SUSAN DYNARSKI
*Massachusetts Institute of Technology*

JONATHAN GRUBER
*Massachusetts Institute of Technology*

# *Can Families Smooth Variable Earnings?*

THE LABOR MARKET in the United States is marked by considerable year-to-year variation in individual earnings.[1] In theory, variation in the earnings of family heads need not be a source of welfare loss to families. Families can rely on their own savings, the labor supply of other family members, and government tax and transfer programs to smooth this variation, so that family consumption remains unchanged. In practice, however, these sources of consumption smoothing may be far from adequate. This issue takes on particular salience in the United States today, because of a substantial increase in the instability of earnings over the past twenty years. As Peter Gottschalk and Robert Moffitt have shown, and as we confirm below, earnings variation has trended upward since the early 1970s. We estimate that over the period 1970–91, earnings variation has grown by a striking 76 percent.

The purpose of this paper is to assess the completeness and the sources of smoothing of idiosyncratic earnings variation. We use two survey data sets with information on income and consumption to estimate the relationship between variation in the earnings of household heads and variation in their families' consumption. In our analysis, we employ an instrumental variables (IV) strategy designed to deal with

We are grateful to Gary Burtless, Robert Moffitt, James Poterba, Steven Pischke, and seminar participants at the Massachusetts Institute of Technology and the Brookings Panel meeting for helpful comments, and to Weiyang Cheong and Farhan Zaidi for research assistance. Jonathan Gruber acknowledges financial support from the National Institute on Aging and the National Science Foundation.

1. See, for example, Gottschalk and Moffitt (1994).

the important problem of measurement error in earnings changes calculated from these data.

We find that families are fairly well able to smooth consumption in the face of variation in the heads' earnings. Such variation has a relatively small effect on nondurables expenditure, with 10 cents or less of each dollar change in the head's earnings being reflected in consumption. There is a somewhat larger effect on durables expenditures, however, with each dollar of earnings change corresponding to a 17 cent change in durables purchases. We also find a larger consumption response to earnings changes induced by changes in wage rates than to those induced by changes in hours of work. And we find evidence of an asymmetric response to earnings changes, with earnings reductions producing a larger effect than earnings increases, particularly for durables expenditures.

These findings raise the question of how families are able to smooth their consumption. We consider two sources of smoothing: offsetting income flows, the most important of which come from the government tax and transfer system, and self-insurance through saving. We find that smoothing through these two channels is roughly equal, with each dollar of earnings change for the head resulting in a 35 to 50 cent offsetting change in other souces of family income and a 25 to 40 cent change in saving.

We then consider the particular effect of unemployment—a large, plausibly exogenous, source of earnings variation. Overall, the consumption effects of unemployment-induced earnings loss are fairly similar to those due to year-to-year earnings variation. But in the case of unemployment, the government plays a somewhat larger consumption smoothing role, with 50 to 55 cents of each dollar in earnings loss compensated by increased transfers and reduced taxes; only about one-quarter of the unemployment-induced earnings loss is smoothed through saving. We also document considerable heterogeneity in the ability of families to smooth earnings variation arising from unemployment, which is consistent with the skewed nature of wealth holding in the United States. For low-education (and low-wealth) households, loss of earnings through unemployment has a much stronger effect on consumption.

Finally, we turn to time-series evidence on the relationship between earnings instability and consumption instability. We find that earnings

variation rises sharply during recessions, whereas consumption variation is much less cyclical. This finding is consistent with the substantial smoothing seen in the microdata. But we conclude with a puzzle: there have been parallel rises in the instability of consumption and of earnings. While the estimates are somewhat sensitive to the definition of consumption, these twin time-series trends are at odds with the microdata evidence of considerable consumption smoothing.

The paper proceeds as follows. We first motivate our analysis by revisiting the time-series evidence on earnings variation. Next, we situate our analysis in the context of the related literature. We then introduce the data and discuss the empirical issues involved in estimating the extent of consumption smoothing. We then present our estimates of the ability of families to smooth year-to-year variation in earnings and examine the sources of consumption smoothing. Next, we focus on the specific case of unemployment, modeling consumption smoothing, sources of insurance, and heterogeneity in the response to unemployment-induced earnings losses. Finally, we present time-series evidence on the relationship between earnings instability and consumption instability and discuss our overall conclusions.

## Motivation: Rising Earnings Instability

The starting point for our analysis is the striking findings of two papers by Gottschalk and Moffitt.[2] Using data from the Michigan Panel Study of Income Dynamics (PSID), the authors document a rise in the transitory component of the earnings of household heads in the 1980s. They find that this transitory component rose by 42 percent from the 1970s to the 1980s. We begin our analysis by revisiting this question, extending their analysis in a number of ways.

The PSID is a longitudinal survey that has been carried out continuously since 1968, following the same sample of families and their "split-offs" over time. The original sample consisted of a nationally representative cross-section of families and a subsample of those in poverty; in the analysis below, we use both samples in order to increase the precision of the estimates. Throughout the analysis, we weight our

2. Moffitt and Gottschalk (1993); Gottschalk and Moffitt (1994).

tabulations and regressions by the PSID sample weights in order to reproduce a nationally representative sample. We obtain very similar findings when we use the nationally representative sample alone, with or without sample weights.

Our analysis of earnings instability differs from that of Gottschalk and Moffitt along several dimensions. First, our sample includes all male heads aged twenty to fifty-nine who are not full-time students; Gottschalk and Moffitt focus only on whites in this age range. Second, we examine total labor earnings, rather than just wages and salaries. Third, we add five years of data on labor income, extending the sample through the 1992 PSID, which includes earnings data for 1991. Fourth, we do not divide the period into two halves (Gottschalk and Moffitt use 1970–78 and 1979–87), but rather, consider the entire period in one regression framework.

Finally, we use a different framework for modeling transitory income. There are a number of options for modeling the transitory component, none of which is fully satisfactory. Gottschalk and Moffitt focus primarily on an individual fixed effects model, where transitory income is defined as the deviation from individual average earnings (after absorbing a general age-earnings profile). We use a differences model, measuring deviations from the previous year's earnings as transitory. These approaches are very similar; indeed, with two observations they are identical. However, both suffer from having low power in distinguishing permanent shifts in earnings prospects from transitory changes.

To attempt to discriminate better between permanent and transitory changes, we also estimate models that include individual fixed effects in the differences specification. In this way, we allow for a specific growth path for each individual and only label as transitory deviations from that growth path. With this fixed effect, we hope to absorb any change in earnings that results from permanent changes in the head's tastes for work or leisure. This method is related to another approach used by Gottschalk and Moffitt, who include person-specific age-earnings profiles in one of their models.

Our basic analysis proceeds as follows. We first estimate a differences model over the period 1970–91. The model controls for time (with year dummies) and a number of family characteristics: a quartic in age; education (our three categories are high school dropout, high

**Figure 1. Variation of Earnings, 1970–91**[a]

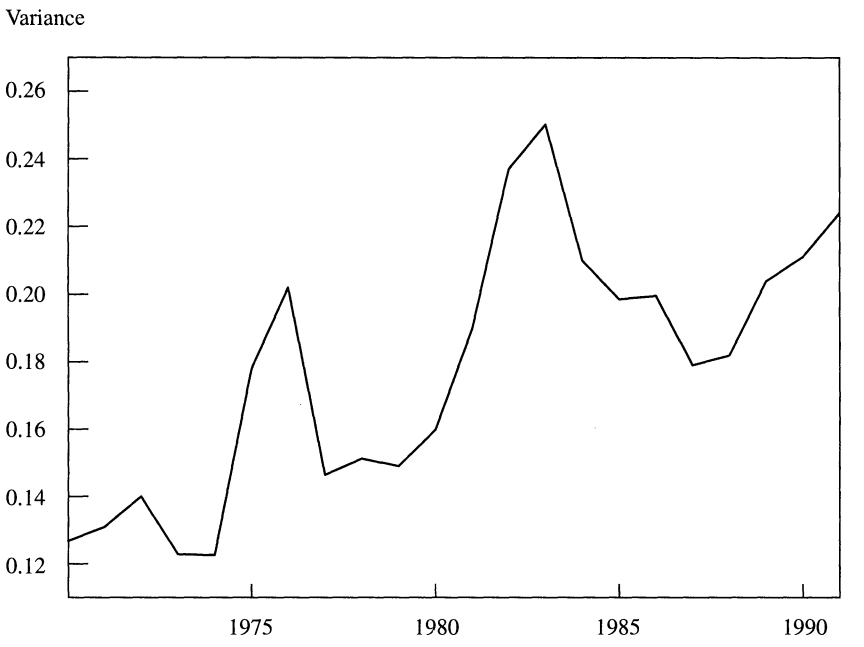


Source: Authors' calculations based on data from the Michigan Panel Study of Income Dynamics (PSID).

a. Figure plots, by year, the mean of the squared residuals produced in regressing changes in log earnings of the head on the family characteristic controls and year dummies described in the text and used throughout the paper.

school graduate, and college graduate); marital status and change in marital status; change in family size; change in the proportion of family size that is children; and change in family "food needs," a PSID measure that is a function of family size and the age of family members. These family composition controls are important for the consumption regressions reported below and are therefore included in each estimation. We include them in our earnings regressions in order to use consistent models across the different dependent variables.

We use the mean of the squared residuals for each year as the measure of aggregate transitory variation in that year. The results for labor earnings are shown in figure 1, which plots aggregate transitory variation against time. There are two findings of note. First, earnings variation has a strong countercyclical component; it peaks in the recessions of the mid–1970s, the early 1980s, and the early 1990s. Second, earnings variation has a strong upward trend: after the first two recessions,

**Figure 2. Variation of Earnings, Individual Fixed Effects Included, 1970–91**[a]


Source: Authors' calculations based on data from the PSID.
a. Produced by the same method as figure 1, except that this equation adds a person-specific fixed effect.

the variation returns to a level much higher than that before the recession. Over the entire period, the variation of earnings of heads rose by 76 percent.

