

NBER WORKING PAPER SERIES

RELATIVE WAGE MOVEMENTS AND
THE DISTRIBUTION OF CONSUMPTION

Orazio Attanasio
Steven J. Davis

Working Paper No. 4771

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1994

For many helpful conversations and comments, we thank John Cochrane, Lars Hansen, Costas Meghir, Kevin M. Murphy, Rob Townsend, Guglielmo Weber and workshop participants at MIT, Stanford, the University of Southern California, the University of Chicago, the Federal Reserve Bank of Cleveland, Boston College, and the 1993 NBER Summer Institute working group on Consumption Behavior. Davis gratefully acknowledges research support by the National Science Foundation. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

RELATIVE WAGE MOVEMENTS AND
THE DISTRIBUTION OF CONSUMPTION

ABSTRACT

We analyze how relative wage movements across birth cohorts and education groups during the 1980s affected the distribution of household consumption. The analysis integrates the labor economics literature on time variation in the wage structure with the consumption insurance literature. In contrast to previous tests of consumption insurance, we examine the impact of systematic, publicly observable shifts in the hourly wage structure. To circumvent the extreme scarcity of longitudinal data with high quality information on both consumption and labor market outcomes, we draw upon the best available cross-sectional data sources to construct synthetic panel data on consumption, labor supply and wages.

We find that low-frequency movements in the cohort-education structure of pre-tax hourly wages drove large changes in the distribution of household consumption. The results constitute a spectacular failure of the consumption insurance hypothesis, and one that is not explained by existing theories of informationally constrained optimal consumption allocations. We also develop a procedure for assessing the welfare consequences of deviations from full consumption insurance and, in particular, from the failure to insulate the consumption distribution from relative wage shifts across cohort-education groups. For a coefficient of relative risk aversion equal to two, fully insulating households from group-specific endowment variation would raise welfare by an amount equivalent to a uniform 2.7% consumption increase.

Orazio Attanasio
Department of Economics
Stanford University
Stanford, CA 93305
and NBER

Steven J. Davis
Graduate School of Business
University of Chicago
Chicago, IL 60637
and NBER

I. Introduction

The U.S. economy underwent pronounced, persistent movements in the structure of relative wages during the 1980s. This paper analyzes 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, examines time variation in the structure of relative wages. 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. ¹ The main research thrust in this literature 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. The prevailing belief, announced 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 this belief, much less subject it to any serious evaluation – 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., Mace (1991), Cochrane (1991), Townsend (1994) and Altonji et al. (1992). This literature exploits a simple and rigorous theory of consumption allocations that delivers, under plausible restrictions on preferences, strong

¹ Levy and Murnane (1992) provide an extensive survey of research on recent U.S. wage structure developments, and 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.

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 can serve more broadly as tools for investigating and interpreting the extent, pattern, and nature of departures from full insurance consumption outcomes.

Previous empirical research in the consumption insurance literature typically relates 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 individual-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), and Atkeson and Lucas (1992). These theories offer no explanation for why relative consumption growth rates depend on publicly observed shocks under an optimal allocation. Rather, they suggest perfect insurance against all publicly observed shocks. We return to this point in the conclusion, 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 of individuals and households facilitates the use of cross-sectional data sets that offer comprehensive, high quality information on either consumption or earnings, but not 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 lon-

gitudinal 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 upon the maintained assumption that the social planner can freely transfer leisure across households (Cochrane, 1991 and 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 female labor supply. A similar strategy would be difficult to implement in work that focuses on idiosyncratic components of individual earnings capacity because of censored wage observations for jobless persons.

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 we use the theory to derive estimating equations suitable for synthetic panel data. Section III treats several specification issues that pertain to life-cycle and demographic factors and preference nonseparabilities. All of our empirical specifications relate time variation in the structure of wages among observationally distinct groups of workers to the distribution of consumption among households headed by these workers.

Section IV takes up econometric issues related to the error structure implied by the theory, the error structure induced by mismeasurement and sampling variation, and the choice between level and difference specifications. We discuss potential pitfalls of least squares estimation on the synthetic panel data, and we develop an instrumental variables strategy designed to overcome these pitfalls.

Section V describes the data and the construction of the synthetic panels. Our syn-

thetic panels contain annual observations from 1980 to 1990 on five-year birth cohorts crossed by four educational attainment categories. We calculate birth-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 VI presents a descriptive analysis of relative wage and consumption movements during the 1980s. Focussing on households with an adult male, we describe relative consumption movements across cohort-education groups. We also describe relative wage movements among men and their wives. In line with previous work, we document pronounced relative wage movements across education groups and birth cohorts. Furthermore, measured wages show large absolute declines among several cohorts of less-educated workers. These sharp relative wage movements across observationally distinct groups of workers provide considerable leverage for estimating the impact of wage changes on the consumption distribution.

Section VI also offers a pictorial analysis of the nexus between relative wage and consumption movements across birth-education groups. 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 relative wage movements across birth cohorts within education groups.

Section VII 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 also indicate that changes in men's pre-tax relative wage rates were the dominant driving force

behind the (large) changes in the distribution of household consumption. Higher frequency comovements between relative wages and consumption show much weaker evidence against the consumption insurance hypothesis and strong evidence against the extreme alternative of no consumption smoothing.

Our finding that low-frequency relative wage movements influence the distribution of household consumption may be unsurprising to many readers, but we view the magnitude of the departure from the consumption insurance null as remarkable. There are many mechanisms that potentially help insulate the household consumption distribution from persistent shocks to the structure of pre-tax hourly wages among men: offsetting labor supply responses by men, offsetting movements in relative wages or labor supply by their wives and other household members, the progressivity of the income tax, changes in the tax structure such as the 1986 increase in the earned income tax credit, public welfare programs such as food stamps, altruistically motivated intra- or inter-generational transfers that cut across birth cohorts or education groups, and consumption smoothing through own holdings of real and financial assets. At least for the U.S. experience during the 1980s, these smoothing mechanisms poorly insulated the distribution of household consumption from shifts in the structure of pre-tax hourly wages among cohort-education groups of men.

To assess the welfare consequences of these shifts in the household consumption distribution, section VIII develops a procedure for quantifying the extent of departure from an optimal consumption allocation. For plausible degrees of risk aversion, our calculations reveal 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 of two, compensating households for the type of cross-group consumption risk they faced during the 1980s requires a uniform 2.7% consumption increase in all states and dates.

II. Theory and Empirical Specifications

A. Testable Implications of the Consumption Insurance Hypothesis

Diamond (1967) and Wilson (1968) derive the risk-sharing properties of Pareto-optimal consumption allocations in full information environments. Scheinkman (1984) was apparently the first to emphasize the empirical implications of optimal risk-sharing behavior for the distribution of consumption growth. Like Mace (1991), Cochrane (1991) and Townsend (1994), we use the theory of optimal risk sharing to formulate testable hypotheses and guide our empirical analysis of relative consumption movements.

A full insurance Pareto-optimal consumption allocation maximizes the discounted weighted sum of individual households' utility functions:

$$\max_{C^j(S^t)} \sum_{j=1}^J \lambda^j \sum_{t=1}^{\infty} \sum_{S^t} (\rho^j)^t \pi(S^t) U[C^j(S^t), \delta^j(S^t)], \quad (1)$$

where $C^j(S^t)$ is consumption by household j at time t when the state of the world is S^t ; λ^j is household j 's Pareto weight, which reflects the initial distribution of wealth; $(\rho^j)^t$ is household j 's discount factor at time t ; $\pi(S^t)$ denotes the probability of state S^t ; and $\delta^j(S^t)$ indexes arbitrary cross-sectional and time variation in household preferences. Feasibility requires that aggregate consumption be less than or equal to the aggregate endowment (or aggregate production) at each date and in each state of the world.

The first-order conditions for Pareto-optimal consumption allocations are

$$(\rho^j)^t \lambda^j U_c[C^j(S^t), \delta^j(S^t)] = \mu(S^t), \quad \text{for all } S^t \text{ and } j = 1, \dots, J, \quad (2)$$

where $\mu(S^t)$ equals the Lagrange multiplier associated with the feasibility constraint, divided by the probability of state S^t . The set of first-order conditions given by (2) embody the implications of full insurance for consumption patterns across individuals, time periods and states of nature. In more concise notation, we write these first-order conditions as

$$(\rho^j)^t \lambda^j U_c(C_t^j, \delta_t^j) = \mu_t. \quad (2)$$

Full insurance implies identical growth rates across households in the marginal utility of consumption. To see this point, first note that the multiplier μ_t in (2) varies with aggregate consumption but is constant across households for a particular state and date. Hence,

given aggregate consumption and the Pareto weights, individual households' consumption allocations do not vary with their endowments. Second, since the Pareto weights are time invariant, we can use observations at two points in time to remove the household fixed effects:

$$\rho^j \frac{U_c[C_{t+1}^j, \delta_{t+1}^j]}{U_c[C_t^j, \delta_t^j]} = \frac{\mu_{t+1}}{\mu_t}, \quad j = 1, \dots, J. \quad (3)$$

Equation (3) says that the growth between t and $t + 1$ in the discounted marginal utility of consumption is the same for all households.

Thus, under the theory, measured consumption growth can vary across households for only three reasons: (i) utility functions differ across households;³ (ii) households experience idiosyncratic preference shocks; and (iii) observed household consumption growth contains idiosyncratic measurement error components. In particular, any variable that is cross-sectionally uncorrelated with preference variation and measurement error in consumption growth is, under the theory, also uncorrelated with the cross-sectional distribution of consumption growth.

To test this implication of full consumption insurance, Cochrane (1991) runs regressions of the form

$$\log(C_{t+1}^j/C_t^j) = \alpha + \beta X_{t+1}^j + \epsilon_{t+1}^j, \quad j = 1, \dots, J, \quad (4)$$

and tests the null hypothesis that $\beta = 0$. Here, X_{t+1}^j is a shock to household j 's endowment or earnings capacity that is assumed to be uncorrelated with measurement error in consumption growth and preference variation across households.

Cochrane derives (4) under the assumption that household j exhibits constant relative risk aversion γ^j and a multiplicative preference shock b_t^j :

$$U[C_t^j, \delta_t^j] \equiv U[C_t^j, b_t^j \gamma^j] = b_t^j \frac{(C_t^j)^{1+\gamma^j}}{1+\gamma^j}. \quad (5)$$

³ Differences in intertemporal substitution elasticities and discount rates imply differences in the volatility and average rates of consumption growth. If preferences are nonhomothetic, utility functions effectively differ across households at different wealth levels.

Under this preference specification, conditions (3) become

$$\rho^j \frac{b_{t+1}^j (C_{t+1}^j)^{\gamma^j}}{b_t^j (C_t^j)^{\gamma^j}} = \frac{\mu_{t+1}}{\mu_t}, \quad j = 1, \dots, J. \quad (6)$$

Taking logs and adding a measurement error ξ_{t+1}^j to the growth of log consumption yields

$$\log\left(\frac{C_{t+1}^j}{C_t^j}\right) = \frac{1}{\gamma^j} \left[\log\left(\frac{\mu_{t+1}}{\mu_t}\right) - \log\left(\frac{b_{t+1}^j}{b_t^j}\right) - \log(\rho^j) \right] + \xi_{t+1}^j. \quad (7)$$

Hence, if X_{t+1}^j is cross-sectionally independent of preference shifts ($\log(b_{t+1}^j/b_t^j)$, γ^j , and ρ^j) and of measurement error (ξ_{t+1}^j), then consumption insurance implies that X_{t+1}^j will have no explanatory power in regressions of the form (4). Cochrane also considers counterparts to (4) that are suitable under more general specifications for preferences and preference variation. We take up these matters in the context of our synthetic panel specifications below.

