Review of Income and Wealth
Series 44, Number 2, June 1998

LABOR INCOME INDICES DESIGNED FOR USE IN CONTRACTS PROMOTING INCOME RISK MANAGEMENT

BY ROBERT J. SHILLER

Yale University

AND

RYAN SCHNEIDER

Sanford C. Bernstein, Inc.

We propose that labor income indices be used to define settlements in many contracts, such as labor contracts, indexed bonds, or income securities. We discuss the issues in producing labor income indices for such uses, and develop prototype indices using U.S. data from the Panel Study of Income Dynamics (PSID). People are grouped by a clustering algorithm based on an estimated transition matrix between jobs, by education level, and by skill category. The groupings are defined so that few people move between them. For each grouping we generate a labor income index (1968–87) using a hedonic repeated-measures regression method. The indices show substantial variability through time, confirming the potential importance of the use of such indices in contracts. There is also substantial variability across groupings, as for example between the agriculture/labor grouping and other groupings, confirming the importance of using the grouping indices rather than aggregate indices in contracts.

Well-constructed indices of labor income for groupings of people who share job opportunities or skills could have important use in defining settlements or payments in contracts that have income risk management (the reduction of uncertainty about future real incomes) as part or all of their goals. As far as we have been able to determine, labor income indices have never been used for this purpose. However, contracts settled on such labor income indices might someday be extremely important: they could reduce extraneous income inequality and have efficiency benefits as well.

I. PROPOSED INDEXATION OF LABOR CONTRACTS

An important use of labor income indices in contracts would be to index wages and salaries in labor contracts. Such indexing might be used in contracts offered to individual employees by their employers, whether or not the contracts are part of collective bargaining. The use of labor income indices would then replace the use of the consumer price index (CPI) cost-of-living allowance (COLA) clauses in some labor contracts today for income risk management, to reduce the uncertainty about future real incomes specified in the contracts. At the initial signing of a contract the employee would have a wage or salary specified for this year, with a provision that in subsequent years the wage or salary would be

Note: The authors thank Vassilis Hajivassiliou, John Hartigan, Lawrence Katz, Alvin Klevorick, and Elisabetta Magnani for their help and comments, and Gary Solon for data. This research was supported by a grant from the National Science Foundation.

adjusted by a formula related to the change in the labor income index for the occupational or other grouping into which that person is placed, as well as other possible factors.

There is a distinct advantage to indexing labor contracts this way, rather than indexing them to the CPI. Contracts indexed to the CPI will eventually, as time goes on, tend to create widening disparities between the employee's contract income and the employee's potential income in other jobs. The discrepancy arises in part because not all those in other jobs will have their pay indexed to the CPI, in part because others are renegotiating their contracts from time to time, in part because others are changing jobs, possibly moving to new industries or regions. Such discrepancies tend to create incentives for employees to leave their current job (if the alternative income is higher) or for employers to layoff or create incentives for current employees to leave (if the alternative income is lower).

Existing contracts that require the firm to pay employees a constant real wage (defined by a consumer price index) or income for a long time may be suboptimal, as pointed out by Gray (1976) and Fischer (1985): the contract allows no adjustment to market conditions. What has tended to happen with such contracts is that the contract fixes a real wage rate, rather than a real income level, and adjustments are made through layoffs. COLA clauses, which grew widely in significance for labor contracts through the 1970s, declined dramatically in importance following the "great recession" of 1981–82, when many firms claimed that, despite their options to layoff employees, they were put into financial distress by these clauses, see Gay (1984). Had labor income indices been available for contract settlement then, the consumer price index might have been replaced or modified using these, rather than merely deleted or downweighted in contract formulae.

The income-indexed labor contracts are perhaps more responsive to the concerns of employees than are CPI-indexed contracts, since fairness is a paramount consideration. Employees sometimes appear not as concerned with the abstract concept of preservation of standards of living. Given this, labor unions might prefer to sign a contract that fixed changes in wages in the out years to those of some such reference group, if better indices of wages or income of such reference groups were available.[1]

The advantage to contracts indexed to a labor income index accrues to both employee and employer. Since labor-management bargaining is inherently about ratios to market value, the presence of any zero-mean random shock to the agreed-upon ratio is damaging, *ex ante*, to both sides of the bargain. Defining contract settlements in terms of an index that does not correlate well with the market value of labor is like throwing a random element into the contract, like introducing a coin toss to determine subsequent incomes. Surely neither side would want that.

The process of labor contract definition is slow and difficult, involving many compromises, and so it is difficult to predict what kinds of indexation will survive

[1]A question is sometimes raised about what would happen if every labor contract were indexed to a labor income index; what, then, would tie down the aggregate income level? However, labor contracts are regularly renegotiated and might also have clauses in them for regular adjustments tied to variables other than a labor income index. This simple question does, however, suggest that there may ultimately be some general equilibrium issues for contract designers to consider when the day comes that these contracts are commonplace.

this process. Given the failure of COLA clauses based on consumer price indices to become a standard in labor contracts, it is worth exploring whether labor income indices might be more useful.

II. Other Proposed Uses of Labor Income Indices

Labor income indices might also be used in contracts between firms and their suppliers. It would seem logical that firms signing long-term contracts to supply their products would want to put into the contract some protection against variations in the cost of their employees (as well as in other costs). Then the firm could count on making, in effect, a contracted profit without risk.

Insurance contracts might also make use of labor income indices. Disability insurance logically would provide for replacement of labor income of the grouping to which the claimant belongs, thereby insulating the person against only disability risks and not insuring against other risks that result from changing circumstances of the labor market relative to the contracted payout. Linking the payout to a labor income index rather than to a CPI would tend to diminish the moral hazard problem that arises when disabled people may lose an incentive to return to work if the income they could earn has fallen in real value. Similarly, life insurance policies could promise a payout as a proportion of a labor income index.

Child support and alimony payment schedules after divorce could be indexed each year to a labor income index for that year for the appropriate grouping of the payer. This would be much better than indexing payments to the CPI, since, with the labor income indexing, payments to be made would more closely match the ability to pay.

