Relative Wage Movements and the Distribution of Consumption
Author(s): Orazio Attanasio and Steven J. Davis
Source: *The Journal of Political Economy*, Vol. 104, No. 6 (Dec., 1996), pp. 1227-1262
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/2138938
Accessed: 23/04/2009 04:09

Relative Wage Movements and the Distribution of Consumption

Orazio Attanasio

University College, London, Institute for Fiscal Studies, and National Bureau of Economic Research

Steven J. Davis

University of Chicago and National Bureau of Economic Research

We analyze how relative wage movements among birth cohorts and education groups affected the distribution of household consumption and economic welfare. Our empirical work draws on the best available cross-sectional data sets to construct synthetic panel data on U.S. consumption, labor supply, and wages during the 1980s. We find that low-frequency movements in the cohort-education structure of pretax hourly wages among men drove large changes in the distribution of household consumption. The results constitute a spectacular failure of between-group consumption insurance, a failure not explained by existing theories of informationally constrained optimal consumption behavior. A welfare analysis indicates that the cost of between-group consumption variability is larger than the cost of aggregate consumption variability by two orders of magnitude.

I. Introduction

The U.S. economy underwent pronounced, persistent movements in the structure of relative wages during the 1980s. This paper analyzes

For many helpful conversations and comments, we thank John Cochrane, Lars Hansen, Costas Meghir, Kevin M. Murphy, Rob Townsend, Guglielmo Weber, two anonymous referees, and workshop participants at Boston College, Chicago, the Federal Reserve Bank of Cleveland, Massachusetts Institute of Technology, the NBER, Rochester, Stanford, University of California, Los Angeles, and University of Southern California. Davis gratefully acknowledges research support by the National Science Foundation. This paper was completed while Attanasio was at the University of Bologna, Italy.

[*Journal of Political Economy*, 1996, vol. 104, no. 6]

how relative hourly wage movements across birth cohorts and education groups affected the distribution of household consumption and economic welfare. The analysis partially integrates two distinct literatures and contributes to each.

One literature, large and very active, seeks to describe and explain time variation in the wage structure, but much of the motivation clearly derives from the perceived welfare consequences of changes in relative wages and overall earnings inequality.[1] The prevailing belief, articulated by Card (1991), seems to be that measured changes in the structure of income and labor earnings closely parallel changes in the distribution of household welfare. Few papers in this literature even question that belief, much less subject it to serious evaluation (the paper by Cutler and Katz [1991] is a notable exception).[2]

In sharp contrast, research on risk sharing in consumption stresses the variety of explicit and implicit insurance mechanisms that insulate the distribution of consumption from shocks to the distribution of earnings capacity and income. The empirical branch of this second literature investigates whether observed consumption outcomes conform to the implications of risk-sharing models (e.g., Cochrane 1991; Mace 1991; Altonji, Hayashi, and Kotlikoff 1992; Townsend 1994). This literature exploits a simple and rigorous theory of consumption allocations that, under plausible restrictions on preferences, delivers strong implications for the cross-sectional distribution of consumption growth. While this literature focuses on hypotheses that many economists view as implausible a priori, the empirical specifications that emanate from the theory serve more broadly as tools for investigating and interpreting the extent, pattern, and nature of departures from consumption insurance.

Previous empirical research in the consumption insurance literature typically relates the variation in idiosyncratic aspects of individual earnings capacity and income to individual or household consumption behavior. In contrast, we examine the impact of systematic, publicly observable shifts in the hourly wage structure on the distribution of household consumption. A focus on this type of wage structure variation offers four advantages.

First, systematic relative wage movements across large groups of workers are uncorrelated with idiosyncratic components of individ-

[1] Goldin and Margo (1992) and Katz and Murphy (1992) provide two of the more comprehensive investigations of relative wage movements in the United States during recent decades. Levy and Murnane (1992) survey the extensive body of research on recent U.S. wage structure developments. Davis (1992) examines patterns of relative wage movements across several countries.

[2] Some sociologists (Mayer and Jenks 1991) also question whether U.S. wage and income statistics accurately portray changes in the level and distribution of economic well-being.

ual-level preference shifts (e.g., changes in health status or household composition). Since these preference shifts plausibly affect both individual earnings capacity *and* household marginal utility, they potentially induce false rejections of consumption insurance hypotheses. Cochrane's (1991) analysis is especially clear on this point.

Second, a focus on publicly observable relative wage movements means that our evidence against consumption insurance cannot be rationalized by theories that stress the role of unobserved shocks in an informationally constrained optimal consumption allocation (e.g., Green 1987; Townsend 1988; Phelan and Townsend 1991; Atkeson and Lucas 1992). These theories offer no explanation why relative consumption growth rates depend on publicly observed shocks. Indeed, under plausible conditions that we spell out below, these theories predict no effect of publicly observed shocks on relative consumption growth rates. We return to this matter in Section VII, where we suggest how these theories can be modified to deliver informationally constrained optimal consumption allocations more consistent with our empirical findings.

Third, a focus on relative movements across observationally distinct groups facilitates the use of cross-sectional data sets that offer comprehensive, high-quality information on either consumption or earnings, but not on both. In particular, we show how to implement tests of the sort advocated by Cochrane (1991) using synthetic panel data rather than longitudinal data on individuals. Given the extreme scarcity of longitudinal data sources with high-quality information on both earnings and consumption, our ability to draw on cross-sectional data sets is an attractive feature of our empirical strategy.

Fourth, the focus on relative wage movements across groups of households enables us to devise suitable tests of consumption insurance in the face of nonseparable preferences between consumption and leisure *and* imperfect transferability of leisure across households. Given nonseparable preferences, previous consumption insurance tests rest on the maintained assumption that the social planner can freely transfer leisure across households (Cochrane 1991; Townsend 1994). We explain why this maintained assumption is unattractive and how its violation leads to false rejections of the consumption insurance hypothesis. We circumvent the need to maintain this problematic assumption by examining the response of relative household consumption to relative wage movements among groups of men with inelastic labor supply, while controlling for women's labor supply.

Section II develops our basic approach to analyzing the nexus between systematic relative wage movements and the distribution of consumption. We review the theory of consumption insurance and use it to derive regression specifications suitable for synthetic panel data. We also consider several specification issues that pertain to life cycle and demographic factors, preference nonseparabilities, and uninsurable idiosyncratic risk. We conclude the section with a discussion of econometric issues pertinent to our work.

Section III describes the data and the construction of the synthetic panels. Our synthetic panels contain annual observations from 1980 to 1990 on 5-year birth cohorts crossed by four educational attainment categories. We calculate cohort-education group means for household consumption data from the Consumer Expenditure Survey (CEX), and we calculate group means for individual wage and labor supply data from the March Annual Demographic Files of the Current Population Survey (CPS). The CEX and CPS represent the best large-scale, cross-sectional data sets on consumption and labor market outcomes, respectively, for the U.S. economy.

Section IV presents a descriptive analysis of relative wage and consumption movements across cohort-education groups during the 1980s. Pronounced movements in the structure of pretax hourly wages provide considerable leverage for gauging the impact of relative wage changes on the consumption distribution. Simple scatterplots reveal a spectacular failure of the consumption insurance hypothesis to account for relative wage and consumption comovements at low frequencies. The scatterplots also reveal that the rejection of the consumption insurance hypothesis arises in connection with relative wage movements across education groups and across birth cohorts within education groups.

Section V carries out formal tests of the consumption insurance hypothesis and characterizes the impact of systematic relative wage movements on the consumption distribution. The results provide strikingly sharp rejections of the consumption insurance hypothesis. Specifications that emphasize low-frequency comovements provide greater support for an extreme alternative to the consumption insurance hypothesis: relative consumption growth rates equal relative wage growth rates. For the decade as a whole, the results indicate that changes in the structure of pretax hourly wages among men were the dominant driving force behind the (large) changes in the distribution of household consumption. Higher-frequency comovements show much weaker evidence against the consumption insurance hypothesis and strong evidence against the extreme alternative of no consumption smoothing.

The magnitude of the departure from between-group consumption insurance is remarkable. Many mechanisms help insulate the consumption distribution from persistent shocks to the structure of pretax hourly wages among men: offsetting labor supply responses by men, their wives, or other household members; the progressivity of the income tax; the 1986 increase in the earned income tax credit; public welfare programs such as food stamps; and private transfers that cut across birth cohorts or education groups. At least for the U.S. experience during the 1980s, these smoothing mechanisms largely failed to insulate the consumption distribution from shifts in the pretax structure of wages.

Section VI calculates the welfare costs of these shifts in the household consumption distribution by comparing them to outcomes under full between-group insurance. For plausible degrees of risk aversion, our calculations indicate that large consumption variations are required to compensate households for the observed shifts in the consumption distribution across cohort-education groups. For example, with a relative risk aversion coefficient of two, compensating households for the type of between-group consumption risk they faced during the 1980s requires a uniform 2.7 percent consumption increase in all states and at all dates. This welfare cost of relative consumption variability among cohort-education groups is larger than the cost of aggregate consumption variability, as calculated by Lucas (1987), by two orders of magnitude.

II. Empirical Implications of Consumption Insurance

A. *Theory and Derivation of Synthetic Panel Specifications*

According to the consumption insurance hypothesis, explicit and implicit mechanisms for sharing consumption risks equalize the growth rate in the marginal utility of consumption across individuals and groups of individuals. This condition can be derived from the first-order condition for a central planner who, given a fixed set of Pareto weights, allocates resources under uncertainty across individuals and over time.

