salaries within jobs in a manner similar to what is used in the ECI. Here, the evidence for nominal wage rigidity is much less

visible; there are many more negative wage change observations, and the LSW asymmetry measure is reduced sharply.¹⁵

Be this as it may, it is debatable whether the relevant unit of observation, from the firm's perspective, is the job or the individual. Returning to our previous example, if this most-tenured worker receives higher pay for some reason other than that her productivity is higher, then her retirement really would reduce the firm's unit labor costs.¹⁶ In that case, the normal attrition of high-tenure employees could provide the firm with flexibility that is missed by focusing on the wage behavior of individuals.

4. Results

This section addresses two questions: Does downward nominal wage rigidity exist? If so, do firms manipulate benefits to achieve more flexibility in total compensation? As discussed in the methodology section, the implications of nominal wage rigidity are most clearly seen in the distribution of wage changes. If workers resist pay cuts and wages are rigid, we expect the distribution of wage changes to be positively asymmetric with a massing of observations at zero and a dropoff below zero. More important, if wages are nominally rigid then the asymmetry and massing at zero should become more pronounced in low-inflation times.

15. Another piece of evidence on this point comes from unpublished BLS data from the National Compensation Survey between 1997 and 1998. Unlike the ECI, the NCS data record the number of individuals per job, so we can restrict our attention to jobs with the same number of individuals in successive interviews. Because these jobs *could* contain precisely the same individuals in the two interviews--in which case they would be free of composition changes--they ought to be more representative of data on individuals than are those jobs where the number of individuals changed. Indeed, jobs in the NCS with the same number of individuals in successive interviews had considerably fewer negative wage-change observations (4 percent versus 17 percent) and a larger spike at zero wage change (20 percent versus 11 percent), again suggesting that the averaging of wages in the ECI job-level data tends to diminish the extent of downward nominal wage rigidity. (We are grateful to Mike Lettau for generating these figures for us.)

16. For a discussion of the literature on how pay moves with tenure/experience/productivity, see Barth (1997) and the references cited therein.

4.1 What is the extent of downward nominal wage rigidity?

Chart 3 and table 2 present summary distributions and statistics for our data. The data are year-to-year changes from 1981 to 1998 in the log of wages and salaries for all private industry jobs in the ECI sample with no imputed data. The top panel of chart 3 displays the distribution of wage changes for all years along with a curve representing a continuous normal distribution with the same mean and standard error as our sample. As can be seen, the distribution of wage changes displays evidence of nominal wage rigidity--it has a clear spike at zero and a faster dropoff on the left side.

The first row of table 2 shows that 14-1/2 percent of all wage changes are negative and 18 percent are zero over the period from 1981 to 1998. As measured by the LSW statistic discussed above, the distribution of wage changes is positively and significantly asymmetric with 13 percent more observations in the upper tail than in the lower.¹⁷ Compared to changes in wages and salaries as measured by the PSID, the micro-ECI data have fewer observations below zero, a much larger concentration at zero, and greater asymmetry.¹⁸

While the evidence above is suggestive of downward nominal wage rigidity, the true test is to see how the distribution varies with inflation. As a first pass, the second and third rows of table 2 contain summary statistics for periods of high and low inflation. It is clear that the spike at zero and the dropoff below zero are more pronounced in the low inflation period: The proportion of observations of exactly zero wage change is 22 percent in the low inflation period versus 11 percent in the high inflation period. The LSW asymmetry measure declines from almost 15 percentage points to below 4 percentage points when moving from the low to the high inflation years. As a more detailed pass, chart 4 shows the distribution of wage and salary changes for each year in our sample. Again, one can see evidence of less truncation in 1981 and 1982, the years in our sample with the highest inflation.

17. Using the approach discussed in footnote 5 above, we calculate approximate standard errors of only about 0.25 percentage point on the LSW statistics, making all of the estimates in table 2 strongly statistically significant.

18. Of course, this apparently larger degree of truncation relative to the PSID could simply reflect the larger median wage change in the PSID data (because the PSID includes the high-inflation 1970s).

Chart 3
Distribution of Changes in Various Compensation Measures
(All Years)

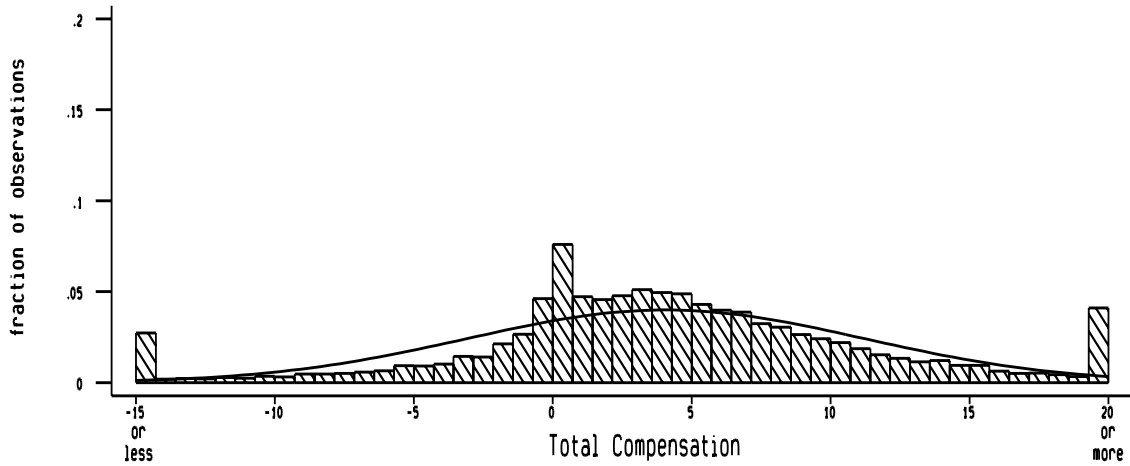



Table 2
Distribution of Wage and Salary Changes

	# Obs.	Median	Standard Deviation	% of Obs.		LSW Asymmetry
				< 0	= 0	
All Years	92,862	3.4	10.8	14.4	17.9	12.6
Low Inflation Years ¹	15,917	2.1	10.2	17.6	22.2	14.6
High Inflation Years ²	17,247	6.9	9.7	8.7	10.6	3.6
Memo: PSID (Wages and salaries, 1971-1988) ³						
Wage and Salary Earners	31,525	6.6	13.0	17.8	8.1	6.8
Wage Earners Only	19,632	6.5	9.5	11.9	10.6	9.7
1. Low inflation years are 1987, 1993, and 1995. 2. High inflation years are 1981, 1982, and 1990. 3. As calculated by Lebow, Stockton, and Wascher (1995).						