Labor income indices might also be used to index payments on bonds. Indeed, Robert Rubin, U.S. Treasury Secretary, has mentioned in a May 16, 1996 press release that the Employment Cost Index, a sort of full-employment labor income index, was one of the candidates under consideration as a measure of "inflation" for the inflation-indexed bonds that the Treasury was planning to create. (In fact the CPI was later chosen.)

There ultimately may also be liquid futures or options markets in such indices, see Shiller (1993a, b). This is especially true if the indices are already used in other contract settlements, such as labor-management or firm-supplier contracts. Firms could use the futures and options markets to swap their wage-bill costs for some other cash flow, and this would be an especially good hedge if their labor contracts were already tied to the labor income index on which the futures or options contracts are based. Employees could also effectively swap their contracted labor income for a more stable income flow using such futures or options markets (either directly or indirectly through intermediaries), and this kind of hedge would work especially well if their labor contract were tied to the index used in the future or options contract.

III. Existing Candidate Indices for Contract Settlement

We do not presently observe labor income indices used in any such contracts. Part of the reason for the absence of such use may be that no published indices designed for this purpose exist.

The Employment Cost Index published by the U.S. Bureau of Labor Statistics would seem useful for many contracts, if the changes in the index really measured well the changes in costs of hiring people that are relevant to individual firms. Labor negotiators sometimes speak of goals for contract package growth rates in terms of a basis spread with the growth rate of the Employment Cost Index, but the Employment Cost Index has never been directly used for risk management contracts, as far as we have been able to determine. The index seems to be used widely by business economists as an indicator of future inflation and is routinely reported in the news as a general indicator of economic conditions. The Employment Cost Index is not based on repeated measures of individuals: for each industry the index is just a fixed-weight Laspeyres index averaging employment costs reported; see O'Conor (1989) and Wood (1982). Should there be a change in the characteristics of people working in the industries or labor types, or a shift of people from one industry type to another, then the Employment Cost Index could be unrepresentative of costs of hiring people with fixed characteristics.

The traditional personal income measures, such as that published by the Bureau of Economic Analysis (BEA) in the United States, are just aggregations of individual incomes without regard for the changing group of people that earn the income or for the changing quality characteristics of this population, including population growth, age distribution, female labor force participation, experience, and education level. The income indices we develop here, which might be regarded as full employment cost indices, may be more suitable for such contracts.

IV. New Indices of Labor Income

In what follows we create indices of individual labor income for use in contracts. Since the intended contracts are those used to manage individual income risk, the indices must be indices of labor income accruing to specific claims that individuals have on income. Creating accurate indices means basing our analysis on the course of labor income of individuals through time, so that our indices follow individual claims on income and not dissimilar claims. It also requires grouping individuals together in such a way that most people do not readily move between groups, so that each index refers to the labor income of a relatively fixed group of people. It also means attempting to control, using hedonic variables, for changes in the characteristics of our sample that identify individual claims, since even when we follow individuals through time there can be a potential for biases in the indices. Biases may arise if the changing individual characteristics indicate that changed individual income is not indicative of changed income opportunities or indicate that the composition of the sample has changed through time. Our indices, constructed with such controls, might be interpreted as indices of labor income of fully-employed people representative of the grouping in which we have placed them, after their student years and before retirement.

To construct these indices we use the Panel Study of Income Dynamics (PSID), a survey conducted annually starting in 1968 by the Survey Research Center of the Institute of Social Research at the University of Michigan.

We use a clustering algorithm based on a method of Hartigan (1975) to define groupings using the transition matrix among occupation-industry categories. We also define groupings in terms of education level and skill category (defined along lines suggested in Reich, 1992). To produce indices, we apply a modification of the hedonic repeated measures regression technique (Shiller, 1993a, b). This technique infers labor-income changes for people in a grouping only from changes in labor income that individuals in that grouping actually experienced.[2]

Defining Groupings of People

We use three main methods of defining groupings of people for which labor income indices will be created: a method based on cluster analysis of PSID occupation-industry categories using the estimated transition matrix between occupation-industry categories, a method based only on education levels, and a judgmental method based on the Reich (1992) classification of jobs into three skill categories.

The first method, using cluster analysis, finds clusters of occupation-industry categories between which there are few transitions. The PSID reports 48 intertemporally consistent occupation-industry categories from 1968 to 1987, see Table 1. From our sample of all heads of households and spouses who have been in the sample for at least two years, we computed a 48×48 estimated transition matrix P. The ij-th element of this matrix is the fraction of the observations on individuals who were observed in occupation-industry category i in which the individual was in category j in the subsequent observation (almost always the next year); we interpret the elements as estimated probabilities.[3]

TABLE 1

OCCUPATIONS AND INDUSTRIES

The base codes that the PSID data use are:	
Occupations:	Industry:
(1) Professional/Technical	(1) Natural ResourcesProduction
(2) Manager/Official/Proprietor	(2) Construction
(3) Sales/Clerical	(3) Manufacturing
(4) Craftsman/Foreman	(4) Transportation/Public Utilities
(5) Operatives	(5) Wholesale/Retail Trade
(6) Laborer, Farmer, Service Worker	(6) Finance/Services
	(7) Public Administration
	(8) Agriculture

Using the transition matrix computed from 72,876 occupation-industry category responses in the PSID, the categories are grouped using a clustering algorithm, a modified version of the improved leader algorithm of Hartigan (1975).

[2]The method has some analogy to that produced by Solon, Barsky, and Parker (1994).

[3]When a person has two jobs, the occupation-industry category was defined as that of the main job, for which the person worked more hours; if hours were equal in the two jobs, the category was defined as that of the job with the higher earnings. For each individual, the number of observations in occupation industry category i is the number of years in which that individual was observed to be working in that industry and for which we also have a subsequent observation in the PSID on the occupation-industry category for that individual.

