

Labor Income Indices Designed for Use in Contracts Promoting Income Risk Management*

Robert J. Shiller
Yale University

Ryan Schneider
McKinsey & Co.

Abstract

Labor income indices are created for groupings of individuals, using data from the Panel Study of Income Dynamics. 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 relatively few people move between them. For each of the groupings, we generate a labor income index using a hedonic repeated-measures regression methodology. Similarities between pairs of indices and between indices and individual labor incomes are described. It is argued that indices like those presented here might someday be used in settlement formulae in contracts promoting income risk management.

Introduction

Labor income indices could have important use in many contracts that have income risk management as part or all of their goals. They could be used to define future payments in contracts between management and labor, between suppliers of goods and their purchasers, and between insurance companies and their clients. Moreover, we could have international markets for contracts that represent claims on labor income flows, contracts traded on futures, options and securities exchanges, that would allow direct hedging of risks to standards of living.

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 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.

In this paper we create indices of individual labor income for use someday in such 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 on income that individuals have. 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. This is the most comprehensive U.S. panel income data set available, with over 20 years of data on income, labor market and personal characteristics necessary to create a consistent index of personal income and control for the changing characteristics of the population. Using this data set we can control for population growth and changes in the ‘quality’ of the labor force in the economy by explicitly including personal characteristics in our analysis.

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. The repeated measures approach allows us to make maximum use of our sample, as entry/exit from the sample does not disqualify observations. Our hedonic repeated measures methods are somewhat analogous to methods already used to take account of

composition bias to investigate the relation of the real wage to the business cycle, see Solon–Barsky and Parker [1994], and Bils [1985], but our design allows us to make use of a larger part of the sample and to control better for characteristics of individuals.

Creating these indices involves substantial judgment, in terms of the choice of individual characteristics to use as hedonic variables in the analysis. To some extent, our indices are perhaps less sensitive to choice of hedonic variables than are hedonic-regression-based indices that do not adhere to the repeated measures format, since changes in our indices are based only on changes in incomes of individual people. Still, the hedonic variables have a substantial impact on the indices. Because hedonic variables are important, different researchers tackling the same problem might arrive at rather different indices. Indeed, this sensitivity of hedonic indices to the choice of hedonic variables has been considered a drawback by some index number constructors, who seek a more “objective” way of constructing indices. But it is not a drawback per se that indices depend on hedonic variables; the dependency only underscores the importance of doing our job carefully. Our labor income indices are not intended to capture some already-popular commonly-understood definition of aggregate income, or to be used without careful attention to their definition. Our indices exist as prototypes of indices whose sole purpose would be to define settlements in risk management contracts, and their ability to function successfully in these contracts is the sole criterion for their success. Because of the difficulties in deciding on a specification of a labor income index, we present ours here as a sort of first step in a discussion, recognizing that there are likely to be many avenues for future improvement of the indices.

We consider in the next two sections some risk management functions that our labor income indices might fulfil. Although we are not entirely sure of the form risk management might take, these functions are fundamental to our analysis, since they are the motivation for the grouping methods and the index number construction methods that are described next. We will then present indices by grouping for the US 1968–1987 and conclude with some observations on the behavior of these indices and their closeness to individual income movements.

Risk Management: Indexation of Labor Contracts

An important use of labor income indices would be to index wages and salaries in labor contracts. Such indexing might be used in contracts offered 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. 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 adjusted by a formula related to the change in the labor income index for the 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 lay off or create incentives for current employees to leave (if the alternative income is lower). If these potential discrepancies could be eliminated by indexing contracts to labor income indices, then contracts could be longer-term. Then, less time could be spent on renegotiation of contracts and, moreover, the contracts might fulfill more of a risk-management function.

Existing contracts that require the firm to pay employees a constant real wage or income for a long time to its employees may be suboptimal, as pointed out by Gray [1976] and Fischer [1985]: the contract allows no adjustment to market conditions. In fact it is generally impossible for all labor contracts to guarantee real income very far into the future, since the aggregate resources to provide such income may not be available to contracting firms. 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]. Similar experiences leading to deindexation occurred in Italy, Belgium and other countries around the same time, see Williamson [1985]. 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. Much of the argument between labor and management is couched in terms of issues of fairness and comparisons of wage rates with those of others, see for example Bewley [1995]. 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.