For a more formal examination, table 3 presents results of our two preferred tests for downward nominal wage rigidity, the correlation of the LSW statistic with measures of inflation, and the Kahn test. The first two rows show the results of OLS regressions of the LSW asymmetry measure for each year's distribution on a measure of inflation, the unemployment rate, and a constant. The unemployment rate is included to control for possible effects of the business cycle on the shape of the wage-change distribution. Two measures of inflation are used: the June-to-June change in the log of the experimental CPI, and the June-to-June change in the log of the median of the ECI wage-change distribution itself. If the asymmetry in the wage change distribution were due to nominal wage rigidity, we would expect the LSW statistic to fall as inflation rises and the distribution is shifted further away from zero. This would imply a negative coefficient on the inflation measure. As seen on table 3, we find that the coefficients on the inflation measures are negative and significant. For example, a 1 percentage point increase in the median of the distribution implies a 1.6 percentage point decline in asymmetry.

Chart 4
Distributions of Changes in Wages & Salaries by Year

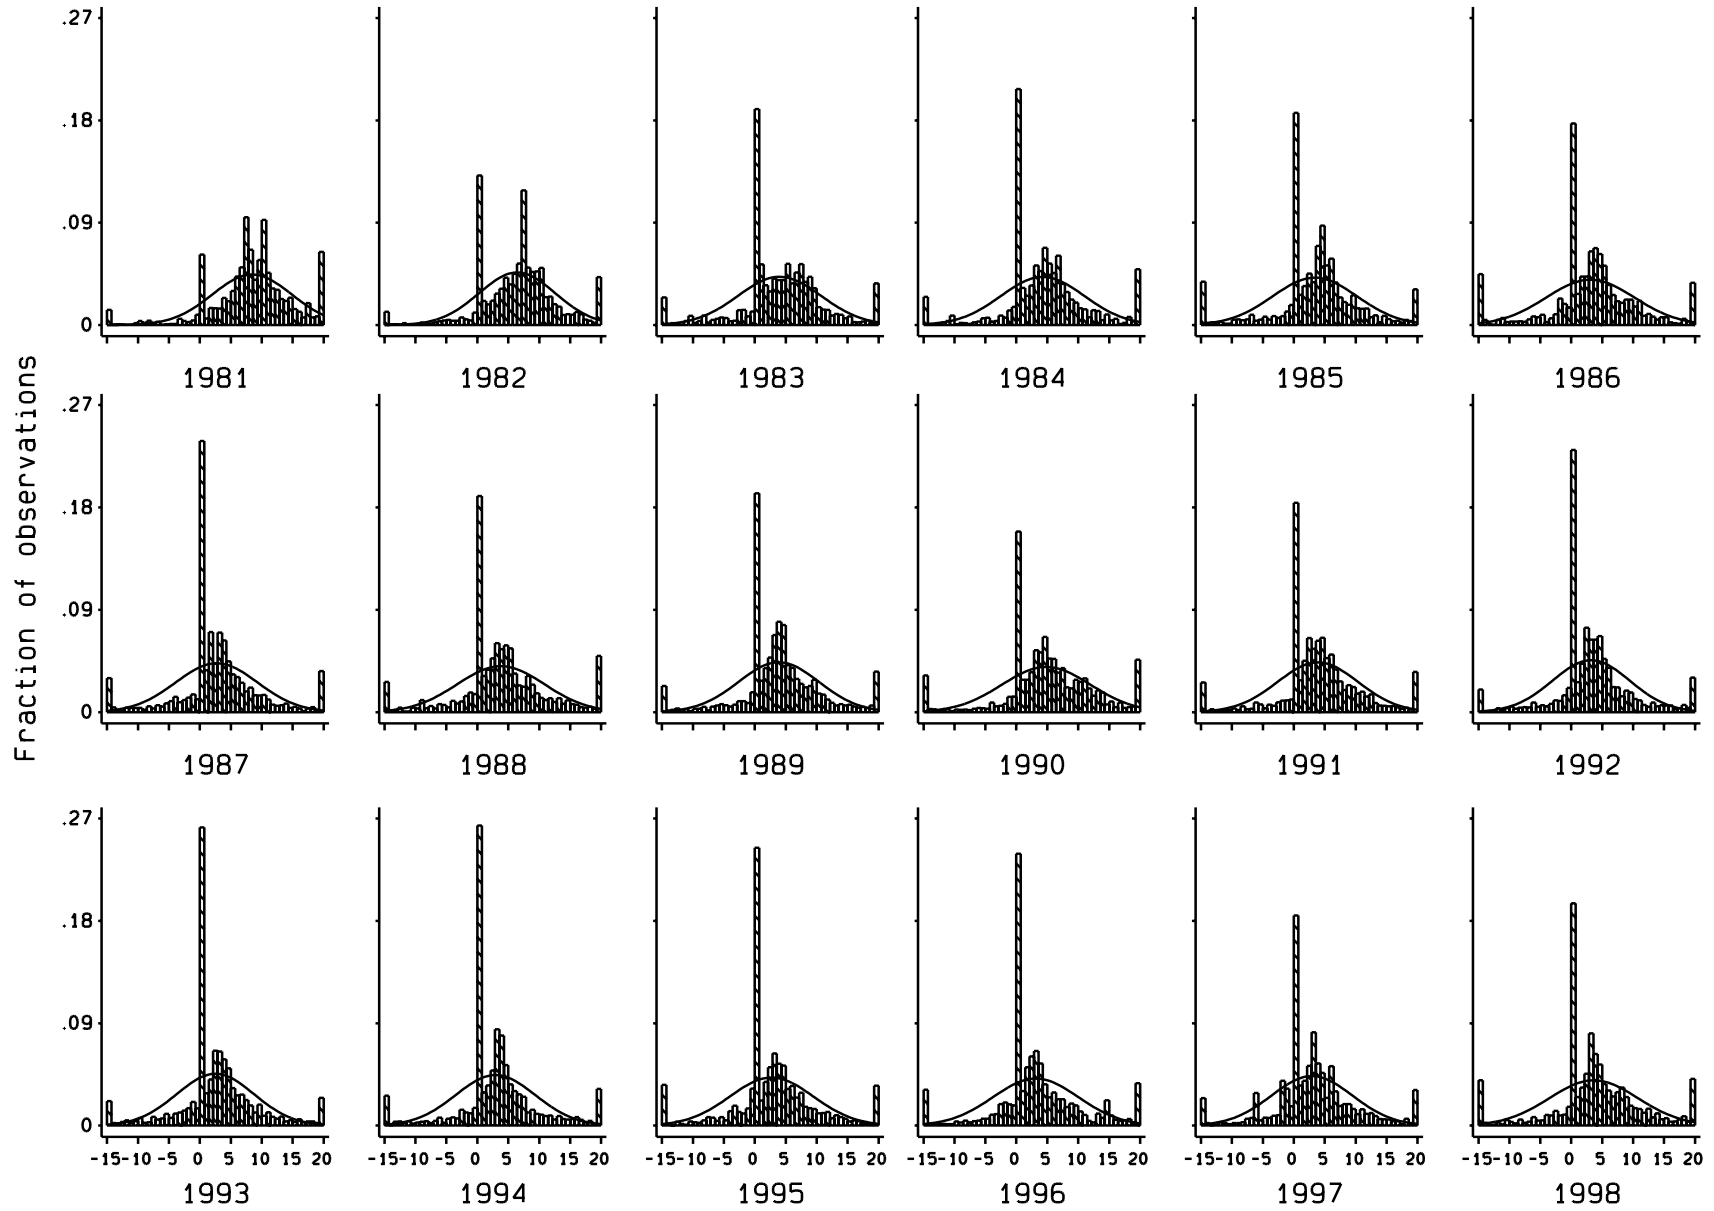


Table 3
Does the shape of distribution of wage and salary growth vary with inflation?

	Wages and Salaries	Memo: PSID (stayers)	
		Wage and Salary Earners	Wage Earners only
LSW Asymmetry tests¹			
Coeff. on median of distribution	-1.54* (.39)	-.18 (.30)	-.88* (.28)
Coeff. on inflation ²	-1.03* (.43)	-.33 (.24)	-.73* (.24)
Kahn tests³			
<i>n</i>	-.47* (.03)	.04 (.03)	-.47* (.03)
<i>z</i>	9.98* (.89)	5.15* (.22)	4.43* (.38)
<p>Note. * indicates significance at the 5% level.</p> <p>1. OLS regression of the LSW asymmetry measure on an inflation measure, the unemployment rate, and a constant. Data for each year are weighted by the number of observations used in creating that year's wage change distribution. PSID results are taken from Lebow, Stockton, and Wascher (1995).</p> <p>2. Twelve-month percent changes (June to June) in experimental CPI.</p> <p>3. Results are estimated using SUR in nonlinear least squares. In each equation, the fraction of compensation changes in the region between <i>r</i> and <i>r</i>-1 percentage points below the median is regressed on a constant and a set of dummy variables, with the equations put through a logistic transformation. Data for each year are weighted as in 1. PSID results are taken from Kahn (1997) and include dummies for bars immediately above and below zero.</p>			

While the correlation of the LSW asymmetry measure with inflation provides evidence of downward nominal wage rigidity, recall from our earlier discussion that we could also obtain a negative correlation between asymmetry and inflation if the underlying distribution of wage changes were skewed for other reasons. Because of this, we concentrate on the results of the Kahn tests, shown in the remainder of table 3.

As discussed in the methodology section, the Kahn test is robust to the shape of the underlying distribution. The main assumption of this test is that, absent downward nominal wage rigidity, the shape of the distribution of wage changes does not vary systematically with inflation. If this is true, then in terms of histogram bars, a bar a given distance from the median should contain *roughly* the same amount of observations in any year. However, if firms find it difficult to give wage cuts, then in times when the range of the bar includes only values below zero, the number of observations in the bar should drop. The *n* coefficient is an