In his empirical implementation, Cochrane uses data from the Panel Survey 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 her data. In other related work, Altonji et al. (1992) reject the hypothesis of complete risk sharing among extended family members, and Lewis (1993) investigates whether the pattern of aggregate consumption growth rates among countries conforms to the full insurance hypothesis.

B. Empirical Specifications Suitable for Synthetic Panel Data

Whereas Cochrane, Mace, Townsend and others use specifications like (4) 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 the advantages of a focus on this type of wage structure variation. Here, we concentrate on deriving synthetic panel analogs to (4) that exploit this type of wage variation.

Let C_{it}^j and W_{it}^j denote consumption and the wage rate for the j th household in group i at time t . Our empirical analysis considers households defined by five-year birth cohorts crossed with four educational attainment categories. Here, we simply refer to them as groups.

After partitioning the data into I groups, take logs in (2), and average over the sample of group- i households at time t to obtain

$$\hat{V}_{it} \equiv \frac{\sum_{j \in i(t)} \log[U_c(C_{it}^j, \delta_{it}^j)]}{\#i(t)} = \log \mu_t - t \frac{\sum_{j \in i(t)} \log \rho^j}{\#i(t)} - \frac{\sum_{j \in i(t)} \log \lambda^j}{\#i(t)} \quad (8)$$

for $i = 1, \dots, I$ and $t = 1, \dots, T$. Here, and below, the hat symbol denotes a sample mean quantity. It should be understood that the time index does not necessarily run over the same periods for all groups.

Rewrite the sample average first-order conditions as

$$\hat{V}_{it} = \log \mu_t - t \bar{R}_i - \bar{L}_i + \epsilon_{it}, \quad i = 1, \dots, I, \quad t = 1, \dots, T, \quad (9)$$

where \bar{R}_i and \bar{L}_i are population counterparts to the corresponding sample means in (8), and where ϵ_{it} is an error term due to differences between population and sample means. This formulation emphasizes that the sample mean of log marginal utility, \hat{V}_{it} , differs from the corresponding population mean due to finite sampling from a heterogeneous population. In practice, \hat{V}_{it} also differs from the population mean due to measurement error in variables that determine the marginal utility of consumption. Under the theory, the population quantity depends on a fixed time effect ($\log \mu_t$) common to all groups, a time-invariant group effect (\bar{L}_i), and a third term that varies linearly with time ($t \bar{R}_i$). If the distribution of time discount factors ρ^j is identical across groups, then the third term becomes the

same linear function of time for all groups. If this distribution differs across groups, then the linear time trends also differ across groups.

Given a parametrization of the utility function, equation (9) leads to a "levels" regression specification for group-averaged data. Alternatively, we can derive a first-difference specification from

$$\hat{V}_{it} - \hat{V}_{is} = (\log \mu_t - \log \mu_s) + (t - s)\bar{R}_i + \epsilon_{it} - \epsilon_{is}, \quad (10)$$

where, as before, the error terms reflect deviations between population and sample means and imperfect measurement of variables that determine the marginal utility of consumption.

When preferences have the isoelastic form (5), the sample mean \hat{V}_{it} becomes

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

where v_{it} is an error term due to deviations between sample and population means. Combining this equation with (9) yields

$$\frac{\sum_{j \in i(t)} \gamma^j \log(C_t^j)}{\#i(t)} = \log \mu_t - t\bar{R}_i - \bar{L}_i - \bar{B}_{it} + \epsilon_{it} - v_{it} \quad (11)$$

Now consider a regression of the sample mean 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:

$$\log \widehat{C}_{it} = \alpha_t + g_i + \beta X_{it} + \epsilon_{it}, \quad i = 1, \dots, I, \quad t = 1, \dots, T. \quad (12)$$

Comparing (11) and (12), the consumption insurance hypothesis implies $\beta = 0$ in this regression specification, provided that X_{it} also satisfies a list of auxiliary statistical assumptions. In particular, conditional on control variables in the regression equation, we require that

- (i) X_{it} is independent of measurement errors in log consumption;

- (ii) X_{it} is independent of cross-group differences in the distribution of the time discount factors ρ^j ;
- (iii) X_{it} is independent of cross-group differences in the distribution of intertemporal substitution parameters γ^j ; and
- (iv) X_{it} is independent of variation in the mean preference disturbances \bar{B}_{it} .

Restrictions (ii)-(iv) say that X_{it} must be independent of any preference variation not captured by other controls in the regression specification. ⁴ Assumption (ii) fails if, coincidentally, groups with higher time discount factors happen to experience unusually large relative wage increases during the sample period. Similarly, (iii) fails if groups with greater intertemporal substitution elasticities happen to experience unusually large relative wage increases. While we cannot rule out such coincidental relationships between preference variation and relative wage movements, they would constitute an extremely implausible interpretation of a consistent pattern of rejections of the null hypothesis, $\beta = 0$. In contrast, in a specification that includes only time and group fixed effects as controls, both theoretical and empirical considerations point to plausible reasons for a systematic violation of (iv). We take up this matter in section III.

The restriction (i) on the measurement error process for X_{it} can fail because of systematic changes over time in the accuracy of the sampling frame for particular groups, misclassification of individuals or households among groups, or the noise induced by finite sampling from groups with heterogeneous populations. Similar measurement problems arise in most econometric work. Our reliance on different cross-sectional data sets to construct left and right side variables in (12) greatly diminishes the likelihood of nonzero correlation between the dependent variable and measurement error in X_{it} .

Turning to the first-difference specification, equation (11) implies

$$\frac{\sum_{j \in i(t)} \gamma^j \log(C_i^j)}{\#i(t)} - \frac{\sum_{j \in i(s)} \gamma^j \log(C_i^j)}{\#i(s)} =$$

⁴ We could dispense with (ii) by introducing group-specific linear time trends in (12). This modification to the regression specification is unappealing, because it would soak up the persistent components of variation in candidate regressors X_{it} , i.e., the persistent components of relative wage movements in the empirical analysis.

$$(\log \mu_t - \log \mu_s) + (s - t)\bar{R}_i - (\bar{B}_{it} - \bar{B}_{is}) + (\epsilon_{it} - \epsilon_{is}) - (v_{it} - v_{is}) \quad (13)$$

Hence, the consumption insurance hypothesis implies $\beta = 0$ in the regression specification

$$\log \widehat{C}_{it} - \log \widehat{C}_{is} = \alpha_t + \beta X_{it} + e_{it} - e_{is} \quad i = 1, \dots, I \text{ and } t = (t - s), \dots, T, \quad (14)$$

provided that X_{it} satisfies a list of statistical assumptions analogous to (i)-(iv).⁵

III. 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 measured relative wage movements not be coincidentally correlated with omitted factors that drive a wedge between measured consumption growth and marginal utility growth. This section treats potentially important sources of coincidental correlation and their consequences for the empirical analysis.

A. Demographics and Life-Cycle Effects

One important source of coincidental correlation involves systematic life-cycle variation in earnings capacity and in the marginal utility of consumption. Earnings capacity increases with experience over much of the life cycle. For a partially overlapping segment of the life cycle, household-level consumption requirements also increase due to increases in family size and age of children. Consequently, the typical cohort can be expected to exhibit a coincidental correlation between relative wage growth and relative consumption growth over much of its life cycle. To control for this correlation, the empirical specifications include high-order polynomials in the head's age plus measures of family size and composition. In terms of the formal theory, these regression controls reflect components of the individual and group preference shifters, b_i^t and \bar{B}_{it} .

⁵ Both Mace and Townsend consider econometric specifications that follow from exponential preferences. Aggregating across household as above, one can easily show that exponential preferences lead to versions of (12) and (14) in which the level of consumption replaces the log of consumption.

B. Preference Nonseparability with Nontransferable Leisure

Consider the implications of relative wage movements when preferences are nonseparable between consumption and leisure. To develop the argument, suppose that wage movements reflect underlying disturbances to an aggregate production technology with diminishing marginal factor products. Assume that nonmarket time is *not* transferable across households, and that preferences exhibit diminishing marginal utility of leisure. Then, if time devoted to market activity increases the marginal utility of consuming market goods, groups with growing relative wages also experience relative consumption growth under a Pareto-optimal allocation. It follows that, even under a Pareto-optimal allocation, β exceeds zero in specifications like (12) and (14) 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.

To derive a synthetic panel specification that conditions on leisure, consider the generalization of (5),

$$U[C_t^j, \bar{H} - L_t^j, \delta_t^j] = v_t^j (C_t^j)^{1+\gamma^j} (\bar{H} - L_t^j)^{1+\phi^j}, \quad (5')$$

where \bar{H} denotes the time endowment in hours, and L_t^j denotes hours worked.⁶ This specification implies that the sample mean log marginal utility function, \hat{V}_{it} , contains the additional term $\sum_{j \in i(t)} (1 + \phi^j) \log(\bar{H} - L_t^j) / \#i(t)$. The appropriate level and difference regression specifications now become

$$\widehat{\log C_{it}} = \alpha_t + g_i + \psi \log(\widehat{\bar{H}} - L_{it}) + \beta X_{it} + e_{it}, \quad \text{and} \quad (12')$$

$$\widehat{\log C_{it}} - \widehat{\log C_{is}} = \alpha_t + \psi [\log(\widehat{\bar{H}} - L_{it}) - \log(\widehat{\bar{H}} - L_{is})] + \beta X_{it} + e_{it} - e_{is}. \quad (14')$$

Thus, when preferences are nonseparable and leisure is nontransferable, one might test the consumption insurance hypothesis by investigating whether $\beta = 0$ in regressions

⁶ The parameter restrictions $b < 0$ and $\gamma, \phi < -1$ imply that nonmarket time and consumption expenditures are substitutes, as presumed in the text.

like (12') and (14') that condition on leisure. ⁷ ⁸ Unfortunately, when preferences are nonseparable, using a wage measure as the X_i variable in (12') or (14') gives rise to an identification problem. Recall that the social planner equates the marginal rate of substitution between consumption and leisure to the wage for households at an interior solution. This intratemporal first-order condition has the same form as (12') when we substitute the wage for X_{it} . Likewise, the first difference of this condition has the same form as (14') when we substitute the first difference of the wage for X_{it} . Consequently, we cannot use relative wage movements to test the consumption insurance hypothesis in (12') and (14') without additional information on the substitutability between consumption and leisure. We achieved identification in (12) and (14) by relying upon the assumption of separable preferences to exclude leisure.

We suggest two approaches to the identification problem that arises with nonseparable preferences. The first consists in focussing on a sample for which male labor supply is believed to be highly inelastic. This consideration is one reason we restrict the sample to households headed by adult males between 23 and 59. We also report results for samples restricted to households headed by males between 30 and 55 and to households who have at least a high school education. ⁹ If we believe that male labor supply is highly inelastic, and

⁷ Using (12') or (14') to test the consumption insurance hypothesis under nonseparable preferences requires, in addition to (i)-(iv), a further auxiliary assumption: X_{it} is independent of cross-group differences in the distribution of the ϕ^j . We can dispense with this assumption by allowing for group-specific slope terms on the leisure controls in (12') or (14').

⁸ Specifications like (12) and (14) can also lead to an upwardly biased estimate of β , if preferences entail a nonseparability between household consumption and wife's leisure. Suppose we test the hypothesis that $\beta = 0$ in (12) or (14) using the husband's wage for X_{it} . If relative wage movements for husbands and wives covary positively, and if reductions in the wife's leisure increase the marginal utility of household consumption, then β exceeds zero under a Pareto-optimal allocation. Section VI below shows that husbands' and wives' relative wage movements exhibit strong positive covariation in our sample.