Employees may be concerned that the 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 may be a temporary drop in real wages until others in the labor market adjust to the new inflation.¹ This consideration might suggest that employees would want to base contracts on indices of labor incomes of people who are currently renegotiating their wages. We find, however, that such indices are conceptually difficult to construct, since 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. Alternatively, those who are concerned about sluggishness of the index in representing current market conditions might add, to contracts indexed in terms of indices of labor income, some clause that attempts to correct for sluggishness. Such a clause might specify that if CPI inflation picks up suddenly above a specified threshold level, then employees will be allowed in that year to borrow according to a formula from their employer

¹The hypothesis that wages may lag behind prices, the “wage lag hypothesis,” was the subject of much discussion in the early to mid part of this century. This early discussion became mired in confusions about inferring causality and identifying directions of lags, see Kessel and Alchian [1960] and Cargill [1969].

against their future income. But, such provisions may well be considered unnecessary. *Ex ante*, sluggishness of labor income indices itself does not bias a contract in favor of either management or labor, since surprises, such as CPI surprises, can be in either direction.

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 this process. Even if one side of the bargaining, the labor side, would prefer to tie future wage increases to the CPI, labor may accept indexation to labor income indices if the management prefers the latter. 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.

Risk Management: Other 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. The firm's profits are determined by the spread between the product price and its costs, and it may be useful for that firm to share the risks of such spread with the other side of the contract; this is especially true if the other side consists of many different firms. The contract probably would not be written in terms of the actual cost of the supplying firm's own employees, since this would create a moral hazard problem; the firm would lose the incentive to contain employment costs. The use of such labor income indices would be especially useful in hedging for the supplier firm if it had labor contracts tied to such indices.

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. Covering the latter risks may add significantly to the cost of the insurance. 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; moral hazard problems in turn create costs to providing indexed disability insurance. Similarly, life insurance policies could promise a payout as a proportion of a labor income index.

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. Such an arrangement, in which contract wages are set in terms of an index of labor income of people in the same labor category as the employee and a subsequent hedge is made of this labor income in liquid hedging markets is a better way to manage risks than for labor to try to extract some promise from the employer to guarantee real incomes. We have argued above, such promises are not likely to be compatible with a free full-employment labor market equilibrium.

We cannot be sure just which uses of indices will ultimately be appealing; we do not have enough information about the situations facing contract designers to be able to extract the most suitable application. It should suffice to say, however, that indices ought to be produced and made available, so that people for whom they are useful can discover them. Innovation with new contract forms is a fairly rare event, and so it is important for us to prepare the way for such events to happen. Most people will not even consider using indices in their contract definition unless they have some indication that good such indices exist: the creation of the indices is the natural first step.

Existing Indices

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 the economists we contacted at the Bureau of Labor Statistics knew, perhaps because of problems in its construction or lack of specificity, not enough types of indices being produced.² 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.

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. We use job transitions in this way since we hypothesize that individuals who can readily exchange jobs will face similar income risks. The PSID reports 48 intertemporally consistent occupation-

²We were told that the employment cost index was not designed to settle contracts. A search of recent newspaper stories in NEXIS produced no evidence of contracts settled in terms of the employment cost index.

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 to category j the next time we observe this individual (almost always the next year); we interpret the elements as estimated probabilities.³ Clearly, the elements in each row of the transition matrix sum to one, by construction.

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]. This algorithm makes no use of the order in which the data are entered, it is not a spatial method; indeed our occupations and industries have no logical spatial positioning.

There are very many ways to define clusters; note that with our 48 occupation-industry pairs and just two clusters, there are over 10^{14} possible assignments. We decided on a simple clustering method that has some intuitive sense to it.