Let μ_t denote the Lagrange multiplier associated with the aggregate feasibility constraint at time t. Then the planner's first-order condition can be written as

$$(\rho^j)^t \lambda^j U_c(C_t^j, \delta_t^j) = \mu_t, \quad j = 1, \ldots, J, \tag{1}$$

for all states of the world at t, where C^j denotes consumption of individual j, δ^j denotes arbitrary preference shocks, ρ^j is a discount factor, and $U_c(\cdot, \cdot)$ denotes the marginal utility function. The time-invariant Pareto weights λ^j are equivalent to individual fixed effects, and they can be eliminated by considering the ratio of the first-order conditions at two points in time:

$$\rho^j \frac{U_c(C^j_{t+1}, \delta^j_{t+1})}{U_c(C^j_t, \delta^j_t)} = \frac{\mu_{t+1}}{\mu_t}, \quad j = 1, \ldots, J. \tag{2}$$

According to (2), any variable that is cross-sectionally uncorrelated with preference variation and measurement error in consumption growth is also uncorrelated with the cross-sectional distribution of consumption growth.

Assuming isoelastic utility with multiplicative shocks and a multiplicative error in measuring consumption implies a log-linear form for (2). In particular, let $U(C^j_t, \delta^j_t) \equiv U(C^j_t, b^j_t\gamma_j) = b^j_t(C^j_t)^{1+\gamma^j}/(1 + \gamma^j)$, and let ξ_{t+1} denote the measurement error in the log consumption change, so that (2) becomes

$$\log\left(\frac{C^j_{t+1}}{C^j_t}\right) = \frac{1}{\gamma^j}\left[\log\left(\frac{\mu_{t+1}}{\mu_t}\right) - \log\left(\frac{b^j_{t+1}}{b^j_t}\right) - \log(\rho^j)\right] + \xi^j_{t+1}. \tag{3}$$

Hence, if a variable X^j_{t+1} is cross-sectionally uncorrelated with preference variation ($\log[b^j_{t+1}/b^j_t]$, γ^j, and ρ^j) and measurement error (ξ^j_{t+1}), then consumption insurance implies that X has no explanatory power for the cross-sectional distribution of consumption growth. Obvious candidates for X that are likely to be important under interesting alternatives to the consumption insurance null hypothesis include shocks to the individual's endowment and earning capacity.

In his empirical implementation, Cochrane (1991) uses data from the Panel Study of Income Dynamics to examine the response of household food consumption to changes in the household head's health and employment status. He finds that long-term illness and job loss indicators imply rejection of the consumption insurance hypothesis under the maintained assumption that these variables are uncorrelated with preference variation and measurement error in the dependent variable. Townsend (1994) reports evidence against the consumption insurance hypothesis using household-level panel data for three Indian villages. Mace (1991) finds only mild evidence against the consumption insurance hypothesis using household-level data from the CEX, but Attanasio and Weber (1992) and Nelson (1994) show that careful treatment of measurement issues leads to sharp rejections of the consumption insurance hypothesis in Mace's data.

Whereas Cochrane, Mace, Townsend, and others use specifications suggested by (3) to examine the impact of individual-level endowment shocks on the distribution of household consumption growth, we focus on systematic relative wage movements across observationally distinct groups of households. The Introduction outlines some advantages of a focus on this type of wage structure variation. Here, we concentrate on deriving synthetic panel specifications that exploit this type of wage variation.

After partitioning households into groups indexed by i, take logs in (1) and average over the sample of group i households at time t to obtain

$$\hat{V}_{it} \equiv \frac{\sum_{j \in i(t)} \log[U_c(C_t^j, \delta_t^j)]}{\#i(t)} = \log \mu_t - t\frac{\sum_{j \in i(t)} \log \rho^j}{\#i(t)} - \frac{\sum_{j \in i(t)} \log \lambda^j}{\#i(t)}, \tag{4}$$

where $\#i(t)$ denotes the number of group i members who are sampled at time t. We can rewrite the sample average first-order conditions as

$$\hat{V}_{it} = \log \mu_t - t\overline{R}_i - \overline{L}_i + \epsilon_{it}, \quad i = 1, \ldots, I, t = 1, \ldots, T, \tag{5}$$

where $\overline{R}_i$ and $\overline{L}_i$ are population counterparts to the sample means in (4), and the error term ϵ_{it} arises because of finite sampling from heterogeneous populations.

Given a parameterization of preferences, (5) leads to a "levels" regression on a synthetic panel constructed from group-averaged data. In particular, when preferences take the isoelastic form specified above, the sample mean $\hat{V}_{it}$ becomes

$$\hat{V}_{it} = \frac{\sum_{j \in i(t)} \log(b_t^j)}{\#i(t)} + \frac{\sum_{j \in i(t)} \gamma^j \log(C_t^j)}{\#i(t)} \equiv \overline{B}_{it} + \frac{\sum_{j \in i(t)} \gamma^j \log(C_t^j)}{\#i(t)} + v_{it}.$$

Combining this equation with (5) yields

$$\frac{\sum_{j \in i(t)} \gamma^j \log(C_t^j)}{\#i(t)} = \log \mu_t - t\overline{R}_i - \overline{L}_i - \overline{B}_{it} + \epsilon_{it} - v_{it}, \tag{6}$$
$$i = 1, \ldots, I, t = 1, \ldots, T.$$

Now consider a regression of the sample mean of log consumption on a full set of time and group fixed effects, plus a variable X_{it} that captures time variation in relative group endowments:

$$\widehat{\log C_{it}} = \alpha_t + g_i + \beta X_{it} + e_{it}, \quad i = 1, \ldots, I, t = 1, \ldots, T. \tag{7}$$

Comparing (6) to (7), we see that consumption insurance implies $\beta = 0$, provided that X_{it} satisfies a list of auxiliary statistical assumptions. In particular, conditional on control variables in the regression equation, we require that X_{it} be uncorrelated with (i) measurement errors in log consumption, (ii) group differences in the mean of the time discount rate, (iii) group differences in the distribution of preference parameters γ^j, and (iv) variation in the mean preference disturbances $\overline{B}_{it}$. Our baseline empirical specifications take X_{it} to be the mean of the log of pretax hourly wages among men in group i at time t.

Proceeding in a similar manner, one finds that consumption insurance also implies $\beta = 0$ in first-difference regression specifications of the form

$$\begin{gathered}\widehat{\log C_{it}} - \widehat{\log C_{is}} = \alpha_t + \beta(X_{it} - X_{is}) + e_{it} - e_{is}, \\ i = 1, \ldots, I, t = (t - s), \ldots, T,\end{gathered} \tag{8}$$

provided that $X_{it} - X_{is}$ satisfies a list of statistical assumptions analogous to points i–iv.

B. *Specification Issues*

Consumption insurance implies that relative wage movements have no effect on the distribution of marginal utility growth. Testing this implication requires the maintained assumption that relative wage movements not be correlated with omitted factors that drive a wedge between marginal utility growth and the growth of measured consumption. It is therefore important to control for any determinants of group-level differences in marginal utility growth that are also correlated with group differences in relative wage growth.

Systematic life cycle variation in earnings capacity and household consumption requirements is one factor likely to lead to such a correlation in our synthetic panel data. Earnings capacity increases with age over much of the life cycle, but so do household-level consumption requirements because of increases in family size and age of children. For this reason, all our empirical specifications contain high-order polynomials in the household head's age, and many include additional controls for family size and composition. In terms of the formal theory, these regression controls reflect components of the individual and group preference shifters, b_t^j and $\overline{B}_{it}$.

Another potentially important factor leading to a correlation between measured consumption growth and relative wage movements involves preferences that are nonseparable between consumption and leisure. To see the point, suppose that wage movements reflect underlying disturbances to an aggregate production technology with diminishing marginal factor products. Assume that nonmarket time is *not* directly transferable across households and that preferences exhibit diminishing marginal utility of leisure. Then as long as time devoted to market activity increases the marginal utility of market-produced goods, groups with growing relative wages also experience relative consumption growth under a Pareto-optimal allocation.[3] It follows that, even under a Pareto-optimal allocation, β exceeds zero in specifications like (7) and (8) that fail to condition on labor supply. Consequently, the omission of labor supply (or leisure) controls can lead to false rejections of the consumption insurance hypothesis.

Since changes in health status, employment status, and income are also correlated with hours worked and expenditures on time-saving goods, the consumption insurance tests carried out by Cochrane, Mace, and Townsend suffer from the same potential specification problem and the same potential for bias against the null hypothesis. Hence, one can construe their evidence against consumption insurance as a failure of the maintained hypothesis that leisure is freely transferable through some extra-market institution. An analogous point applies to Abel and Kotlikoff's (1994) test of the intergenerational altruism hypothesis in cohort data.

We address the issue of consumption-leisure nonseparability in two respects. To treat nonseparability between consumption and women's leisure, we include the group mean of log female leisure hours. Such a regression specification can be derived from a nonseparable generalization of the isoelastic preferences we adopted above. In this case, the female leisure measure is appropriately intepreted as controlling for shifts in the marginal utility of household consumption.

Treating nonseparability between consumption and men's leisure in specifications that include men's wages is more difficult. The difficulty arises because an equation that includes consumption, male leisure, and male wages is observationally equivalent to the intratemporal first-order condition governing the consumption-leisure choice. Such an equation, therefore, cannot be used to test the consumption insurance hypothesis. To treat this problem, we follow two related identification strategies. One strategy simply restricts the sample to households headed by men with a low labor supply elasticity (i.e., well-educated, prime-age men). A second strategy posits nonsepara-

[3] See Sec. III*C* and app. A in Attanasio and Davis (1994) for a more extended discussion of this point.

ble preferences and then computes a range of marginal utility functions implied by a plausible set of values for male labor supply elasticities. This strategy substitutes alternative expressions for the marginal utility of consumption into the left side of regression specifications like (7) and (8).