estimate of the fraction by which the number of observations in a bar drops when it falls below zero. Because downward nominal wage rigidity also implies a pile-up of observations at zero, the Kahn test includes a constraint that the number of observations "removed" from the bars below zero must be reflected in the massing at zero. However, because we might expect a massing at zero for other reasons--rounding, long-term contracts, etc.--the Kahn test also estimates a z coefficient, which represents the amount of the massing at zero due to factors other than nominal wage rigidity.

The lower section of table 3 presents our results from the standard Kahn test as described in equation (1).¹⁹ The n coefficient is significantly negative and about $-1/2$, implying that histogram bars are reduced in half when they fall below zero. This estimate of n provides strong evidence of nominal wage rigidity and suggests that if wages were totally flexible, the fraction of observations below zero would be double what we observe. In comparison with Kahn's results using both wage and salaried workers in the PSID, we find much more significant evidence for nominal wage rigidity. Interestingly, our results are the same as Kahn's for hourly-wage earners alone.²⁰

4.2 Do firms use benefits to achieve greater flexibility?

While we find strong evidence of downward nominal rigidity in wages and salaries, it is possible that firms vary benefits--which may be less salient or visible to workers, and are more difficult to value and compare--to achieve more flexibility in total compensation. The detailed benefits information in the micro-ECI data allow us to test this theory.

19. The model is estimated using nonlinear least squares. We note three features of this estimation, none of which affects the results qualitatively. First, following Kahn, we use SUR to account for correlation in the errors across equations. Second, we weight the data for each year by the number of observations used in generating that year's wage change distribution. Third, because the dependent variables are limited--heights of histogram bars cannot be negative--we perform a logistic transformation on each equation. That is, for the first equation in (1), $PROP2 = f(\cdot)$, we actually estimate $\ln[PROP2/(1-PROP2)] = \ln[f(\cdot)/(1-f(\cdot))]$.

20. This result raises the question of whether measurement error of some kind in the salary data from the PSID is obscuring the true underlying rigidity of salaries. Unfortunately, as mentioned above, we cannot separate wage and salary jobs in our data to test this further.

As discussed above, we break the total compensation measure into three categories in order of diminishing employer control over the benefits granted: wages and salaries, total compensation excluding legally required benefits, and total compensation. Chart 3 and table 4 show the basic summary data for these compensation measures. As can be seen in the chart, there appears to be a shortage of negative changes in compensation, just as with wages and salaries. However, this shortage diminishes as benefits are added. Even more noticeably, the spike at zero drops when using the more inclusive measures of compensation. Surprisingly, these changes in the shape of the distribution appear to come equally from firm-controlled benefits (such as nonproduction bonuses and insurance) and legally required benefits.