⁹ Juhn, Murphy and Topel (1991) present evidence that less-educated men exhibit greater labor supply elasticity.

we control for female leisure, then a finding that *male* wage growth enters the equation for consumption growth constitutes convincing evidence against the consumption insurance hypothesis.¹⁰

Alternatively we can ask the question: can we explain the observed comovements among relative wages, consumption and hours worked appealing to nonseparability between consumption and wages? Our second approach to the identification problem entails calibrating the marginal rate of substitution between consumption and leisure to construct the marginal utility of consumption. The parameters of the MRS are chosen to yield values for labor supply elasticities consistent with existing evidence. We can then estimate equations like (12) or (14) with the (log of) the marginal utility of consumption (rather than log consumption) on the left hand side of the regression, and test the hypothesis that relative wages enter these equations. A rejection of the null would be an indication that the consumption insurance hypothesis cannot be salvaged by appealing to nonseparability between consumption and leisure.

C. Leisure Transferability via Market Transactions

Our emphasis on the potential importance of conditioning on leisure when testing the consumption insurance hypothesis differs from previous research. Cochrane (1991) and Townsend (1994), in particular, argue that free transferability of leisure across households implies the suitability of their consumption insurance tests in the face of nonseparable preferences. This argument fails when the social planner's optimal leisure allocation across households is implemented through market transactions. These transactions for time-saving goods and services are recorded as additional consumption expenditures in the data. Examples include prepared food, personal services that cover laundry and other household chores, and expenditures on certain forms of transportation. Aside from time-saving goods and services, more expensive clothing and other market goods are often complementary inputs for market work activities. These observations bear out a problematic aspect of

¹⁰ Altug and Miller (1990) consider preferences defined over consumption, husband's leisure and wife's leisure, but they assume separability between consumption and male leisure.

previous consumption insurance tests that fail to condition on leisure.

Appendix A analyzes a model of leisure transferability via market transactions. The model posits two types of market-produced commodities: goods that generate utility directly and time-saving goods. Households combine time-saving goods and nonmarket time to produce effective leisure. Shocks to an aggregate production technology drive relative wage movements. The social planner chooses hours worked in the market sector and quantities of the two market-produced goods to implement a Pareto-optimal allocation. As before, an optimal allocation entails identical growth rates across households in the marginal utility of consumption.

In this model, the mechanism for transferring leisure across households induces a non-separability between leisure and *measured* consumption, even if the underlying preferences are separable. In addition, this mechanism induces positive covariation across households between relative wage growth and relative growth in measured consumption. Hence, if expenditures on time-saving goods are the mechanism by which an optimal allocation is implemented, then testing the consumption insurance hypothesis with specifications like (12) and (14) involves a bias in favor of the alternative hypothesis $\beta > 0$. 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 specification problem and 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 (1992) test of the intergenerational altruism hypothesis in cohort data.

In principle, one could address the bias caused by time-saving goods by associating them with particular subsets of commodities observed in the data. One could then estimate (12) and (14) for consumption measures based only on utility-producing goods. Given a correct identification of time-saving goods (and suitable separability assumptions), this specification would provide an unbiased estimate of β under the consumption insurance null. As a practical empirical matter, however, it would be difficult to create a satisfactory partition of observed commodities into utility-producing and time-saving goods. For this

reason, we pursue the approaches set out in the preceding section. We interpret the model with nonseparable preferences and nontransferable leisure as a reduced form of the model with leisure transferability via market transactions.

Our analysis of wages, time-saving goods, and effective leisure can be viewed as an application of the general theory of household commodity production (Becker, 1965, especially section IV). Juster and Stafford (1991) summarize several pieces of empirical evidence that support the importance of the wage-induced substitution effects stressed in our analysis.

D. Other Preference Nonseparabilities

Our empirical investigation focuses on the consumption of nondurable goods and services. If preferences are not separable between these components of consumption and other components, then the left side of equations (12) and (14) mismeasure the marginal utility of consumption. While this mismeasurement constitutes a potential source of bias, we do not think it is a serious concern for two reasons. 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.¹¹ 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 in section II.B 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.

IV. Econometric Issues

A. The Error Structure and Estimation Methods

¹¹ 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.

Under the null hypothesis of full insurance against publicly observable shocks, the equation errors in (12) and (14) arise only from preference variation and measurement error. The expectational errors for the uninsurable aggregate shocks are captured by time effects. Hence, if the X_{it} in (12) and (14) were measured without error – and maintaining the auxiliary assumptions about preference variation – OLS would deliver consistent coefficient estimates.

Sampling variation alone, however, induces measurement error in the sample mean quantities that appear as regressors in (12) and (14). Of course, sampling variation is unlikely to be the only source of measurement error in the constructed regressors. To address the potential inconsistency of the OLS estimator for β caused by measurement error, we devise an instrumental variable strategy for estimating (12) and (14).

To develop an appropriate IV strategy, represent the regression equation involving the unobserved, true variables as

$$\Delta_k \bar{C}_{it} = \text{controls} + \beta \Delta_k \bar{W}_{it} + \epsilon_{it}, \quad (15)$$

where $\Delta_k X_{it} \equiv X_{it} - X_{i,t-k}$ for $k > 0$, with $\Delta_0 X_{it} \equiv X_{it}$, and where \bar{C} and \bar{W} denote true values of log consumption and log wages, respectively. C and W represent the corresponding observed quantities. Under the null, the *only* reasons to instrument for $\Delta_k W_{it}$ are sampling variation and measurement error in W_{it} .

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 (15). Since 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

$$W_{it} = \bar{W}_{it} + u_{it}, \quad (16)$$

where the error u_{it} satisfies $\text{Cov}(u_{it}, \bar{W}_{i,t-k}) = 0$ 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 $i \neq i'$. Given these assumptions, we describe an IV estimation strategy designed to perform well, while minimizing the loss of observations due to instrumenting. For reasons that will become apparent, we treat the levels, one-year difference, and k -year difference specifications separately.

Before describing our IV estimators, one further observation is in order. Under the null hypothesis of full insurance against observable relative wage movements, and given the postulated measurement-error structure for (16), lags and leads of group means are valid instruments for the contemporaneous wage measure in (15).

Instrument Choice for the Levels Specification

This case fits the classical measurement error model. Thus, if we estimate (15) with $k = 0$ by OLS, we will obtain a downward biased estimator for β . To address this problem, consider the following instrument Z_{it} for W_{it} :

$$Z_{it} = \begin{cases} W_{i,t+1} & \text{if } W_{i,t-1} \text{ is unobserved;} \\ (1/2)(W_{i,t+1} + W_{i,t-1}) & \text{if both } W_{i,t+1} \text{ and } W_{i,t-1} \text{ are observed;} \\ W_{i,t-1} & \text{if } W_{i,t+1} \text{ is unobserved.} \end{cases} \quad (17)$$

The appeal of this scheme is three-fold. First, no observations are lost due to instrumenting. Second, the averaging over the immediate past and immediate future values of the wage measure reduces the noise component in the instrument that arises from sampling variation and measurement error. 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.

Instrument Choice for One-Year Difference Specification

Suppressing the i subscript for expositional convenience, we have

$$\Delta_1 W_t = (\bar{W}_t + u_t) - (\bar{W}_{t-1} + u_{t-1}).$$

It follows that the once-lagged value of $\Delta_1 W_t$ is *not* a valid instrument for $\Delta_1 W_t$ in the regression,

$$\Delta_1 C_{it} = \text{controls} + \beta \Delta_1 W_{it} + e_{it},$$

because both $\Delta_1 W_t$ and $\Delta_1 W_{t-1}$ involve the error u_{t-1} .

Thus, pursuing the same approach as before, consider the following instrument for $\Delta_1 W_t$:

$$Z_t = \begin{cases} \Delta_1 W_{t+2} & \text{if } \Delta_1 W_{t-2} \text{ is unobserved;} \\ (1/2)(\Delta_1 W_{t+2} + \Delta_1 W_{t-2}) & \text{if both } \Delta_1 W_{t+2} \text{ and } \Delta_1 W_{t-2} \text{ are observed;} \\ \Delta_1 W_{t-2} & \text{if } \Delta_1 W_{t+1} \text{ is unobserved.} \end{cases} \quad (18)$$

While this instrumenting scheme shares the first two virtues described above for the levels scheme, it suffers from a serious shortcoming. True *changes* in wages contain much less persistence through time than true levels. Thus, Z_t may not be highly correlated with $\Delta_1 \bar{W}_t$. In practice, IV estimates based on (18) produced enormous standard errors.

In light of this shortcoming, consider an alternative instrument that brackets the time interval of the true change. Let

$$Z_t = \Delta_3 W_{t+1}. \quad (19)$$

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, (19) entails the loss of only one observation per group.

Instrument Choice for k-Year Difference Specifications ($k \geq 2$)

In this case, we could pursue the basic idea behind (17), and define an instrument that involves averaging immediate lags and leads of k -year differences. However, the benefits of two-sided averaging are less pronounced than in the levels case. Hence, we use the simple instrumenting scheme,

$$Z_t = \Delta_k W_{t-1}. \quad (20)$$

This scheme entails no loss of observations (except possibly for young cohorts), since our wage data extends further back in time than our consumption data.

B. Estimating the Covariance Matrix

The regression residuals in (12) and (14) 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 k th-order moving average term, as evident from (14) or (14'). 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. Appendix B supplies details.

V. 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 Consumer Expenditure Survey (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 who share at least two of three major expense categories: food, housing, and other living expenses. The CEX also contains information on labor market outcomes for individual household members. 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. A quarterly interview elicits information about expenditure patterns during each of the preceding three months. Information about income and labor market outcomes refers to the twelve months preceding the interview. Our investigation uses CEX data for calendar years 1980 to 1990.

The Annual Demographic Supplement to the Current Population Survey (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 to 60,000 households per year. Our investigation uses CPS data for calendar years 1975 to 1990.

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 empirical analysis restricts attention to households with a male head (all married

couples, plus single male consumers and a few households with a single male parent). The only exception occurs when we use wage figures for all women. 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 the all women sample. In 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 five-year bands. The oldest cohort contains persons born between March 1925 and February 1930, where the cutoff month is chosen to maximize conformity between the CEX and CPS.¹² We consider four educational attainment categories: less than twelve years of schooling, exactly twelve years, more than twelve but less than sixteen years, and sixteen or more years of schooling. Our synthetic panel groups result by crossing these four education categories with the five-year birth cohorts.

C. Consumption and Wage Measures

Our consumption measure equals household expenditures on nondurable goods and services. We exclude expenditures on durables, health, education and housing. The main motivation for excluding these components is to avoid treating dynamics and other problems connected with durability. In addition, the CEX includes only out-of-pocket health expenditures, not insurance payments. We deflate 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 that we effectively applied the same deflator to all groups.¹³

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

¹² 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 three-month intervals in the CEX. Uncertainty about the exact date of birth in the CEX implies a theoretical rate of misallocation to five-year birth cohorts of approximately one percent under our procedures. In practice, missing and erroneous age responses in the CEX generate a larger misallocation rate.

¹³ The standard deviation across groups of the 1980-to-1990 log change in the price indexes is about one-half of one percent.

those with incomplete income responses.¹⁴ 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, consumption observations 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 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 GDP deflator for personal consumption expenditures. We excluded persons who were students or in the military for at least part 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% 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. Namely, the CPS enables us to construct wage regressors with much smaller sampling error than we could obtain from CEX-based wage measures. The sampling error component of the CEX-based consumption measures is a less serious concern, because consumption is the dependent variable in our regression specifications.

VI. Relative Wage and Consumption Behavior

¹⁴ It is standard practice among BLS statisticians to exclude CEX observations with incomplete income responses, when computing means, on the basis of data quality.