The clustering algorithm makes a first pass through the non-diagonal data to assign C "leader" pairs—the basis of clusters—pairs having estimated transition probabilities in both directions [$p(j, i)$ & $p(i, j)$] that are above the initial threshold value $\tau(0)$ specified in advance. Our notion of clusters requires that the initial leader pairs show transition probabilities above the threshold both ways; we want the initial clusters to represent stable occupation-industry pairs, not one-way transitions. The algorithm then uses V passes through the data to assign the remaining (48–2C) occupation-industry categories to the C clusters based on the transition probabilities $p(i, \{c\})=\Sigma(j\in c)p(i, j)$, $c=1, \ldots, C$, above a threshold $\tau(v)$, $v=1, \ldots, V$. When an occupation-industry category has a $p(i, \{c\})$ above the threshold for more than one cluster, it is assigned to the cluster for which $p(i, \{c\})$ is highest. The threshold $\tau(v)$ is lowered by the program as each pass, $v=1, \ldots, V$, through the data occurs. The process continues until all elements have been assigned.[4]

This algorithm is invariant to changes in the input order of the data and only requires the initial specification of the threshold value $\tau(0)$—the algorithm automatically adjusts the threshold value down in very small increments to assign the remaining data. Based on a $\tau(0)$ value of 0.025, there are $C=7$ clusters in the data that are used to generate indices of labor income for hedging purposes.

The 48 initial occupation-industry categories and the 7 job clusters obtained using the modified leader algorithm are presented in Table 2, along with names we judgmentally assigned to these clusters after viewing the results of the cluster analysis. As a summary measure of our success, Table 3 presents the 7×7 transition matrix for the job clusters. Table 3 also presents the proportion changing occupation-industry categories within the job cluster.

Note that the clustering was quite successful in the sense that most off-diagonal elements of the transition matrix in Table 3 are 0, those that are non-zero are quite small. Of course, there are still transitions between the job clusters which may tend to compromise the use of the indices for contract settlement. The extent of compromise will depend on the similarity of the changes in the labor income indices between clusters for which the transition probabilities are higher.

The second method of defining groupings of people is very simple. A person's grouping was defined only in terms of education level, without any reference to the occupation-industry category. There are three education categories: did not graduate from high school, graduated from high school, and graduated from college. The transition matrix for these three categories is shown in Table 4. It is evident that transitions are fairly rare, and thus that education level appears to be another good way of grouping people for contracting purposes.

The third method of defining groupings of people is based on Reich's (1992) attempt to group worker occupations in the U.S. economy into skill categories, within which, he argued, people are likely to have similar income trends. His analysis is largely descriptive, dividing workers into those engaged in symbolic-analytic services (SAS), routine production services (RPS), in-person services (IPS), and others. There is neither empirical analysis nor results in Reich (1992)

[4]There are pathological data sets for which not all occupation-industry categories would be assigned to a cluster, but that problem does not arise in our analysis.

TABLE 2

RESULTS OF CLUSTER ANALYSIS

Cluster (Grouping) Names:
A: Professional/Technical
B: Production
C: Services
D: (Public/Private) Works I
E: (Public/Private) Works II
F: Trade/Labor
G: Agriculture/Labor

	1 Prof/Tech	2 Manager Officer/ Proprietor	3 Sales/ Clerical	4 Craftsman/ Foreman	5 Operative	6 Laborer/ Farmer/ Service Worker
1 Natural Resources Production	A	E	B	B	G	G
2 Construction	A	D	E	D	F	F
3 Manufacturing	A	B	B	B	B	G
4 Transportation, Public Utilities	A	D	D	F	G	G
5 Wholesale/Retail Trade	A	F	F	F	F	G
6 Finance/Services	A	C	C	C	C	G
7 Public Administration	A	E	E	E	F	G
8 Agriculture	G	G	G	G	G	G

TABLE 3

TRANSITION MATRIX FOR GROUPINGS ASSIGNED USING CLUSTERING METHOD

	To: A	B	C	D	E	F	G	Sample Size (Initial cluster)
A	0.99905	0.00095	0.00000	0.00000	0.00000	0.00000	0.00000	14,893
B	0.02130	0.94550	0.03320	0.00000	0.00000	0.00000	0.00000	16,859
C	0.01907	0.01907	0.87600	0.04330	0.04256	0.00000	0.00000	9,228
From: D	0.02992	0.04189	0.01151	0.82512	0.03018	0.03811	0.02327	6,882
E	0.00000	0.00000	0.00000	0.03313	0.89728	0.02281	0.04726	3,241
F	0.00000	0.00000	0.00000	0.00000	0.04230	0.94765	0.01004	11,216
G	0.00000	0.00000	0.00000	0.00000	0.00000	0.01241	0.98760	10,557

Note: Estimated probability of transferring to a different occupation-industry category within the same job cluster:

0.06780 0.05896 0.08713 0.07115 0.06324 0.08978 0.01241

to compare our results with, and Reich describes the skill categories only very broadly and by example. Although it is difficult to assign individuals to these skill categories based only on information in the PSID, to provide another method of producing groupings we attempted a judgmental assignment of the PSID

TABLE 4

TRANSITION MATRIX FOR EDUCATION CATEGORIES (GROUPINGS)

		To:			
		No High School	High School	College	Sample Size
	No High School	0.9937	0.0063	0.00	16,027
From:	High School	0.00	0.9982	0.0018	37,471
	College	0.00	0.00	1.00	19,378

TABLE 5

GROUPING BASED ON SKILL CATEGORIES

Skill Classifications (Reich, 1992):

S = Symbolic/Analytic Services (SAS)
I = In-Person Services (IPS)
R = Routine Production Services (RPS)

	1 Prof/Tech	2 Manager Officer/ Proprietor	3 Sales/ Clerical	4 Craftsman/ Foreman	5 Operative	6 Laborer/ Farmer/ Service Worker
1 Natural Resources Production	S	I	R	R		
2 Construction	S	S	I	R	R	R
3 Manufacturing	S	S	I	R	R	R
4 Transportation, Public Utilities	S	S	I	R	R	R
5 Wholesale/Retail Trade	S	I	I	R	I	R
6 Finance/Services	S	S	I	R	S	R
7 Public Administration	S	S	I	R	I	R
8 Agriculture	S					