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. 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 $p(i, \{c\}) = \sum_{j \in c} p(i, j)$, $c = 1, \dots, C$, above a threshold $\tau(v)$, $v = 1, \dots, V$. 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, $v = 1, \dots, V$, through the data occurs. The process continues until all elements have been assigned.⁴

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

³For 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.

⁴There 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.

stable occupation-industry pairs, not one-way transitions. The additional members of the clusters are assigned only by reference to transitions into the cluster since we want clusters that are attractors, not repellers. 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.⁵

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 nonzero 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. The education level might be a better indicator of the kinds of jobs an individual is suited for than any clustering of occupation-industry categories. 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 risk management 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

⁵With $\tau(0) = 0.050$ there are only four groups, but these four are in effect four of the seven initial clusters, with the other three clusters now subsets of the four chosen with this higher threshold value.

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] 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 occupation-industry categories to his three skill categories. This assignment is presented in Table 5, and it is important to note that 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. The percentages of individuals from our sample we placed in each grouping correspond ex post roughly (within 5%) to the percentages Reich posits for the U.S. economy. 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

Our goal is to create indices whose changes through time reflect changes in labor income that many people in a grouping share, an income that is relevant to the earning power of the representative person in that grouping, and so that changes in the representation in the sample of people used to construct the index does not cause spurious changes in the index through time. We select as hedonic variables indicators of the kind of claims on labor income that people have. 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 proves 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 associated with the very income risks 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. Substantial judgment is necessary in choosing hedonic variables for our purposes.

To appreciate the importance of hedonic variables, consider our manhours variable. Over our sample period there has been an increase in female labor force participation rates, 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 increasing participation rate would engender settlements in risk management contracts even if no person saw a change in wages or income opportunities.

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, $t = 1, \dots, T$, 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 an 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):

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

The \mathbf{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

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

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.

An alternative setup for the x matrix might be attractive if we wanted to assume that switches between grouping 1 and grouping 2 are windfalls to people in grouping 1 (we are supposing that the second grouping in this example is the higher wage grouping), and represent real increases in their standards of living. This setup might be suggested by the view of interindustry wage differentials put forth by Krueger and Summers [1988], and Topel and Ward [1993], who argued that job-to-job transitions are important sources of wage growth.

On balance, we have decided that it is best to work on the assumption that when people move between groupings, the increase in pay is not a windfall but a compensation for different working conditions. Some people oscillate between groupings, suggesting that the move was not a windfall to them; if we used a formulation to estimate income indices that treated the wage increase as an increase in the grouping corresponding to the first income observation, then we might tend to see a steady rise through time in our index number for grouping 1 and decline for grouping 2, even if there were no changes in incomes in either grouping.

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. We have already argued that manhours may be particularly important: suppose that over this time period there has been an increase in female participation rates; without taking account of these increases, there would tend to be an upward bias in our index as a measure of the income opportunity risks we want to hedge.

Adapting the \mathbf{x} matrix to account for manhours alone, changing H to one from zero, we have equation (2):

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

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 manhours. In such a formulation, if the hedonic variable (or, more generally, any hedonic variable if there are more than one) 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 nonstochastic 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 \mathbf{x} corresponding to this hedonic variable (corresponding to the last four columns in equation (2)).

We assume that the variance matrix $\mathbf{\omega}$ of the transformed errors is diagonal with variances along the diagonal proportional to the interval between measures. This means that we are assuming that individual log income deviations from the log

income predicted by the regression is a random walk.⁷ The GLS estimate is:

$$\mathbf{B} = (\mathbf{x}'\boldsymbol{\omega}^{-1}\mathbf{x})^{-1}\mathbf{x}'\boldsymbol{\omega}^{-1}\mathbf{y} \quad (3)$$

Using the coefficients defined by (3), defining from \mathbf{B} for grouping g the $(1 + H) \times 1$ vector $\mathbf{B}_{g,t}$ of regression coefficients corresponding to time t , $t = 0, \dots, T-1$, ($\mathbf{B}_{g,t}$ 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 \mathbf{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 $I_{g,t}$, $t = 0, \dots, T-1$, $g = 1, \dots, G$:

$$I_{g,t} = I_{g,t-1} + \mathbf{x}_{t-1}(\mathbf{B}_{g,t} - \mathbf{B}_{g,t-1}) \quad (4)$$

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 (\mathbf{x}_{t-1}) multiplied by $(\mathbf{B}_{g,t} - \mathbf{B}_{g,t-1})$. We defined \mathbf{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 education, race and sex. Of course specialized indices that used other values for some hedonic variables, e.g., an index of for low-education people in a job grouping, might also be useful for contract settlement.