A. Movements in Relative Wages

Table 2 summarizes movements in men's real hourly wages by birth cohort and education group.¹⁵ 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. Several types of information about the wage structure can be read from the table. Looking across a single row 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 twelve years of schooling suffered a real wage decline of eleven 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 twelve years of schooling, real wages fell between 1980 and 1990 for all birth cohorts. Real wages declined by about ten percent for cohorts that satisfy the age restrictions throughout the decade. Real wages also declined over the decade for all cohorts of high school educated men, although the declines were typically more modest. Among men with some post-secondary education, real wages increased substantially for the youngest cohorts and modestly for older cohorts. Finally, among college-educated men real wages rose for all cohorts, including rapid real wage gains for the youngest cohorts. To pick out an extreme comparison, 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 twelve years of schooling.

Table 2 shows more modest between-cohort movements in relative wages within most education groups. However, among men with more than twelve years of schooling, the younger cohorts experienced notably more rapid wage growth than the older cohorts. The between-cohort variation in wage growth is especially pronounced among college educated men.

Patterns of relative wage movements among birth-education groups of women are

¹⁵ 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.

broadly similar during the 1980s. Table 3 highlights this similarity, summarizing the covariation of wages between men and women in the same birth-education group and between husbands and wives. The table reports, for example, regressions of five-year changes in mean wages for men on the corresponding changes in mean wages for women in the same birth-education group. The men/women regressions classify men and women into groups based on their separate characteristics, whereas the husbands/wives regressions classify both sexes into groups based on the husband's characteristics. Table entries in parentheses report results for specifications that include cohort fixed effects and thereby isolate the between-education component of the covariation in men's and women's wage movements. The large positive slope coefficients and the high R^2 values in the table bear out the strong similarity of relative wage movements among men and women during both halves of the decade.

In summary, wages for highly educated men rose sharply during the 1980s relative to wages of their less educated contemporaries. These relative wage movements involved real wage increases for men with a college education and real wage declines for men with twelve or fewer years of schooling. 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. Finally, the wage growth patterns among men were largely reinforced by the relative wage movements among women and wives. The large relative wage movements among birth-education groups of men, and the reinforcing pattern of movements among women, indicate that our synthetic panel data offer ample leverage for testing the consumption insurance hypothesis and estimating the effects of systematic relative wage movements on the household consumption distribution.

B. Movements in Relative Consumption

Table 4 summarizes movements in real household consumption by birth cohort and educational attainment of the male head. Although the consumption data are notably noisier than the wage data, some key patterns are discernible. 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 5 condenses the information on real consumption growth by educational attainment status. This table reports simple means of the cohort-specific real consumption changes in Table 4. The results reveal 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% 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.

Although not shown here, we have also examined year-to-year movements in real household consumption by cohort-education groups. The more detailed examination reveals that consumption declines among the less educated are concentrated in the early 1980s. After 1983 or 1984, consumption levels appear fairly flat for most cohorts of the least educated and slowly rising for most cohorts of the high school educated. Among the college educated, consumption by the younger cohorts rises strongly throughout the decade.

These descriptive results constitute a *prima facie* case against the consumption insurance hypothesis. While we have not placed household consumption on an adult equivalence basis, it is hard to believe that *exogenous* differences in the growth of household consumption requirements explain the sharp differences in consumption growth rates by education category. It is also noteworthy that real household consumption falls absolutely for some groups while rising absolutely for others. This pattern, whereby consumption moves in opposite directions for different groups, points towards the likely untenability of the consumption insurance hypothesis for any stable parametrization of preferences under which agents care about consumption risk.

C. Comovements between Relative Wages and Consumption

Figure 1 plots the 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.

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 ten-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% of the considerable variation 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 endowment shocks. Comparing Figures 1 and 2, it is the persistent movements in relative endowments (i.e., low frequency relative wage variation) that drives relative consumption movements.

Figure 3 shows a scatterplot that corresponds to the levels specification (12). The plotted values are residuals from regressions on a quartic polynomial in age plus a maximally linearly independent set of year and group fixed effects. Given that these residuals represent deviations from group-year means for an (unbalanced) panel that contains as many as eleven annual observations per group, we also interpret this scatterplot as de-

picting the low-frequency covariation between relative wages and consumption. The least squares regression line in Figure 3 exhibits a large, positive slope and is tightly estimated, thereby confirming the failure of the consumption insurance hypothesis in Figure 2.

We know from Figure 2 that the covariation in relative wages and consumption across education groups generates a sharp rejection of the consumption insurance hypothesis. Figure 4 shows that the between-cohort covariation within education groups also generates a sharp rejection of the hypothesis. In this figure, we plot by education group the same residuals that appear in Figure 3.¹⁶ 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 birth-education groups. Our evidence is 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 inclusion of additional controls, women's wages, nonseparable preferences, alternative samples, and instrumental variables estimation.

VII. Econometric Results

A. The Synthetic Panel Specifications

This section reports the results of estimating several versions of equations (12) and (14) using synthetic panel data. We typically begin with a benchmark specification that contains year effects, a polynomial in age, plus group fixed effects for the levels specifica-

¹⁶ That is, we constrain the age polynomial to be the same for all education groups. This procedure reflects our belief in small exogenous differences across education groups in the shape of life-cycle consumption requirements.

tions. We then consider alternative specifications that add wage and leisure measures for women (or wives) and controls for family size and composition. We report OLS and IV results for a sample of all households with a 23-to-59 year old male head and for a sample of married couples only. In addition to the levels specification (12), we estimated k -year difference specifications (14) for all k from 1 to 9. We report a selected set of the difference specifications to illustrate the general pattern of results.

The wage measure for women represents an alternative source of endowment shocks to the household. The leisure controls for women reflect our concern for a potentially important nonseparability between household consumption and the leisure of women or wives. 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.

Our family size and composition controls are the (mean of) log family size, the number of adults, the number of children under three years of age, and the number of other children. These variables are 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 using the schemes described in section IV.B. 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

Table 6.1 contains results for the level specifications. The first two lines of Table 6.1 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 IV estimates are considerably larger – 0.81 and 0.59.¹⁷ 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 lines add female wage measures. As expected from Table 3, the collinearity between male and female wage movements inflates the standard errors on the individual slope coefficients, especially under IV estimation. The OLS results reiterate the benchmark results, but the individual IV coefficients are too imprecisely estimated to draw any inferences.

Lines 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 preference nonseparability, but inclusion of this control does not mitigate the size or statistical significance of male wage effects.

Finally, the remaining lines 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 IV estimates for this coefficient are less precise when we include the composition variables.

C. Results for the difference specifications

Tables 6.2-6.4 report results for the one-year, five-year and eight-year differencing intervals. The results for two-year differencing intervals, like the results for one-year intervals, show no consistent effects of relative wage movements on relative consumption. For

¹⁷ The IV results are virtually unaffected when we use only lagged wage values to construct instruments.

differencing intervals of three or more years, we find sizable and statistically significant departures from the consumption insurance hypothesis, as illustrated by tables 6.3 and 6.4. 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 IV 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 three or more years closely parallel results for the level specifications.

D. Can Nonseparability Salvage the Consumption Insurance Hypothesis?

As stressed in section III, 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 6 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 the separability assumption.

We proceed by considering nonseparable preference specifications of the form (5'), which can easily be generalized to account for female leisure. This specification is quite simple and probably inadequate for capturing many important features of the interaction between consumption and labor supply. However, we aim only 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 0.3 and 0.6 and an uncompensated wage elasticity (evaluated at sample means) between -0.3 and 0.3.¹⁸ Corresponding values for the compensated elasticity lie between 0.21 and 0.82. These figures are roughly consistent with available evidence

¹⁸ 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

(Pencavel, 1986).

Given values for γ and ϕ , we used the implied marginal utility of consumption, $\log(C_i^j) + \frac{1+\phi}{\gamma} \log(\bar{H} - L_i^j)$, as the dependent variable in regression specifications otherwise identical to the ones reported in Tables 6. The results were highly similar to Tables 6.1-6.4, and they are available upon request. From these 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, 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 greater supply elasticity, we also considered samples further restricted to exclude groups with fewer than twelve years of schooling.

Tables 7 reports results for a 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.) Due to smaller elasticity. These elasticities satisfy

$$\xi = \frac{M_c H - 1}{M_L - w M_c} \frac{w}{H},$$

$$\eta = -\frac{1}{M_L - w M_c} \frac{w}{H}, \quad \text{and}$$

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

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 $c = 23,000$, $w = 11$, $H = 2080$, and $L = 5000 - 2080$. For instance, $\gamma = -3$, $\phi = -4$ implies $\beta = -0.037$, $\eta = 0.55$ and $\psi = 0.5$.

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. Even the annual difference specifications often show statistically significant evidence against the null hypothesis when estimated by instrumental variables.

E. Assessing the Comovement between Relative Wages and Consumption

We interpret the results in Tables 6 and 7 as a spectacular rejection of the hypothesis of consumption insurance against publicly observable relative wage movements. Regardless of how one interprets these results, it is useful to precisely quantify the fraction of relative consumption variation explained by relative wage variation. Because of sampling errors in the synthetic panel variables, R^2 values in the fitted regressions lie below one, even when the true relationship between consumption and wages exhibits a perfect fit. To address this difficulty, we estimate an upper bound on the R^2 value in the presence of sampling error. We then compare the actual R^2 values for the synthetic panel regressions to the estimated upper bounds to assess the closeness of the consumption-wage relationship.

Suppose that (classical) measurement error and sampling variation are the only source of the equation error in regressions (12) and (14). Suppose further, for the sake of simplicity, that these errors are homoscedastic. Under these conditions, the R^2 value in the level specification is given by

$$R^2 = 1 - \frac{\sigma_v^2}{\sigma_v^2 + \sigma_c^2} \quad (21)$$

where

$$\sigma_v^2 = \beta^2 \frac{\sigma_x^2}{N_x} + \frac{\sigma_y^2}{N_y}, \quad (22)$$

σ_y and σ_x are the within-cell variances of log consumption and X_{it} in equation (12), N_y and N_x are the cell sizes of the samples used to compute average consumption and average X , and σ_c is the variance of log consumption *across cells* not explained by the regression controls. Equation (22) reflects the orthogonality of the measurement errors in consumption and wages, as implied by their construction from independent samples. For

the difference specifications theoretical bounds can be derived in a similar fashion, but taking into account the MA structure induced by the measurement error in the levels.

Table 8 reports actual and estimated upper bounds on R^2 for the simplest specifications that include only male wages, after regressing out the effects of the age polynomial and year effects plus group effects for the level specification. The actual R^2 values equal the fraction of residual consumption variation accounted for by men's wages. The upper bound values are given by (21) and (22) for the level specifications and by their analogs for the difference specifications. We compute (21) and (22) separately for each cell, evaluating at the IV estimates of β (reported in Table 6) and the cell-specific values of σ_x/N_x and σ_y/N_y . We compute σ_c as the cross-cell variance in the consumption measure after regressing out the controls. The upper bound values reported in Table 8 are averages of the cell-specific values.

The results in Table 8 make two points. First, sampling variation reduces the attainable R^2 value by 3-18%, depending on the length of the differencing interval. Thus sampling variation in the CPS- and CEX-based means induces a nontrivial, but modest, deterioration in the potential goodness of fit. Second, male relative wage movements alone explain a large fraction of variation in relative household consumption movements at longer differencing intervals. The table suggests that, over the decade as a whole, variation in men's pre-tax hourly wages across cohort-education groups was the single most important factor driving changes in the cross-group distribution of household consumption (for the population encompassed by our selection criteria).

VIII. Welfare Implications

How large are the welfare losses implied by differences among cohort-education groups in consumption growth rates? How large are the welfare losses implied by the failure to insulate the consumption distribution from relative wage shifts among cohort-education groups? In this section, we propose and implement a procedure for addressing these questions. Our welfare calculations do not reflect uninsured events that occur prior to 1980 or prior to age 23. Nor do they reflect idiosyncratic components of consumption and

Appendix B: Estimating the Covariance Matrix

The asymptotic variance-covariance matrix of $\hat{\beta}$ in equation (12) estimated by instrumental variables is given by

$$AsyV(\hat{\beta}) = P_{xz} X' Z (Z' Z)^{-1} Z' \Omega Z (Z' Z)^{-1} Z' X P_{zz}, \quad (B.1)$$

where Z is the matrix of instruments, X is the regressor matrix, and $\Omega = E[\epsilon' \epsilon]$.