In figure 2 we show the same labor earnings graph, but control for an individual fixed effect. The pattern is almost identical to that of figure 1, with the exception of a large drop in the last year that appears only in figure 2. Over the entire period, the rise in variation is almost identical in the two models. Thus, regardless of our specification, we confirm the conclusion of Gottschalk and Moffitt: transitory earnings variation has risen dramatically in the United States over the past twenty years.[3]

3. One potential problem with our fixed effects specification is that we posit that the individual effect is fixed over a very long period; some heads are in our sample for all twenty-three survey years. As a result, if individuals move to very different earnings trajectories later in their careers, the shifts will not be captured by this permanent

William Dickens, however, raises the key issue of the welfare implications of this finding.[4] If there is full consumption smoothing of variable earnings, there is potentially little welfare cost to this increased instability. But if individuals are not able to smooth their consumption over periods of high and low earnings, this increased variation may cause large welfare losses. Thus the bulk of our analysis is devoted to an assessment of the completeness of consumption smoothing of variable earnings. We first investigate this in a cross-sectional context; in the final section, we examine the time-series trend in consumption variation.

## Background

The analysis of the following sections is closely related to two different strands of the literature on consumption behavior. The first tests the hypothesis of *full consumption insurance*. This benchmark is met when mechanisms for pooling risk, either within or across families, equalize the growth rate of the marginal utility of consumption across households. As demonstrated by John Cochrane, Angus Deaton, and Robert Townsend, the full consumption insurance hypothesis implies that the growth rate of consumption will depend only on the growth rate of aggregate resources and changes in household preferences (for example, as a result of aging or changes in family size).[5] Therefore this hypothesis implies that the growth in each household's consumption will not depend on changes in household resources that are uncorrelated with shifts in preferences, once time-series changes in endowments have been taken into account.

This theory has been tested by estimating a model of growth in consumption against growth in income, controlling for aggregate resources. If there is full consumption insurance, then idiosyncratic

component. We have replicated our findings using instead a rolling average model, whereby we estimate our fixed effects model over periods of six years (five differences) and take the average of the residuals for the individual over that period. This model allows the permanent component for each individual to evolve over time. The substantial upward trend remains under this approach, although one necessarily loses the cyclical variation.

4. Dickens (1994).

5. Cochrane (1991); Deaton (1992a); Townsend (1994).

variation in family resources should not be reflected in family consumption. Barbara Mace, using consumption data from the Consumer Expenditure Survey (CEX) of the Bureau of Labor Statistics, finds that one cannot reject the proposition of full consumption insurance. But Julie Nelson points out a host of problems with Mace's implementation and finds that once these issues are corrected, full consumption insurance is strongly rejected. Cochrane confirms this rejection in data from the PSID.[6]

Cochrane also notes, however, that this may not be a strong test, for two reasons. First, changes in income may be correlated with changes in preferences. For example, a desire for more leisure and fewer consumption goods could be manifested as correlated falls in income and consumption, biasing the analysis toward a rejection of the theory of full insurance. Second, there may be significant measurement error in income in survey data, which would bias toward finding consumption smoothing. To deal with these two sources of bias, Cochrane suggests replacing income with a series of measures of plausibly exogenous changes in the individual's environment: involuntary job loss, illness, strike days. He finds a strong rejection of full consumption insurance using these measures. Studies by Paul Burgess and others, Mark Dynarski and Steven Sheffrin, Martin Browning and Thomas Crossley, and Gruber find significant effects of unemployment on consumption as well. Gruber also shows that consumption responds to the generosity of unemployment insurance, using this exogenous source of income variation to further reject the full consumption insurance hypothesis.[7]

The present study is also closely related to microdata tests of Milton Friedman's *permanent income* hypothesis.[8] A key prediction of this

6. Mace (1991); Nelson (1994); Cochrane (1991). Additional tests of the benchmark of full consumption insurance are carried out by Altug and Miller (1990), who cannot reject this benchmark. Hayashi, Altonji, and Kotlikoff (1996) expand Altug and Miller's study to test for the presence of consumption insurance from others (as opposed to self-insurance) and strongly reject this proposition. In addition, there is a growing literature on consumption smoothing in developing countries; see, for example, Deaton (1992a), Paxson (1992), Townsend (1994, 1995), Morduch (1995), and Gertler and Gruber (1997).

7. Burgess and others (1981); Dynarski and Sheffrin (1987); Browning and Crossley (1996); Gruber (1997).

8. Friedman (1957). There is also a large literature on macrodata tests of the permanent income hypothesis, but testing for consumption insurance in a macro context is meaningless, since the key test is a comparison of consumption changes across persons.

hypothesis is that only permanent variation in income should be reflected in consumption, whereas transitory variation is absorbed through saving or dissaving. Robert Hall and Frederic Mishkin test this hypothesis by statistically decomposing income into its permanent and transitory components.[9] They find that transitory income changes do predict changes in consumption, which is consistent with the permanent income hypothesis only for very high interest rates. Further work in this framework considers the implications of measurement error, additional years of data, and the modeling of liquidity constraints.[10] The early literature does not produce very strong evidence against the permanent income hypothesis, on net; more recent articles find stronger rejections.[11]

There are two important distinctions between the full consumption insurance and the permanent income hypotheses.[12] First, the permanent income hypothesis draws a sharp distinction between transitory and permanent variation in income; the latter should be reflected in consumption decisions, while the former should not. But under the full consumption insurance hypothesis, neither transitory nor permanent idiosyncratic variation should be reflected in consumption; any change in resources, relative to aggregate shifts, should be smoothed through interpersonal transfers. The second key difference between these hypotheses is their treatment of self-insurance, as distinct from insurance by others. The permanent income hypothesis focuses on the use of self-insurance (through saving and dissaving) to smooth transitory income changes. But "'self-insurance' through accumulation of assets is an *alternative* to consumption insurance, not a mechanism for implementing it."[13] Therfore the full consumption insurance hypothesis focuses on the role of interpersonal transfers.

This paper navigates a course between these two streams in the consumption behavior literature. First, we do not claim to achieve a

---

See Deaton (1992a) and Browning and Lusardi (1996) for superb reviews of the microdata and macrodata tests of the permanent income hypothesis.

9. Hall and Mishkin (1982).

10. See Altonji and Siow (1987) on measurement error, Mariger and Shaw (1993) on additional data, and Hayashi (1985) and Zeldes (1989) on liquidity constraints.

11. See Deaton (1992b) on the earlier literature and Browning and Lusardi (1996) on more recent articles.

12. See Cochrane (1991) and Hayashi, Altonji, and Kotlikoff (1996) for further discussion of these points.

13. Cochrane (1991, p. 960).

clean distinction between transitory and permanent variation in income. We follow the full consumption insurance hypothesis in specifying a simple model regressing change in consumption on change in earnings, rather than attempting a more formal decomposition into permanent and transitory earnings variation. We do condition on the set of covariates given above, which capture life-cycle changes in preferences (that is, due to age and family structure). And in the PSID sample, we are able to include an individual fixed effect in the differences specification, so that we are examining only deviations from a person-specific growth path. But even this relatively rich framework does not capture idiosyncratic permanent changes, such as a sudden promotion with a salary raise that causes a deviation from trend growth.

This is a problem in all previous analyses relating change in consumption to change in income. It is impossible to capture the theoretically appropriate measure of transitory income variation, so the analysis must rely on some empirical proxy. A number of studies take a reduced-form approach that is similar to ours. Another common approach to the problem is to model structurally the transitory component of the income process.[14] This solution consists of posing a particular autocorrelation structure for the transitory component of earnings and then using the covariance structure of earnings to fit this model. If the model that is fit is appropriate, then this structural approach is a more efficient means of identifying the transitory component of earnings; but it may yield a misleading inference if the model is inappropriate. In particular, this approach is very sensitive to the assumptions on the covariance structure of measurement error, relative to other transitory variance in earnings.

Rather than impose this set of structural assumptions, we follow the reduced-form specification adopted by the full consumption insurance literature. In a further attempt to separate permanent from transitory changes, we also consider the effect of earnings variation arising through change in hours of work (most likely transitory) and change in wages (most likely permanent). Regardless of the theoretical label that one attaches to it, the rising variation of earnings changes documented in the previous section merits examination. Has this earnings variation been reflected in consumption variation?

The second compromise between these two literatures is in our def-

14. See, for example, Abowd and Card (1989) and Moffitt and Gottschalk (1993).

inition of consumption insurance. We do not follow the strict interpretation of the full consumption insurance hypothesis in focusing only on the role of interpersonal transfers; rather, we focus on smoothing from any source, including self-insurance. Our tests are therefore more precisely an assessment of full consumption smoothing, rather than of interfamily insurance, per se. We do, however, extend our analysis to specify the sources of consumption smoothing: we measure the extent to which consumption smoothing is due to transfers from others (including the government) and to self-insurance, through saving and dissaving.

Relative to past work on tests of either the full consumption insurance or the permanent income hypotheses, our analysis offers four contributions. First, like Joseph Altonji and Aloysius Siow in the context of testing the permanent income hypothesis, but unlike the literature testing full consumption insurance, we introduce an instrumental variables strategy to deal with the potentially important problem of measurement error in earnings variation.[15] This is an important consideration in our context, since white noise measurement error in earnings changes would bias toward zero the earnings coefficient in our regressions.

Second, we do not simply test for deviations from the admittedly extreme benchmark of full consumption insurance, but actually quantify those deviations. How large is the effect of earnings variation on consumption variation? In the limit, with a large data set and a precise estimation strategy, one could reject full insurance with very small deviations. But with concave utility, the welfare loss from imperfect consumption smoothing rises nonlinearly with the deviation from full smoothing. Thus it is important to assess whether the deviations from full smoothing are empirically meaningful.[16]

Third, unlike all previous literature on these hypotheses, we focus not on variation in family income, but on variation in the earnings of the head alone. This has two advantages: by focusing on prime-age

15. Altonji and Siow (1987).

16. It is important to note that we approximate family living standards by consumption, thereby ignoring the benefits of leisure. Even if a family's consumption drops because of a reduction in the work hours of the head, its standard of living may in fact rise, due to the extra leisure that the head is enjoying. To the extent that leisure is negatively correlated with income, this lowers further the welfare cost of earnings variation. However, if one source of consumption smoothing is increased labor supply by the spouse, this will off-set the leisure gains of the head.

male heads, we are able to mitigate concerns about endogeneity through joint determination of consumption and hours of work; and we are able to assess how other forms of family income provide insurance against the earnings variation of the head. That is, we are able to go beyond an assessment of the completeness of consumption smoothing to investigate the sources of smoothing: to what extent does smoothing occur through transfers from others (primarily the government), and to what extent through self-insurance (primarily saving and dissaving)? Fourth, we explore the sensitivity of our findings to two different data sets, the PSID and the CEX.