Note that if all repeated-measure pairs were exactly one period apart (there are no people who left and then reentered the sample), if we are estimating only a single grouping, if all hedonic variables are constant through time for each individual, and if we defined \mathbf{x}_{t-1} in (4) by taking the average of the rows of \mathbf{x} that correspond to

⁷Recent work by Gottschalk and Moffitt [1994] stresses the increasing importance of transitory, rather than permanent, earnings changes. Taking account of the transitory changes might seem to suggest some negative correlation between successive errors corresponding to the same individuals. However, much of the transitory component of earnings changes is captured by the independent variables in our regression, and so we would expect to see less such correlation in our error terms than are in the earnings figures themselves.

people who were in the sample both at $t-1$ and t and taking the columns of this average corresponding to time $t-1$, then our chain index would reduce to an index whose changes from $t-1$ to t are nothing more than the average change in log labor income for all people observed both at time $t-1$ and t . That this is so can be seen by rearranging the \mathbf{x} matrix, as exemplified by equation (2) with the two groupings merged into one and with the first $h(i, j)$ column deleted (to prevent multicollinearity when the hedonic variables are constant through time for each individual), so that the regression would estimate first differences of the coefficients through time and so that it is block diagonal. We have reordered both rows and columns, the rows to keep observations of each time period rather than each individual together, and the columns for the same purpose; we are here restoring the order we saw in equation (2), although now the dependent variable is time differences in log labor income.

$$\mathbf{x} = \begin{bmatrix} 1 & h(1,1) & 0 & 0 & 0 & 0 \\ 1 & h(3,1) & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & h(1,2) & 0 & 0 \\ 0 & 0 & 1 & h(3,2) & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & h(1,3) \\ 0 & 0 & 0 & 0 & 1 & h(2,3) \\ 0 & 0 & 0 & 0 & 1 & h(3,3) \end{bmatrix} \quad \mathbf{y} = \begin{bmatrix} y(1,1) & -y(1,0) \\ y(3,1) & -y(3,0) \\ y(1,2) & -y(1,1) \\ y(3,2) & -y(3,1) \\ y(1,3) & -y(1,2) \\ y(2,3) & -y(2,2) \\ y(3,3) & -y(3,2) \end{bmatrix} \quad (5)$$

Then, since \mathbf{x} is now block diagonal, our regression procedure could be simplified so that it consists of running a separate regression for each time period, each regression estimating the change in the coefficients between the two periods, that is, the t th regression estimating $\mathbf{B}_{g,t} - \mathbf{B}_{g,t-1}$. Since a regression plane with a constant term always passes through sample means, and since there is effectively a constant term in this regression, if we took \mathbf{x}_{t-1} in (4) to be the average of the rows in the t th bloc, then our fitted value would be nothing other than the average growth rates of labor incomes of all individuals in the sample in both $t-1$ and t .

That the index would reduce in this special case to such a simple average of log income changes for individuals does not mean that it would be a bad index; it would still be a repeated measures index, and would still be an improvement over simple averages of incomes. This special case, moreover, does not apply to our

regressions. Notably, our manhours variable (log of number of work hours) is not constant through time; other variables also change through time. We are doing something important here in controlling for manhours, making our index conform to full employment as defined by our current sample. Once we have broken block diagonality for one variable, then the coefficients of all variables are affected. Block diagonality is destroyed also by our practice of estimating coefficients for a number of groupings at the same time.