Regarding another separability issue, our empirical analysis examines consumption of nondurable goods and services. If preferences are not separable between these and other components of consumption, then the left sides of (7) and (8) mismeasure the marginal utility of consumption. For two reasons, we do not think that this mismeasurement is a serious concern. First, the different groups of households in our synthetic panels are unlikely to experience very different patterns of relative price movements for, say, durable goods.[4] Second, even if the relative prices of durables vary systematically across groups in our synthetic panel, there is no apparent reason why these omitted price movements would be correlated with the relative prices of leisure (i.e., relative wages) that we use to test the consumption insurance hypothesis. In other words, the auxiliary statistical assumption iv specified above is unlikely to fail. If we were to condition on the consumption services derived from durable goods, we would needlessly dilute the power of our tests.

Finally, consider the implications of uninsurable idiosyncratic consumption risk. In deriving our regression specifications, we averaged the planner's first-order conditions (1) to obtain (4) and (5). In doing so, we derived an implication of consumption insurance across groups in a world without idiosyncratic consumption risk. As an alternative, consider the implications of cross-group consumption insurance in a model with idiosyncratic risk. In particular, suppose that informational or other constraints limit the planner to transfers that are contingent only on group membership.[5] In this case, optimal risk sharing among groups entails $(\overline{\rho}^i)^t\overline{\lambda}^i EU_c^i(C_t^j, \delta_t^j) = \mu_t$ for all i. In words, the planner equates the expected marginal utility of consumption across groups up to a proportionality factor that depends on the mean of the discount factors and Pareto weights.

By comparing the sample-average counterpart of this risk-sharing

[4] We have exploited the highly detailed information on household consumption in the CEX to construct group-specific price indexes over nondurables and services. These indexes revealed only trivial cross-group variation in inflation rates, and so we did not pursue the matter.

[5] For example, suppose that individual-level log wages contain three components: an exogenous random component common to the group, an exogenous random component specific to the individual, and a component that increases with the individual's effort. Assume also that the individual-specific components of the wage cannot be separately observed or inferred, and the effort is costly to supply. The resulting moral hazard problem in effort supply hampers insurance against within-group, but not between-group, wage risk.

condition to (5) or (6), one obtains an expression for the regression specification error introduced by within-group consumption risk. This error involves the difference between the mean of log marginal utility and the log of mean marginal utility. For example, with isoelastic preferences the specification error in the levels regression (7) becomes

$$\frac{\sum_{j \in i(t)} \log[b_t^j (C_t^j)^{\gamma j}]}{\# i(t)} - \log\left[\frac{\sum_{j \in i(t)} b_t^j (C_t^j)^{\gamma j}}{\# i(t)}\right]. \tag{9}$$

Thus, in the presence of uninsurable idiosyncratic consumption risk, our interpretation of the null hypothesis $\beta = 0$ as a test of cross-group consumption insurance requires that X_{it} be uncorrelated with (9) (conditional on other regression controls). As far as we can see, the theory of consumption insurance offers no reason for a systematic relationship between (9) and, say, relative wage movements of the sort that would violate this requirement.

We can also adopt regression specifications that derive directly from the model with uninsurable idiosyncratic consumption risk. When we work with the between-group risk-sharing condition that holds in the presence of uninsurable within-group risk, the theoretical levels regression corresponding to (6) becomes

$$\log\left[\frac{\sum_{j \in i(t)} b_t^j (C_t^j)^{\gamma j}}{\# i(t)}\right] = \log \mu_t - t \log \bar{\rho}^i - \log \bar{\lambda}^i + \text{sampling error}. \tag{6'}$$

This theoretical regression serves as the basis for alternative empirical specifications that differ from (7) and (8) only in the construction of the dependent variable.

C. *Econometric Issues*

Under the null hypothesis, the equation errors in (7) and (8) arise only from preference variation, measurement error, and (possibly) specification error. The expectational errors for uninsurable aggregate shocks are captured by time effects. To address the potential inconsistency of the ordinary least squares (OLS) estimator for β caused by sampling variation and other sources of measurement error in X_{it}, we use instrumental variables estimation.

To develop suitable instruments, first represent the regression equation involving the unobserved, true variables as

$$\Delta_k \tilde{C}_{it} = \text{controls} + \beta \Delta_k \tilde{W}_{it} + \epsilon_{it}, \tag{10}$$

where $\Delta_k X_{it} \equiv X_{it} - X_{i,t-k}$ for $k > 0$, with $\Delta_0 X_{it} \equiv X_{it}$, and $\tilde{C}$ and $\tilde{W}$ denote true values of log consumption and log wages, respectively; C and W represent the corresponding observed quantities.

Because we rely on different data sources to construct the C and W variables, sampling variation and measurement error in W are uncorrelated with the equation error in the regression model (10). Since, in addition, our regression specifications entail differenced quantities (or year and group fixed effects in the levels case), a reasonable error model for the log wage measure is

$$W_{it} = \tilde{W}_{it} + u_{it}, \tag{11}$$

where the error u_{it} satisfies $\text{cov}(u_{it}, \tilde{W}_{i,t-k}) = 0$ for all k, $\text{cov}(u_{it}, C_{i,t-k}) = 0$ for all k, $\text{cov}(u_{it}, u_{i,t-k}) = 0$ for $k \neq 0$, and $\text{cov}(u_{it}, u_{i',t-k}) = 0$ for $i \neq i'$. We select our instruments to perform well under these assumptions, while minimizing the loss of observations caused by instrumenting. Note that, given the measurement error structure postulated for (11), lags *and* leads of group means can be valid instruments for the contemporaneous wage measure in (10) under the null hypothesis.

In this setup, our levels specification fits the classical measurement error model. Thus if we estimate (10) with $k = 0$ by OLS, we obtain a possibly downward-biased estimator for β. To address this problem, consider the following instrument Z_{it} for W_{it}:

$$Z_{it} = \begin{cases} W_{i,t+1} & \text{if } W_{i,t-1} \text{ is unobserved} \\ \tfrac{1}{2}(W_{i,t+1} + W_{i,t-1}) & \text{if both } W_{i,t+1} \text{ and } W_{i,t-1} \text{ are observed} \\ W_{i,t-1} & \text{if } W_{i,t+1} \text{ is unobserved.} \end{cases} \tag{12}$$

There are three reasons why this scheme is appealing. First, no observations are lost because of instrumenting. Second, averaging over the immediate past and immediate future values of the wage measure reduces the noise component in the instrument. Third, since true relative wages change slowly over time, we anticipate a high correlation between the true current wage and the measured wages in adjacent years.

Since measurement error and sampling variation in the *levels* induce a moving average error structure in differenced data, certain lags (and leads) are no longer valid instruments. Hence, we adopt different instrumenting schemes for the difference specifications. For the 1-year difference specification ($k = 1$), we use an instrument that brackets the time interval of the true change. In particular, we instrument $\Delta_1 W_t$ by $Z_t = \Delta_3 W_{t+1}$. This instrument involves u_{t+1} and u_{t-2} but not u_t and u_{t-1}. Since our CPS-based wage measures extend further back in time than the CEX-based consumption measures, this instrument entails the loss of only one observation per group. For the multiyear difference specifications ($k \geq 2$), we instrument the time t wage change by $Z_t = \Delta_k W_{t-1}$. This scheme entails no loss of observations (except possibly for young cohorts) since our wage data extend further back in time than our consumption data.

With regard to standard errors, the regression residuals in (7) and (8) are likely to be characterized by heteroscedasticity and (for the difference specifications) by autocorrelation. Heteroscedasticity arises from variation across year-group cells in the extent of measurement error, the number of sampling units, and the degree of within-group heterogeneity. Differencing over k-year intervals induces a kth-order moving average term. We use a robust method (à la White) to form a consistent estimator for the covariance matrix in the levels specification. In the difference specifications, we use a more elaborate covariance matrix estimator that is consistent under heteroscedasticity and autocorrelation. Attanasio and Davis (1994, app. B) supply details.

III. The Data

A. *The Consumer Expenditure Survey and the Current Population Survey*

Our empirical analysis draws on two large-scale, public-use micro data sources. The CEX gathers information on income, demographic characteristics, and expenditure patterns of consumer units. A consumer unit is a group of individuals living in the same household who are related or share at least two of three major expense categories: food, housing, and other living expenses. Since 1980, the CEX has been carried out on a continuous basis with monthly rotation, surveying approximately 5,000 households per year. Barring attrition, each CEX household is surveyed for four consecutive quarters. The quarterly interview elicits information about expenditure patterns during each of the preceding three months. Information about income and labor market outcomes refers to the 12 months preceding the interview. Our investigation uses CEX data for calendar years 1980–90.

The Annual Demographic Supplement to the CPS gathers information on household income, demographic characteristics, and the labor market outcomes of individual members. The survey, carried out in March, elicits information about income and labor market outcomes for the preceding calendar year. The March CPS files contain this information for roughly 40,000–60,000 households per year. Our investigation uses CPS data for calendar years 1975–90.

The large CPS sample sizes constitute a major advantage to combining information from these two data sources. By constructing wage and leisure measures from the CPS rather than the CEX, we greatly reduce the sampling error component in our synthetic panel regressors. Aside from considerations of sample size, CPS income and earnings data are superior to the corresponding CEX data (Cutler and Katz 1991).

B. *Forming Synthetic Panel Groups*

Our analysis considers two main samples: one restricted to households with a male head (married couples, single males, and a few households with a single male parent) and another sample further restricted to married couples. In the CEX data, we form synthetic panel groups based on the birth year and educational attainment of the male household head. We follow the same practice in the CPS, except when constructing female wage and leisure measures for the less restrictive sample. For that case only, we define groups of women in terms of their own birth year and educational attainment.