This observation is also supported by the statistics in table 4. Measuring compensation including all firm-provided benefits except those required by law increases the fraction of negative observations by 5 percentage points and reduces the LSW asymmetry measure by 4 percentage points; including legally required benefits adds another 3 percentage points to the proportion of wage cuts and reduces asymmetry by the same amount. The fraction of observations at exactly zero change also falls sharply when broadening the definition of compensation. Note, however, that the proportion of total observations that are near zero (between -1 and 1) jumps sharply when moving from wage and salaries to total compensation excluding legally required and then to total compensation. We will use this fact shortly.

Table 4
Various Compensation Measures in Private Industry

	Median	Standard deviation	% of Observations			LSW Asymmetry
			< 0	= 0	near 0	
Wages and salaries	3.4	10.8	14.4	17.9	5.4	12.6
Total Comp ex. required benefits	3.6	10.9	19.3	7.6	10.7	8.5
Total Compensation	3.7	10.9	22.3	0.9	14.4	5.3

Table 5 presents results of our more direct test of downward nominal wage rigidity, the Kahn test, for the various measures of compensation. The results for the standard Kahn test are shown in the upper part of the table. The first column contains the coefficients shown earlier for wages and salaries. The second column contains results when the measure of compensation including firm-controlled benefits is used. While still significantly negative, the size of the n coefficient drops markedly. Taken literally, the results in the second column suggest that about 34 percent of potential salary cuts are suppressed rather than 47 percent as shown in the first column. The results for total compensation are even more dramatic: The fraction of suppressed wage cuts falls to 17 percent.

These results from the standard Kahn test for total compensation are misleading, however. Recall from table 4 that as benefits were added, the proportion of observations of exactly zero change fell sharply while the proportion near zero rose, suggesting that many of the additional negative observations are small ones, perhaps providing just a little flexibility. But as we discussed in section 2, the Kahn test may understate the extent of rigidity in this situation, because the test includes a constraint that the censored wage cuts must be piled up onto the bar that includes zero. We therefore loosen this constraint by estimating the “modified” Kahn test shown by equation (2), which allows observations removed from the bars below zero to be added to one of three bars--the bars immediately above and below zero as well as the bar that includes zero. If the truncated observations in fact are added to the zero bar, then the n coefficient should be little changed by this modification. But if a sizable portion of truncated observations instead are added to the bars near zero, the modified Kahn test should generate a more negative--and more accurate-- n parameter.

The second half of table 5 presents the results of this modified test. For wages and salaries and total compensation excluding required benefits, we estimate that almost all of the truncated observations are added to the histogram bar that includes zero.²¹ As a result,

21. Our estimation constrains the proportion of truncated observations added to each of the three histogram bars to lie between zero and one.

Table 5
Kahn Tests for Various Compensation Measures in Private Industry

	Wages & Salaries	Total Comp. Excl. Required	Total Compensation
Standard Kahn test (equation 1)			
<i>n</i>	-.47* (.03)	-.34* (.02)	-.17* (.02)
<i>z</i>	9.98* (.89)	3.55* (.58)	1.95* (.33)
Modified Kahn test (equation 2)			
<i>n</i>	-.47* (.03)	-.35* (.03)	-.30* (.02)
<i>z</i>	9.98* (.89)	3.82* (.63)	1.94* (.30)
Fraction of truncated observations:			
Added to zero bar (θ)	1.00 (.00)	.91 (.05)	.53 (.04)
Added to bar above zero ($1-\theta-\delta$)	.00 (.00)	.08 (.13)	.30 (.06)
Added to bar below zero (δ)	.00 (.00)	.01 (.14)	.17 (.07)

loosening the constraint only slightly alters the *n* coefficient for these compensation measures. The results are markedly different for total compensation, where many of the truncated observations are added to the bars above and below zero, and the *n* coefficient rises in absolute value from -.17 to -.30--statistically indistinguishable from the coefficient for total compensation excluding required benefits. Using this modified test suggests that legally required benefits add small changes at and around zero but do not increase the flexibility of benefits beyond what firm-controlled benefits adds.

This result accords well with our intuition. It seems reasonable to suppose that, when faced with downward nominal wage and salary rigidity, firms achieve greater flexibility by altering benefits. However, it is harder to believe that firms rely to an equal extent on variations in their payments of legally required benefits to achieve flexibility. First of all, the

cost of these benefits are out of firms' control. Second, firms presumably would like to target individual employees or specific groups of employees when giving pay cuts, whereas changes to legally required benefits are almost always across-the-board.

Of course, the basic message of table 5 is that, while not eliminating downward rigidity, the inclusion of benefits substantially lessens it; the Kahn coefficient is reduced by more than one third, from 0.47 to 0.30. Given that firms are achieving greater flexibility by varying benefits within their control, it is interesting to know which benefits are driving the results. Table 6 presents summary statistics and Kahn test results for wages and salaries combined one at a time with selected benefits. (The Appendix table presents results for the full collection of benefits.) The top row shows the benchmark results for just wages and salaries, and the following rows show results for wages and salaries plus a particular benefit. Interestingly, no benefit in itself has much effect on the results, indicating that one must examine combinations of benefits to explain the added flexibility they provide. Nevertheless, a few benefits do have effect on the Kahn n parameter that are at least noticeable (although these do not have a noticeable effect on the n parameter from the modified Kahn test). One is nonproduction bonuses, which, given anecdotal evidence about flexible "pay for performance," would be expected to provide some flexibility for firms, despite the fact that these bonuses, on average, only add up to about 1.6 percent of total compensation in the ECI. A few other, less performance-specific benefits also seem to play a role. Health insurance,

Table 6
Adding One Benefit at a Time to Wages and Salaries

	Percent of total comp.	% of Obs.		LSW Asym.	Kahn test n	Modified Kahn n
		< 0	= 0			
Wages & salaries	72.5	14.4	17.9	12.6	-.47	-.47
Non-prod. bonuses	72.5 + 1.6	16.3	14.7	11.7	-.40	-.45
Health insurance	72.5 + 5.3	16.6	9.8	9.8	-.40	-.46
Paid holidays	72.5 + 2.3	15.2	16.3	12.0	-.42	-.46
Pension & saving	72.5 + 2.9	17.0	14.6	10.8	-.43	-.47

paid holidays, and employer contributions to pension and saving plans--which together account for more than 10 percent of total compensation--all make small contributions to reducing the Kahn n parameter.