The hedonic repeated regression index number construction technique that we have described, and illustrated with equation (2), lets us avoid shrinking panel data due to attrition when compared with other methods, such as that used by Solon, Barsky and Parker [1994]. For example, the PSID has cumulative attrition of 44% in the past 20 years (though only 2–4% per year). If we restricted our sample to only individuals who were in the sample each year, we would have only 335 individuals who could be used for analysis. Instead, since the PSID has grown to over 35,000 individuals measured, the repeated measures technique gives us a sample size of 72,876 complete observations, leading to more efficient estimators.

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

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

The education variable was excluded in the regressions for the education categories. Note that the White/Nonwhite, 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 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. There are thus 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/nonwhite, 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.

The work hours variable is included since, as discussed above, we want our labor income indices to represent earning power, and not be influenced by aggregate changes in the tastes for hours worked. We could have constrained the log hours worked variable to have a coefficient of one; we could do this by making our dependent variable the log of income per manhour rather than income itself. But, we decided not to do this since we feel that there is a good chance that work done at reduced hours is for a lower wage (or implicit wage) than work done at full time. We felt that it is important that our index be a sort of best prediction about what people could earn if they chose to work full time.

Given that we have log hours worked as a hedonic variable, there is a question whether we should include as an additional hedonic variable the status as unemployed for part of the year. We decided to include this variable, for much the same reason that we did not constrain the coefficient of the log of hours worked to be one. When a person reports that he or she was unemployed for part of the year, there is the presumption that part of the labor income earned was under extraordinary circumstances, for which the compensation should be unusual. We of course recognize that there will be some question whether we want to leave the variable out; one might regard status as unemployed as one of the vicissitudes that we want to allow contracts to insure against.

Experience, years at the firm, 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 abundantly 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 for education brings in the risk that we are not allowing income to feed back into education. There is a question whether we want to include for observations at time t the education level at time t , or rather the education level at a given age, or when the individual enters our sample. If we did the latter, then we would have an education level that does not change for individuals through time; it would allow us to control for different kinds of claims (representing aptitude or motivation, say) at the beginning of the working years. But, should we control for time-varying education levels of individuals? If we do not do that, then our risk-management contracts will be controlling for the uncertainty about changes in education levels. We judge that unexpected changes in labor income due to unexpected changes in education levels are not something that people will want to ensure against. Education is bought at some expense, an expense that is not counted against income, and so increases (decreases) in education and corresponding increases in income may not be increases (decreases) in welfare.

Before concluding this section, let us describe some difficulties we have in choosing hedonic variables and interpreting the results. We wish to interpret our results as controlling for the effect on each hedonic variable on labor income, and yet we cannot be sure that the hedonic variables themselves are causally responsible for the observed effect on labor income. For example, we include education as a hedonic variable, however, we cannot be sure that we are really controlling for education and not something else that is correlated with education. It is conceivable that education levels are correlated with good work attitudes, and that it is these attitudes, rather than education levels, that influence labor income. If for some reason the correlation between education levels and work attitudes were different in the PSID sample than in the general population, then our change index would not be properly taking account of education levels.

It is important to realize what might have happened had we included as a hedonic variable any variable that is the same function of time for all individuals

up to an additive constant term that may differ across individuals. Age is such a variable; it is equal to time itself plus the individual's age in the base year, time 0. It would be useful in general to control for possible changes in the age mix of our sample, but doing so in our framework introduces a multicollinearity problem. The problem created by multicollinearity can be described as that we cannot learn how labor income would behave if we could somehow stop people from aging: since everyone ages at the same rate, the effect of a year of aging itself cannot be determined, it cannot be distinguished from any other time-varying effect on income. If we delete the age variable for $t = 0$, to destroy multicollinearity, then we still have $T - 1$ columns in our \mathbf{x} matrix corresponding to age variables, and we can estimate the effect of age at time t on the change in labor income from age 0 to t . But we could add any constant to all coefficients of age and subtract the same coefficient times time from each coefficient of a time dummy, without affecting fit; there is thus an arbitrary trend in the resulting index for constant age.