Birth cohorts are defined in terms of 5-year bands. The oldest cohort contains persons born between March 1925 and February 1930; the cutoff month is chosen to maximize conformity between the CEX and CPS.[6] We consider four educational attainment categories: fewer than 12 years of schooling, exactly 12 years, more than 12 but fewer than 16 years, and 16 or more years of schooling. Our synthetic panel groups result by crossing these four education categories with the 5-year birth cohorts.

C. *Consumption and Wage Measures*

Our consumption measure equals household expenditures on nondurable goods and services. We exclude expenditure on durables, health, education, and housing. The main motivation for excluding these components is to avoid treating dynamics and other issues connected with durability. In addition, the CEX includes only out-of-pocket health expenditures, not insurance payments. We deflate

[6] The CPS records age at the March survey date. The CEX records age at each quarterly interview. Hence, increments in reported age between interviews enable us to bound birth dates within 3-month intervals in the CEX. Uncertainty about the exact date of birth in the CEX implies a theoretical rate of misallocation to 5-year birth cohorts of approximately 1 percent under our procedures. In practice, missing and erroneous age responses in the CEX generate a higher misallocation rate.

group consumption expenditures using group-specific price indexes that we constructed from the detailed expenditures data in the CEX. It turns out, however, that cross-group variation in inflation rates is tiny, so we effectively applied the same deflator to all groups.[7]

In constructing the consumption measure, we exclude nonurban households, those residing in student housing, those with a male head older than 59 or younger than 23, and those with incomplete income responses.[8] We exclude nonurban households because the CEX did not sample them in 1982 and 1983. By excluding young heads (and those residing in student housing), we minimize migration to higher educational attainment categories as a cohort ages. By excluding old heads, we minimize the impact of retirement and retirement choices on our sample.

Since CEX households are interviewed on a staggered basis, data for a typical household straddle two adjacent calendar years. Our annual consumption measures weight each household in proportion to the number of monthly observations that fall into the calendar year. For example, a household last interviewed in July of 1990 contributes six monthly expenditure observations to each of the 1989 and 1990 consumption measures.

We measure hourly earnings from the CPS, computed as annual earnings divided by the product of weeks worked and usual hours per week. We converted to real wages using the gross domestic product deflator for personal consumption expenditures. We excluded persons who were students or were in the military for at least part of the year, those who failed the age restriction described above, those who reported self-employment as their primary source of earnings, and those who earned less than 75 percent of the minimum wage. We imputed an estimate of the conditional mean earnings for top-coded individuals using the same procedure as Katz and Murphy (1992). In constructing labor supply measures, we include the self-employed.

Table 1 displays cell count summary statistics for the wage and consumption measures used in our study. The large cell counts for the CPS-based wage measures bear out one important advantage of our empirical strategy: the CPS enables us to construct wage regressors with much smaller sampling error than we could obtain from CEX-based wage measures.

[7] The standard deviation across groups of the 1980–90 log change in the price indexes is about 0.5 percent.

[8] It is standard practice among Bureau of Labor Statistics statisticians, when computing means, to exclude CEX observations with incomplete income responses on the basis of data quality.

TABLE 1

CELL COUNT SUMMARY STATISTICS FOR GROUP-LEVEL CPS AND CEX DATA

Variable Type	Source	Number of Cells	Mean Cell Count	Minimum	Maximum
Wages:					
Men	CPS	420	902	227	2,291
Women	CPS	420	762	192	2,056
Husbands	CPS	420	695	190	1,576
Wives	CPS	420	509	148	1,166
Consumption:					
All	CEX	288	499	137	1,071
Married	CEX	288	414	117	896

NOTE.—Groups are defined by crossing 5-year birth cohorts with four educational attainment categories. Each cell corresponds to one annual observation on a group. The number of cells equals the total number of annual group-level observations that are admissible under our sample selection criteria. For CPS data, an admissible cell is one in which all men are between 23 and 59 years of age. The CPS samples of wives are restricted to women with husbands between 23 and 59 years of age. The CEX samples are restricted to households with a male head or husband of female head between 23 and 59 years of age. Other selection criteria are defined in the text. The number of admissible cells and the cell count summary statistics for CPS (CEX) data pertain to the 1975–90 (1980–90) sample period. The cell count equals the number of nonmissing observations for the indicated variable type.

IV. Relative Wage and Consumption Behavior

A. *Movements in Relative Wages*

Table 2 summarizes movements in men's hourly wages by birth cohort and education group.[9] Each panel describes real wage movements for all cohorts in a particular education category relative to the 1980 value for the 1945–50 birth cohort in the same education category. Looking across a single row of this table traces out the cross-sectional age profile of wages for the indicated year and education category. Looking down a column traces out the evolution of real wages for a particular cohort-education group. For example, the 1950–55 cohort of men with fewer than 12 years of schooling suffered a real wage decline of 11 log points between 1980 and 1990. Comparisons across columns reveal differences in the evolution of real wages between cohorts and between cohort-education groups.

The most pronounced relative wage movements in table 2 involve differences by educational attainment. Among men with fewer than 12 years of schooling, real wages fell (by about 10 percent) between 1980 and 1990 for all birth cohorts. Real wages declined more modestly over the decade for cohorts of high school–educated men. Among men with some postsecondary education, real wages increased substantially for the youngest cohorts and modestly for older

[9] Here, and throughout the paper, we use only the observations for which all members of the group satisfy the age restrictions in the indicated year.

TABLE 2

Real Hourly Wage Movements for Men by Birth Cohort and Education Group: Log Deviations from the 1980 Value for Own Education Group and 1945–50 Cohort

	5-Year Birth Cohort								
Year	1960–65	1955–60	1950–55	1945–50	1940–45	1935–40	1930–35	1925–30	1920–25
	Less than High School Education (0–11 Years)								
1975				−.06	.06	.09	.15	.16	.14
1980			−.07	.00	.08	.10	.15	.14	
1985		−.18	−.06	−.06	.01	.06	.09		
1990	−.29	−.19	−.19	−.09	−.03	−.02			
	High School Education								
1975				−.06	.02	.04	.09	.10	.10
1980			−.12	.00	.04	.07	.13	.13	
1985		−.21	−.12	.00	.04	.08	.06		
1990	−.31	−.18	−.13	−.07	.02	−.03			
	Some Postsecondary Education (13–15 Years)								
1975				−.10	.05	.14	.13	.12	.19
1980			−.15	.00	.12	.09	.14	.14	
1985		−.17	−.06	.08	.09	.13	.12		
1990	−.26	−.09	−.02	.04	.08	.12			
	College Education (16+ Years)								
1975				−.15	.10	.22	.26	.28	.27
1980			−.18	.00	.17	.19	.23	.28	
1985		−.11	.03	.17	.25	.29	.30		
1990	−.16	.03	.14	.17	.24	.23			

cohorts. Finally, among college-educated men, real wages rose for all cohorts, including rapid real wage gains for the youngest cohorts. An example of an extreme comparison is that wages for the 1950–55 birth cohort of college-educated men rose by 44 log points relative to wages of their contemporaries with fewer than 12 years of schooling.

Between-cohort relative wage movements are modest for less educated men. However, among men with more than 12 years of schooling, younger cohorts experienced notably more rapid wage growth than older cohorts. Between-cohort variation in wage growth is especially pronounced among college-educated men.

Relative wage movements among cohort-education groups of women during the 1980s show patterns similar to the ones experienced by men. This similarity holds whether men and women are grouped according to their individual characteristics or as husband-wife pairs. See Attanasio and Davis (1994) for evidence.

In summary, wages for highly educated men rose sharply during the 1980s relative to wages of their less educated contemporaries. Between-cohort variation in wage growth was modest among less educated men and pronounced among the most educated men, with more rapid wage growth for younger cohorts. These large relative wage movements among cohort-education groups of men, and the reinforcing pattern of movements among women, indicate that our synthetic panel data offer ample leverage for testing consumption insurance hypotheses and for estimating the effects of systematic relative wage movements on the household consumption distribution.

B. *Movements in Relative Consumption*

Although the consumption data are notably noisier than the wage data, some key patterns emerge. Among the least educated, real consumption declines sharply over part or all of the decade for the 1945–50 and earlier cohorts. Among the high school–educated, real consumption also declines sharply for several cohorts. In contrast, among the more educated, real consumption shows modest to sharp gains for all but the oldest cohorts.

Table 3 reports average (across cohorts) real consumption growth during the 1980s by educational attainment status. The table reveals a sharp, systematic pattern of rapid relative consumption growth for more educated groups during both halves of the decade. The least educated experienced a striking 15 percent decline in real household consumption from 1980 to 1985 and essentially no change over the remainder of the decade. The college-educated experienced reasonably strong consumption growth over both halves of the decade. It is noteworthy that real household consumption falls absolutely for some

TABLE 3

CONSUMPTION GROWTH BY EDUCATIONAL ATTAINMENT OF MALE HEAD, 1980–90: SIMPLE AVERAGES OF CHANGES IN THE MEAN OF LOG(Consumption) FOR 5-YEAR COHORTS

	Educational Attainment of Male			
Time Interval	<12 Years	High School	Postsecondary	College
1980–85	−.15	−.06	−.05	.07
1985–90	.01	.03	.06	.10

NOTE.—The sample is restricted to cohorts for which the male head (or husband of female head) is between 23 and 59 years of age throughout the indicated time interval. The consumption measure equals household expenditures on nondurable goods and services, as defined in the text.

groups while rising absolutely for others. This pattern, whereby consumption moves in opposite directions for different groups, points toward the likely untenability of the consumption insurance hypothesis for any stable parameterization of preferences under which agents care about consumption risk.