The dimension of Ω is $\sum_{j=1}^N T_j$, where T_j is the number of periods over which group j is observed. The structure of the residuals implies that Ω is block diagonal. Blocks on the main diagonal represent the variance-covariance matrix of the residuals of each group.²¹ These matrices are diagonal for the level specification. They include nonzero entries along the main diagonal and the two k th bands around the main diagonal for the k -year difference specification.

Our estimator for the Ω matrix accommodates heteroscedasticity of unknown form. The estimator for the variance-covariance matrix is obtained by substituting $Z' \Omega Z$ with $P_0 + P_1$, where

$$P_0 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{\tau}^{T_j} z_{j,\tau} z'_{j,\tau} u_{j,\tau}^2, \quad \text{and}$$

$$P_1 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{t=k}^{T_j} [z_{j,t} z'_{j,t-k} u_{j,t} u_{j,t-k} + z_{j,t-k} z'_{j,t} u_{j,t} u_{j,t-k}].$$

T_j is the number of periods over which group j is observed, N is the number of groups, and k equals the differencing interval. $P_1 = 0$ for the level specification.

The main problem with this formula is that it does not guarantee that the estimated $Z' \Omega Z$ is positive definite. We divide the second sum in the expression for P_1 by T_j rather than by $T_j - k$ to alleviate this problem. This sort of problem commonly arises in the estimation of variance-covariance matrices that allow for autocorrelation. Following Fuller's (1985) suggestion, we downweight the component P_1 of the variance-covariance matrix. Our procedure is asymptotically equivalent to the uncorrected one (as T_j increases).

²¹ Ω is block diagonal under the assumption that the residuals are contemporaneously uncorrelated across groups. This assumption holds under the null, if shocks to preferences are uncorrelated across groups.

first year such that $\chi_{it} = 1$. The term ψ_{t+1} reflects the growth in aggregate consumption between t and $t + 1$. It satisfies $\psi_{1980} = 1$ and

$$\sum_{i=1}^I \omega_{i,t+1} C_{i,t+1} = \sum_{i=1}^I \omega_{i,t+1} C_{i,t+1}^*$$

$$= \sum_{i=1}^I \omega_{i,t+1} \left[((1 - \chi_{it}) + \chi_{it}(1 - \chi_{i,t+1})) C_{i,t+1} + \chi_{it} \chi_{i,t+1} \psi_{t+1} \left(\frac{b_{it}}{\rho b_{i,t+1}} \right)^{1/\gamma} C_{it}^* \right].$$

In words, ψ_{t+1} is a scaling factor on the consumption of continuing groups chosen to equate actual and optimal per capita consumption. This procedure effectively uses each group's initial relative consumption to set the Pareto weights in the planner's problem.

Next, define the welfare level achieved under the optimal allocation,

$$V^*(C^*) = U(C^*, 1).$$

We can now calculate the consumption variation required to compensate households for the failure to insulate the consumption distribution from *all* sources of group-level consumption risks as the value of $\bar{\lambda}^O$ that satisfies

$$U(C, \bar{\lambda}^O) = V^*(C^*).$$

Likewise, we can calculate the consumption variation required to compensate households for the failure to insure against relative wage variation among groups as the value of $\bar{\lambda}^W$ that satisfies

$$U(\hat{C}, \bar{\lambda}^W) = V^*(C^*).$$

Here, $(\bar{\lambda}^O - 1)$ denotes the uniform (across groups, dates and states) percentage increase in actual consumption that would be required to bring average welfare to the same level achieved by the optimal allocation. Similarly, $(\bar{\lambda}^W - 1)$ denotes the uniform percentage increase in consumption required to compensate households for the failure to insulate the consumption distribution from relative wage movements.

In carrying out these 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. If one posits

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

then a regression of the form

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

delivers the estimated taste shifters $\hat{b}_{it} = \exp \left[-\gamma \sum_{j=1}^4 \hat{\delta}_j (\text{age}^j)_{it} \right]$. As anticipated, these estimated consumption requirement profiles are concave with respect to age.

To randomize over alternative paths, let \check{C} be the consumption path obtained by reversing the observed relative movements: $\check{C} = C - (C^* - C)$. This path implies the same aggregate consumption but alters the distribution, so that winners become losers and vice-versa. Finally, define λ^O by

$$0.5[U(C, \lambda^O) + U(\check{C}, \lambda^O)] = V^*,$$

and analogously for λ^W . Thus, the compensating consumption variations, $\lambda^O - 1$ and $\lambda^W - 1$, involve a randomization over two alternative consumption growth paths: the one that actually occurred during the 1980s, and the "opposite" path.

Figure 5 displays the computed values of λ^O and λ^W for various coefficients of relative risk aversion. The figure reveals large welfare losses associated with the failure to insulate the household consumption distribution from group-specific endowment variation. As the coefficient of relative risk aversion ranges from -9.5 to -0.5, the consumption variation required to compensate households from all sources of group-level risk ranges from a staggering 24.8% to .6%. The consumption variation required to compensate households for the risk associated with men's pre-tax relative wage movements ranges from 10.8% to .4%. For a coefficient of relative risk aversion equal to two, the compensating consumption variations equal 2.7% for all shocks and 1.8% for men's relative wage movements only.

Thus, for relative risk aversion equal to two, a household with an annual consumption of \$40,000 would have to be compensated by more than \$1000 (in each period) to achieve the same level of welfare as implied by full insurance against group-level endowment variation. This calculation ignores any potential gains associated with pooling consumption risks within groups. The large welfare losses reported in Figure 5 indicate that there are powerful impediments to insurance against publicly observable endowment variation among cohort-education groups.

IX. 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, and 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 conditioning on women's leisure, household size and composition, and samples restricted to men with inelastic labor supply.

In our view, the magnitude of the covariance between relative wages and consumption constitutes a spectacular failure of the consumption insurance hypothesis. This hypothesis is not even remotely consistent with the evidence developed here. Our calculations suggest that the observed departures from optimal consumption allocations carry large welfare costs. In addition, our evidence against the consumption insurance hypothesis involves publicly observed endowment 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 across 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 endowment shocks. The link between consumption and observable endowment 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 across 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 towards 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 birth-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 one- and two-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 observationally distinct workers, carry unimportant welfare consequences.

To close, we remark upon one other direction for future research. Many advanced and middle income economies experienced large relative wage movements among observationally 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 consumption insurance. 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 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.

Appendix A: A Model of Leisure Transferability via Market Transactions

Consider a model in which time-saving goods allow for the transferability of leisure across households. Represent the underlying preferences over true consumption and leisure by the strictly concave function $U(C, L)$. Household j faces a time constraint

$$L^j = g^j(\bar{H} - H^j, G^j),$$

where \bar{H} denotes the time endowment, H^j denotes hours of market work activity, and G^j denotes expenditures on time-saving goods produced in the market. The function $g^j(\bar{H} - H^j, G^j)$ represents the j th household's technology for transforming nonmarket time and expenditures on time-saving goods into effective leisure, L .

We adopt the following assumptions regarding the household's technology for producing leisure: (i) Effective leisure increases with nonmarket time and with expenditures on time-saving goods. (ii) Nonmarket time and time-saving goods are substitutes in the production of leisure, so that $g_{12} \leq 0$. (iii) The rate at which time-saving goods can be substituted for nonmarket time diminishes with greater expenditures on these goods, so that iso-leisure curves are convex. (iv) $g(\bar{H} - H, G) \leq \bar{H} - H$ for all G .

Market goods C and G are produced from time inputs supplied by I types of workers. The aggregate production function exhibits diminishing marginal factor returns and is separable across factor types. Hence, we can write the "wage" function for individual j in group i as $W^j(H^i, S^t)$, where $H^i \equiv \sum_{j \in i} H^j$, and $dW^j/dH^i < 0$ for all j and i . The dependence of the wage function on S^t captures technological or other disturbances that drive relative wage movements.

Under these assumptions, the Pareto problem becomes

$$\max_{C^j(S^t), H^j(S^t), G^j(S^t)} \sum_{j=1}^J \lambda^j \sum_{t=1}^{\infty} \sum_{S^t} \rho^t \pi(S^t) U[C^j(S^t), g^j(\bar{H} - H^j(S^t), G^j(S^t))],$$

subject to the aggregate feasibility constraint,

$$\sum_j H^j(S^t) W^j(H^i, S^t) \geq \sum_j [C^j(S^t) + G^j(S^t)], \quad \text{for all } S^t, \text{ and}$$

$$C^j \geq 0, \quad 0 \leq H^j \leq \bar{H}, \quad \text{and } G^j \geq 0 \quad \text{for all } j \text{ and } S^t.$$

If households have identical technologies for producing leisure, the first-order conditions for an interior solution to this problem become

$$H^j : \quad \rho^t U_L(C^j, L^j) g_1(\bar{H} - H^j, G^j) = W^j(H^i, S^t) \mu_t \lambda^j, \quad (\text{A.1})$$

$$G^j : \quad \rho^t U_L(C^j, L^j) g_2(\bar{H} - H^j, G^j) = \mu_t \lambda^j, \quad (\text{A.2})$$

$$C^j : \quad \rho^t U_C(C^j, L^j) = \mu_t \lambda^j, \quad (A.3)$$

for all j and at all dates and states. In particular, at an interior solution for hours worked and time-saving goods, each household satisfies the tangency condition,

$$\frac{g_1(\bar{H} - H^j, G^j)}{g_2(\bar{H} - H^j, G^j)} = W^j(H^j, S^t). \quad (A.4)$$

It follows from convexity of the iso-leisure curves that the ratio of nonmarket time to expenditures on time-saving goods moves inversely with the wage.

This analysis delivers three useful insights. First, it is apparent that the technology for transferring leisure across households induces a nonseparability between measured consumption expenditures, $C + G$, and measured leisure, $\bar{H} - H$, even when the true underlying preferences are separable.

Second, the optimal allocation entails complete consumption insurance in the sense that the marginal utility of consumption, U_C , is equated across households. Thus, the imperfect transferability of leisure across households does not upset the consumption insurance property of the optimal allocation.

Third, although the optimal allocation retains the consumption insurance property, relative levels of *measured* consumption covary positively with relative wages. To see this point, consider an increase in j 's (relative) wage that leaves aggregate production and μ_t unchanged. The wage increase raises the (opportunity) cost of producing leisure for j , so that L_j declines and U_L^j rises. Both the scale effect associated with the decline in L_j and the substitution effect reflected in (A.4) imply a decline in $\bar{H} - H^j$ and, hence, an increase in j 's hours of work.¹⁹ In addition, a substitution effect ($g_{12} \leq 0$) and a price effect (rise in U_L^j) both increase j 's optimal allocation of time-saving goods. Thus, relative levels of measured consumption expenditures, $C + G$, covary positively with relative wages under the consumption insurance allocation.²⁰

¹⁹ Condition (A.1) is directly analogous to the wage-equals-marginal-revenue-product condition for a firm with a downward sloping product demand schedule. Leisure corresponds to output, U_L corresponds to the demand schedule, and nonmarket time corresponds to the factor input subject to a price increase.

²⁰ If C and L are utility complements, then the decline in L induces a decline in C under the optimal allocation. The conclusion in the text presumes that such an effect, if present, is smaller than the rise in G . Such a high degree of complementarity between C and L strikes us as implausible. In any case, extreme complementarity between C and L would create a bias towards acceptance of the consumption insurance hypothesis in our empirical work. It cannot account for our rejections of the hypothesis.