TABLE 6

TRANSITION MATRIX FOR GROUPINGS BASED ON SKILL CATEGORIES

	SAS	IPS	RPS	Sample Size (Initial Grouping)
SAS	0.9360	0.0512	0.0218	28,489
IPS	0.0761	0.8431	0.0808	16,695
RPS	0.0312	0.1066	0.8622	27,692

occupation-industry categories to his three skill categories. This assignment is presented in Table 5. Some industry-occupations (especially those in agriculture or mining) appear to have no place in the Reich scheme and so are unassigned. For the most part, we have interpreted professional or managerial jobs as symbolic-analytic services, sales jobs as in-person services, and craftsmen, operatives and laborers as routine production services. There are a few exceptions to this rule; for example, we have classified managers and operatives in "finance/services" as in-person services, based on our impression of the work these people do. Estimated transition matrices between these groupings are shown in Table 6. Estimated transition probabilities are higher in general than with either our improved-leader algorithm clusters or our educational groupings. Hence, we feel that our skill grouping method is the least successful of our three grouping methods, the lack of success possibly due to our need to rely on inadequate information about skills in the PSID.

Repeated Measures Design

The basic idea of our index number construction method is very simple: changes in the index should be estimated only from changes in individual incomes. Hedonic variables enter our method only to ensure that there are no sample composition effects on our indices, and to ensure that our index represents normal full employment income. Even though the PSID, on which our empirical work is based, attempts to follow all members of families through time, its composition changes dramatically over years because of births, deaths and participation rate changes as well as defection from the sample. Moreover, the people in our various groupings also change through time. The index that the method produces may be thought of as indicating actual changes in income that *individuals* in various groupings experience, where the individuals are standardized, in terms of a number of characteristics, to be representative of their grouping.

We select as hedonic variables indicators of the kind of claims that people have on labor income. For example, we pick race and sex as hedonic variables; these do not change for individuals and so changes in racial or sexual composition of our sample indicates changes in the representativeness of the sample. We omit from our list of hedonic variables any variables that change stochastically for individuals if these changes may plausibly be associated with the very income changes that we want to represent. For example, we would omit the individual's tax bracket, since this changes in response to income changes. Unfortunately, there is sometimes ambiguity whether a hedonic variable is a good indicator of individual claims on labor income. For example, education level may reflect innate ability or motivation, and thus be a good indicator of a claim on labor income, but education may also change in response to income changes.

To appreciate the importance of hedonic variables, consider our manhours variable. Over our sample period there has been an increase in the female labor force participation rate, related to the changing societal attitudes towards women's working. Without taking account of this increase, there would tend to be an upward bias in our index as a measure of the labor income risks we want to hedge. The increased total income due to the increasing participation rate would affect settlements in risk management contracts even if no person saw a change in wages or income opportunities.[5]

[5]Our method will have a selection bias problem if those who drop out of the labor force do so because their individual earnings capability has changed relative to that of those who stay in, and if the importance of this effect changes through time. In future work, one might combine the methods used here with explicit economic modeling of the participation decision, as in Haveman and Buron (1993).

Repeated Measures Index Construction Method

The hedonic-repeated measures method that we use is described in Shiller (1993a, b), where it is shown that the method may be thought of as ordinary hedonic regression augmented with dummy variables indicating each individual in the sample, though those dummy variables do not appear explicitly in the formulation as it is presented here. The present application must be modified from that described before only in that there are multiple indices, one for each grouping, and occasional shifts of subjects between groupings.

The index number construction method is based on a generalized least squares (GLS) regression. Each observation of the dependent variable in the regression is the change in the log of labor income between successive observations of income for an individual; these observations on changes in log income are arrayed into a column vector y. The consecutive observations on labor income that are the basis of y are not always one year apart, since there are some gaps in our data on labor incomes. Given that there are G groupings of people (G indices to be produced) and H hedonic variables, the matrix x of independent variables for our regression will be constructed by first constructing a matrix with the same number of rows as y but with $(G+H)T$ columns, and then deleting columns as necessary to prevent multicollinearity; at least one column (which we will take to be the first) will have to be deleted. Before deletion of columns, for each grouping and for each hedonic variable there are T columns. In the t-th such column, $t=1, \ldots, T$, for any row, the element is zero unless t corresponds either to the time period of the first observation on labor income, or to the time period of the second observation on labor income. If it is the first observation, then the element of that column is minus the hedonic variable for the date of the first observation of labor income for the individual (or minus one if the column corresponds to a grouping and that is the grouping of that individual in the time period of the first income observation corresponding to that row). If it is the second observation, then the element of that column is the hedonic variable for the date of the second observation on labor income for the individual (or plus one if the column corresponds to a grouping and that is the grouping of that individual in the time period of the second observation on income corresponding to that row).

For example, let us suppose that there are four time periods, periods 0, 1, 2, and 3 ($T=4$), that there are only two groupings ($G=2$), that there are no hedonic variables ($H=0$), and that there are only three individuals in our sample. Suppose also that the first individual is always in grouping one, the second individual always in grouping two, but that the third individual moves from grouping two to grouping one between periods 1 and 2. Our method of defining our groupings was supposed to assure that there are not very many moves between groupings, but still there are such moves, and we include such a move in our example to illustrate how we handle them. Defining $y(i, j)$ as the log labor income of individual i[6] in time period j, we have equation (1). The

[6]Top-coded values were replaced with actual values using data courtesy of Gary Solon.

$$
(1) \qquad x=\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 1 \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \qquad y=\begin{bmatrix} y(1,1) & -y(1,0) \\ y(1,2) & -y(1,1) \\ y(1,3) & -y(1,2) \\ y(2,3) & -y(2,2) \\ y(3,1) & -y(3,0) \\ y(3,2) & -y(3,1) \\ y(3,3) & -y(3,2) \end{bmatrix}
$$

x matrix has seven, $(G+H)T-1$, columns, three for the first grouping, and four for the second grouping; we have already deleted the first column of the original x matrix. The first occupation has columns (columns one through three) corresponding to periods 1, 2, and 3; there is no zero period since there would be multicollinearity in the x matrix if we had included such a column. The deletion of the first column will not have any effect on our ability to produce index numbers since the (log) index will be set to zero (the log of 1) in period 0 anyway. The second grouping has four columns since we need to account for the spread between the incomes in the two groupings, to account for income changes of people who switch between the groupings. Note how our method handles switches between groupings. Suppose that grouping 2 is a higher-wage grouping than is grouping 1; the wage increase is not treated as an increase in the income in any grouping. The move from grouping one to grouping 2 is what breaks the collinearity of the x matrix, and allows us to estimate all four values, including the period-0 value for the index for grouping 2.