C. *Comovements between Relative Wages and Consumption*

Figure 1 plots annual differences of mean log consumption against mean log wages for the cohort-education groups in our sample. The plotted values are residuals from regressions on a cubic polynomial in age and year fixed effects. The year effects control for the uninsurable aggregate shocks in our sample, and the age polynomial represents a crude control for systematic life cycle variation in household consumption requirements. The theory developed in Section II delivers this specification when preferences are isoelastic and separable between consumption and leisure (see eq. [8]).

The scatterplot in figure 1 conforms well to the implications of the consumption insurance hypothesis or any other theory that predicts smoothing of high-frequency earnings variation. The least-squares regression slope is mildly positive but insignificantly different from zero. As indicated by the low R^2 value, there is no apparent relationship between year-to-year relative wage and consumption movements among cohort-education groups.

Figure 2 shows an analogous scatterplot for 10-year first differences in our sample. This figure reveals a remarkably close relationship between low-frequency relative wage and consumption movements during the 1980s. The slope coefficient equals .92 with a standard error of .12. Cross-group differences in men's relative wage growth explain a remarkable 82 percent of the considerable variation

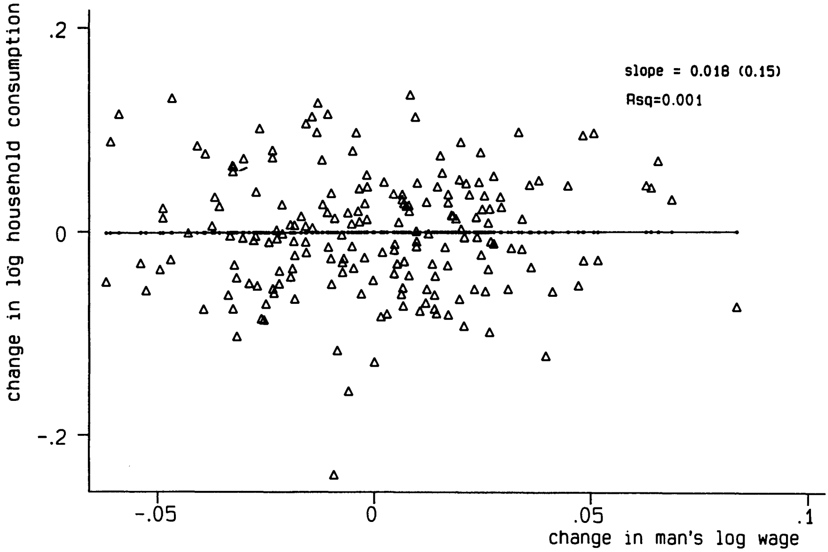

FIG. 1.—Household consumption vs. man's wage, annual log change residuals, 1981–90. Groups are defined by four-way education crossed with 5-year birth cohorts. Plotted values are residuals from regressions on year effects and a cubic in age.

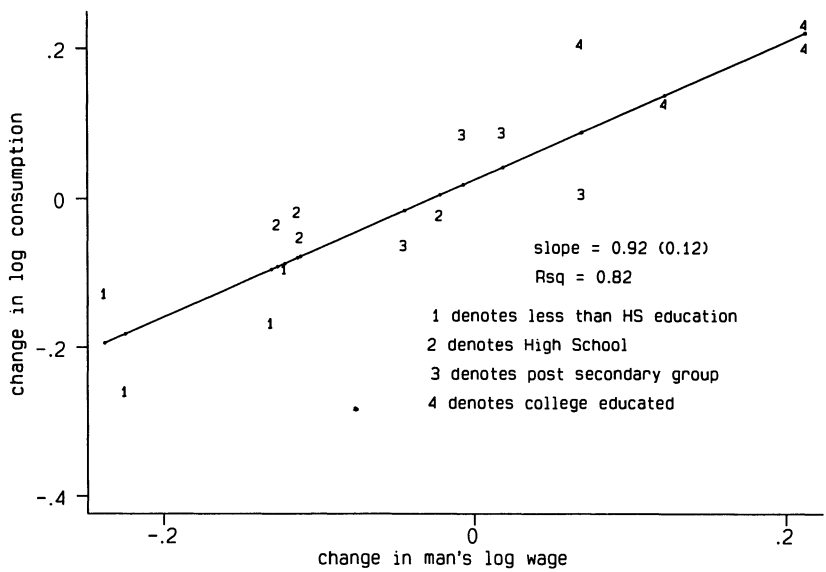

FIG. 2.—Household consumption vs. man's wage, 1980–90 log change residuals. Groups are defined by four-way education crossed with 5-year birth cohorts. Plotted values are residuals from regressions on a cubic in age.

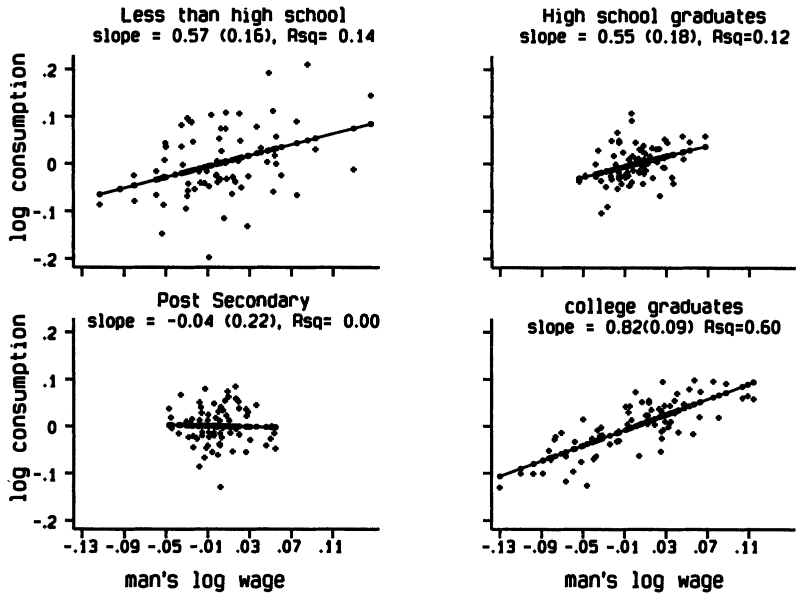

FIG. 3.—Household log consumption vs. man's log wage by education category, annual observations, 1980–90. Plotted values are residuals from regressions on a quartic in age plus year and group fixed effects.

in relative consumption growth. The labels on the individual points in the scatterplot indicate that the between-education components of relative wage and consumption movements drive the regression line.

This second scatterplot points to a spectacular failure of the consumption insurance hypothesis with respect to publicly observable, systematic components of relative wage variation. Indeed, the evidence is highly favorable to an extreme alternative hypothesis under which relative consumption growth equals relative wage growth. The sharp contrast between figures 1 and 2 also highlights another important advantage of our empirical strategy. By drawing on cross-sectional data sets to construct long synthetic panels, we can identify persistent components of earnings variation. Comparing figures 1 and 2 shows that the persistent components of relative wage movements drive relative consumption movements.

Figure 3 shows scatterplots that correspond to the levels specification (7). In this figure, we plot by education group the residuals from regressions on a quartic polynomial in age plus a maximally linearly independent set of year and group fixed effects.[10] We know from

[10] We constrain the age polynomial to be the same for all education groups, which reflects our belief in small exogenous differences across education groups in the shape of life cycle consumption requirements.

figure 2 that the covariation in relative wages and consumption between education groups generates a sharp rejection of the consumption insurance hypothesis. Figure 3 tells us that the between-cohort covariation within education groups also produces a sharp rejection of the hypothesis. Three of the four education groups generate large and significantly positive slope coefficients on relative wages. The large cross-cohort variation in wage growth among the college-educated generates an especially stark rejection of the consumption insurance hypothesis.

In summary, these scatterplots provide visually compelling evidence that the household consumption distribution is poorly insulated from persistent relative wage movements among cohort-education groups. The scatterplots are consistent with the view that the consumption distribution is well insulated from transitory relative wage movements, perhaps through the smoothing mechanisms envisioned by life cycle and permanent income theories. The more formal econometric investigation carried out below shows how these results are affected by considerations related to measurement error, the inclusion of additional controls, women's wages, nonseparable preferences, uninsurable idiosyncratic risk, and alternative samples.

V. Econometric Results

A. *The Synthetic Panel Specifications*

This section reports the results of estimating several versions of equations (7) and (8) on synthetic panel data. We typically begin with a benchmark specification that contains the mean log pretax hourly wage among men, year effects, a polynomial in age, plus group fixed effects for the levels specifications. We then consider alternative specifications that add wage and leisure measures for women (or wives) and controls for family size and composition.[11] We report OLS and instrumental variables results for a sample of all households with a 23–59-year-old male head and for a sample of married couples only. In addition to the levels specification (7), we estimated k-year difference specifications (8) for all k from one to nine. We report a selected set of the difference specifications to illustrate the general pattern of results.

Our family size and composition controls are (the mean of) log family size, the number of adults, the number of children under 3 years of age, and the number of other children. These variables are

[11] We experimented with several measures of female leisure without much effect on the results. The reported results use CPS-based means of the logarithm of annual leisure hours, defined as 52 times 126 minus annual hours worked.

intended to capture life cycle preference variation that differs over time and across groups. On a priori grounds, it is not clear that these controls belong in the specification. While we believe that consumption requirements vary exogenously over the life cycle, time and group *variation* in the shape of life cycle consumption requirements could be driven by variation in relative wages. Thus the inclusion of these controls stacks the deck in favor of the consumption insurance hypothesis. In any case, their inclusion represents an easy and flexible alternative to an adult equivalence scheme.