There has been other recent work on changes in the distribution of income and consumption that is somewhat related to our analysis. Studies by Thesia Garner and by David Cutler and Lawrence Katz show that the widening of the income distribution in the United States since the early 1970s is reflected in a widening distribution of consumption.[17] Daniel Slesnick notes, however, that this finding is sensitive to the choice of equivalence scale and finds that consumption inequality has been falling over this period.[18] Cutler and Katz and also Orazio Attanasio and Steven Davis compare changes in mean consumption across groups (delineated by some combination of education, occupation, and age) with changes in mean income and find a strong relationship: groups whose wages suffered from the recent widening of the U.S. income distribution also suffered in terms of consumption.[19]

The finding that secular interpersonal increases in wage inequality are reflected in consumption, however, has little bearing on families' ability to smooth idiosyncratic intertemporal income variation. Indeed, Attanasio and Davis find that if they use higher frequency changes across groups, there is evidence of full consumption smoothing.

## Data and Empirical Issues

The key constraint in carrying out an analysis that compares the variation of income and of consumption in the United States is the

17. Garner (1993); Cutler and Katz (1991).

18. Slesnick (1992). There is also related work on shifts in the distribution of family income, which has widened more quickly over the past twenty years than has the distribution of individual income; see, for example, Karoly and Burtless (1995).

19. Attanasio and Davis (1996).

quality of the available consumption data. Whereas high-quality income data are collected by a variety of sources, consumption data for a nationally representative sample are available from only two: the Panel Study of Income Dynamics and the Consumer Expenditure Survey.[20] This paucity of data, as well as the limitations of each of these surveys (noted below), have motivated studies at a higher level of aggregation, matching data from consumption in the CEX to higher quality data on income from other surveys, such as the PSID and the Current Population Survey (CPS) of the Bureau of Labor Statistics.[21] But such a strategy averages out the idiosyncratic variation in earnings that is of greatest interest, from our perspective.

The PSID collects a set of high-quality indicators of earnings and labor force attachment. In addition, in most years the survey has collected data on two subcomponents of consumption. The first is food: respondents are asked how much their family "usually" spends on food at home and away from home, as well as how much of their food is paid for by food stamps. The second is housing: families are asked for their expenditures on rent or mortgage payments in each year.[22] These two variables provide a very incomplete measure of consumption; using the more complete data of the CEX, we calculate that food and housing expenditures amount to only 34 percent of total consumption expenditures for the median family. Counterbalancing this key limitation of the PSID is the strong advantage that it is the only multiple year, nationally representative, longitudinal database on consumption expenditures available in the United States.

The CEX, by contrast, is the only survey in the United States that collects a complete inventory of consumption data. It has been collected on a regular basis since 1980.[23] This survey contains information on

20. Other longitudinal surveys focusing on particular population groups provide some consumption data; for example, Ohio State University's National Longitudinal Survey of Youth, for young adults, and for older persons, the Social Security Administration's Retirement History Longitudinal Survey and the National Institute on Aging's Health and Retirement Study.

21. See, for example, Cutler and Katz (1991), Attanasio and Davis (1996), Garcia, Lusardi, and Ng (1997), and Lusardi (1996).

22. In some years, data on utilities payments are also collected, but since these data are more frequently missing, we do not use them in this analysis.

23. There were periodic surveys before 1980, most recently in 1972–73. There are difficult issues of data comparability across these earlier surveys and the 1980s surveys; see Cutler and Katz (1991) for a discussion.

expenditure on several hundred separate categories of goods, collected for up to four consecutive quarters, and household demographic characteristics are collected at each interview as well. In addition, in the first and fourth interviews, information is gathered on labor force behavior and income.[24] Thus the CEX provides a short panel, with only two observations on both consumption and income. This is potentially an important disadvantage, in that it limits one to simple difference models in identifying the transitory component of earnings.

By using both of these data sources, we attempt to "triangulate" the estimates of interest. The CEX has data on food and housing consumption, which allow one to confirm estimates from the PSID. One can go beyond the PSID's consumption measures in the CEX. And one can move beyond simple differences models in the PSID.

*Sample and Definitions of Variables*

This section describes our samples from the Panel Study of Income Dynamics and the Consumer Expenditure Survey and discusses the construction of the earnings and consumption variables.

THE PANEL STUDY OF INCOME DYNAMICS SAMPLE. For the microdata analysis, we use PSID data only for 1976 onward, in order both to exploit the higher quality labor earnings data available in that period and to avoid the problem of missing data on consumption expenditure in the period 1973–75. Our sample consists of all male family heads who are between the ages of twenty and fifty-nine and are not full-time students. Our key dependent variable is consumption, for which we use three measures: food, housing, and the sum of food and housing. While the last is the broadest measure, it suffers from the fact that housing expenditures are missing in three years of our sample period (1982 and 1988–89), while food expenditures are only missing in two years (1988–89). Each of the components of food and housing is deflated by the item-specific annual consumer price index (CPI), and the real components are summed. Our key independent variable, earnings, is measured as the sum of wages and salaries, bonuses and overtime, and professional income.[25]

24. In theory, employment status and income are updated each quarter; in practice, for all but a very small share of observations these variables change only at the last interview.

25. This is not the same earnings variable as that used in the time-series analysis

One important issue with the PSID data is their timing. PSID interviews are carried out early in the year: 35 percent occur in February or March, another 40 percent in April, and another 20 percent in May. Earnings data collected at interview date $t$ refer to the preceding year; that is, earnings data collected in April 1977 refer to calendar year 1976. The timing of the consumption question, which asks about "usual spending," is less obvious, and is the source of some disagreement among researchers who use the PSID. Hall and Mishkin argue that it refers to the previous year as well, while Stephen Zeldes argues that it refers to the time of the interview.[26] In either case, given that we are interested in the relationship between consumption and earnings over a calendar year, it seems most appropriate to match consumption from the April interview with labor earnings from the previous year (as collected at the same interview). If Hall and Mishkin are correct, this will be timed appropriately; if Zeldes is correct, it will be somewhat mistimed, but even so, consumption from shortly after the year has ended should be more tightly related to earnings last year than should consumption from the beginning of that year. Moreover, this most closely matches the only timing that one can use with the CEX, as described below.[27]

THE CONSUMER EXPENDITURE SURVEY SAMPLE. Our CEX data covers the period 1980–93, and our sample is once again male heads who are aged twenty to fifty-nine and are not full-time students.[28] The CEX provides data on expenditures on a variety of consumption items. We measure consumption by expenditure level; that is, we make no effort to amortize durables purchases, but rather, treat them as consumption in the year of purchase.[29] In addition, we provide results for a variety

above, since nonwage labor income is bracketed before 1976. In the time-series analysis, we relied on the PSID's constructed "labor income" variable, including imputed labor income from business activities, which is defined consistently throughout our sample period. The results of the aggregate analysis are very similar if we simply use wage income, which is also defined consistently throughout the period.

26. Hall and Mishkin (1982); Zeldes (1989).

27. Therefore the timing convention for our time-series analysis is to refer to data from interview date $t$ as coming from year $t-1$, which is the timing of the labor data.

28. In the PSID, married males are almost always automatically assigned to be the family head. In the CEX, the wife in a couple is sometimes assigned to be the head. In these cases, our observation is the male in the couple.

29. Rather than follow the CEX convention, which includes some tax payments in total expenditure, we redefine the concept to exclude taxes.

of subcomponents of consumption. Each subcomponent is deflated by the relevant (monthly) subcomponent of the CPI.

In the CEX, as in the PSID, timing is an important issue. The change in consumption is a difference across three quarters. As a result, one must control for seasonality. We do so by including the month of the interview as a control variable. The change in earnings is the change in annual earnings reported three quarters apart, so that the two reports overlap somewhat. In addition, the consumption data overlap somewhat with the income data: the first quarter's consumption data are from the year included in the last quarter's (annual) income report. Unfortunately, given the structure of the CEX, there is little that one can do to address these limitations.

CENSORING OUTLIERS. A potential source of concern in both of these data sets is outliers. To the extent that the large outliers in the consumption variables are real changes in consumption, it is important that they be included, as they may reflect instances of particularly poor consumption insurance. Yet to the extent that they are simply coding or reporting errors, it is inappropriate to include them. Also, there are changes over time in the top coding of the income variables, which affect the underlying measured variation.

We deal with both of these problems by censoring the top and bottom 1 percent of the distribution of our consumption and income changes in each year. These cutoffs are sufficient to ensure that the results are not affected by changes in top codes.[30] By censoring, we incorporate the information from these large changes, but do not allow them to have an undue influence on the regression estimates.

Table 1 presents the means of the key variables in our two data sets. The samples are very consistent. The PSID sample is slightly younger and less well educated, but earnings and consumption are very close to those of the CEX sample. The fact that the overlapping consumption data are so similar speaks to the reasonably high quality of the available PSID consumption information.

30. In addition, we impose a consistent top code of $100,000 on the nominal values of the head's wage income and his wife's labor income.