The example shown in equation (1) showed no hedonic variables (only the constant term). It is important to include some hedonic variables, since there may be changes both in the composition of the sample and in the "prices" of characteristics of labor. For an example, let us adapt the x matrix to account for man-hours alone, changing H to one from zero, we have equation (2).

$$
(2) \qquad x=\begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & -h(1,0) & h(1,1) & 0 & 0 \\ -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -h(1,1) & h(1,2) & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & -h(1,2) & h(1,3) \\ 0 & 0 & 0 & 0 & 0 & -1 & 1 & 0 & 0 & -h(2,2) & h(2,3) \\ 0 & 0 & 0 & -1 & 1 & 0 & 0 & -h(3,0) & h(3,1) & 0 & 0 \\ 0 & 1 & 0 & 0 & -1 & 0 & 0 & 0 & -h(3,1) & h(3,2) & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & -h(3,2) & h(3,3) \end{bmatrix}
$$

There are 11 columns in this matrix: 3 $(T-1)$ columns including constant dummies corresponding to grouping one, 4 (T) columns including constant dummies corresponding to grouping 2, and 4 (T) columns corresponding to the single hedonic variable man-hours. In such a formulation, if the hedonic variable is constant through time for each individual (as with a variable such as race or sex) then a column would have to be dropped for the variable (for each such hedonic variable), since the sum of the columns corresponding to the variable would otherwise be zero. Moreover, if the hedonic variable behaves as a non-stochastic function of time the same for all individuals up to an additive constant term that may differ across individuals, then we will also have to drop a column for the variable (for each such variable), since the columns corresponding to the variable would otherwise be collinear with the columns corresponding to the constant term. The sum of the columns corresponding to the constant term [corresponding to the first 7 columns in equation (2)] each multiplied by $(f(t)-f(0))$ where t is the corresponding time period equals the sum of the T columns in x corresponding to this hedonic variable [corresponding to the last four columns in equation (2)].

We assume that the variance matrix ω of the transformed errors is diagonal with variances along the diagonal proportional to the interval between measures. Thus, we are assuming that individual log income deviations from the log income predicted by the regression is a random walk. The GLS estimate is:

$$B=(x'\omega^{-1}x)^{-1}x'\omega^{-1}y. \tag{3}$$

The GLS method under this assumption about serial correlation of errors provides efficient estimates and F statistics corrected for serial correlation. Using the coefficients defined by (3), defining from B for grouping g the $(1+H)\times 1$ vector $B_{g,t}$ of regression coefficients corresponding to time t, $t=0,\ldots,T-1$, ($B_{g,t}$ is the coefficient of the constant dummy corresponding to index g and time t, followed by the coefficients of the H hedonic variables corresponding to time t; values are zero corresponding to a coefficient that is omitted from the regression) and defining the corresponding $1\times(1+H)$ vector x_{t-1} of constant and hedonic variables (the numeral 1 followed by the values of the hedonic variables at time $t-1$ to be used for the index), we derive from our estimated regression coefficient vector a chain index $I_{g,t}$, $t=0,\ldots,T-1$, $g=1,\ldots,G$:

$$I_{g,t}=I_{g,t-1}+x_{t-1}(B_{g,t}-B_{g,t-1}) \tag{4}$$

where $I_{g,0}$ equals 0. (In reporting the index in the tables below we take its antilog and multiply by 100.) Using a chain index keeps the index relevant for the "standardized" individual in each time period. Now the index change in each year is last year's mean quality vector (x_{t-1}) multiplied by $(B_{g,t}-B_{g,t-1})$. We defined x_{t-1} for each job grouping in terms of the average hedonic variables of that grouping in time period $t-1$ using the values in our sample. However, we replaced the sample values, using census data, with U.S. averages for each year, for education, race and sex. Note that since x_{t-1}, and not x_t, multiplies $B_{g,t}$ in the formula (4), this formula controls for changes through time in hedonic variables even though the formula involves values of hedonic variables that change through time.[7]

Choice of Hedonic Variables

Our list of hedonic variables was chosen to allow our method to take account of time variation of individual characteristics that might cause their labor incomes to be spurious indicators of earning power, that is of the market full-time income for individuals in the grouping, and to take account of possible

[7]Further interpretations of the role of hedonic variables in our indices are drawn out in Shiller and Schneider (1995).

unrepresentativeness of the PSID sample.[8] The hedonic variables that we included in our analysis are:

Employment:	Log of Number of Work Hours in the Year
Personal:	White/Non-white, Sex
Job Status:	Unemployed, Self-Employed, Retired
Human Capital:	Education, Experience, Years in Current Job

The education variable was excluded in the regressions for the education categories. Overtime hours are included in the employment variable, as well as hours at a second job. Note that the White/Non-white, Sex, Unemployed, Self-Employed, and Retired variables are 0–1 dummy variables. The unemployment variable is 1 if the individual is unemployed for more than 3 months in the year. The education variable is years of education, so that 12 corresponds to high school graduate and 16 to college graduate. The experience variable is the total months worked since entering the labor force; it is calculated by cumulating employment in each year from the date when the respondent reported entering the labor force. It differs from age minus a constant because different people start work at different ages, and because different people have different spells out of the labor force or spells unemployed. The years at current job variable, which the PSID obtains directly by asking respondents, also differs from age minus a constant. There are thus only two variables, as well as the constant, for which we must drop columns of the x matrix: we must drop columns for the constant term, White/Non-white, and Sex. For the seven groupings, $G=7$, $H=9$, and $T=20$, we have $(G+H)T-3=317$ columns in the x matrix. For the skill categories, G drops from 7 to 3, and so we have 237 columns, while with the education categories H is also reduced, to 8, so that there are 217 columns.