The number of available observations varies with the control set, the estimation method, and the differencing interval. To facilitate comparability, the reported results use the largest sample available across all specifications and control sets, given a particular differencing interval. This leads to a loss of observations for some specifications; results were not affected by these minor sample changes.

In carrying out instrumental variables estimation, we instrument the wage measures as described in Section II*C*. We instrument the leisure and demographic variables in an analogous manner. In the difference specifications, we also use the group fixed effects as instruments for the demographic variables.

B. *Results for the Level Specifications*

Panel A of table 4 contains results for the level specifications. The first two rows of table 4 show, for the benchmark specification, an OLS coefficient estimate on male wages of .65 in the male-headed sample and .42 in the married couples sample. As expected in the presence of measurement error, the corresponding instrumental variables estimates are considerably larger: .81 and .59.[12] These slope coefficients are precisely estimated and provide strong evidence against the hypothesis of full insurance against publicly observable relative wage movements. The estimates are closer to unity—indicating that relative wage changes translate one for one into relative consumption changes—than to zero.

The next two rows add female wage measures. As expected, the collinearity between male and female wage movements inflates the standard errors on the individual slope coefficients, especially under instrumental variables estimation. The OLS results reiterate the

[12] The estimates are virtually unaffected when we use only lagged wages to construct instruments, a result that may appear contrary to the permanent income hypothesis. However, in view of the panel length and the persistent nature of relative wage movements during the 1980s, this result is driven by the cross-sectional correlation between relative wages and consumption. The permanent income hypothesis places no restrictions on such cross-sectional correlations, as discussed in Deaton (1992, pp. 146–48).

TABLE 4

SYNTHETIC PANEL REGRESSIONS

		Households with a 23–59-Year-Old Male Head			Married Couples Only		
Estimation Method	Household Controls*	Man's Wage	Woman's Wage	Woman's Leisure	Husband's Wage	Wife's Wage	Wife's Leisure
		A. Levels Specification ($N = 288$)					
OLS	No	.653 (.072)	. . .	. . .	.419 (.063)	. . .	. . .
IV	No	.814 (.079)	. . .	. . .	.592 (.076)	. . .	. . .
OLS	No	.616 (.094)	.063 (.105)	. . .	.380 (.084)	.070 (.098)	. . .
IV	No	1.760 (.640)	−1.392 (.941)	. . .	7.807 (25.534)	−10.036 (35.248)	. . .
OLS	No	.712 (.077)	. . .	−.654 (.315)	.477 (.068)	. . .	−.633 (.295)
IV	No	1.068 (.111)	. . .	−2.162 (.450)	.885 (.107)	. . .	−2.321 (.433)
OLS	Yes	.476 (.070)	. . .	. . .	.462 (.061)	. . .	. . .
IV	Yes	.824 (.245)	. . .	. . .	1.272 (.665)	. . .	. . .
OLS	Yes	.500 (.070)	. . .	−.347 (.289)	.488 (.061)	. . .	−.330 (.292)
IV	Yes	.898 (.245)	. . .	−1.343 (.899)	1.250 (.665)	. . .	.914 (4.346)
		B. Annual Difference Specification ($N = 192$)					
OLS	No	.019 (.157)	. . .	. . .	−.057 (.124)	. . .	. . .
IV	No	.463 (.269)	. . .	. . .	.370 (.235)	. . .	. . .
OLS	No	−.006 (.161)	.093 (.135)	. . .	−.114 (.126)	.240 (.124)	. . .

IV	No	−.356	1.276	. . .	.807	−.772	. . .
		(1.882)	(2.774)		(1.541)	(2.578)	
OLS	No	−.003	. . .	.661	−.067	. . .	.568
		(.158)		(.567)	(.124)		(.520)
IV	No	.564	. . .	−1.539	.444	. . .	−1.229
		(.310)		(1.400)	(.278)		(1.075)
OLS	Yes	−.011	. . .	. . .	−.063	. . .	. . .
		(.134)			(.118)		
IV	Yes	.270	. . .	. . .	.571	. . .	. . .
		(.394)			(.271)		
OLS	Yes	−.028	. . .	.521	−.075	. . .	.610
		(.134)		(.483)	(.118)		(.490)
IV	Yes	.257	. . .	−.575	.447	. . .	−.910
		(.394)		(1.004)	(.284)		(.160)
		C. 8-Year Difference Specification (N = 36)					
OLS	No	.699	. . .	. . .	.541	. . .	. . .
		(.129)			(.144)		
IV	No	.701	. . .	. . .	.523	. . .	. . .
		(.133)			(.157)		
OLS	No	.686	.020	. . .	.428	.173	. . .
		(.286)	(.404)		(.307)	(.416)	
IV	No	1.434	−1.139	. . .	2.409	−2.763	. . .
		(2.426)	(3.796)		(7.153)	(10.560)	
OLS	No	.910	. . .	−1.991	.778	. . .	−2.131
		(.132)		(.625)	(.147)		(.668)
IV	No	.976	. . .	−2.444	.843	. . .	−2.507
		(.135)		(.668)	(.145)		(.682)
OLS	Yes	.869	. . .	. . .	.720	. . .	. . .
		(.173)			(.155)		
IV	Yes	.897	. . .	. . .	.690	. . .	. . .
		(.142)			(.131)		
OLS	Yes	.865	. . .	−1.385	.773	. . .	−1.407
		(.173)		(.648)	(.155)		(.726)
IV	Yes	.866	. . .	−1.283	.737	. . .	−1.324
		(.142)		(.473)	(.131)		(.564)

* The household controls are log of family size, the number of adults, the number of children under 3, and the number of other children.

benchmark results, but the individual instrumental variables coefficients are too imprecisely estimated to draw any inferences.

Rows 5 and 6 add controls for female leisure to the benchmark specification. The results indicate that increases in female leisure reduce household consumption expenditures, consistent with our prior views about the nature of the preference nonseparability, but this control does not mitigate the size or statistical significance of male wage variables.

Finally, the remaining rows add family size and composition variables, with and without the controls for female leisure. Once again, the coefficient on male wages is little affected, although the instrumental variables estimates for this coefficient are less precise when we include the composition variables.

C. *Results for the Difference Specifications*

Panels B and C of table 4 report results for the 1-year and 8-year differencing intervals. The results for 2-year differencing intervals, like the results for 1-year intervals, show no consistent effects of relative wage movements on relative consumption. For differencing intervals of 3 or more years, we find sizable and statistically significant departures from the consumption insurance hypothesis, as illustrated by panel C. Relative wage movements over these longer differencing intervals are associated with large relative consumption movements. These patterns in the estimation results emerged for all sets of controls we considered, with the exception of instrumental variables estimates of specifications that include female wages. This latter specification yielded imprecisely estimated slope coefficients on male and female wage measures. In short, results obtained for differencing intervals of 3 or more years closely parallel results for the level specifications.

D. *Can Nonseparability Salvage the Consumption Insurance Hypothesis?*

As stressed in Section II, one cannot test the hypothesis that men's relative wage movements have no effect on the distribution of marginal utility growth while simultaneously controlling for nonseparability between men's leisure and household consumption, unless one brings additional identifying information to bear. The tests in table 4 achieve identification by maintaining strong separability between consumption and men's leisure. We now investigate whether the previous evidence against the consumption insurance hypothesis hinges on this separability assumption.

We proceed by considering nonseparable preferences of the form

$$U(C_t^j, \bar{L} - H_t^j, \delta_t^j) = b_t^j (C_t^j)^{1+\gamma^j} (\bar{L} - H_t^j)^{1+\phi^j},$$

where $\bar{L}$ denotes the time endowment in hours, and H_t^j denotes hours worked.[13] We seek to assess whether a plausible degree of complementarity between consumption and men's labor supply can rationalize the covariance between men's relative wages and household consumption. "Plausible" means consistent with available evidence on male labor supply elasticities. In particular, we experimented with values of γ and ϕ that correspond to an intertemporal substitution elasticity between .3 and .6 and an uncompensated wage elasticity (evaluated at sample means) between $-.3$ and .3. Corresponding values for the compensated elasticity lie between .21 and .82.[14] These figures are consistent with available evidence (Pencavel 1986).

Given values for γ and ϕ, we used the implied marginal utility of consumption,

$$\log(C_t^j) + \frac{1 + \phi}{\gamma} \log(\bar{L} - H_t^j),$$

to construct the dependent variable in regression specifications otherwise identical to the ones reported in table 4, and we obtained highly similar results. We infer that our evidence against the consumption insurance hypothesis cannot be rationalized by a plausible degree of complementarity between consumption and men's labor supply.

As another check on the potential importance of nonseparable preferences, we considered samples restricted to men with relatively inelastic labor supply. To the extent that labor supply is inelastic,

[13] The parameter restrictions $b < 0$ and $\gamma, \phi < -1$ imply that nonmarket time and consumption expenditures are substitutes, as presumed in the text.