**Table 1. Mean Characteristics of PSID and CEX Respondents**[a]

Units as indicated

| *Characteristic* | *PSID data* | *CEX data* |
|---|---|---|
| Distribution by educational attainment[b] | | |
| High school dropout | 0.20 | 0.16 |
| High school graduate | 0.55 | 0.55 |
| College graduate | 0.25 | 0.29 |
| Age | 38.9 | 40.5 |
| | (10.5) | (9.9) |
| Married[b] | 0.79 | 0.84 |
| Family size | 3.12 | 3.32 |
| | (1.53) | (1.56) |
| Expenditures | | |
| Food | 4,367 | 3,891 |
| | (2,370) | (2,534) |
| Housing | 3,590 | 3,302 |
| | (3,347) | (3,237) |
| Food plus housing | 8,036 | 7,193 |
| | (4,467) | (4,535) |
| Nondurables | . . . | 18,894 |
| | | (13,902) |
| Total | . . . | 23,481 |
| | | (19,616) |
| Income | | |
| Head's earnings | 23,288 | 23,330 |
| | (19,005) | (16,577) |
| Wife's earnings | 6,364 | 10,216 |
| | (8,840) | (11,897) |
| Transfer payments | | |
| Government | 721 | 369 |
| | (2,225) | (1,525) |
| Total | 1,100 | 1,276 |
| | (3,189) | (3,994) |
| Taxes paid | 5,499 | 4,995 |
| | (6,809) | (6,711) |
| *Summary statistic* | | |
| *N* | 59,323 | 48,368 |

Source: Authors' calculations based on data from the PSID and the Bureau of Labor Statistics's Consumer Expenditure Survey (CEX).

a. All entries except those for educational attainment are the mean values for the given characteristics, with standard deviations in parentheses. All mean values except for age, being married, and family size, are expressed in dollars.

b. Proportion of sample in each category.

*Regression Framework*

Our basic regression framework for estimating the ability of families to smooth variation in the earnings of the head is

$$\Delta C_i = \alpha + \beta \Delta Y_i + \mathbf{X}_i' \boldsymbol{\delta} + \epsilon_i, \tag{1}$$

where $C_i$ is consumption for family $i$, $Y_i$ is the earnings of the head of family $i$, and $\mathbf{X}_i$ is a vector of the year dummies and family characteristic control variables described above. In the CEX, we replace the food needs variable with a more detailed set of controls for the age structure of the household: change in shares of family that are children under five years old, children over five, and adults over sixty-four.

We estimate this model in levels, rather than the log form traditionally used to test the full consumption insurance and permanent income hypotheses, because we are interested in estimating precisely how much consumption changes for each dollar change in earnings.[31] It is also of interest to examine this relationship in elasticity form, in order to compare the relative response of different types of consumption to earnings changes. Thus we also calculate elasticities, evaluated at the regression mean.

There are three important empirical concerns with interpreting $\beta$ as a causal effect of earnings variation in equations such as (1). The first is measurement error. The PSID validation study finds that 15 to 30 percent of the cross-sectional variation in earnings is measurement error. Using matched data from the CPS and social security records, John Bound and Alan Krueger find that 20 to 25 percent of the variation in first differences is due to measurement error.[32]

A classic solution to the problem of measurement error is to employ an instrumental variables strategy. To be valid in this instance, the instrument must satisfy two conditions: it must be correlated with the change in earnings and uncorrelated with the error term in equation 1. An instrument that meets both of these requirements is an independent measure of change in earnings. We are able to construct such a measure by using information available in both the PSID and the CEX. Each

31. Although the first microdata test of the permanent income hypothesis (Hall and Mishkin, 1982) estimated the model in levels, most subsequent approaches have used logs, which is consistent with constant relative risk aversion utility functions.

32. See Duncan and Hill (1985) on the PSID, and Bound and Krueger (1991).

survey has variables that measure the number of hours worked in the previous year and the current wage rate.[33] The product of these, which we hereafter refer to as imputed earnings, is an independent measure of earnings in the previous year.[34] Since we are interested in change in earnings, we use the change in imputed earnings as the instrument for the change in reported earnings.

The two-equation system that we estimate by instrumental variables is

$$\Delta Y_i = a + b\Delta Z_i + \mathbf{X}_i'\mathbf{d} + \mu_i \tag{2}$$

$$\Delta C_i = \gamma + \beta\widehat{\Delta Y}_i + \mathbf{X}_i'\mathbf{d} + \nu_i, \tag{3}$$

where $\Delta Z_i$ is the instrument, $\mathbf{X}_i$ is the vector of covariates of equation 1, and $\widehat{\Delta Y}_i$ is the change in earnings predicted by equation 2.

This instrumental variables strategy will be valid if the measurement error in change in imputed earnings is uncorrelated with the measurement error in change in reported earnings. Note that the instrument will be invalidated only by correlated error in the changes, not in the levels. For example, if a family always underreports earnings, hours, and wages, the instrument is valid, because the correlated measurement errors will be eliminated when one takes differences. But if a family underreports earnings, hours, and wages in one year only and not in

33. Our PSID measure of hours is hours worked in the previous year. Our CEX hours measure is usual weekly hours worked times weeks worked in the previous year. Our PSID wage measure is actual hourly pay for hourly workers, and for salaried workers, reported current salary normalized by forty hours per week. We construct our CEX wage measure by dividing weekly pay by weekly hours of work. The CEX collects data on gross pay for the last pay period and the frequency of pay. Eighty-seven percent of our CEX sample are paid weekly or biweekly and another 10 percent are paid monthly, so that this earnings figure closely approximates earnings at the interview date. Most of the remaining 3 percent of respondents do not report frequency of pay, so we set their wages to missing. Note that the wage and hour variables are reported independently of annual earnings in both the PSID and the CEX, and thus constitute an independent measure of earnings.

34. For example, for an individual interviewed in April 1977, we construct imputed earnings with the product of hours worked in 1976 and wage in April 1977. In the PSID sample, we could instead use the wage rate at the point of the previous interview (in this example, April 1976). Doing so yields very similar results for our wage times hours instrument. We use the timing convention described above for consistency with the CEX, where it is the only approach possible, since we only have two observations on each individual.

other years, differencing will not eliminate the correlated measurement errors in reported and imputed earnings. With such a pattern of misreporting, measurement error is not eliminated from equation 3 and our estimated $\beta$ would be biased toward zero.

The second concern is endogeneity through planned coincident changes in consumption and labor supply, as highlighted by Cochrane.[35] The direction of the bias to $\beta$ from endogeneity is not obvious and depends on the complementarity or substitutability of leisure and consumption. This endogeneity should be mitigated, to some extent, by our use of a prime-age male sample, for which full-time work is the general activity, so that changes in labor supply may be largely exogenous. But we cannot rule out that year-to-year variation in hours is determined by the same factors that drive consumption decisions.

The third concern is that equation 1 assumes that the consumption smoothing process is linear. In fact, features such as liquidity constraints could give rise to a nonlinear process: individuals may be able to smooth losses that are less than existing wealth holdings but unable to borrow in order to smooth variation that exceeds ex ante asset levels. For this reason, different sources of earnings variation could have quite different effects, in terms of dollar change in earnings: smaller, higher frequency changes in earnings may influence consumption much less than larger, lower frequency shocks.

In order to address all three concerns, we use a second instrument for estimating equation 1: unemployment shocks. In particular, we create a dummy variable for becoming unemployed for more than one month. This measure is plausibly exogenous to consumption decisions and represents a major change in the earnings prospects of the head; on average, a head who becomes unemployed faces a 30 percent reduction in earnings. In addition, any measurement error in this indicator is likely to be independent of error in reported earnings. Thus, by comparing the effects of year-to-year changes in earnings with the effects of unemployment shocks, we can assess the sensitivity of our findings to endogeneity and nonlinearity in consumption smoothing. We find that year-to-year earnings variation and reductions in earnings due to unemployment have similar effects, suggesting that these empirical prob-

35. Cochrane (1991).

lems do not account for our finding of fairly comprehensive consumption smoothing.

## Smoothing Variation in Earnings

This section discusses our findings on the ability of families to smooth year-to-year variation in earnings.

### *Results from the Panel Study of Income Dynamics*

Table 2 shows the results from estimating equation 1 in the PSID by ordinary least squares (OLS). As discussed above, OLS estimates of $\beta$ are likely to be biased toward zero by measurement error in the change in earnings. We present these results for comparison with the previous literature, which relies largely on OLS. We show results for three dependent variables: food, housing, and the sum of food and housing. We also report the elasticities implied by these regressions.

There is a highly significant effect of income changes on consumption changes, concordant with the previous literature rejecting full insurance. But the coefficient is quantitatively very small. Each dollar rise in the earnings of the head increases consumption of food by only 1 cent, and of food and housing by only 1.7 cents. Food consumption appears to be more elastic with respect to income changes than is housing consumption; a 10 percent rise in the earnings of the head raises food consumption by 0.53 percent, while it raises housing consumption by only 0.46 percent. The resulting elasticity of food plus housing consumption with respect to income is only 0.051.

The control variables show that the growth rate of food consumption declines and the growth rate of housing consumption rises with age. The growth rate of food consumption rises with education, but there is no significant effect of education on growth in housing consumption. There is a small negative effect on consumption growth of being married, but a large positive effect of becoming married. There are also strong effects of change in family size and change in food needs.

Table 3 reports instrumental variables estimates of equation 1 in the PSID, using as the instrument the change in imputed earnings (esti-

**Table 2. Estimating Consumption Smoothing Using PSID Data, OLS Estimation[a]**

| | *Dependent variable* | | |
|---|---|---|---|
| *Independent variable* | *Food* | *Housing* | *Food plus housing* |
| Change in earnings of head | 0.0099 | 0.0068 | 0.0171 |
| | (0.0014) | (0.0017) | (0.0025) |
| Age | −402.9 | 404.3 | 50.0 |
| | (228.8) | (240.0) | (385.2) |
| Age$^2$ | 16.17 | −16.10 | −1.15 |
| | (9.04) | (9.41) | (15.09) |
| Age$^3$ | −0.279 | 0.268 | 0.070 |
| | (0.154) | (0.159) | (0.255) |
| Age$^4$ | 0.0017 | −0.0016 | 0.0001 |
| | (0.0010) | (0.0010) | (0.0016) |
| Change in family size | 421 | 52 | 463 |
| | (37) | (28) | (55) |
| Change in children-to–family size ratio | 81 | −32 | −47 |
| | (119) | (115) | (187) |
| Change in food need[b] | 0.0805 | 0.0330 | 0.1173 |
| | (0.0358) | (0.0194) | (0.0524) |
| Dummy variables | | | |
| Black | −21 | −26 | −34 |
| | (30) | (27) | (46) |
| High school graduate | 19 | −1 | 23 |
| | (24) | (23) | (37) |
| College graduate | 48 | −18 | 31 |
| | (30) | (30) | (47) |
| Married | −13 | −68 | −72 |
| | (26) | (31) | (46) |
| Change in marital status[c] | 311 | 573 | 869 |
| | (60) | (80) | (114) |
| Implied elasticity[d] | 0.053 | 0.046 | 0.051 |
| *Summary statistic* | | | |
| *N* | 43,812 | 34,311 | 27,484 |

Source: Authors' calculations based on data from the PSID.

a. The dependent variable is the annual change in expenditures in the given category. Regressions include the independent variables listed plus a full set of year dummies. The sample period is 1976–92; regressions involving food exclude 1988–89, and regressions involving housing exclude 1982 and 1988–89. Standard errors are shown in parentheses.

b. A PSID variable; see text for details.

c. Becoming married equals 1; becoming single equals −1; no change equals 0.

d. Elasticity implied by the coefficient on head's earnings, evaluated at variable means.