Experience, years in current job, self-employed status and retired status are variables that change through time for individuals, reflecting decisions that they make. For example, a person who decides to become self-employed may be choosing a lower labor income, and possibly a future income path that also grows more slowly, because of perceived advantages to self-employment, such as personal choice of hours of work. If, let us say, the number of self-employed persons changes through time, due to changes in taste for self-employment, then we would see a spurious reduction in the growth rate of our indices if we did not include the self-employment variable. Such a spurious reduction in the growth rate of the index would cause risk management contracts based on the index to force unnecessary payments; for example, an employer hedging employment cost risk might find that the decline in the index due to rise in self-employment would cause contract settlements unmatched by any declines in that employer's wage bill.

It is not entirely clear whether we want to include education as a hedonic variable. Education should be included as a hedonic variable, at least to control for possible changes in the representativeness of the sample, but then controlling

[8]The PSID is not a true random sample of the population; specifically, it excludes working adults who are living with their parents, and who are thus neither head of household nor spouse. When we use census data for hedonic variables in constructing indices, we are partially controlling for unrepresentativeness of our sample.

for education brings in the risk that we are not allowing income to feed back into education.

There is possibly another use of hedonic variables, out of a concern that a labor income index might be expected to respond sluggishly to new market conditions, since most others in their grouping may not negotiate new wages or switch jobs for months or years. Should there be a sudden unexpected pickup in inflation, there could possibly be a temporary drop in real wages until others in the labor market adjust to the new inflation (although the literature on this wage lag hypothesis does not appear to confirm that this tends to happen). Possibly some hedonic variables indicating extent of labor market involvement could be used. We did not do so here since we judged that we cannot be sure who is really renegotiating wages and in what ways those who appear to be renegotiating wages are atypical of the market.

Our list of hedonic variables is limited by data availability, and we do not have any assurance that no important hedonic variables have been omitted. There could be unobserved changes through time in the composition of the sample in terms of such unmeasured hedonic variables. Our repeated measures design will assure that omission of such variables will not cause any direct bias in our indices due to the changing weight given to people with different values of these unobserved variables, but will not protect against biases caused if the "price" of these unobserved characteristics also changes through time.

V. Results of Index Number Construction

Table 7 shows the repeated measures chain indices of labor income from the seven job clusters, and Table 8 shows the repeated measures chain indices of labor income for the education and skill (Reich) categories. Due to the large sample size, the standard errors for these index values are generally low: typically about 1 percent or 2 percent, occasionally 3 percent or 4 percent.[9] The relation between these indices and the Employment Cost Index and Personal Income is described in Schneider (1995).

Our indices are somewhat confirming of trends noted from simpler analyses. Much has been made in the literature of the fact that college graduates or professionals have seen an income improvement relative to high-school graduates or laborers in the 1980s. Our indices confirm this: It can be seen from the tables here that there was an uptrend through the 1980s in the ratio of the cluster A (professional/technical) to cluster G (agriculture/labor) indices, in the ratio of college graduates to high school graduates indices, and in the ratio of the symbolic-analytic-services to routine production services indices. Our log college graduate index rose relative to our log high school graduate index by 0.130 between 1979 and 1987. Katz and Murphy (1992) found somewhat similar results using their average wage data from the Current Population Survey. They did not lump people

[9]The validity of these standard error estimates hinges upon the validity of our assumption that error terms (in the GLS-transformed regression, with variables rescaled by dividing by the square root of the interval between measures) are serially uncorrelated. In fact, the serial correlation coefficient for estimated residuals was only 0.066 for the job clusters regression, 0.085 for the educational levels regression, and 0.089 for the skill categories regression.

with some college in with high-school graduates as we did; they found that the college graduates showed an improvement 1979–87 in log average weekly wage relative to the log average weekly wage of people with 12 years schooling of 0.117, and relative to the average weekly wage of people with 13–15 years of schooling of 0.062.[10] However, our indices do not confirm the steady downtrend in the return to a college education over the 1970s that others have reported. Our log college graduate index rose relative to our log high school graduate index by 0.013 between 1971 and 1979. Katz and Murphy found that the log average weekly wage for college graduates fell relative to the log average weekly wage for those who had 12 years of schooling by 0.115 from 1971 to 1979; and fell relative to the log average weekly wage for those who had 13–15 years schooling by 0.067.[11]

The cluster A, (agricultural/labor) also showed much less income growth over the sample period 1968 to 1987 than did any other cluster. This result may help us interpret why Haveman and Buron (1993) found that the fraction of individuals who could not be able to escape poverty through their own effort grew by 38 percent from 1973 to 1988.

TABLE 7

CHAIN INDICES OF LABOR INCOME: 1968 = 100

	Job Cluster (Grouping):						
Year	A	B	C	D	E	F	G
1968	100	100	100	100	100	100	100
1969	114.2	104.4	106.7	99.6	110.3	95.0	101.2
1970	127.8	116.3	105.5	108.3	124.4	118.3	104.4
1971	144.2	123.7	121.9	106.5	119.8	122.7	113.6
1972	160.3	136.3	133.8	122.4	124.6	129.9	119.1
1973	164.4	142.9	150.0	123.3	137.6	139.2	121.5
1974	168.4	156.1	151.8	134.9	144.4	150.9	132.6
1975	183.4	174.8	171.2	149.0	162.2	179.4	145.5
1976	196.6	188.0	187.9	158.6	180.2	198.6	159.9
1977	213.4	200.5	189.8	172.3	201.6	205.4	169.3
1978	226.6	207.5	201.5	185.6	211.3	203.6	176.6
1979	254.0	219.6	228.4	199.4	222.7	218.3	195.7
1980	262.9	236.5	256.6	211.7	247.3	237.7	208.7
1981	296.3	256.0	280.1	231.4	269.9	259.7	228.2
1982	312.2	278.4	305.6	244.6	277.0	274.4	229.1
1983	346.8	301.8	322.4	241.1	276.9	288.6	227.8
1984	375.4	346.5	359.8	276.5	309.2	306.7	244.5
1985	398.9	361.2	372.0	301.1	335.3	327.8	266.8
1986	421.8	378.2	380.0	324.2	340.3	339.1	282.9
1987	453.2	379.3	368.7	320.2	341.4	362.9	289.6

VI. How Well Do the Indices Represent Individual Incomes?

It is important to assess how well the indices for each grouping capture the movements in income that individuals in that grouping experience. In doing this,

[10]These numbers were inferred from data in Katz and Murphy (1992), page 40, Table I, column 3, rows 5, 6, and 7.