[14] Denote the labor supply elasticities with respect to wage movements by ξ for the uncompensated elasticity, η for the compensated elasticity, and ψ for the intertemporal elasticity. These elasticities satisfy

$$\xi = \frac{M_c H - 1}{M_L - wM_c}\left(\frac{w}{H}\right),$$

$$\eta = -(M_L - wM_c)^{-1}\left(\frac{w}{H}\right),$$

and

$$\psi = \frac{\gamma}{1 + \gamma + \phi},$$

where $M = U_L/U_c$, $M_c = \partial M/\partial c$, $M_L = \partial M/\partial L$, H is labor supply, L is leisure, c is consumption, and w is the wage. We evaluated ξ and η at $c = 23{,}000$, $w = 11$, $H = 2{,}080$, and $L = 5{,}000 - 2{,}080$. For instance, $\gamma = -3$ and $\phi = -4$ implies $\beta = -.037$, $\eta = .55$, and $\psi = .5$.

there is little scope for complementarity between work and consumption to drive variation in household consumption. A large empirical literature holds that prime-age men supply labor inelastically with respect to the wage. Hence, we considered samples restricted to men between 30 and 55. Since some evidence suggests that less educated men exhibit a greater supply elasticity, we also considered samples further restricted to exclude groups with fewer than 12 years of schooling. These exclusions limit the scope for work-consumption complementarities and sample selectivity (i.e., correlation between relative wage movements and incidence of household head status) to influence the results.

Table 5 reports results for the sample restricted to all men between 30 and 55 who have at least a high school education. (Results for the corresponding sample of married couples and the sample that includes the least educated men are highly similar.) Owing to smaller sample sizes, the coefficient estimates are slightly less precise than before, but the effects of male wages on consumption strongly confirm the previous results. If anything, the null hypothesis is even more consistently rejected.

E. *Can Idiosyncratic Consumption Risk Explain the Lack of Between-Group Insurance?*

Section II outlined a model of between-group consumption insurance in the presence of uninsurable idiosyncratic risk. The alternative model led to regression specifications similar to the ones estimated in tables 4 and 5, except for the construction of the dependent variable. In particular, given a value for γ, equation (6′) instructs us to measure the dependent variable as the log of the sample mean of C^{γ}. We did so for values of γ in the set $\{-.5, -1, -2\}$. Using the alternative dependent variables, we then reestimated the specifications in table 4 on the sample of married couples and obtained very similar results. The results based on the alternative dependent variables do not imply that idiosyncratic consumption risk is unimportant. They do imply that, by itself, the presence of such risk cannot account for our rejection of between-group consumption insurance.

VI. Welfare Implications

How large are the welfare costs of observed differences in consumption growth among cohort-education groups? How large are the welfare losses implied by the sensitivity of the consumption distribution to relative wage movements among cohort-education groups? We address these questions by computing the percentage consumption vari-

TABLE 5

Synthetic Panel Regressions, Restricted Sample: Households with a 30–55-Year-Old Male Head Who Has at Least a High School Education

		Levels Specification (N = 144)			8-Year Difference Specification (N = 18)		
Estimation Method	Household Controls*	Man's Wage	Woman's Wage	Woman's Leisure	Man's Wage	Woman's Wage	Woman's Leisure
OLS	No	.764	. . .	. . .	.800	. . .	. . .
		(.096)			(.147)		
IV	No	.981	. . .	. . .	.796	. . .	. . .
		(.086)			(.139)		
OLS	No	.778	−.050	. . .	.349	1.034	. . .
		(.104)	(.147)		(.306)	(.629)	
IV	No	1.278	−1.049	. . .	.150	.444	. . .
		(.208)	(.643)		(.602)	(1.196)	
OLS	No	.859	. . .	−.978	1.112	. . .	−2.661
		(.104)		(.457)	(.170)		(.992)
IV	No	1.310	. . .	−2.880	1.290	. . .	−3.718
		(.122)		(.648)	(.132)		(1.519)
OLS	Yes	.556	. . .	. . .	.666	. . .	. . .
		(.086)			(.265)		
IV	Yes	1.507	. . .	. . .	1.699	. . .	. . .
		(.742)			(.138)		
OLS	Yes	.577	. . .	−.206	.908	. . .	−2.746
		(.742)		(.409)	(.265)		(1.035)
IV	Yes	2.412	. . .	−6.163	1.908	. . .	−2.746
		(.742)		(12.845)	(.138)		(.564)

* The household controls are log of family size, the number of adults, the number of children under 3, and the number of other children.

ation—uniform across groups, states, and dates—required to compensate households for not living in a world with between-group consumption insurance. We tailor the calculations so that they are unaffected by events that occur prior to 1980 and prior to age 23. To isolate the role of group-level risk, we carry out the calculations in terms of group means.

The first step is to compute the consumption allocation that would have prevailed under between-group insurance. To that end, let C_{it} denote the observed consumption path from 1980 to 1990, and let C^*_{it} denote the path implied by the risk-sharing condition (2). Also, define the welfare criterion function as

$$U(C, \lambda) = \begin{cases} \lambda^{1+\gamma} \sum_{i=1}^{I} \sum_{t=1}^{T} \omega_{it} \rho^{t-1} b_{it} C_{it}^{1+\gamma} (1+\gamma)^{-1}, & \gamma < 0, \gamma \neq -1 \\ \sum_{i=1}^{I} \sum_{t=1}^{T} \omega_{it} \rho^{t-1} b_{it} (\log C_{it} + \log \lambda), & \gamma = -1, \end{cases}$$

where the index i runs over all groups that satisfy $23 \leq \text{age} \leq 59$ at time t, and ω_{it} equals the fraction of such households in group i at time t.

The path C^*_{it} is constructed recursively so that the change in marginal utility for all groups is the same and the aggregate resource constraint is satisfied. Given isoelastic preferences, the former condition implies

$$C^*_{i,t+1} = \psi_{t+1} \left(\frac{b_{it}}{\rho b_{i,t+1}} \right)^{1/\gamma} C^*_{it} \quad \forall i, \tag{13}$$

where ψ_{t+1} is a scaling factor that reflects aggregate consumption growth and the consumption allocation of new groups. In particular, ψ_{t+1} is equal to actual aggregate consumption growth of the groups that satisfy the age condition both at t and at $t + 1$. Then C^*_{it} is set equal to C_{it} for $t = 1980$ and for the groups that satisfy the age requirement for the first time at t and according to (13) otherwise. It is easily verified that these calculations for C^*_{it} solve a planner's problem that respects the realized path for aggregate consumption and the initial relative group consumption levels. In carrying out the calculations, we set $\rho = (1.02)^{-1}$, and we estimate the multiplicative taste shifters from the shape of the average age-consumption profile during the 1980s.[15]

[15] If one posits

$$b_{it} = \exp\left[-\gamma \sum_{k=1}^{4} \delta_k (\text{age}^k)_{it} \right],$$

We specify the stochastic process for the evolution of relative consumption as a two-point distribution defined in terms of the actual path and its opposite (relative to the planner's allocation): $\check{C}_{it} = C_{it} - (C^*_{it} - C_{it})$. This opposite path implies the same aggregate consumption as the observed path but alters the distribution, so that winners become losers and vice versa. In effect, we treat the 1980s as one realization of a symmetric two-point distribution. This approach is simple, transparent, and reasonable in the absence of an obvious way to calibrate ex ante uncertainty about group consumption risk from data for a single decade.

Given paths for C_{it}, C^*_{it}, and $\check{C}_{it}$, we compute the compensating variation as the value of λ^A that solves $.5[U(C, \lambda^A) + U(\check{C}, \lambda^A)] = U(C^*, 1)$. Conceptually, the calculation of λ^A corresponds to an experiment in which agents are probabilistically assigned to different cohort-education groups. Prior to assignment, agents are asked what uniform percentage increase in consumption would bring them to the same expected utility level as the consumption insurance allocation.

To compute the welfare losses implied by the sensitivity of the consumption distribution to relative wage movements, we proceed in an entirely analogous manner. We replace the time path for C_{it} in the preceding calculations by a path for $\hat{C}_{it}$, generated as predicted values from a fitted regression. Let λ^W denote the compensating variation associated with $\hat{C}_{it}$. The values for λ^W discussed below are based on the full-sample instrumental variables estimate of how men's hourly wages affect consumption, as reported in the second row of table 4 (panel A).

Table 6 reports λ^A and λ^W for various degrees of risk aversion. The table also reports the welfare cost of aggregate consumption variability, per the calculation of Lucas (1987). Lucas calculates the consumption variation required to compensate a representative agent with isoelastic preferences for aggregate consumption risk.[16] Com-

---

then a regression of the form

$$\log C_{it} = \alpha_t + g_i + \sum_{k=1}^{4} \delta_k(\text{age}^k)_{it} + e_{it}$$

delivers the estimated taste shifters

$$\hat{b}_{it} = \exp\left[-\gamma \sum_{k=1}^{4} \hat{\delta}_k(\text{age}^k)_{it}\right].$$

As anticipated, these estimated consumption requirement profiles are concave with respect to age.

[16] Lucas's calculation presumes complete markets for sharing consumption risks. For related analyses that consider the interaction between aggregate variability and idiosyncratic consumption risk, see İmrohoroğlu (1989) and Atkeson and Phelan (1994).

TABLE 6

Welfare Costs of Aggregate and Between-Group Consumption Variability

Coefficient of Relative Risk Aversion	Consumption Variation Required to Compensate for: Aggregate Consumption Risk (1)	Between-Group Consumption Risk: λ^A: All Group-Level Risk (2)	Between-Group Consumption Risk: λ^W: Men's Hourly Wages Only (3)
−.5	.00004	.0061	.0041
−1.0 (log utility)	.00008	.0126	.0084
−2.0	.00017	.0267	.0176
−5.0	.00042	.0851	.0501
−10.0	.00845	.2591	.1135

Note.—Col. 1 reports the uniform proportionate increase in consumption required to compensate a representative agent for aggregate consumption risk, per the calculation of Lucas (1987). This column reflects a calibration to the variability of postwar U.S. data on detrended log consumption (see Lucas for details). Cols. 2 and 3 report the uniform proportionate increase in consumption required to compensate households for between-group consumption risk, per the calculations described in the text. Groups are defined as 5-year birth cohorts crossed with four educational attainment categories.

paring across columns in table 6 reveals that the welfare cost of between-group consumption risk is larger than the cost of aggregate consumption risk by two orders of magnitude. The sensitivity of the consumption distribution to pretax hourly wages among men accounts for about two-thirds of the overall welfare cost of between-group consumption risk.