5. Potential Effects of Downward Nominal Wage Rigidity and the Micro-Macro Puzzle

In section 4, we presented evidence of substantial downward nominal wage rigidity in the micro ECI data, which raises the question of the impact of this rigidity on key macroeconomic variables. At least in theory, firms are thought to reduce employment levels in the face of worker resistance to nominal pay cuts. Therefore, we begin by getting a rough sense of the impact of our estimated degree of rigidity on aggregate wage inflation and equilibrium unemployment (the NAIRU) in a very simple model of how this rigidity affects labor market behavior.

Beginning with the potential effects of rigidity on wage inflation, we first generate counterfactual wage change distributions assuming no downward nominal wage rigidity. This is easily done using the results of the modified Kahn test in table 5. There, we found that for wages and salaries, 47 percent of observations in the histogram bars to the left of zero were truncated and piled up in the bars containing or next to zero. To construct our counterfactual distribution, we move the truncated observations back into the bars to the left of zero. Note that by doing this, we are assuming that downward nominal wage rigidity has no other effects on the distribution of wage changes--an assumption to which we will return later.

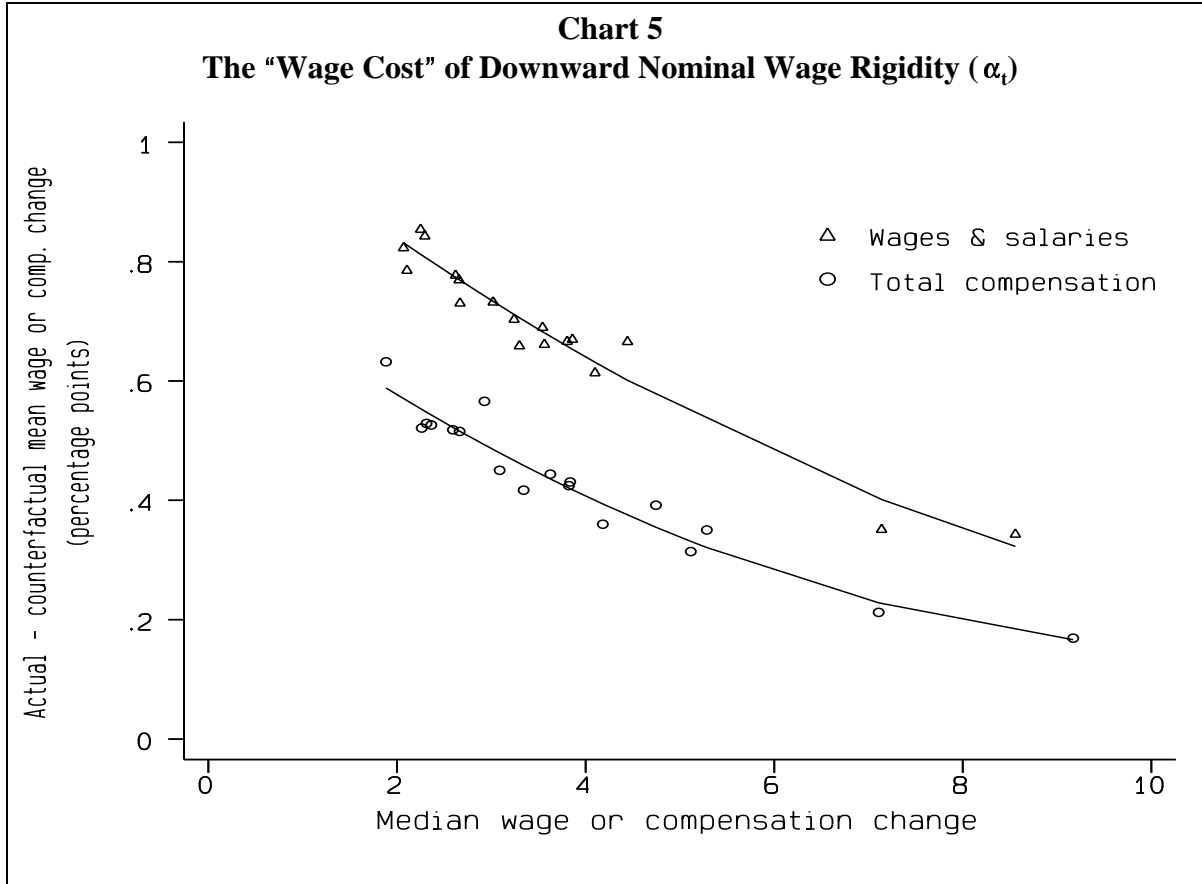
The mean wage change in this counterfactual distribution will, by construction, be smaller than in the actual distribution whenever n --the parameter from the Kahn test that measures the degree of truncation--is estimated to be negative. The difference between the mean wage changes in the two distributions also will depend on inflation. For any given n , the mean wage change in the counterfactual distribution will be only a little smaller than in the actual distribution when inflation (and the median of the distribution) is high and the downward rigidity is binding for only a small share of jobs. But the mean wage change will be a great deal smaller in the counterfactual distribution when inflation is low and the rigidity is assumed to bind for a larger share of jobs.

Although we assumed in the Kahn test that all histogram bars used in the regressions are affected equally by downward nominal rigidity, the regressions only include bars extending to wage cuts of about 10 percent.²² In constructing our counterfactual distributions, then, the question becomes how should we treat larger wage cuts than that. It seems unlikely that worker resistance to pay cuts causes firms to truncate extremely large wage cuts at zero, and this argues for treating the region far in the left tail as unaffected by downward rigidity. Toward that end, we assumed that n is as estimated for all histogram bars up to some cutoff, and that n equals 0 for larger wage cuts. We calculated results for cutoff points of -10 percent, -25 percent and -50 percent.²³

The results of this exercise are summarized in chart 5 for both wages and salaries and total compensation. The chart relates the median wage change in each year's distribution to the "wage cost" of rigidity--the difference between mean wage changes in the actual and counterfactual distributions, which we will call $\alpha_{\vartheta} = \alpha(n, \dot{p}_{\vartheta})$ in what follows. Clearly, under the assumptions of this stylized exercise, the wage cost is larger in years in which inflation and, hence, the median of the distribution are lower. As would be expected, given that the shape of the wage change distribution is peaked and not uniform, the relationship is slightly nonlinear. We have drawn a quadratic regression line through these points. When comparing the line for total compensation with that for wages and salaries, we see a smaller effect of downward nominal rigidity on average compensation changes than on average wage changes, reflecting the smaller estimate of n for total compensation (-.30) compared with wages and salaries (-.47). The curves are constructed using -25 percent as the lowest wage cut subject to truncation; but the more negative is the assumed cutoff point, the more pronounced will be the difference between the means of the actual and counterfactual distributions.