**Table 3. Comparing OLS and IV Estimates of Consumption Smoothing Using PSID Data[a]**

| | *OLS method* | | *IV method* | |
|---|---|---|---|---|
| *Dependent variable* | *Earnings coefficient* | *Implied elasticity*[b] | *Earnings coefficient* | *Implied elasticity*[b] |
| *Basic equations* | | | | |
| Food | 0.0099 | 0.053 | 0.0381 | 0.205 |
| | (0.0014) | | (0.0066) | |
| Housing | 0.0068 | 0.046 | 0.0241 | 0.163 |
| | (0.0017) | | (0.0075) | |
| Food plus housing | 0.0171 | 0.051 | 0.0581 | 0.174 |
| | (0.0025) | | (0.0111) | |
| *Equations with individual fixed effects*[c] | | | | |
| Food | 0.0098 | 0.053 | 0.0378 | 0.203 |
| | (0.0016) | | (0.0078) | |
| Housing | 0.0068 | 0.046 | 0.0269 | 0.182 |
| | (0.0019) | | (0.0090) | |
| Food plus housing | 0.0168 | 0.050 | 0.0595 | 0.179 |
| | (0.0028) | | (0.0133) | |

Source: Authors' calculations based on data from the PSID.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Instrumental variables (IV) column instruments change in earnings with change in imputed earnings, as described in the text.

b. Evaluated at variable means.

c. These equations add a person-specific fixed effect identifying each individual in the sample.

mated using wage at the interview and hours worked last year).[36] The OLS coefficients (from table 2) are included for comparison.

We find that instrumenting significantly raises the effect of income changes on consumption changes. Each dollar of income rise in this IV specification leads to an increase of 3.8 cents in food consumption and of 5.8 cents in food plus housing consumption. In elasticity terms, this is an elasticity of food consumption of 0.205, and of food plus housing consumption of 0.174. While significant, these effects remain fairly small. Thus, while one can reject full consumption insurance, the deviations from that benchmark are not substantively significant.[37]

36. The imputed earnings instrument has significant explanatory power. The $R^2$ of the first stage is 0.192, and the $t$ statistic on the instrument is 20.94.

37. Note that our food plus housing coefficient need not equal the sum of the food and housing coefficients. This is because (1) the sample is changing across these regressions, due to missing data on food or housing expenditure; and (2) we censor separately the food, housing, and food plus housing variables. For this second reason, in the CEX

As discussed above, this is only one possible specification for identifying the transitory component of individual earnings. Another approach would be to include individual fixed effects in this changes model, measuring as transitory only deviations from a person-specific growth path. We show the results of this approach in the lower panel of table 3. In fact, these results are almost identical to those in the first panel. These findings imply that our CEX differences results would be robust to the inclusion of fixed effects. Therefore for the rest of the paper we present only the results of estimating equations that do not include individual fixed effects. We have replicated all of our PSID results with the inclusion of fixed effects and find that they are very similar in every case.[38]

*Results from the Consumer Expenditure Survey*

Table 4 reports the basic consumption smoothing estimates from the CEX. We show only the coefficients of interest, but these regressions include all the variables in equation 1 described above. We present results for both the OLS and the IV models, using imputed earnings as an instrument. The table includes each of our consumption subcategories in the CEX.

We begin by measuring the effect on food and housing consumption. The effects in the CEX are very similar to those in the PSID, albeit somewhat smaller. The OLS estimates indicate that a dollar change in earnings leads to a 0.4 cent change in food consumption and a 0.5 cent change in food plus housing consumption. The IV estimates yield a

sample the results for total consumption expenditure need not equal the sum of the results for nondurables and durables expenditure.

38. Lawrence Katz has pointed out to us that the tripling of the estimates between the OLS and the IV specifications is inconsistent with a simple measurement error explanation. As noted above, the measurement error in earnings changes in the PSID is estimated to be roughly 20 to 25 percent of the variation. This implies that our coefficient estimate should rise by a factor of only 1.25 when we instrument, if we are simply purging measurement error. One possible explanation for the larger rise, suggested to us by John Abowd, is that the types of events that induce variation in imputed earnings are somewhat different from the types of events that induce variation in actual earnings. If variations in imputed earnings were perceived as more permanent than movements in actual earnings, and if permanent changes were reflected more strongly in consumption (as given by the permanent income hypothesis), instrumenting would lead to a rise in the estimated coefficient. It seems unlikely that this process is playing a large role here, however, since imputed earnings is simply an independent measure of actual earnings.

**Table 4. Estimating Consumption Smoothing Using CEX Data[a]**

| | *OLS method* | | *IV method* | |
|---|---|---|---|---|
| *Dependent variable* | *Earnings coefficient* | *Implied elasticity[b]* | *Earnings coefficient* | *Implied elasticity[b]* |
| Food | 0.0044 | 0.026 | 0.0238 | 0.142 |
| | (0.0015) | | (0.0064) | |
| Housing | 0.0008 | 0.006 | 0.0042 | 0.029 |
| | (0.0012) | | (0.0048) | |
| Food plus housing | 0.0053 | 0.017 | 0.0339 | 0.109 |
| | (0.0022) | | (0.0089) | |
| Nondurables | 0.0344 | 0.042 | 0.0893 | 0.110 |
| | (0.0058) | | (0.0237) | |
| Durables | 0.0329 | 0.169 | 0.1727 | 0.888 |
| | (0.0112) | | (0.0047) | |
| Total consumption | 0.0680 | 0.067 | 0.2440 | 0.240 |
| | (0.0130) | | (0.0550) | |
| Utilities | 0.0004 | 0.005 | 0.0017 | 0.021 |
| | (0.0006) | | (0.0026) | |
| Clothing | 0.0058 | 0.693 | 0.0111 | 0.177 |
| | (0.0013) | | (0.0053) | |
| Entertainment | 0.0073 | 0.102 | 0.0013 | 0.018 |
| | (0.0022) | | (0.0091) | |
| Vehicle maintenance and fuel | 0.0034 | 0.030 | 0.0016 | 0.014 |
| | (0.0017) | | (0.0072) | |
| Home services | 0.0028 | 0.076 | 0.0224 | 0.604 |
| | (0.0015) | | (0.0063) | |
| Alcohol and tobacco | 0.0011 | 0.047 | 0.0044 | 0.188 |
| | (0.0004) | | (0.0015) | |
| Medical care and insurance | 0.0013 | 0.051 | 0.0022 | 0.086 |
| | (0.0008) | | (0.0034) | |
| Other insurance | −0.0007 | −0.013 | −0.0087 | −0.166 |
| | (0.0010) | | (0.0039) | |
| Contributions to others | 0.0016 | 0.048 | 0.0088 | 0.266 |
| | (0.0008) | | (0.0033) | |
| *Summary statistic* | | | | |
| $N$[c] | 19,155 | | 12,875 | |

Source: Authors' calculations based on data from the CEX.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes, including a full set on month dummies. Standard errors are shown in parentheses. IV column instruments change in earnings with change in imputed earnings, as described in the text. The sample period is 1980–93.

b. Evaluated at variable means.

c. For equation involving total consumption.

**Table 5. Estimating Consumption Smoothing: Food at Home versus Food Away from Home[a]**

| | *PSID data* | | *CEX data* | |
|---|---|---|---|---|
| *Dependent variable* | *Earnings coefficient* | *Implied elasticity[b]* | *Earnings coefficient* | *Implied elasticity[b]* |
| Food at home | 0.0225 | 0.156 | 0.0106 | 0.092 |
| | (0.0054) | | (0.0042) | |
| Food away from home | 0.0215 | 0.540 | 0.0070 | 0.158 |
| | (0.0034) | | (0.0035) | |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument change in earnings with change in imputed earnings, as described in the text. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

b. Evaluated at variable means.

coefficient on food plus housing of 0.034 and an elasticity of 0.109. Once again, we find that consumption smoothing is fairly complete.

The main advantage of the CEX data is that they allow one to expand the measures of consumption expenditure. That is, one can estimate a set of "dynamic Engel curves," tracing out how the consumption of different categories of goods changes with transitory variation in income.

The fourth row of table 4 considers the effects on total nondurables expenditures. Using our IV specification, we find that a dollar change in earnings results in an increase in nondurables expenditures of 9 cents; the elasticity is 0.110. It is striking that this elasticity is actually lower than the estimated elasticity for food consumption, given that the Engel curve for food is generally assumed to be relatively flat.

We address this point further in table 5, which decomposes food spending in both data sets into spending on food at home and away from home. We find that the absolute effect on spending on food away from home is roughly as large as that on food at home, despite the fact that spending on food away from home is only 20 to 25 percent of total food expenditure. As a result, the relatively large elasticity for food expenditures is driven by spending on food away from home; spending on food at home has an IV elasticity of only between 0.092 (in the CEX) and 0.156 (in the PSID), while spending on food away from home has an IV elasticity of between 0.158 and 0.540. Thus the elasticity of actual food consumption is low; it is the elasticity of food

preparation services that drives the fairly large effects on total food expenditures.

The fifth row of table 4, however, shows a relatively large response for durables expenditures. Despite the fact that durables are only 11 percent of the consumption bundle on average, changes in earnings have a larger absolute effect on durables spending than on nondurables spending, according to the instrumental variables estimates. This is reflected in an elasticity that is eight times larger for durables than nondurables. Because of the high elasticity for durables, as the next row shows, the effect of earnings changes on total consumption is fairly large: for each dollar change in earnings, total consumption changes by 24 cents.