The welfare cost of between-group consumption risk is also large in an absolute sense. According to table 6, a household with annual consumption of $40,000 and relative risk aversion of two would have to be compensated by $1,080 per year to achieve the same level of welfare as implied by an allocation with between-group insurance. (This calculation ignores any potential gains associated with pooling consumption risks within groups.) The large size of the welfare losses associated with between-group consumption risk implies that there are powerful impediments to insurance against publicly observable shocks to the distribution of earnings capacity among cohort-education groups.

VII. Concluding Remarks

We began this paper by observing that the U.S. economy underwent pronounced and persistent movements in the structure of relative wages during the 1980s. We have shown that relative wage movements among birth cohort–education groups of men drove large changes in the distribution of household consumption. Among the less educated, real household consumption fell sharply for most cohorts during the early 1980s, paralleling their sharp declines in real wages. Among the college-educated, especially for the younger cohorts, real wages and real household consumption rose throughout the decade. Our econometric analysis shows that the close alignment between men's relative wage movements and relative household consumption movements continues to hold after consideration of alternative samples and alternative specifications motivated by consumption-leisure nonseparability and idiosyncratic consumption risk.

In our view, the magnitude of the covariance between relative wages and consumption constitutes a spectacular failure of the hypothesis of between-group consumption insurance. The hypothesis is not even remotely consistent with the evidence developed here. Our calculations indicate that the observed departures from between-group risk sharing involve large welfare costs. In addition, our evidence against the consumption insurance hypothesis involves publicly observed shocks. Indeed, the sharp decline in relative and real wages among the less educated has been a major public policy concern in recent years. Hence, our findings cannot be rationalized as a consequence of unobserved shocks in environments with informationally constrained insurance.

One potential line of explanation for our results stresses the interaction between publicly observable shocks and private information about individual attributes that relate to the acquisition of human capital. The costs of and expected returns to education, for example, are likely to vary greatly among individuals. If differences in the net returns to education (or other acquired skills) are private information, then the optimal consumption allocation may vary with publicly observable shocks. The link between consumption and observable shocks arises to call forth further human capital acquisition by those agents who are best positioned to augment the stock of needed skills. This line of explanation is in the spirit of existing theories of informationally constrained optimal consumption allocations, although we are not aware of research in this tradition that models the connection between observable shocks and private information about the returns to human capital acquisition.

A second potential line of explanation for our results stresses the difficulties of devising and maintaining institutions that share consumption risks among broad social groups. While informational problems may underlie these difficulties, they may also reflect the absence of suitable mechanisms for enforcing risk-sharing agreements that are Pareto improving ex ante and the infeasibility of articulating complete risk-sharing contracts. The barriers to devising, articulating, and enforcing optimal contracts in private settings suggest that similar problems hamper political and social risk-sharing compacts. The scale of political, social, and even military resources deployed to alter the distribution of consumption also suggests that actual consumption allocations deviate sharply from allocations constrained only by production technologies and private information.

We hope that future research discriminates between these two lines of explanation for the impact of publicly observable endowment shocks on the distribution of household consumption. There seems ample scope for both theoretical and empirical research directed toward this issue.

The empirical results in this paper largely confirm the view that animates much research on the earnings distribution by labor economists. As we noted in the Introduction, this literature typically takes for granted that measured changes in the structure of income and earnings closely parallel changes in the distribution of household welfare. The close alignment between relative wage and consumption movements among cohort-education groups supports this view, but at least two caveats are in order. First, we devised our empirical strategy to maximize the connection between relative wage and consumption movements. In particular, our focus on households with a nonelderly adult male head omits the groups that are most insulated from changes in the earnings structure. Second, we adduced only modest evidence that relative wage movements over 1- and 2-year intervals affect the consumption distribution. In this regard, our results are consistent with the view that short-term changes in the earnings distribution, even when they involve large groups of workers, carry unimportant welfare consequences.

To close, we remark on one other direction for future research. Many advanced and middle-income economies experienced large relative wage movements among distinct groups of workers during the 1970s and 1980s (Davis 1992). Several of these countries offer cross-sectional data sets with information on consumption expenditures and labor earnings comparable to the information contained in the U.S. Consumer Expenditure Survey and Current Population Survey. These data sets provide the grist for synthetic panel analyses of relative wage and consumption comovements in several countries. Developing this line of research would provide an empirical basis for quantifying and interpreting cross-country differences in the extent of risk sharing. Much existing research considers cross-country differences in labor market institutions, tax structures, and income maintenance programs with an eye toward their distributional consequences (e.g., Card and Freeman 1993), but the focus typically falls on income rather than on consumption outcomes. Given the variety and complexity of private, public, market, and extra-market institutions that play risk-sharing roles, it would be useful to supplement the existing style of research with more direct evidence on how endowment shocks affect the consumption distribution.

References

Abel, Andrew B., and Kotlikoff, Laurence J. "Intergenerational Altruism and the Effectiveness of Fiscal Policy—New Tests Based on Cohort Data." In *Savings and Bequests,* edited by Toshiaki Tachibanaki. Ann Arbor: Univ. Michigan Press, 1994.

Altonji, Joseph G.; Hayashi, Fumio; and Kotlikoff, Laurence J. "Is the Extended Family Altruistically Linked? Direct Tests Using Micro Data." *A.E.R.* 82 (December 1992): 1177–98.

Atkeson, Andrew G., and Lucas, Robert E., Jr. "On Efficient Distribution with Private Information." *Rev. Econ. Studies* 59 (July 1992): 427–53.

Atkeson, Andrew G., and Phelan, Christopher. "Reconsidering the Costs of Business Cycles with Incomplete Markets." In *NBER Macroeconomics Annual 1994,* edited by Stanley Fischer and Julio J. Rotemberg. Cambridge, Mass.: MIT Press, 1994.

Attanasio, Orazio, and Davis, Steven J. "Relative Wage Movements and the Distribution of Consumption." Working Paper no. 4771. Cambridge, Mass.: NBER, June 1994.

Attanasio, Orazio, and Weber, Guglielmo. "Consumption Growth and Excess Sensitivity to Income: Evidence from U.S. Micro Data." Manuscript. Stanford, Calif.: Stanford Univ., 1992.

Card, David. "Comments and Discussion." *Brookings Papers Econ. Activity,* no. 2 (1991), pp. 62–64.

Card, David, and Freeman, Richard B., eds. *Small Differences That Matter: Labor Markets and Income Maintenance in Canada and the United States.* Chicago: Univ. Chicago Press (for NBER), 1993.

Cochrane, John H. "A Simple Test of Consumption Insurance." *J.P.E.* 99 (October 1991): 957–76.

Cutler, David M., and Katz, Lawrence F. "Macroeconomic Performance and the Disadvantaged." *Brookings Papers Econ. Activity,* no. 2 (1991), pp. 1–61.

Davis, Steven J. "Cross-Country Patterns of Change in Relative Wages." In *NBER Macroeconomics Annual 1992,* edited by Olivier J. Blanchard and Stanley Fischer. Cambridge, Mass.: MIT Press, 1992.

Deaton, Angus. *Understanding Consumption.* Oxford: Clarendon, 1992.

Goldin, Claudia, and Margo, Robert A. "The Great Compression: The Wage Structure in the United States at Mid-century." *Q.J.E.* 107 (February 1992): 1–34.

Green, Edward J. "Lending and the Smoothing of Uninsurable Income." In *Contractual Arrangements for Intertemporal Trade,* edited by Edward C. Prescott and Neil Wallace. Minneapolis: Univ. Minnesota Press, 1987.

İmrohoroğlu, Ayşe. "Cost of Business Cycles with Indivisibilities and Liquidity Constraints." *J.P.E.* 97 (December 1989): 1364–83.

Katz, Lawrence F., and Murphy, Kevin M. "Changes in Relative Wages, 1963–1987: Supply and Demand Factors." *Q.J.E.* 107 (February 1992): 35–78.

Levy, Frank, and Murnane, Richard J. "U.S. Earnings Levels and Earnings

Inequality: A Review of Recent Trends and Proposed Explanations." *J. Econ. Literature* 30 (September 1992): 1333–81.

Lucas, Robert E., Jr. *Models of Business Cycles.* New York: Blackwell, 1987.

Mace, Barbara J. "Full Insurance in the Presence of Aggregate Uncertainty." *J.P.E.* 99 (October 1991): 928–56.

Mayer, Susan E., and Jenks, Christopher. "Recent Trends in Economic Inequality in the United States: Income vs. Expenditures vs. Material Well-Being." Manuscript. Chicago: Univ. Chicago, 1991.

Nelson, Julie A. "On Testing for Full Insurance Using Consumer Expenditure Survey Data." *J.P.E.* 102 (April 1994): 384–94.

Pencavel, John H. "Labor Supply of Men: A Survey." In *Handbook of Labor Economics,* vol. 1, edited by Orley C. Ashenfelter and Richard Layard. Amsterdam: Elsevier Sci., 1986.

Phelan, Christopher, and Townsend, Robert M. "Computing Multi-period, Information-Constrained Optima." *Rev. Econ. Studies* 58 (October 1991): 853–81.

Townsend, Robert M. "Information Constrained Insurance: The Revelation Principle Extended." *J. Monetary Econ.* 21 (March/May 1988): 411–50.

———. "Risk and Insurance in Village India." *Econometrica* 62 (May 1994): 539–91.