22. This is because only histogram bars that are above zero in some years and below zero in others are used in identifying n , and the range of median wage inflation is about 10 percentage points.

23. Altonji and Devereux (1998) estimate a model with such a cutoff point. Using PSID data, they estimate cutoffs ranging from -25 to -50 percent. It is unclear how those figures, based on data on individuals, relate to the ECI data on jobs.



Taking the logic of this exercise a bit further, we can calculate the implicit effect of our estimated wage costs on the NAIRU. It is typically assumed that in a world with no downward nominal wage rigidity, there is only a short-run relationship between inflation and unemployment. This relationship is often captured using a Phillips Curve, which relates compensation growth (\dot{w}) to expected inflation \dot{p}^e , (usually proxied for by lags of inflation), trend productivity growth ($\dot{\pi}$), and a measure of slack such as the unemployment rate (U), as well as other variables (Z), given the parameters a , β , and γ :

$$\dot{w}_e = a + \dot{p}_e^e + \dot{\pi}_e + \beta U_e + \gamma Z_e + \varepsilon_e. \quad (3)$$

When combined with a simple markup equation relating price inflation to the change in wages and trend productivity, $\dot{p}_e = \dot{w}_e - \dot{\pi}_e$, the estimated parameters allow one to calculate a

NAIRU, defined as the unemployment rate that is consistent with no expectational errors ($\dot{p}^e = \dot{p}$) in the absence of shocks ($Z = \varepsilon = 0$):

$$NAIRU_{\theta} = \frac{a}{-\beta} . \quad (4)$$

In this model, the NAIRU is constant. However, if downward nominal wage rigidity has significant effects on wage costs--which, as we saw in chart 5, would become larger as inflation falls--then it would require higher unemployment to offset this wage cost and maintain stable inflation. In this case, inflation and unemployment would be negatively correlated in the long-run, and the NAIRU would no longer be constant.

More specifically, suppose that $\dot{w}_{\theta}^{\oplus} = \dot{w}_{\theta}^{\kappa\oplus} + \alpha_{\theta}$, so that the actual wage change we observe under the assumption of downward nominal rigidity ($\dot{w}_{\theta}^{\oplus}$) is equal to wage change assuming no such rigidity ($\dot{w}_{\theta}^{\kappa\oplus}$) plus the positive wage cost term α_{θ} . In this case, estimating equation (3) using actual wage inflation is incorrect because the wage data used are affected by downward nominal wage rigidity, whereas the equation is specified assuming a world without such rigidity. Instead, we can rewrite the Phillips Curve consistent with downward rigidity as follows:

$$\dot{w}_{\theta} = a + \alpha_{\theta} + \dot{p}_{\theta}^{\ominus} + \pi_{\theta} + \beta U_{\theta} + \gamma Z_{\theta} + \varepsilon_{\theta} . \quad (5)$$

The time-varying NAIRU derived from equation (5) is:

$$NAIRU_{\theta} = \frac{a + \alpha_{\theta}}{-\beta} . \quad (6)$$

Equation (6) shows that downward nominal wage rigidity in this simple model has two implications for the NAIRU: First, except at very high inflation rates when α_{θ} is close to zero, the NAIRU under rigidity is always above that in the flexible world. Second, the NAIRU in equation (6) increases when inflation falls and the wage cost of downward rigidity

rises. To calculate how the NAIRU varies when inflation falls from, say, 10 percent to zero, we calculate the following relation:

$$NAIRU(\dot{p} = 0) - NAIRU(\dot{p} = 10) = \frac{\alpha(n, 0) - \alpha(n, 10)}{-\beta} . \quad (7)$$

We estimate the Phillips Curve from equation (5) using the measure of α_t for total compensation from chart 5. As shown in the first column of table 7, when we constrain α_t to enter the Phillips curve with a coefficient of unity, we estimate β to equal -0.37. (We will return to the rest of table 7 shortly.)

Plugging this estimate of β into equation (7), we generate estimates from this model of the impact on the NAIRU of moving from 10 percent inflation to zero, broken down into stages to illustrate the nonlinearity of the relationship. Results are shown in table 8.²⁴ Focusing first on the results for wages and salaries using a 25 percent cutoff, this simple framework predicts that a reduction of inflation from 10 percent to 2 percent would lead to a 1.4 percentage point increase in the NAIRU, and a further reduction to zero inflation would add another 0.6 percentage point to the NAIRU. Moving to total compensation, we know that the degree of truncation falls from about 50 percent to about 30 percent when benefits are included; and table 8 shows that as a result, a reduction from 10 percent inflation to zero is predicted to add a total of 1.4 percentage points to the NAIRU--about one-third less than would be implied by looking only at the results for wages and salaries. Of course, all of these figures are a little larger when using a 50 percent cutoff, and a little smaller when using a 10 percent cutoff.

24. Our calculations assume that the median wage change moves one-for-one with inflation plus productivity growth, taken to be 1.2 percent per year--the average since the late 1970s. That is, we assume that when inflation is zero, the median of the wage-change distributions shown in chart 5 is at 1.2 percent.

Table 7
Phillips Curve Estimates

	(1)	(2)	(3)
Coefficient on the Unemployment rate (β)	-.37 (.23)	-.40 (.24)	-.39 (.22)
Coefficient on the "Wage Cost" α	1.00	-.52 (1.54)	—
Adjusted R-squared	.933	.934	.939

Note. The dependent variable is the growth rate of the published ECI for total compensation in private industry, from 1981 to 1998. Standard errors are in parentheses. Along with α and a demographically weighted unemployment rate, the equation contains the change in the unemployment rate, three lags of the rate of change of the nonfarm business deflator (with sum constrained to unity), a ten-year moving average of trend productivity growth (with coefficient constrained to unity), and variables to capture changes in the minimum wage and employer contributions for social insurance.

Table 8
How Large is the Micro-macro Puzzle?