This finding confirms the notion that durables are, to some extent, savings; individuals absorb income variation, in part, by adjusting the timing of durables purchases. This raises the important and difficult question of how to combine the responses of durables and nondurables in computing the effective change in total consumption. On the one hand, if one simply adds these responses, one overstates the welfare implications of a drop in earnings, since the loss in utility caused by a decrease in durables purchases is presumably smaller than that caused by an equal decrease in other expenditures. On the other hand, by revealed preference, there is some welfare loss from delayed durables purchases. So the true change in consumption, as opposed to measured expenditures, lies somewhere between the effect on nondurables expenditures and the effect on total expenditures; exactly where in this range is a function of how one values the utility loss from altering the timing of durables purchases.

The remaining rows of table 4 examine the effect of earnings variation on the subcomponents of nondurables expenditures. Focusing on the IV model, the elasticities are largest for clothing, home services (repair and maintenance), alcohol and tobacco, and contributions to others. For clothing and home services, as for durables, the welfare loss from delayed purchase may be relatively small. For contributions to others, the estimates are consistent with an altruistic model, which posits that transfers between households will rise with the donor's income.[39]

39. See Cox (1987).

The elasticity of consumption with respect to income variation is smallest for entertainment, medical care and insurance, vehicle maintenance and fuel, and utilities. The estimate for other insurance purchases (for example, life and automobile insurance) is actually negative. This may reflect delayed automobile purchases, which are included in the durables category. But it is also consistent with a falling absolute coefficient of risk aversion; as individuals get richer, they feel less need for a given dollar of insurance.

To summarize, the results from both the PSID and the CEX show that there is a highly significant effect of income variation on consumption variation. This effect is quantatively small for nondurables purchases; of each dollar of earnings change, less than 10 cents is reflected in nondurables expenditures. The effect becomes more sizable, however, when one values durables purchases at their expenditure value. But even in this case, individuals are smoothing over 75 cents of each dollar of earnings change.

### *Changes in the Quantity and Price of Labor*

One important issue raised by the discussion above (and by the permanent income hypothesis literature) is the differential impact of transitory and permanent income variation. Under the permanent income hypothesis, changes in earnings that are perceived as relatively permanent will have larger effects on consumption than will changes that are perceived as relatively transitory. As noted above, it is difficult to decompose earnings changes convincingly into their permanent and transitory components, but one crude approximation is to decompose year-to-year variation in labor earnings into two component sources of variation: changes in hours worked (the quantity of labor), and changes in wage rates (the price of labor). For our prime-age male sample, changes in hours are likely to be transitory, perhaps resulting from overtime or vacation. But changes in wages are likely to be relatively permanent, reflecting promotions or new job matches.

This contention is supported by previous work on the dynamics of earnings and hours. Henry Farber and Robert Gibbons find that wage residuals approximately follow a martingale process for younger workers, while John Abowd and David Card report substantial mean rever-

**Table 6. Estimating Consumption Smoothing: Changes in Hours versus Changes in Wage[a]**

| | *Coefficient on head's earnings* | | | |
|---|---|---|---|---|
| | *Change in hours worked* | | *Change in wage* | |
| *Dependent variable* | *PSID data* | *CEX data* | *PSID data* | *CEX data* |
| Food | 0.015 | 0.027 | 0.092 | 0.017 |
| | (0.005) | (0.008) | (0.016) | (0.010) |
| Housing | 0.010 | 0.005 | 0.040 | 0.004 |
| | (0.005) | (0.006) | (0.015) | (0.007) |
| Food plus housing | 0.016 | 0.031 | 0.128 | 0.030 |
| | (0.008) | (0.011) | (0.025) | (0.014) |
| Nondurables | . . . | 0.057 | . . . | 0.094 |
| | | (0.030) | | (0.037) |
| Durables | . . . | 0.091 | . . . | 0.207 |
| | | (0.058) | | (0.073) |
| Total consumption | . . . | 0.148 | . . . | 0.294 |
| | | (0.069) | | (0.085) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Change in earnings is instrumented either with the change in hours worked or with the change in hourly wage, as indicated. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

sion in hours of work.[40] We also find more persistence in wage changes, as compared with hours changes, in our PSID data.[41]

One can distinguish between these two sources of variation in a straightforward manner, by separating imputed earnings (wage times hours) into its two components: change in hours worked and change in the hourly wage rate. Separate instrumental variables estimations of equation 1 will then demonstrate the differential effects of these two sources of income variation on consumption.

The results of using these two instruments in the PSID are presented in table 6. The results show a striking difference between the effects of these two sources of variation. Earnings variation due to changes in hours has a relatively small effect on consumption, with each dollar of hours-induced earnings change producing only a 1.6 cent change in

40. Farber and Gibbons (1996); Abowd and Card (1989).

41. More specifically, regressing change in wages or hours on lagged changes, we find that the sum of the coefficients on the first three lags is −0.98 for hours, but only −0.75 for wages.

food plus housing consumption. But each dollar change in earnings due to change in wages corresponds to a 12.8 cent change in food plus housing consumption. That is, the effect of wage-induced change in earnings is eight times as large as the effect of hours-induced change in earnings. This is consistent with the notion that wage changes are relatively permanent, and that permanent changes have a larger effect on consumption.[42]

The results for the decomposition into hours and wages are not as robust across data sets as are our earlier findings, however. Table 6 also shows this decomposition for the CEX. For food and housing, the effect of hours-induced variation in earnings is actually larger than that of wage-induced variation. For nondurables and durables, the effect of wage-induced variation in earnings is roughly twice as large as that of hours-induced variation. Thus, while the direction is in accordance with the PSID, the magnitudes are not comparable. The reason for this inconsistency across data sets is not clear. Both hourly wage measures have limitations. The PSID measure is more precise for hourly workers, but less precise for salaried workers (since it is normalized by forty hours, rather than reflecting actual hours). Our PSID findings are fairly similar for hourly and salaried workers, however, so differential accuracy does not seem to be an important problem.

Nevertheless, the broad picture that is painted by these findings is consistent across the two data sets. There is clearly a much larger effect of wage variation than of hours variation. Yet individuals can smooth over 70 cents of each dollar in earnings changes due to changes in wages. Thus even in the case of wage variation, the vast majority of earnings variation is smoothed.

### *Upward and Downward Movements in Earnings*

Another natural extension of this analysis is to consider asymmetric responses to upward and downward movements in earnings. If individuals face liquidity constraints and have low savings, they may be less able to smooth downward changes in earnings. We investigate this

42. An alternative interpretation is that the measurement error in hours is more highly correlated with the measurement error in earnings than is the measurement error in wages. But it seems unlikely that this would account for coefficient differences of the magnitude observed in the PSID.

**Table 7. Estimating Consumption Smoothing: Earnings Increases versus Decreases[a]**

| | *Coefficient on head's earnings* | | | |
|---|---|---|---|---|
| | *Earnings increases only* | | *Earnings decreases only* | |
| *Dependent variable* | *PSID data* | *CEX data* | *PSID data* | *CEX data* |
| Food | 0.035 | 0.018 | 0.041 | 0.035 |
| | (0.010) | (0.035) | (0.011) | (0.011) |
| Housing | 0.022 | −0.007 | 0.026 | 0.017 |
| | (0.011) | (0.007) | (0.012) | (0.008) |
| Food plus housing | 0.046 | 0.018 | 0.070 | 0.056 |
| | (0.017) | (0.013) | (0.018) | (0.016) |
| Nondurables | . . . | 0.071 | . . . | 0.109 |
| | | (0.035) | | (0.041) |
| Durables | . . . | 0.051 | . . . | 0.346 |
| | | (0.069) | | (0.081) |
| Total consumption | . . . | 0.142 | . . . | 0.392 |
| | | (0.081) | | (0.095) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument change in earnings with change in imputed earnings, as described in the text. The first two columns give coefficients on upward movements in predicted earnings (obtained from the first-stage regression); the final two columns show coefficients on downward movements. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

proposition in our instrumental variables framework by estimating our first-stage equation to generate predicted earnings, and dividing predicted earnings into its upward and downward components. We then include these components separately in our consumption model.[43]

The findings of this decomposition are presented in table 7. In the PSID, there is only slight evidence of a stronger effect of downward movements in earnings than of upward movements. For food plus housing, the coefficient on downward movements is roughly 50 percent larger than that on upward movements. In no case are the coefficients statistically distinguishable. In the CEX, the gap is similar for nondurables; while the effect on downward movements is roughly 50 percent larger, the estimates are not significantly different from each other. But for durables purchases, we do find a very large gap, indicating that the large effects for durables shown in table 4 appear to be driven mainly

43. This is equivalent to IV estimation done in two steps. In this context, our standard errors are somewhat understated, since we do not account for the two-step nature of the estimation (whereas our direct IV estimates correct the standard errors automatically). Given the very strong first-stage relationship between actual and imputed earnings, this understatement should be small.

by responses to downward movements in earnings. Thus families appear to smooth upward changes in earnings largely through other sources of insurance or saving, whereas downward changes are mainly reflected in reduced durables purchases.

The analysis thus far, particularly in the CEX, shows that earnings variation arising from most sources is reflected in families' consumption in only a very limited way. This naturally raises the question of *how* families smooth consumption against swings in the earnings of the head.

## Sources of Consumption Insurance

There are three possible sources of insurance against idiosyncratic variation in the earnings of the head: nonhead earnings, nonlabor income, and saving. The key component of nonhead earnings is the labor supply of other family members; that is, other family members can increase their labor supply when the earnings of the head fall.

There are several sources of nonlabor income. One is increased receipts from government social insurance programs, such as unemployment insurance or workers' compensation. Another source is transfers from others to the family when income is transitorily low, which would be consistent with the altruistic model that finds support in the results above. A final source of nonlabor income is the tax system. A family's net income is automatically insulated from movements in the head's earnings through taxation: lower earnings imply lower payments of income and payroll taxes, so that the net fall in resources is smaller than the gross fall in the head's earnings.