[11]*Loc. cit.* column 2, rows 5, 6, and 7.

TABLE 8

CHAIN INDICES OF LABOR INCOME BY EDUCATIONAL LEVEL AND SKILL CATEGORIES

	Educational Level				Skill Category	
Year	NHS	HS	COLL	SAS	RPS	IPS
1968	100	100	100	100	100	100
1969	105.8	104.1	109.3	115.5	98.8	104.5
1970	107.3	115.0	116.2	145.7	125.1	118.1
1971	112.7	122.0	129.2	147.5	131.2	127.7
1972	115.4	126.3	135.1	153.8	140.3	130.5
1973	124.8	133.6	136.5	157.2	147.1	150.1
1974	136.1	149.1	154.3	166.4	151.6	147.9
1975	144.7	157.7	163.7	174.5	156.6	151.9
1976	145.5	162.2	174.6	206.6	158.0	158.6
1977	150.9	165.2	187.3	218.8	176.2	170.1
1978	163.9	182.1	198.0	210.3	177.5	194.2
1979	178.7	195.9	210.1	216.4	199.3	197.9
1980	191.7	216.2	223.6	238.1	219.1	209.7
1981	188.1	218.9	239.1	250.5	229.8	218.3
1982	205.2	229.5	249.6	251.7	237.1	221.4
1983	194.6	238.6	266.9	259.7	237.2	240.4
1984	215.9	244.2	283.1	279.8	250.7	254.9
1985	217.7	240.7	292.6	292.9	255.7	265.7
1986	238.1	250.1	295.3	312.6	273.3	286.9
1987	238.9	265.7	324.5	332.3	291.5	307.0

it should be borne in mind that there are certain kinds of income movements that we do not intend our indices to cover, namely income movements that occur for an individual because of changes in one of the hedonic variables that we included; for example, income changes that occur because the person drops out of the labor force or retires.

We first produced a set of all one-year and five-year log labor income changes that we can construct for individuals from our sample. For each individual who stayed in our sample for all twenty years, we obtained 19 one-year changes and 15 five-year changes. Individuals who were in our sample for less than the full twenty years produced fewer observations for us.

We separated these observations into groupings by the grouping of the individual at date of the first observation of income, disregarding the grouping at the second date. For each grouping, we compute the regression residuals in the regression (with constant term) of the change in log labor income on the change over the same interval in the hedonic variables with the change in the index over the same interval excluded from the list of independent variables (A), and the regression residuals with the change in the index (B) included. We included the seven hedonic variables that might change through time for each of the seven job cluster groupings and skill groupings, and the six such hedonic variables for the education groupings, since education was omitted as a hedonic variable for these groupings.

The statistic we show in Table 9 for each grouping and for both time intervals is one minus the ratio of the estimated variance of the B regression residuals to the estimated variance of the A regression residuals. This statistic may be interpreted as a measure of the success of our indices in capturing the individual labor income movements that are not explained by the special factors that are represented by our hedonic variables. The final row in the table was computed analogously to the other rows, but with all the groupings put together and the CPI in place of the labor income indices.

The ability of our indices to capture individual income movements as revealed in Table 9 is substantial, though by no means perfect. For one year changes, close to half of the variance of individual log labor income changes not explained by changes in the hedonic variables is explained by the indices, in each job cluster grouping and education grouping. The CPI explains only 20 percent, the skill groupings do only a little better. The five-year changes are perhaps more interesting, as they relate to larger income movements. For the seven job-cluster groupings and the education groupings the indices explain over half of the movement in labor income, as high as 74 percent; the CPI explains 41 percent. It is clear from this table that some of our labor income indices explain most of the variance of five-year changes in log labor income, and that the indices do substantially better than the CPI.

TABLE 9

PROPORTION OF UNEXPLAINED VARIANCES OF LABOR INCOME CHANGES EXPLAINED BY INDICES

Grouping	1-Year Changes	5-Year Changes
A	0.51	0.73
B	0.47	0.60
C	0.42	0.54
D	0.39	0.50
E	0.44	0.52
F	0.43	0.53
G	0.49	0.74
NHS	0.41	0.59
HS	0.39	0.51
COLL	0.50	0.70
S	0.27	0.48
I	0.26	0.46
R	0.21	0.43
CPI	0.20	0.41
No. Observations	64,314	23,912

Note: Figures shown are 1-B/A where B is the estimated variance of the residual in a regression of changes in log labor income on the changes in the hedonic variables and the change in the log labor income index and A is the estimated variance of the estimated residual in the same regression with the change in the log labor income index omitted. The row marked CPI shows results for all individuals where the log of the consumer price index (all urban consumers, annual average) from the U.S. Bureau of Labor Statistics is substituted for the log labor income index.

VII. HOW MUCH INCOME VARIATION THROUGH TIME DO THE INDICES REVEAL?

Table 10 shows a matrix with three different measures of the variability of the indices and their differences. The standard deviations of five-year changes in the log indices are shown along the diagonal computed from the 15 observations we have on five-year changes for each index. These numbers (which take account of the serial correlation of residuals in the GLS method) show how variable through time are the individual indices.

The standard deviation of differences between the five-year change in the log of the row index and the contemporaneous five-year change in the log of the column index are shown above the diagonal, standard deviations computed using fifteen observations for each pair. These numbers show how differently each index behaves through time.