From our regression, we could define an index for age 25 constant through time, but this index would not tell us the effects of a fountain of youth. If we handle multicollinearity by deleting an age column, then the estimated index would effectively attribute all income differences across age at the time period for which an age variable column was deleted to differences in the individuals who happened to be in the age groups, rather than to age itself. The arbitrary trend in the index may not affect the use of the index for risk management purposes, but it may make the interpretation of the index less clear.

Alternatively, we could produce from the regression an index for a specific cohort, basing the index at time t on the fitted value at t of the regression not for a fixed age but for the age of the cohort at time t . With this alternative, the index is not contaminated by any arbitrary linear trend, because the values used to produce fitted values are subject to the same multicollinearity that was in the original data. So, in this case the multicollinearity yields no ambiguity to the resulting indices.

There are no such issues with our estimates, since the only hedonic variables that are the same function of time for all individuals up to an additive constant that may differ across individuals are sex and race. In these cases the function is a constant. There are other variables related to age, experience and years on the job, which are not the same functions of time (up to an individual specific additive constant) for all individuals. Our choice to include these variables, of course, makes us rely on those transitions that make these functions of time different for different people; if there are not many of these, then we run into the same

ambiguities that we found with the age variable itself.

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. Because of the large sample size, the standard errors for these index values are generally low: typically about 1% or 2%, occasionally 3% or 4%.

The relation between these indices and the Employment Cost Index and Personal Income are described in Schneider [1995]. To check our data, the time path of a simple average of the incomes from the data set used to produce our indices was compared with the time path of Personal Income; there were no significant inconsistencies.

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 .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 with some college in with high-school graduates as we did; they found that the college graduates showed an improvement 1979–1987 in log average weekly wage relative to the log average weekly wage of people with 12 years schooling of .117, and relative to the average weekly wage of people with 13–15 years of schooling of .062.⁸ 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 .013 between 1971 and 1979. Katz and Murphy found that the log average weekly wage for

⁸These numbers were inferred from data in Katz and Murphy [1992], page 40, Table I, Column 3, rows 5, 6, and 7.

college graduates fell relative to the log average weekly wage for those who had 12 years of schooling fell by .115 from 1971 to 1979; and fell relative to the log average weekly wage for those who had 13–15 years schooling by .067.⁹ For settling risk management contracts between established and continuing workers and their employers, our indices, based on changes in income for individuals and controlling for a number of factors, would appear to be better than these simple averages of wages for the purpose of contract settlement; we would not like to see settlements generated in these risk management contracts that are due to the changing composition of the work force.

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, 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. There are seven hedonic variables that might change through time for each of the seven job cluster groupings and skill groupings, and six such hedonic variables for the education groupings, since education was omitted as a hedonic variable. We have computed some numbers, shown in Table 9, that are somewhat analogous to partial *R* squareds, that are computed for one-year and five-year changes in labor incomes. To do this, 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. Thus, we are including changes in labor income that occur at times of changes between groupings. For each grouping, we compute the regression residuals in the regression of the change in log labor income on the hedonic variables with the change in the index excluded from the list of independent variables (A), and the regression residuals with the change in the index (B) included. For the seven job cluster groupings and the skill groupings, there are eight independent variables in each regression when the index is included, seven when it is excluded (there is

⁹*Loc. cit.*, column 2, rows 5, 6 and 7.

always also a constant term); there is one fewer independent variable with the education 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 is explained by the indices, in each job cluster grouping and education grouping. The CPI explains only 20%, 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%, although now the CPI explains 41%. 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 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 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 much differently the different indices behave through time. The p values for an F test that pairs of the indices are the same are shown below the diagonal. These numbers show how statistically significant are the differences between the individual indices. These are p values from ordinary regression F tests, 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, so that there is substantial income risk that could be dealt with using contracts settled in terms of the indices. The results show, moreover, that the differences across indices are big 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.

Conclusion

The Table 9 results show that our indices can go a long way to hedge income movements of individuals, that much of individual income movement 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.

Our labor income indices for job cluster groupings and education level groupings illustrate what might be done someday to create indices for contract settlement. Of course, the indices cannot be used today for such purposes because our data source, the PSID, is woefully out of date; our data stop in 1987. Still, we have learned something from our efforts about the nature of labor income risks and their hedgability.