The third source of consumption insurance is saving. In the standard permanent income hypothesis model, families use saving and dissaving to smooth consumption over changes in the employment status of the head. The results discussed in the previous section clearly reject the strong form of this hypothesis, since families are not using saving to smooth consumption fully. Nonetheless, a substantial proportion of the consumption smoothing that does occur may be through saving behavior.

We use a regression framework to measure the magnitudes of the three souces of family income smoothing for which we have data: wife's

**Table 8. Estimating Sources of Family Income Smoothing[a]**

| | *Coefficient on head's earnings* | |
|---|---|---|
| *Dependent variable* | *PSID data* | *CEX data* |
| Wife's earnings | −0.016 | −0.012 |
| | (0.014) | (0.015) |
| Transfer payments | | |
| Government | −0.122 | −0.046 |
| | (0.007) | (0.004) |
| Total | −0.145 | −0.062 |
| | (0.009) | (0.007) |
| Taxes paid | 0.347 | 0.262 |
| | (0.014) | (0.016) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on change in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument change in earnings with change in imputed earnings, as described in the text. The PSID sample period is as described in table 2, note a, but the regression involving taxes also excludes 1992. The CEX sample period is 1980–93.

earnings, transfer income, and taxes. For each source of income, we estimate an equation of the form

$$\Delta I_i = \gamma \Delta Y_i + \mathbf{X}_i' \boldsymbol{\phi} + \epsilon_i, \tag{4}$$

where $I_i$ is a given source of family income. The coefficient $\gamma$ measures how much a source of family income insurance changes for each dollar change in the head's earnings ($\Delta Y_i$). Measurement error in $\Delta Y_i$ will bias the estimate of $\gamma$ toward zero; we therefore employ an instrumental variables strategy, using imputed earnings as the instrument. These estimates are reported in table 8.

RESPONSE OF WIFE'S EARNINGS. Both the CEX and the PSID measure the wife's labor income over the same time frame as the earnings of the head. The results presented in table 8 indicate that, in general, there is relatively little compensation for lost earnings through the earnings of the wife. In the PSID, the coefficient is small and insignificant. In the CEX, the coefficient is somewhat larger, but still insignificant. Thus, overall, we find that spousal labor supply plays only a small role in the smoothing of year-to-year earnings variation.

TRANSFER INCOME. In both data sets, transfer income is defined as government transfers, retirement income, and other sources of support that are not earned labor income or asset income. As table 1 shows, the mean of this variable is very similar across the two data sets. The results reported in table 8 indicate that transfer income plays a nontrivial role

in consumption smoothing. In the PSID results, transfer income flows off-set 15 cents of each dollar lost due to earnings variation. The basic pattern is similar in the CEX, albeit with smaller coefficients. The table also shows that the majority of the response of transfer income comes through government transfers; private transfers account for, at most, one-quarter of the transfer estimates.

TAXES. A mechanical, but important, source of smoothing of earnings variation is taxation. Indeed, one justification that has been offered for a redistributive tax system is the provision of insurance against income variation.[44] Yet previous analyses of consumption smoothing have largely ignored taxes. The CEX reports the actual tax payments of households over the previous year, while the PSID calculates expected tax payments on the basis of the comprehensive income data in the survey.[45] We include both the income and the payroll tax payments of households. Payroll tax payments are reported in the CEX; for the PSID sample, we calculate them by applying the payroll tax schedule to household earnings.

The results presented in table 8 show a very important role for taxes. In the PSID, changes in tax burdens smooth 35 cents of each dollar of earnings variation. In the CEX, the coefficient is somewhat smaller, indicating smoothing of 26 percent. These coefficients can be interpreted as marginal tax rates.[46]

There are two possible reasons for the larger tax effects in the PSID. First, the PSID contains data for five years before the tax reforms of the early 1980s, which substantially lowered marginal rates (although the effect of the reforms was somewhat counteracted by a rising payroll tax rate over this period). Indeed, if we estimate our PSID models over the same sample period as the CEX, we find that the tax coefficients fall and are more similar to those that we estimate in the CEX, whereas there is no difference in other sources of consumption insurance. Second, self-reported taxes (as in the CEX) may respond less to changes

44. Varian (1980) bases this justification on permanent (cross-person) differences in income, not transitory (within-person) changes, but the same basic principles apply.

45. The PSID does not report tax payments for 1992, so we are unable to use the last year of our data for this exercise.

46. These coefficients are consistent with previous estimates of the average marginal tax rate in the United States. Barro and Sahasakul (1986) estimate year-by-year averages of the sum of federal income and payroll taxes. The average of their estimated rates for 1976 through 1983 (the last year in their data) is 0.348.

in reported earnings than do tax payments calculated directly on the basis of those reported earnings (as in the PSID).

IMPLIED SAVING RESPONSE. To summarize our findings for the three sources of consumption smoothing, we find in the PSID that roughly one-half of earnings changes are off-set by net income flows to the family; in the CEX, where the income data are probably less reliable, the offset is only 32 cents. Most of this offset arises from reduced tax payments, although some comes through increased government transfers.

Taken together with our consumption estimates, these findings imply that 25 to 40 percent of change in earnings from year to year is smoothed by saving. This estimate rises to about 40 to 55 percent if one includes durables expenditures as saving. Thus saving and off-setting sources of family income appear to play an almost equally important role in smoothing consumption. That is, there is an equal role for insurance through others (predominantly the government) and self-insurance through saving.[47]

## Unemployment

Unemployment is a large, arguably exogenous, source of earnings variation. As such, earnings variation due to unemployment may be harder for households to smooth than that arising from the sources explored above, both because unemployment is unplanned and because its effects on earnings are large.[48] In fact, previous analyses have shown that unemployment is associated with a significant decline in consumption. But none have quantified the deviation from full smoothing due to unemployment. For each dollar of earnings loss due to unemployment, by how much does consumption fall?

47. We also investigate the differential response of transfers and taxes to upward and downward movements in earnings. As might be expected, transfers respond somewhat more strongly to downward movements, whereas taxes respond somewhat more strongly to upward movements; this is consistent with the redistributive structures of tax and transfer programs in the United States.

48. As noted above, the coefficient from regressions for the PSID and the CEX of the change in log earnings on our unemployment dummy (controlling for the year dummies and family characteristics used throughout the paper) yields sizable effects: unemployment lowers the earnings of the head by over 30 percent.

In this section we consider the effect of becoming unemployed on consumption. It is possible for unemployment to be somewhat anticipated; for example, individuals may be on regular temporary layoff. Our goal is to draw a sample for which unemployment appears unanticipated. In the PSID, we do so by contrasting the sample of workers who report at interview date $t$ (for example, April 1978) that they had one month or more of unemployment last year (in this example, January 1997 to December 1977) with those who report no unemployment last year. We further condition the sample on both employment at interview date $t-1$ (April 1977) and no unemployment during the year preceding that interview (January–December 1976). The dummy variable for unemployment therefore measures the transition from employment to unemployment for a sample that appears to be regularly employed ex ante, and thus for whom unemployment is not a regular event. In addition, by conditioning on employment at interview date $t-1$, which is in the early spring, we generally capture unemployment spells occurring in the second half of the year, which is closer to the frame of reference for the consumption question.[49]

Our CEX measure is defined analogously. However, the CEX provides information not on weeks of unemployment, per se, but only on weeks without work. We therefore replace the condition of zero weeks of unemployment in the year preceding interview $t-1$ with a condition of more than forty-eight weeks of employment in that year; and the condition of one month of unemployment in the year preceding interview $t$ with a condition of fewer than forty-eight weeks of employment.

We estimate the response of consumption variation to earnings variation, using this dummy variable for unemployment as an instrument. In neither data set are we able to determine whether unemployment is due to voluntary or involuntary job separation. Thus, to the extent that quits are planned and reflected in consumption profiles, our estimates may understate the impact of exogenous job loss.[50] But for prime-age

49. In fact, our unemployment spells may be occurring before April 1977, since the question in April 1978 simply asks about unemployment during the previous year. That is, an individual may have been employed during all of 1976, unemployed during February 1977, and employed again by April 1977. But for individuals who report no unemployment during 1976 and employment in April 1977, it seems likely that any unemployment during 1977 occurred after April.

50. That is, quitters may lower their consumption in the previous period in antici-

**Table 9. Estimating Consumption Smoothing for Earnings Losses due to Unemployment[a]**

| | *Coefficient on head's earnings* | |
|---|---|---|
| *Dependent variable* | *PSID data* | *CEX data* |
| Food | 0.076 | 0.055 |
| | (0.018) | (0.013) |
| Housing | 0.027 | 0.006 |
| | (0.017) | (0.010) |
| Food plus housing | 0.088 | 0.067 |
| | (0.027) | (0.018) |
| Nondurables | . . . | 0.108 |
| | | (0.049) |
| Durables | . . . | 0.129 |
| | | (0.097) |
| Total consumption | . . . | 0.241 |
| | | (0.114) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on decrease in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument earnings loss with an unemployment dummy, as described in the text. The sample is confined to heads who experience an unanticipated period of unemployment in a given year. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

heads, a period of search of more than one month seems most likely to be associated with exogenous job loss, rather than a planned quit.

### *Basic Results*

The results using the unemployment dummy as an instrument are presented in table 9. For most components of consumption, these estimates are very similar to those for general downward movements in earnings, shown in table 7. In the PSID, each dollar of earnings loss due to unemployment reduces food consumption by 7.6 cents, and food plus housing consumption by 8.8 cents; these effects are slightly larger than those in table 7. In the CEX, the pattern is similar. The effects are slightly larger for food and housing, and almost identical for total nondurables consumption; each dollar of unemployment-induced earn-

pation of having low income and, as a result, have a small consumption change when they quit.

ings loss lowers nondurables expenditure by 10.8 cents, as compared to 10.9 cents in table 7.[51]

As in the analysis above, the major impact of earnings variation is on durables expenditures. An employment-induced drop of one dollar in earnings produces a drop of 12.9 cents in durables expenditure. This is considerably smaller than the 34.6 cent drop in durables expenditure in response to general downward movement in earnings shown in table 7. As a result of this differential response of durables expenditures, the effect of earnings variation on total consumption is actually smaller for the case of unemployment (0.241) than for general downward movement in earnings (0.392).