Appendix B: Estimating the Covariance Matrix

The asymptotic variance-covariance matrix of $\hat{\beta}$ in equation (12) estimated by instrumental variables is given by

$$AsyV(\hat{\beta}) = P_{xz} X' Z (Z' Z)^{-1} Z' \Omega Z (Z' Z)^{-1} Z' X P_{zz}, \quad (B.1)$$

where Z is the matrix of instruments, X is the regressor matrix, and $\Omega = E[\epsilon' \epsilon]$.

The dimension of Ω is $\sum_{j=1}^N T_j$, where T_j is the number of periods over which group j is observed. The structure of the residuals implies that Ω is block diagonal. Blocks on the main diagonal represent the variance-covariance matrix of the residuals of each group.²¹ These matrices are diagonal for the level specification. They include nonzero entries along the main diagonal and the two k th bands around the main diagonal for the k -year difference specification.

Our estimator for the Ω matrix accommodates heteroscedasticity of unknown form. The estimator for the variance-covariance matrix is obtained by substituting $Z' \Omega Z$ with $P_0 + P_1$, where

$$P_0 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{\tau}^{T_j} z_{j,\tau} z'_{j,\tau} u_{j,\tau}^2, \quad \text{and}$$

$$P_1 = \frac{1}{N} \sum_{j=1}^N \frac{1}{T_j} \sum_{t=k}^{T_j} [z_{j,t} z'_{j,t-k} u_{j,t} u_{j,t-k} + z_{j,t-k} z'_{j,t} u_{j,t} u_{j,t-k}].$$

T_j is the number of periods over which group j is observed, N is the number of groups, and k equals the differencing interval. $P_1 = 0$ for the level specification.

The main problem with this formula is that it does not guarantee that the estimated $Z' \Omega Z$ is positive definite. We divide the second sum in the expression for P_1 by T_j rather than by $T_j - k$ to alleviate this problem. This sort of problem commonly arises in the estimation of variance-covariance matrices that allow for autocorrelation. Following Fuller's (1985) suggestion, we downweight the component P_1 of the variance-covariance matrix. Our procedure is asymptotically equivalent to the uncorrected one (as T_j increases).

²¹ Ω is block diagonal under the assumption that the residuals are contemporaneously uncorrelated across groups. This assumption holds under the null, if shocks to preferences are uncorrelated across groups.

References

- Abel, Andrew and Kotlikoff, Laurence J. (1992) "Intergenerational Altruism and the Effectiveness of Fiscal Policy - New Tests Based on Cohort Data," forthcoming in Toshiaki Tachibanaki, ed., *Savings and Bequests*, University of Michigan Press.
- Altonji, Joseph G., Hayashi, Fumio and Kotlikoff, Laurence J. "Is the Extended Family Altruistically Linked?" *American Economic Review*, 82, no .5 (December 1992), 1177-1198
- Altug, Sumru and Miller, Robert A. "Household Choices in Equilibrium." *Econometrica*, 58, no. 3 (May 1990), 543-570.
- Atkeson, Andrew G. and Lucas, Robert E. Jr. (1992) "On Efficient Distribution with Private Information," *Review of Economic Studies*,
- Attanasio, Orazio and Weber, Guglielmo (1992) "Consumption Growth and Excess Sensitivity to Income: Evidence from U.S. Micro Data," manuscript, Stanford University.
- Card, David "Comments and Discussion" to Cutler and Katz (1991) *Brookings Papers on Economic Activity*, (1991:2), 62-68.
- Card, David and Richard Freeman (1993) *Small Differences that Matter: Labor Markets and Income Maintenance in Canada and the United States*, Chicago: The University of Chicago Press for the National Bureau of Economic Research.
- Cochrane, John H. "A Simple Test of Consumption Insurance." *Journal of Political Economy*, 99, no. 5 (October 1991), 957-976.
- Cutler, David M. and Katz, Lawrence F. "Macroeconomic Performance and the Disadvantaged," *Brookings Papers on Economic Activity*, (1991:2), 1-74.
- Davis, Steven J. "Cross-Country Patterns of Change in Relative Wages," *NBER Macroeconomics Annual*, 7 (1992), 239-292.
- Diamond, Peter (1967) "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," *American Economic Review*, 57, 759-776.
- Fuller, Wayne A. (1985) : *Measurement Error Models* , New York: John Wiley.
- Goldin, Claudia and Margo, Robert "The Great Compression: The Wage Structure in the United States at Mid-Century." *Quarterly Journal of Economics*, 107, no 1 (February 1992), 1-34.
- Green, Edward (1987) "Lending and the Smoothing of Uninsurable Income," in Prescott and Wallace, eds., *Contractual Arrangements for Intertemporal Trade*. Minneapolis: University of Minnesota Press.
- Juhn, Chinhui, Kevin M. Murphy and Robert H. Topel (1991) "Why Has the Natural Rate of Unemployment Increased over Time?" *Brookings Papers on Economic Activity*, 1991:2, 75-142.
- Katz, Lawrence F. and Murphy, Kevin M. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors," *Quarterly Journal of Economics*, 107, no. 1 (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" *Journal of Economic Literature*, 30, no. 3 (September 1992), 1333-1381.
- Lewis, Karen (1993): "What Can Explain the Apparent Lack of International Consumption Risk Sharing?", Manuscript, University of Pennsylvania.
- Mace, Barbara J. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy*, 99, no. 5 (October 1991), 928-956.
- Mayer, Susan E. and Christopher Jenks (1991) "Recent Trends in Economic Inequality in the United States: Income vs. Expenditures vs. Material Well-Being," unpublished.
- Nelson, Julie (1994) "On Testing Full Insurance Using Consumer Expenditure Survey Data," *Journal of Political Economy*, 102, no. 2 (April), 384-94.
- Pencavel, John (1986) "Labor Supply of Men: A Survey," in Orley C. Ashenfelter and Richard Layard, eds., *Handbook of Labor Economics, I*, Amsterdam: Elsevier Science Publishers.
- Phelan, Christopher and Townsend, Robert "Computing Multiperiod Information-Constrained Optima." *Review of Economic Studies*, (1991).
- Scheinkman, Jose (1984) "General Equilibrium Models of Economic Fluctuations: A Survey of Theory," University of Chicago.
- Townsend, Robert M. (1988) "Information Constrained Insurance: The Revelation Principle Extended," *Journal of Monetary Economics*, 21, no. 2/3 (March/May), 411-450.
- Townsend, Robert M. (1994) "Risk and Insurance in Village India," *Econometrica* 62, no. 3 (May), 539-591.
- White, Halbert (1980) "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica*, 48, 817-838.
- Wilson, Robert (1968) "The Theory of Syndicates," *Econometrica*, 36, 1 (January), 119-132

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	2291
Wages. Women	CPS	420	762	192	2056
Wages. Husbands	CPS	420	695	190	1576
Wages. Wives	CPS	420	509	148	1166
Consumption. All	CEX	288	499	137	1071
Consumption. Married	CEX	288	414	117	896

Notes:

- (1) Groups are defined by crossing five-year birth cohorts with four educational attainment categories.
- (2) 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. CPS samples of wives are restricted to women with husbands between 23 and 59 years of age. CEX samples are restricted to households with a male head or husband of female head between 23 and 59 years of age. See the text for other selection criteria.
- (3) The number of admissible cells and the cell count summary statistics for CPS (CEX) data pertain to the 1975 (1980) to 1990 sample period. The cell count equals the number of nonmissing observations for the indicated variable type.

Table 2

Real Wage Movements for Men by Birth Cohort and Education Group Deviations
from the 1980 Value for Own-Education Group and 1945-50 Cohort

Educational Attainment = Less than High School (0-11 Years)

Year	Five-Year Birth Cohort								
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30	20-25
1975				-0.06	0.06	0.09	0.15	0.18	0.14
1980			-0.07	0.00	0.08	0.10	0.15	0.14	
1985		-0.18	-0.06	-0.06	0.01	0.06	0.09		
1990	-0.29	-0.19	-0.19	-0.09	-0.03	-0.02			

Educational Attainment = High School

Year	Five-Year Birth Cohort								
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30	20-25
1975				-0.06	0.02	0.04	0.09	0.10	0.10
1980			-0.12	0.00	0.04	0.07	0.13	0.13	
1985		-0.21	-0.12	0.00	0.04	0.08	0.06		
1990	-0.31	-0.18	-0.13	-0.07	0.02	-0.03			

Educational Attainment = Some Post-Secondary (13-15 Years)

Year	Five-Year Birth Cohort								
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30	20-25
1975				-0.10	0.05	0.14	0.13	0.12	0.19
1980			-0.15	0.00	0.12	0.09	0.14	0.14	
1985		-0.17	-0.06	0.08	0.09	0.13	0.12		
1990	-0.28	-0.09	-0.02	0.04	0.08	0.12			

Educational Attainment = College (16+ Years)

Year	Five-Year Birth Cohort								
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30	20-25
1975				-0.15	0.10	0.22	0.26	0.26	0.27
1980			-0.18	0.00	0.17	0.19	0.23	0.26	
1985		-0.11	0.03	0.17	0.25	0.29	0.30		
1990	-0.16	0.03	0.14	0.17	0.24	0.23			

Table 3

The Covariance Between Men's and Women's Wage Growth, 1980 to 1990

Regressions Using Five-Year Changes in the Mean Log Wage for Cohort-Education Groups

Time Interval	Sample	Dependent Variable	Regressor	Slope Coeff.	Stand. Error	R ²
1980-85	All	Men	Women	.97(.99)	.22(.20)	.53(.71)
1980-85	Married	Husbands	Wives	.95(1.16)	.29(.24)	.37(.70)
1980-85	All	Women	Men	.54(.64)	.12(.13)	.53(.66)
1980-85	Married	Wives	Husbands	.39(.53)	.12(.11)	.37(.66)
1985-90	All	Men	Women	1.11(.93)	.25(.26)	.52(.70)
1985-90	Married	Husbands	Wives	.78(.80)	.33(.22)	.24(.75)
1985-90	All	Women	Men	.47(.52)	.11(.14)	.52(.60)
1985-90	Married	Wives	Husbands	.31(.62)	(.13)(.17)	.24(.51)

Notes:

- (1) Groups are defined by crossing five-year birth cohorts with four educational attainment classes.
- (2) In the sample of married persons, women are assigned to groups based on the characteristics of their husbands.
- (3) The samples are restricted to groups that satisfy the restriction, $23 \leq \text{age} \leq 59$, for all members over the entire five-year interval.
- (4) The entries in parentheses correspond to a specification that includes cohort fixed effects. The other entries correspond to a specification that includes a constant but no cohort effects.
- (5) Each regression contains twenty observations.