The task that we have set for ourselves, of defining indices that would be useful for risk management contracts, while extremely important, is not easy, and it is not easy to judge how well we have done. Still, we do not have to do a perfect job to make indices that are useful for contract settlement. The ultimate measure of the success of any effort to produce indices for contract settlement would have to be whether people someday find that they are useful in their contracts. That judgment, however, might not come for a long time. In the meantime, we will have to use various imperfect means to measure our success.

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.”

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%, 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 classify well according to the skill categories.

One piece of modest evidence of our success in grouping is that labor income indices in those groupings for which transition probabilities are smaller tend to show larger deviations from the other indices. Groupings A and G, which showed the fewest transitions to other groupings, showed the greatest deviations from the overall index.

Perhaps some of the groupings that we have developed on the basis of the transition matrix clustering method might be grouped together for contract settlement since their labor income indices appear to have moved substantially together in our results. Still, in the absence of many job transitions between the groupings, we cannot be sure that the indices will continue to move together in the future.

We considered changing our methods so that the first step in our procedure defines job clusters directly in terms of individual characteristics in such a way that all individuals in a cluster had similar income movements.¹⁰ But we judged that, with the PSID data set, there is more information about the occupation-industry category transition matrix than there is about the similarity of paths through time of labor incomes. Since we have only twenty years of data on labor income, and since the incomes show substantial low-frequency movement, we think the simpler

¹⁰Defining indices in terms of similarity of income movements is explored by Shiller and Athanasoulis [1995] in another context.

approach, defining groupings only on the occupation-industry transition matrix, may be reasonably close to the best we can do with present data.

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. We would like to see a lot more information collected on the personal characteristics of the individuals, so that grouping or clustering algorithms can be based on much more data. If and when a substantial history of such data becomes available, there will be much more research to do to define better indices for contract settlement.

Table 1
Occupations and Industries

The base codes that the PSID data use are:

Occupations	Industry
1) Professional/Technical	1) Natural Resources Production
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

Table 2
Results of Cluster Analysis

Cluster (Grouping) Names:		Occupation					
		1	2	3	4	5	6
A: Professional/Technical							
B: Production	1	A	E	B	B	G	G
C: Services	2	A	D	E	D	F	F
D: (Public/Private) Works I	3	A	B	B	B	B	G
E: (Public/Private) Works II	4	A	D	D	F	G	G
F: Trade/Labor	5	A	F	F	F	F	G
G: Agriculture Labor	6	A	C	C	C	C	G

Table 3
Transition Matrix for Groupings Assigned Using the Clustering Procedure

		To						Sample size (initial cluster)	
		A	B	C	D	E	F		G
From	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
	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

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
---------	---------	---------	---------	---------	---------	---------

Table 4
Transition Matrix for Education Categories (Groupings)

		To			Sample size
		No high school	High school	College	
From	No high school	0.9937	0.0063	0.00	16,027
	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)

		Occupation					
		1	2	3	4	5	6
Industry	1	S	I	R	R		
	2	S	S	I	R	R	R
	3	S	S	I	R	R	R
	4	S	S	I	R	R	R
	5	S	I	I	R	I	R
	6	S	S	I	R	S	R
	7	S	S	I	R	I	R
	8	S					

Table 6
Transition Matrix for Groupings Based on Skill Categories

				Sample size
	SAS	IPS	RPS	(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

Table 7
Chain Indices of Labor Income: 1968 = 100

Year	Job Cluster (grouping)						
	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

Table 8
Chain Indices of Labor Income by Educational Level and Skill Categories

Year	Educational level			Skill category		
	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

Table 9
Proportion of Unexplained Variances of Labor Income Changes
Explained by Indices

Grouping	1-year changes	5-year changes	Grouping	1-year changes	5-year changes
A	0.51	0.73	NHS	0.41	0.59
B	0.47	0.60	HS	0.39	0.51
C	0.42	0.54	COLL	0.50	0.70
D	0.39	0.50	S	0.27	0.48
E	0.44	0.52	I	0.26	0.46
F	0.43	0.53	R	0.21	0.43
G	0.49	0.74	CPI	0.20	0.41
No.obs.	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.