	Increase in NAIRU in simple model generated by reducing inflation from				
Measure of compensation	10 to 6 percent	6 to 4 percent	4 to 2 percent	2 to 0 percent	10 to 0 percent
	Cutoff at 50 percent				
Wages & salaries	.8	.5	.6	.7	2.6
Total compensation	.3	.3	.5	.6	1.7
	Cutoff at 25 percent				
Wages & salaries	.5	.4	.5	.6	1.9
Total compensation	.2	.3	.4	.5	1.4
	Cutoff at 10 percent				
Wages & salaries	.0	.2	.4	.5	1.1
Total compensation	.0	.2	.3	.4	1.0

5.1 The Micro-Macro Puzzle

Table 8 presents estimates that point to a possibly sizable effect of downward nominal wage rigidity on the NAIRU. The last column suggests that, all other things equal, over the period from 1980 to 1999 when inflation, as measured by the nonfarm business deflator, fell from 10 percent to 1 percent, the NAIRU might have been expected to increase between 3/4 and 2-1/4 percentage points.²⁵ The size of this predicted effect of rigidity seems at odds with actual events. Instead of rising, economists' estimates of the NAIRU have moved down; for example, the time-varying NAIRU estimated in Gordon (1998) has declined with inflation over this time period.

This evidence points to a puzzle--why does the wage rigidity at the micro level not appear to affect the aggregate inflation and unemployment rate? It is a straightforward matter to test the hypothesis implied by the earlier model that the wage cost of rigidity, α_t , affects wage inflation one-for-one. As shown in the second column of table 7, the coefficient on the wage rigidity variable, α_t , is not significantly different from zero. Indeed, when the Phillips curve is estimated without the restriction that the coefficient on α_t be unity, the estimated coefficient is *negative*--though estimated imprecisely. We take this result not as suggesting that rigidity leads to lower unemployment, but that the relationship between inflation and equilibrium unemployment is likely much more complicated than depicted in the simple model used for table 8.

Because a thorough investigation of what could resolve this micro-macro puzzle is beyond the scope of this paper, here we simply lay out several possible explanations.²⁶ Of course, as we have discussed at length, one way firms could offset inflexible wage costs is to adjust benefit costs, and we have found that including benefits along with wages resolves perhaps one-third of the puzzle. Beyond benefits, we see two general links between micro

25. This result stands in contrast to the argument made by Akerlof, Dickens, and Perry (1998) that one should not expect to have seen very much macroeconomic impact over the postwar period, because inflation (combined with productivity growth) has been high enough to avoid the need for very many nominal wage cuts. We have shown that given the extent of rigidity that we find in the ECI, a decline in inflation to 4 percent, and certainly to 2 percent, would be predicted to have had a notable effect on the equilibrium unemployment rate.

26. Several of these issues are discussed as well in Hogan (1997).

rigidity and macro consequences which, if broken, might explain our “micro-macro” puzzle: The link between rigidity and higher compensation costs, and the link between higher compensation costs and lower employment levels.

First of all, whereas we assumed a constant degree of rigidity over our sample period--for example, a constant Kahn n parameter--the rigidity could in fact diminish over time in a low-inflation environment. Either firms could make increasingly greater use of flexible forms of pay, such as bonuses, when faced with rigidity, or workers’ resistance to nominal pay cuts itself could diminish over time.

Also related to the link between nominal rigidity and compensation growth, our calculation of the wage cost of rigidity in chart 5 assumed that downward nominal rigidity affected only the portion of the distribution to the left of zero. However, suppose firms respond to this rigidity by holding down pay increases at the top end of the wage change distribution as well, thereby keeping overall labor costs invariant to downward nominal wage rigidity. Such an adjustment implies that the portion of the distribution above the median also would be sensitive to inflation--becoming less disperse and more concentrated during times of low inflation.²⁷ Although such wage compression could well have employment consequences, the channel would have to be more subtle than the wage cost as measured here.

Firms may have other ways of preventing downward nominal wage rigidity from affecting their compensation costs. Firms may be able to make use of promotions and quits to adjust their labor costs, especially if the tendency of wages to rise with tenure is not reflective of rising productivity (Wilson, 1999). Furthermore, when downward nominal wage rigidity is prevalent, optimizing firms might build knowledge of this rigidity into their wage structure to begin with; for example, firms might offer lower starting wages than they otherwise would have to help reduce the need for nominal wage cuts in the future.

27. In our data, some simple tests suggest that the upper half of the wage change distribution is not sensitive to inflation. In particular, we tested whether the difference between the median of the distribution and the 75th, 90th, 95th, or 99th percentiles are sensitive to inflation and we found little evidence that this was true.

The second link, that between higher overall compensation costs and lower employment, could be broken if firms respond to compensation costs by adjusting on other margins. If the constraint on compensation is temporary, perhaps because the adverse demand shocks for certain types of labor are subsequently reversed, then firms may prefer to smooth through the disturbance by varying hours or allowing profits to decline without reducing employment. And even if the constraint is not viewed as being temporary, if firms are not optimizing fully at all times, their inability to cut wages could spur them to undertake productivity enhancements, thereby validating their higher compensation costs.

Of course, it also is possible that downward nominal wage rigidity has put upward pressure on the NAIRU, but that the impact has been obscured by other factors affecting the NAIRU at the same time (and that are not adequately captured by the Phillips curve). A recent study by Katz and Krueger (1999) explores several hypotheses to explain the apparent decline in the NAIRU over the past twenty years, and finds that the single largest factor has been changing demographics. However, this effect was already taken account of in our calculations, which employ a demographically adjusted unemployment rate. According to Katz and Krueger's estimates, the remainder of the influences that they examine (improvements in job matching, rising incarceration rates, and a reduction in workers' bargaining power) would not have been large enough to offset the hypothesized effects that we show in table 8. But, perhaps, other factors have had some impact. In that regard, one final possibility is that the low inflation itself may have improved the functioning of labor and product markets enough to offset any difficulties arising from downward nominal wage rigidity.

6. Conclusions

We examine the extent of downward nominal wage rigidity using the microdata underlying the BLS's employment cost index--an extensive, establishment-based dataset with detailed information on wage and benefit costs. We find that the number of nominal wage cuts is about half of what would be expected in the absence of such rigidity--stronger evidence of downward nominal wage rigidity than found in previous studies using panel data on individuals. Firms appear able to undo some but not all of this rigidity by varying

benefits: Total compensation displays about one-third less rigidity than do wages alone. Given our estimated amount of rigidity, a simple model predicts that the disinflation over the 1980s would have raised equilibrium unemployment notably--in contrast to the actual behavior of unemployment over this period.