Table 4

Real Household Consumption Movements by Birth Cohort and Education
Households with an Adult Male Head between 23 and 59
Deviations from the 1980 Value for Own-Education Group and 1945-50 Cohort

Educational Attainment = Less than High School (0-11 Years)

Year	Five-Year Birth Cohort							
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30
1980			-0.21	0.00	-0.05	0.09	-0.01	-0.14
1985		-0.51	-0.18	-0.18	-0.17	-0.19	-0.22	
1990	-0.20	-0.37	-0.18	-0.22	-0.23	-0.17		

Educational Attainment = High School

Year	Five-Year Birth Cohort							
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30
1980			-0.17	0.00	0.05	0.18	0.31	0.20
1985		-0.32	-0.16	-0.06	0.09	0.15	0.06	
1990	-0.30	-0.10	-0.06	0.01	0.01	-0.01		

Educational Attainment = Some Post-Secondary (13-15 Years)

Year	Five-Year Birth Cohort							
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30
1980			-0.16	0.00	0.13	0.34	0.31	0.25
1985		-0.21	-0.11	0.06	0.13	0.21	0.10	
1990	-0.21	-0.07	0.00	0.13	0.20	0.12		

Educational Attainment = College (16+ Years)

Year	Five-Year Birth Cohort							
	60-65	55-60	50-55	45-50	40-45	35-40	30-35	25-30
1980			-0.19	0.00	0.25	0.34	0.47	0.42
1985		-0.21	0.01	0.18	0.32	0.36	0.38	
1990	-0.24	0.00	0.17	0.27	0.36	0.39		

Table 5

Consumption Growth by Educational Attainment of Male Head, 1980 to 1990

Simple Averages of Changes in the Mean of Log(Consumption) for Five-Year Cohorts

Time Interval	<i>Educational Attainment of Male</i>			
	< 12 Years	High School	Post HS	College
1980-85	-0.15	-0.06	-0.05	0.07
1985-90	0.01	0.03	0.06	0.10

Notes:

- (1) 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.
- (2) The consumption measure equals household expenditures on nondurable goods and services, as defined in the text.

Table 6.1

Synthetic Panel Regressions - Levels Specification*

Estimation Method	Household Controls**	Households with a 23-59 Year Old Male Head			Married Couples Only		
		Man's Wage	Woman's Wage	Woman's Leisure	Husband's Wage	Wife's Wage	Wife's Leisure
OLS	No	0.653 (0.072)	-	-	0.419 (0.063)	-	-
IV	No	0.814 (0.079)	-	-	0.592 (0.076)	-	-
OLS	No	0.616 (0.094)	0.063 (0.105)	-	0.380 (0.084)	0.070 (0.098)	-
IV	No	1.760 (0.640)	-1.392 (0.941)	-	7.807 (25.534)	-10.036 (35.248)	-
OLS	No	0.712 (0.077)	-	-0.654 (0.315)	0.477 (0.068)	-	-0.633 (0.295)
IV	No	1.068 (0.111)	-	-2.162 (0.450)	0.885 (0.107)	-	-2.321 (0.433)
OLS	Yes	0.476 (0.070)	-	-	0.462 (0.061)	-	-
IV	Yes	0.824 (0.245)	-	-	1.272 (0.665)	-	-
OLS	Yes	0.500 (0.070)	-	-0.347 (0.289)	0.488 (0.061)	-	-0.330 (0.292)
IV	Yes	0.898 (0.245)	-	-1.343 (0.899)	1.250 (0.665)	-	0.914 (4.346)

* Each regression contains 288 observations.

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

Table 6.2

Synthetic Panel Regressions - Annual Difference Specification*

Estimation Method	Household Controls**	Households with a 23-59 Year Old Male Head			Married Couples Only		
		Man's Wage	Woman's Wage	Woman's Leisure	Husband's Wage	Wife's Wage	Wife's Leisure
OLS	No	0.019 (0.157)	-	-	-0.057 (0.124)	-	-
IV	No	0.463 (0.269)	-	-	0.370 (0.235)	-	-
OLS	No	-0.006 (0.161)	0.093 (0.135)	-	-0.114 (0.126)	0.240 (0.124)	-
IV	No	-0.356 (1.882)	1.276 (2.774)	-	0.807 (1.541)	-0.772 (2.578)	-
OLS	No	-0.003 (0.158)	-	0.661 (0.567)	-0.067 (0.124)	-	0.568 (0.520)
IV	No	0.564 (0.310)	-	-1.539 (1.400)	0.444 (0.278)	-	-1.229 (1.075)
OLS	Yes	-0.011 (0.134)	-	-	-0.063 (0.118)	-	-
IV	Yes	0.270 (0.394)	-	-	0.571 (0.271)	-	-
OLS	Yes	-0.028 (0.134)	-	0.521 (0.483)	-0.075 (0.118)	-	0.610 (0.490)
IV	Yes	0.257 (0.394)	-	-0.575 (1.004)	0.447 (0.284)	-	-0.910 (0.160)

* Each regression contains 192 observations.

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

Table 6.3

Synthetic Panel Regressions - Five-Year Difference Specification*

Estimation Method	Household Controls**	Households with a 23-59 Year Old Male Head			Married Couples Only		
		Man's Wage	Woman's Wage	Woman's Leisure	Husband's Wage	Wife's Wage	Wife's Leisure
OLS	No	0.557 (0.115)	-	-	0.485 (0.098)	-	-
IV	No	0.670 (0.133)	-	-	0.572 (0.117)	-	-
OLS	No	0.405 (0.168)	0.237 (0.193)	-	0.333 (0.150)	0.235 (0.178)	-
IV	No	0.913 (0.698)	-0.332 (0.876)	-	1.738 (1.322)	-1.546 (1.685)	-
OLS	No	0.663 (0.130)	-	-0.816 (0.482)	0.644 (0.110)	-	-1.195 (0.421)
IV	No	0.916 (0.180)	-	-1.635 (0.580)	0.903 (0.152)	-	-2.081 (0.509)
OLS	Yes	0.380 (0.116)	-	-	0.528 (0.105)	-	-
IV	Yes	0.445 (0.162)	-	-	0.654 (0.150)	-	-
OLS	Yes	0.426 (0.116)	-	-0.624 (0.462)	0.613 (0.105)	-	-0.813 (0.448)
IV	Yes	0.554 (0.162)	-	-1.108 (0.502)	0.765 (0.150)	-	-1.102 (0.487)

* Each regression contains 108 observations.

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

Table 6.4

Synthetic Panel Regressions - Eight-Year Difference Specification*

Estimation Method	Household Controls**	Households with a 23-59 Year Old Male Head			Married Couples Only		
		Man's Wage	Woman's Wage	Woman's Leisure	Husband's Wage	Wife's Wage	Wife's Leisure
OLS	No	0.699 (0.129)	-	-	0.541 (0.144)	-	-
IV	No	0.701 (0.133)	-	-	0.523 (0.157)	-	-
OLS	No	0.686 (0.286)	0.020 (0.404)	-	0.428 (0.307)	0.173 (0.416)	-
IV	No	1.434 (2.426)	-1.139 (3.796)	-	2.409 (7.153)	-2.763 (10.560)	-
OLS	No	0.910 (0.132)	-	-1.991 (0.625)	0.778 (0.147)	-	-2.131 (0.668)
IV	No	0.976 (0.135)	-	-2.444 (0.668)	0.843 (0.145)	-	-2.507 (0.682)
OLS	Yes	0.869 (0.173)	-	-	0.720 (0.155)	-	-
IV	Yes	0.897 (0.142)	-	-	0.690 (0.131)	-	-
OLS	Yes	0.865 (0.173)	-	-1.385 (0.648)	0.773 (0.155)	-	-1.407 (0.726)
IV	Yes	0.866 (0.142)	-	-1.283 (0.473)	0.737 (0.131)	-	-1.324 (0.564)

* Each regression contains 36 observations.

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

Table 7

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

Estimation Method	Household Controls***	Levels Specification*			Annual Difference Specification**		
		Man's Wage	Woman's Wage	Woman's Leisure	Man's Wage	Woman's Wage	Woman's Leisure
OLS	No	0.764 (0.096)	-	-	0.056 (0.241)	-	-
IV	No	0.981 (0.086)	-	-	1.084 (0.429)	-	-
OLS	No	0.778 (0.104)	-0.050 (0.147)	-	0.035 (0.245)	0.110 (0.230)	-
IV	No	1.278 (0.208)	-1.049 (0.643)	-	0.978 (0.7784)****	0.359 (1.7677)****	-
OLS	No	0.859 (0.104)	-	-0.978 (0.457)	0.014 (0.244)	-	0.884 (0.891)
IV	No	1.310 (0.122)	-	-2.880 (0.648)	1.368 (0.621)	-	-2.452 (1.991)
OLS	Yes	0.556 (0.086)	-	-	0.136 (0.200)	-	-
IV	Yes	1.507 (0.742)	-	-	1.056 (0.576)	-	-
OLS	Yes	0.577 (0.742)	-	-0.206 (0.409)	0.107 (0.200)	-	0.646 (0.736)
IV	Yes	2.412 (0.742)	-	-6.163 (12.845)	1.308 (0.713)	-	-2.724 (2.141)

* Each regression contains 144 observations.

** Each regression contains 93 observations.

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

**** The standard errors for this specification are not corrected for heteroscedasticity or auto correlation because the estimated variance covariance matrix was not positive definite.

Table 7 (continued)

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

Estimation Method	Household Controls***	Five-Year Difference Specification*			Eight-Year Difference Specification**		
		Man's Wage	Woman's Wage	Woman's Leisure	Man's Wage	Woman's Wage	Woman's Leisure
OLS	No	0.457 (0.149)	-	-	0.800 (0.147)	-	-
IV	No	0.768 (0.153)	-	-	0.796 (0.139)	-	-
OLS	No	0.527 (0.184)	0.050 (0.260)	-	0.349 (0.306)	1.034 (0.629)	-
IV	No	1.279 (0.534)	-1.116 (1.087)	-	0.150 (0.602)	1.444 (1.196)	-
OLS	No	0.722 (0.179)	-	-1.107 (0.659)	1.112 (0.170)	-	-2.661 (0.992)
IV	No	1.269 (0.241)	-	-2.683 (1.019)	1.290 (0.132)	-	-3.718 (1.519)
OLS	Yes	0.495 (0.147)	-	-	0.666 (0.265)	-	-
IV	Yes	0.666 (0.143)	-	-	0.699 (0.138)	-	-
OLS	Yes	0.507 (0.147)	-	-0.083 (0.699)	0.908 (0.265)	-	-2.746 (1.035)
IV	Yes	0.976 (0.143)	-	-1.562 (0.865)	0.908 (0.138)	-	-2.746 (0.564)

* Each regression contains 48 observations.

** Each regression contains 18 observations.

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

Table 8
Fraction of Cross-Cell Consumption Variation Explained by Men's Wages

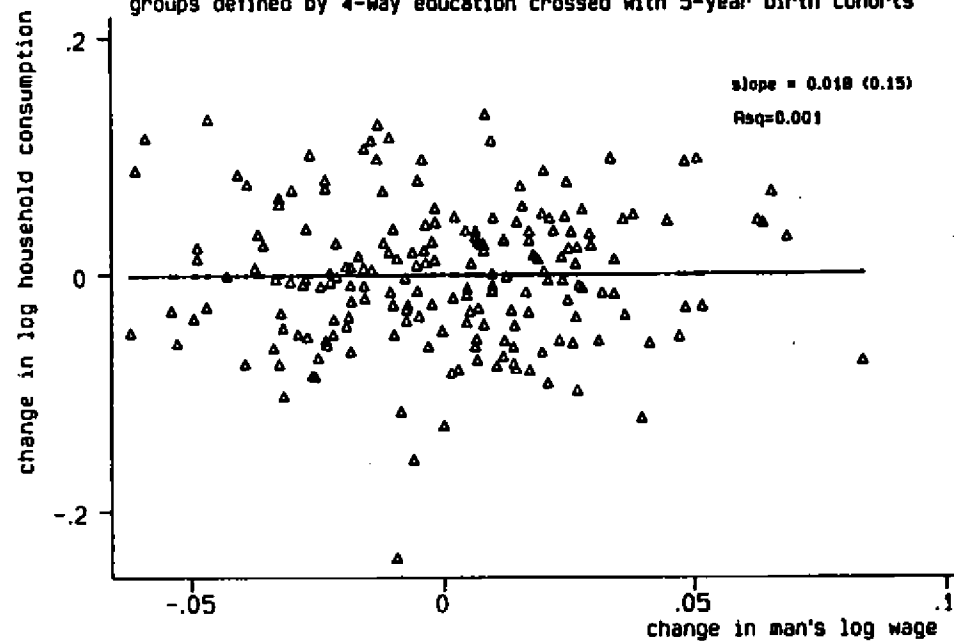
	Length of Differencing Interval									
	0	1	2	3	4	5	6	7	8	9
Actual R^2	.21	.00	.02	.07	.15	.18	.21	.39	.45	.61
Upper Bound	.91	.82	.88	.90	.92	.94	.95	.96	.96	.97

Notes:

- (1) The "actual R^2 " value reports the fraction of residual consumption variation accounted for by men's relative wages. Residuals are obtained from regressions of log consumption (levels or differences) on an age polynomial and year effects plus group effects for the level specifications.
- (2) For a differencing interval of length zero, the "upper bound" values report the results of implementing equations (21) and (22) in the text, evaluating β at the IV estimates reported in Table 6. For differencing intervals greater than zero, upper bound values are computed in a similar fashion that takes into account the moving average structure induced by measurement error in the levels.