Table 10
Standard Deviations and F-tests for Differences Across Indices

Grouping	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
Grouping	NHS	HS	COLL	Grouping	S	I	R
NHS	6.64	6.51	10.85	S	8.13	15.50	10.47
HS	0.08	6.21	10.09	I	0.09	7.84	11.29
COLL	0.00	0.05	8.01	R	0.01	0.08	11.26

Note: Figures above the main diagonal (in %, 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 %) of the indices. Figures below the diagonal are p values for F tests of the null hypothesis that the two indices are the same.

Bibliography

- Bewley, Truman, 1995. "A Depressed Labor Market as Explained by Participants," unpublished manuscript, Yale University.
- Bils, Mark J., 1985. "Real Wages over the Business Cycle: Evidence from Panel Data," *Journal of Political Economy*, 93: 666–689.
- Cargill, Thomas F., 1969. "An Empirical Investigation of the Wage-Lag Hypothesis," *American Economic Review*, 59(5): 806–811, December.
- Fischer, Stanley, 1986. "Wage Indexation and Macroeconomic Stability," in Stanley Fischer, *Indexing, Inflation, and Economic Policy*. Cambridge, MA: MIT Press.
- Gay, Robert S., 1984. "Union Settlements and Aggregate Wage Behavior in the 1980s," *Daily Labor Report*, No. 235, Page D-1, December 6.
- Gottschalk, Peter, and Robert Moffitt, 1994. "The Growth of Earnings Instability in the US Labor Market," *Brookings Papers on Economic Activity*, 217–272.
- Gray, Jo Anna, 1976. "Wage Indexation: A Macroeconomic Approach," *Journal of Monetary Economics*, 2: 221–235.
- Hartigan, J. 1975. *Clustering Algorithms*. New York: John Wiley & Sons.
- Katz, Lawrence F. and Kevin M. Murphy, 1992. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107(1): 35–78.
- Kessel, Reuben and Armen Alchian, 1960. "The Meaning and Validity of the Inflation-Induced Lag of Wages Behind Prices," *American Economic Review*, 50: 43–66.
- Krueger, Alan B. and Lawrence H. Summers, 1988. "Efficiency Wages and the Interindustry Wage Structure," *Econometrica*, 56: 259–293, March.
- O’Conor, Karen, 1989. "Measuring the Precision of the Employment Cost Index," *Monthly Labor Review*, 112(3): 29–36.
- Reich, R., 1992. *The Work of Nations*. New York: Random House.
- Schneider, R., 1995. "New Risk Management Concepts in Derivative Markets: Standardized Assets and Hedging," unpublished Ph.D. dissertation, Yale University.
- Shiller, R., 1993a. "Measuring Asset Values for Cash Settlement in Derivative Markets: Hedonic Repeated Measures Indices and Perpetual Futures," *Journal of Finance*, 48(3): 911–931.
- _____, 1993b. *Macro Markets: Creating Institutions for Managing Society’s Largest Economic Risks*. Oxford: Oxford University Press.
- Shiller, Robert J. and Stefano Athanasoulis, 1995. "World Income Components: Measuring and Exploiting International Risk Sharing Opportunities," National Bureau of Economic Research Working Paper No. 5095, Cambridge, MA, April.
- Solon, Gary, Robert Barsky and Jonathan Parker, 1994. "Measuring the Cyclicalities of Real Wages: How Important is Composition Bias?" *Quarterly Journal of Economics* 109: 1–25.
- Topel, Robert and M. Ward, 1993. "Earnings and Labor Market Experience," *Quarterly Journal of Economics*, reproduced.
- Williamson, John, 1985. *Inflation and Indexation: Argentina, Brazil and Israel*, Institute for International Economics.
- Wood, D. 1982. "Estimation Procedures for the Employment Cost Index," *Monthly Labor Review*, 105: 40–42.