The p values for an F test for pairs of the indices that are the same are shown below the diagonal. These numbers show statistically significant differences between the individual indices. These are p values from ordinary regression F tests (also taking account of serial correlation of error terms via the GLS method), computed for the hedonic repeated measures regressions that produced the indices in Tables 7 and 8; regressions with 317 independent variables for the seven job-cluster groupings; 217 independent variables for the education groupings; and 237 independent variables for the skill groupings. For each F test the null hypothesis was a set of 19 restrictions on the coefficients of the dummy variables corresponding to dates. For each date t from $t = 1$ through $t = 19$ there was a restriction that the t-th coefficient of the first grouping minus the zeroth coefficient for that grouping (if any) equals the t-th coefficient of the second grouping minus the zeroth coefficient for that grouping (if any).

The Table 10 results show that there has been substantial variability in the indices. The substantial income risk shown by the variability could be addressed by using contracts settled in terms of the indices. The results show, moreover, that the differences across indices are large enough that it would pay to write contracts in terms of the individual indices, not just an aggregate index. The p values show that the differences between the indices are usually significant at conventional levels.

VIII. Conclusion

The Table 9 results show that much of individual income movements would be hedgable in terms of contracts constructed using our indices. The Table 10 results show that there has been some success in defining groupings of people, since the income movements are significantly different across the groupings we have defined. Together, these results indicate both the importance and feasibility of income risk management contracts based on such indices.

Our grouping method met with varying success across groupings in terms of the transition matrix. It was most successful in identifying grouping A—professional. Although there were substantial numbers of transitions into this grouping, there were virtually no transitions out of it (less than 1/10 of one percent per year per grouping pair). This is the grouping that, of all our job groupings, showed the highest earnings growth since 1968, much higher than the indices we computed using our interpretation of the Reich classification "symbolic-analytical services," or using the education category "college."

TABLE 10

STANDARD DEVIATIONS AND *F*-TESTS FOR DIFFERENCES ACROSS INDICES

	A	B	C	D	E	F	G
A	7.67	10.99	8.38	13.57	16.35	15.53	14.28
B	0.13	6.77	7.61	9.93	14.50	9.07	12.63
C	0.09	0.09	5.67	9.16	12.25	11.11	9.67
D	0.03	0.05	0.05	7.51	9.24	11.98	5.87
E	0.01	0.01	0.01	0.08	11.87	13.29	7.43
F	0.02	0.02	0.02	0.03	0.02	10.65	11.50
G	0.01	0.01	0.05	0.07	0.07	0.04	8.67
	NHS	HS	COLL				
NHS	6.64	6.51	10.85				
HS	0.08	6.21	10.09				
COLL	0.00	0.05	8.01				
	S	I	R				
S	8.13	15.50	10.47				
I	0.09	7.84	11.29				
R	0.01	0.08	11.26				

Note: Figures above the main diagonal (in percent, for pairs of indices) are standard deviations of differences between five-year growth rates. Figures along the diagonal are standard deviations of five-year growth rates (in percent) of the indices. Figures below the diagonal are p values for F tests of the null hypothesis that the two indices are the same.

Our clustering method met with substantial success in identifying grouping G—labor/agricultural. The proportions of transitions out of category G was only 1.24 percent, though again there were more transitions into this category.

The education groupings also met with substantial success. Transitions between education categories were very low, and there were substantial differences across the indices. The skill groupings as we interpreted them show a transition matrix that we interpret as less diagonal, and therefore the groupings are less successful. Of course, we did not have the information in our occupation-industry categories to knowledgeably classify according to the skill categories.

Our efforts here were circumscribed by the data categories defined in the PSID. We hope that those contemplating further expansions of such panel study data collection efforts will bear in mind the needs of researchers who wish to produce labor income indices for contract settlement, so that, in the future, there can be more success in defining people who share labor income risks.

REFERENCES

Bils, M. J., Real Wages over the Business Cycle: Evidence from Panel Data, *Journal of Political Economy*, 93, 666–89, 1985.

Fischer, S., Wage Indexation and Macroeconomic Stability, in S. Fischer (ed.), *Indexing, Inflation, and Economic Policy*, MIT Press, Cambridge MA, 1986.

Gay, R. S., Union Settlements and Aggregate Wage Behavior in the 1980s, *Daily Labor Report*, No. 235, Page D-1, December 6, 1984.

Gray, J. A., Wage Indexation: A Macroeconomic Approach, *Journal of Monetary Economics*, 2, 221–35, 1976.

Hartigan, J., *Clustering Algorithms*, John Wiley & Sons, New York, 1975.

Haveman, R. and L. Buron, Escaping Poverty through Work: The Problem of Low Earnings Capacity in the United States, 1973–88, *Review of Income and Wealth*, 39(2), 141–57, June 1993.
Katz, L. F. and K. M. Murphy, Changes in Relative Wages, 1963–87: Supply and Demand Factors, *Quarterly Journal of Economics*, 107(1), 35–78, 1992.
O'Conor, K., Measuring the Precision of the Employment Cost Index, *Monthly Labor Review*, 112, No. 3, 29–36, March 1989.
Reich, R., *The Work of Nations*, Random House, New York, 1992.
Schneider, R., New Risk Management Concepts in Derivative Markets: Standardized Assets and Hedging, unpublished Ph.D. dissertation, Yale University, 1995.
Shiller, R., Measuring Asset Values for Cash Settlement in Derivative Markets: Hedonic Repeated Measures Indices and Perpetual Futures, *Journal of Finance*, 48, 911–31, 1993a.
———, *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*, Oxford University Press, Oxford, 1993b.
——— and S. Athanasoulis, World Income Components: Measuring and Exploiting International Risk Sharing Opportunities, National Bureau of Economic Research Working Paper No. 5095, Cambridge MA, April 1995.
——— and R. Schneider, Labor Income Indices Designed for Use in Contracts Promoting Income Risk Management, National Bureau of Economic Research Working Paper No. 5254, 1995.
Solon, G., R. Barsky, and J. Parker, "Measuring the Cyclicality of Real Wages: How Important is Composition Bias? *Quarterly Journal of Economics*, 109, 1–25, 1994.
Wood, D., Estimation Procedures for the Employment Cost Index, *Monthly Labor Review*, 105(5), 40–2, May 1982.