Figure 1

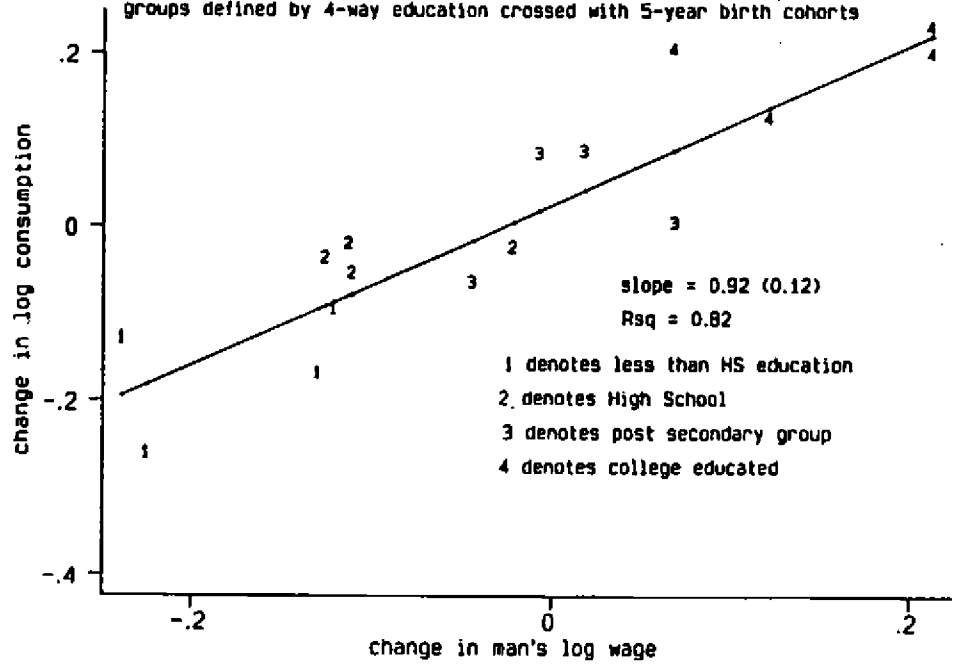
Household Consumption vs. Man's Wage- Annual log change residuals- 1981-1990
groups defined by 4-way education crossed with 5-year birth cohorts



Plotted values are residuals from regressions on a cubic in age and year effects

Figure 2

Household Consumption vs. Man's Wage- 1980-1990 log change residuals
groups defined by 4-way education crossed with 5-year birth cohorts



Plotted values are residuals from regressions on a cubic in age

Figure 3

Household Log Consumption vs. Man's Wage- Annual data from 1980-1990
groups defined by four-way education crossed with 5-year birth cohorts
solid line shows OLS regression with slope 0.63 (0.07), $R^2=0.21$

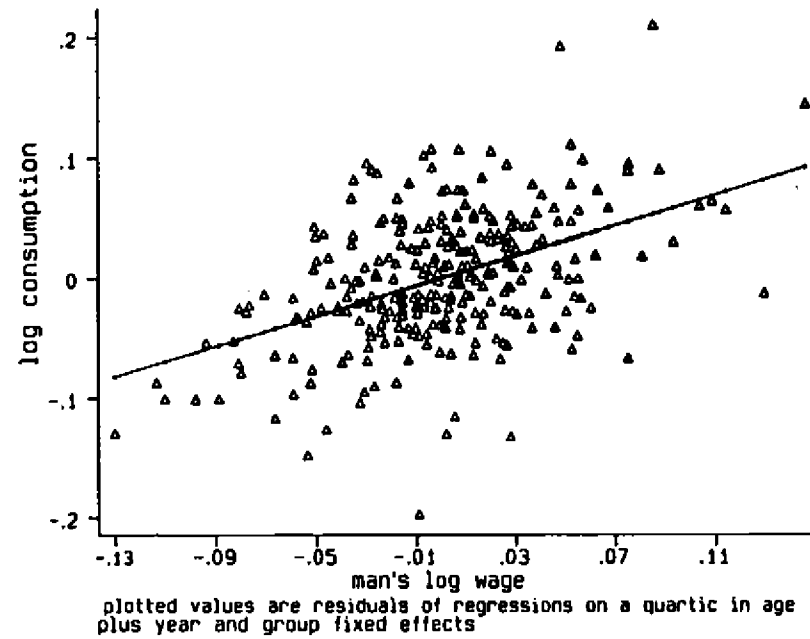


Figure 4

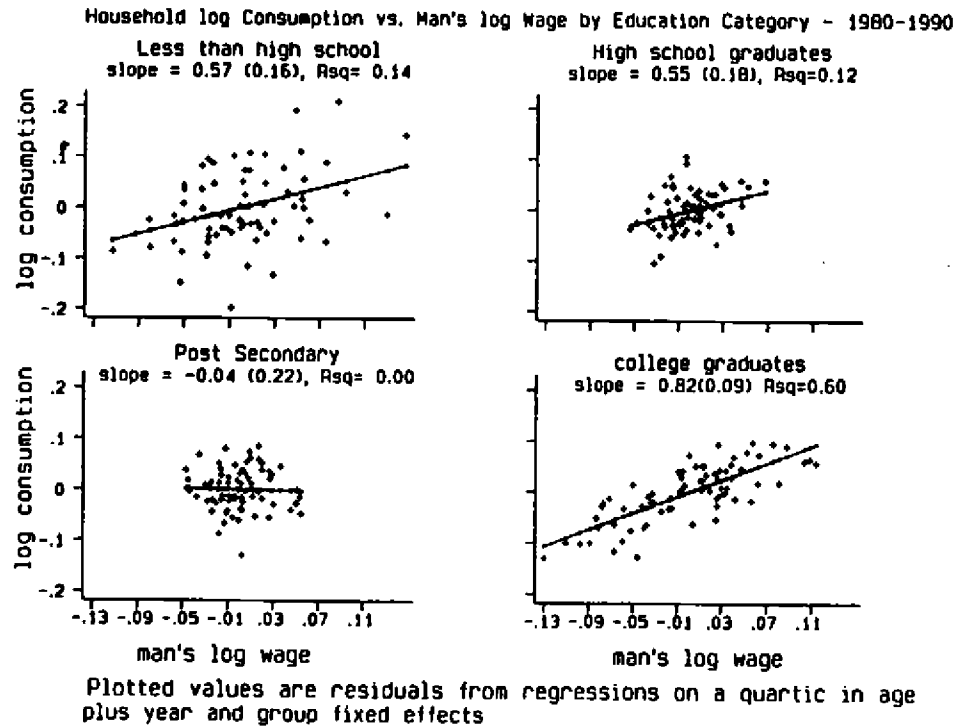
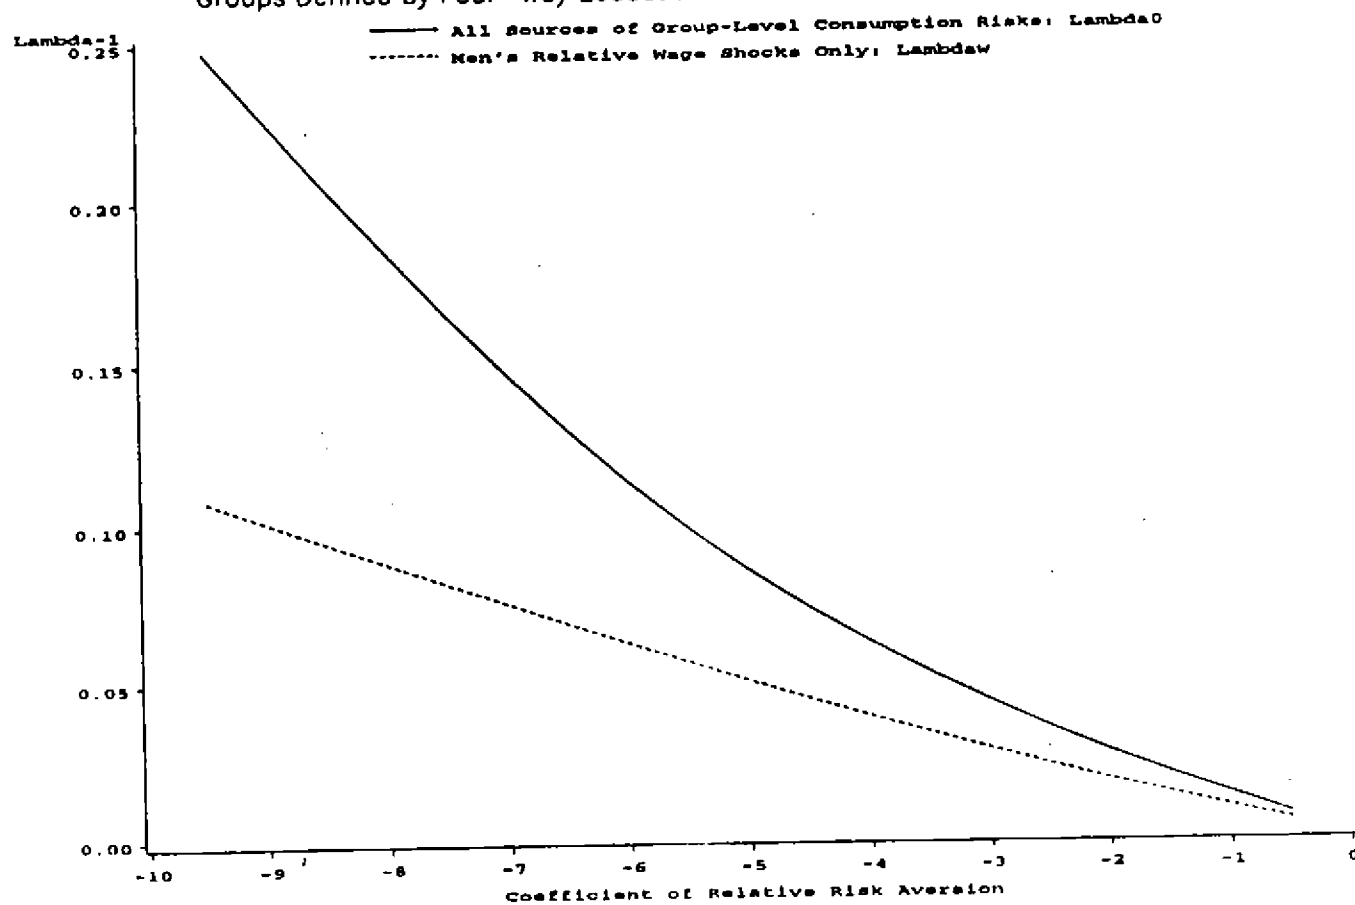


Figure 5
 Welfare Losses Implied by Cross-Group Differences in Marginal Utility Growth
 Groups Defined by Four-Way Education Crossed with 5-Year Birth Cohorts



Plotted values show the uniform percentage increase in consumption required to compensate households for the failure to insure against the indicated shocks.