If downward nominal wage rigidity has any macroeconomic effects they are likely to be more complicated than is commonly depicted. It is possible that rigidity has significant effects on aggregate unemployment that have been masked by other factors. It is also possible that rigidity has had no effect on aggregate unemployment because firms have made other adjustments that could be reflected in other macroeconomic variables such as profits or productivity. Clearly, understanding the macroeconomic effects of downward nominal wage rigidity remains an important question for future research.

A second area for future research is to better understand the relationship between rigidity at the individual and job levels. Resistance to pay cuts is certainly generated at the individual level, but it is unclear whether individual- or job-level rigidity matters more for firms. To the extent that firms wish to single out workers for pay cuts, individual-level rigidity would be more costly. To the extent that pay differences across individuals within a job reflect something other than productivity differences, firms may be able to achieve cost reductions through turnover, say, without cutting any individual's pay, and rigidity at the job level would be more relevant. Of course, if costs at the individual level were more important, the tensions highlighted in section 5 would only be increased, because, as we have noted, rigidity probably is more substantial at the individual level. In either case, better understanding whether firms focus on jobs or individuals may provide insight in how firms adapt to downward nominal wage rigidity.

References

- Akerlof, George, William Dickens, and George Perry, "The Macroeconomics of Low Inflation," *Brookings Papers on Economic Activity* (1996:1), pp. 1-76.
- Altonji, Joseph, and Paul Devereux, "The Extent and Consequences of Downward Nominal Wage Rigidity," mimeo, July 1998.
- Barth, Erling, "Firm-specific Seniority and Wages," *Journal of Labor Economics* 15(3), July 1997, pp. 495-506.
- Bewley, Truman, "Why Don't Wages Fall During Recessions?" manuscript, Yale University, October 1998.
- Card, David, and Dean Hyslop, "Does Inflation Grease the Wheels of the Labor Market?" in *Reducing Inflation: Motivation and Strategy*, Christina Romer and David Romer, eds., National Bureau of Economic Research Studies in Business Cycles, v. 30, 1997.
- Craig, Ben, "Are Wages Inflexible?" *Economic Commentary*, Federal Reserve Bank of Cleveland, April 1, 1995.
- Fortin, Pierre, "The Great Canadian Slump," *Canadian Journal of Economics* 29(3), August 1996, pp. 761-87.
- Freedman, Charles, and Tiff Macklem, "A Comment on The Great Canadian Slump," *Canadian Journal of Economics* 31(3), August 1998, pp. 646-65.
- Gordon, Robert J., "Foundations of the Goldilocks Economy: Supply Shocks and the Time-Varying NAIRU," *Brookings Papers on Economic Activity* (1998:2), pp. 297-346.
- Groschen, Erica, and Mark Schweitzer, "The Effects of Inflation on Wage Adjustments in Firm-Level Data: Grease or Sand?" mimeo, Federal Reserve Bank of New York, December 1996.
- Hogan, Seamus, "What does Downward Nominal-Wage Rigidity Imply for Monetary Policy?" Bank of Canada Working Paper 97-13, June 1997.
- Kahn, Shulamit, "Evidence of Nominal Wage Stickiness from Microdata," *American Economic Review* (December 1997), pp. 993-1008.
- Kahneman, Daniel, Jack L. Knetsch, and Richard Thaler, "Fairness as a Constraint on Profit Seeking: Entitlements in the Market," *American Economic Review* 76 (September 1986), pp. 728-41.

Katz, Lawrence, and Alan Krueger, "The High-Pressure U.S. Labor Market of the 1990s," *Brookings Papers on Economic Activity* (1999:1), forthcoming.

Kimura, Takeshi, and Kazuo Ueda, "Downward Nominal Wage Rigidity in Japan: Is Price Stability Costly?" mimeo, Bank of Japan, May 1997.

Lebow, David, David Stockton, and William Wascher, "Inflation, Nominal Wage Rigidity, and the Efficiency of Labor Markets," Finance and Economics Discussion Series #94-45, Board of Governors of the Federal Reserve System, October 1995.

McLaughlin, Kenneth, "Rigid Wages?" *Journal of Monetary Economics* vol. 34 (1994), pp. 383-414.

McLaughlin, Kenneth, "Are Wage Changes Skewed Away from Nominal Wage Cuts?" mimeo, Hunter College, September 1998.

Shafir, Eldar, Peter Diamond, and Amos Tversky, "Money Illusion," *Quarterly Journal of Economics*, 112(2), May 1997.

Shea, John, "Comment," in *Reducing Inflation: Motivation and Strategy*, Christina Romer and David Romer, eds., National Bureau of Economic Research Studies in Business Cycles, v. 30, 1997.

Tobin, James, "Inflation and Unemployment," *American Economic Review*, vol. 62 (1972), pp. 1-18.

Wilson, Beth Anne, "Wage Rigidity: A Look Inside the Firm," Finance and Economics Discussion Series #99-22, Board of Governors of the Federal Reserve System, April 1999.

Appendix

Adding One Benefit at a Time to Wages & Salaries

	% of Obs.		LSW Asym.	Kahn test n^*	Modified Kahn n^*
	< 0	=0			
W&S	14.4	17.9	12.6	-.47	-.47
W&S plus...					
Overtime	14.9	16.9	12.3	-.46	-.45
Vacation	15.3	16.0	12.0	-.47	-.45
Paid holidays	15.2	16.3	12.0	-.42	-.46
Sick leave	14.9	16.8	12.4	-.47	-.45
Shift differentials	14.5	17.7	12.7	-.48	-.45
Non-production bonuses	16.3	14.7	11.7	-.40	-.45
Severance pay	14.5	17.7	12.6	-.46	-.44
Supplemental unemployment	14.6	17.8	12.6	-.44	-.48
Life insurance	17.3	12.0	9.9	-.46	-.51
Health insurance	16.6	9.8	9.8	-.40	-.46
Accident insurance	16.0	15.0	11.4	-.46	-.51
Pension & saving	17.0	14.6	10.8	-.43	-.47
Social security	16.3	8.2	10.9	-.62	-.65
Federal UI	15.6	14.5	11.6	-.42	-.43
State UI	20.6	4.8	7.0	-.35	-.44
Workers' compensation	20.8	5.1	6.8	-.28	-.38

* All coefficients are significant at the 1 percent level.