Overall, though, we find that this very different instrumental variable produces results quite similar to those uncovered by our basic instrumental variables regressions, shown in tables 3 and 4. Nondurables expenditures change by about 10 cents for each dollar change in earnings and total expenditures change by less than 25 cents. This suggests that the previous findings are not badly biased by measurement error, endogeneity, or nonlinearities in consumption smoothing. But, once again, these results raise the question of how families smooth this large shock to the earnings prospects of the head.

### *Sources of Consumption Insurance*

Table 10 extends the analysis of the sources of family income smoothing to unemployment. This is a particularly interesting case, since here is an explicit government insurance program—unemployment insurance—that is designed to deal with income variation due to unemployment.

In both the PSID and the CEX, we find that there are substantial (net) income flows offsetting the income loss from unemployment; the totals are similar, although the sources are quite different. In the CEX, there is a large offsetting response from spousal earnings to the income loss from unemployment; in the PSID, there is little response from spousal earnings. These CEX findings, however, are at odds with a

51. The similarity of these consumption responses may seem counterintuitive, since unemployment is a rather severe shock to earnings prospects. But, as discussed below, the government explicitly provides insurance against the income shock of unemployment. Thus, while the shock of unemployment may be severe, insurance against that shock is relatively complete.

**Table 10. Estimating Sources of Family Income Smoothing in Response to Unemployment[a]**

| | *Coefficient on head's earnings* | |
|---|---|---|
| *Dependent variable* | *PSID data* | *CEX data* |
| Wife's Earnings | −0.022 | −0.124 |
| | (0.032) | (0.033) |
| Transfer payments | | |
| Government | | |
| Unemployment insurance | −0.220 | −0.072 |
| | (0.020) | (0.005) |
| Total government | −0.232 | −0.104 |
| | (0.024) | (0.008) |
| Total transfers | −0.258 | −0.165 |
| | (0.012) | (0.016) |
| Taxes paid | 0.282 | 0.204 |
| | (0.022) | (0.034) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on change in head's earnings from regression specification like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument earnings loss with an unemployment dummy, as described in the text. The sample is confined to heads who experience an unanticipated period of unemployment in a given year. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

large literature on the *added worker effect*, which finds that there is no strong effect of the husband's unemployment on the labor supply of the wife.[52]

By contrast, the response of transfer income is much larger in the PSID than in the CEX. While it is highly significant in both data sets, in the PSID we find that for each dollar lost in income due to unemployment, transfers rise by 26 cents. Most of this increase is through government transfers, in particular, through income flows from the unemployment insurance program. The finding of a 22 cent rise in unemployment insurance payments for each dollar of earnings lost through unemployment is sensible, in that unemployment insurance receipt rates among the unemployed are around 40 percent and the average replacement rate over this period averages roughly 50 percent.[53] The CEX results for transfers are much smaller. In both data sets, there is also a highly significant and sizeable response of tax payments.

52. See Gruber and Cullen (1996) for a review of this literature and some new evidence that confirms the absence of an added worker effect.

53. See Blank and Card (1991) on receipt rates, and Gruber (1997) on replacement rates.

Overall, we find that about 50 to 55 cents of each dollar of earnings loss due to unemployment is compensated by increased income flows, primarily through taxes and transfers. Taken together with our consumption estimates, this suggests that only 20 to 25 cents of each dollar of earnings loss due to unemployment is reflected in dissaving. Thus for unemployment-induced earnings variation, relative to earnings variation in general, there is a larger consumption smoothing role for transfers and a smaller role for saving.

*Are Our Results Consistent with Wealth Holding?*

One weakness of our analysis is that we do not have direct data on saving; instead, we define the saving response as a residual. This suggests that there is value in confirming our contentions in terms of actual wealth or saving. For example, our findings for unemployment imply that a sizeable share of the resulting income loss is financed by dissaving. However, this is potentially at odds with the well-known fact that most households in the United States have low asset holdings. This fact has particular salience, considering that the households where the head experiences unemployment are largely drawn from the lower end of the income distribution. Moreover, unemployed households are likely to face serious liquidity constraints when trying to borrow to finance the unemployment spell. Thus it is natural to ask whether household asset holdings are plausibly large enough to finance one-quarter of the loss in earnings due to unemployment.

This question is addressed elsewhere by Gruber, using data from the Census Bureau's Survey of Income and Program Participation (SIPP).[54] We use the data from this study, but add sample restrictions to match our PSID and CEX samples (males aged twenty to fifty-nine). The SIPP is a large, nationally representative survey, which follows a sample of roughly 15,000 households for a period of two to three years. Households are interviewed every four months and provide retrospective information on each of the previous four months, including weekly employment status information. At two points in each SIPP panel, households also provide an asset inventory. The advantage of the SIPP data, for our purpose, is that one can match asset holdings shortly before an unemployment spell with the precise income loss from that spell.

54. Gruber (1996b).

We use the sample of households for which the head, within one year of the wealth interview, experiences a job separation that results in unemployment of more than four weeks. We then compare the ex ante wealth holdings of those households to the ex post income loss due to the head's unemployment. The income loss is defined by taking the head's ex ante after-tax weekly earnings and multiplying by weeks of unemployment; we then add back the unemployment insurance benefits received during the spell. This will not account for other offsetting income flows from spousal labor supply or transfers other than unemployment insurance, but our PSID results show that these are small relative to the unemployment insurance and tax effects that are incorporated in the analysis.[55]

The adequacy of wealth holding to smooth consumption in the face of income loss is defined for two different concepts of wealth: gross liquid assets, which include interest-earning assets in banks and other institutions, household equity in stocks and mutual funds, and other assets, such as bonds and checking accounts; and total net worth, which is the sum of all household net assets, liquid and illiquid, including (in addition to those previously mentioned) equity in retirement savings accounts, homes, vehicles, and personal businesses. Which one of these is the appropriate concept for measuring the ability to finance consumption during unemployment is unclear. Liquid assets are easily accessed to finance income loss, whereas illiquid assets may be harder to tap. Nevertheless, there is only a small penalty for drawing down retirement savings, and those who have some equity in their residences may be able to take out a home equity loan, so focusing only on liquid assets may seriously understate the resources available to households.

The results of our calculations are reported in the top panel of table 11.[56] We present results for the median ratio of wealth to lost income, as well as the distribution of the sample across categories defined by wealth-to–income loss ratios as follows: wealth holdings that are less than 10 percent of the expected income loss, less than 25 percent, less

55. We compare wealth holding only to realized ex post income loss, not to ex ante expected income loss. Since the duration of unemployment is endogenous, this approach may bias upward our adequacy calculations, because those with low wealth will have relatively short spells.

56. These results include imputed wealth observations in the SIPP. Since imputation rates are much higher at higher wealth levels, excluding them would potentially skew the sample.

**Table 11. Ratio of Wealth to Unemployment-Induced Income Loss, by Educational Attainment[a]**

| *Wealth as share of income loss* | *Definition of wealth* | |
|---|---|---|
| | *Gross liquid assets* | *Net worth* |
| *Entire sample* | | |
| Median | 0.35 | 6.35 |
| Less than 0.10 | 0.37 | 0.17 |
| Less than 0.25 | 0.46 | 0.20 |
| Less than 0.50 | 0.54 | 0.23 |
| Less than 1.00 | 0.64 | 0.28 |
| *High school dropouts* | | |
| Median | 0.05 | 3.78 |
| Less than 0.10 | 0.55 | 0.21 |
| Less than 0.25 | 0.62 | 0.25 |
| Less than 0.50 | 0.69 | 0.30 |
| Less than 1.00 | 0.76 | 0.36 |
| *High school graduates* | | |
| Median | 0.40 | 7.26 |
| Less than 0.10 | 0.33 | 0.15 |
| Less than 0.25 | 0.43 | 0.18 |
| Less than 0.50 | 0.53 | 0.21 |
| Less than 1.00 | 0.63 | 0.27 |
| *College graduates* | | |
| Median | 1.24 | 7.11 |
| Less than 0.10 | 0.16 | 0.15 |
| Less than 0.25 | 0.23 | 0.16 |
| Less than 0.50 | 0.32 | 0.18 |
| Less than 1.00 | 0.45 | 0.21 |

Source: Author's calculations based on data from the Census Bureau's Survey of Income and Program Participation.

a. Each entry gives either the median wealth-to-income loss ratio or the share of the given subsample whose wealth holdings are less than the given share of their unemployment-induced income loss.

than 50 percent, and less than 100 percent. These categories are defined cumulatively, so that one minus the final row gives the share of the sample whose wealth is higher than their income loss.

Our results are consistent with the notion that assets are used to smooth about one-quarter of the income loss from unemployment. We find that the median household has gross liquid assets equal to 35 percent of its income loss from the unemployment spell, and that about one-half of households have assets greater than one-quarter of their income loss. Thirty-six percent of households have assets that are greater than their entire income loss. If one includes illiquid assets, wealth appears more than adequate to finance the kind of consumption

smoothing that one observes in the PSID and the CEX. Thus our finding is not implausible, even if families cannot borrow.

*Heterogeneity*

Thus far, the analysis has considered the average effect of income variation on all male heads in our PSID and CEX samples. But there is reason to suspect considerable heterogeneity in the ability of families to smooth unemployment shocks. There is marked heterogeneity in wealth holding in the United States. As R. Glenn Hubbard, Jonathan Skinner, and Stephen Zeldes report, the median nonhousing wealth holdings of households headed by thirty- to thirty-nine-year-old high school dropouts is only one-sixth that of households headed by college graduates in the same age range.[57] This suggests that higher income households may be better able to smooth consumption. Yet, to the extent that government social insurance programs are means tested, or have redistributive benefits structures, there may be more scope for consumption smoothing among low-income households.[58] Thus an important question is whether there are differences in the ability to smooth income variation across households of different levels of wealth holding.

In order to examine heterogeneous responses to earnings variation, one must divide the sample by underlying ability to smooth consumption. Using actual wealth holdings for this purpose is problematic, however, for two reasons. First, wealth holding is endogenous to earnings variation. That is, if a family has low asset holdings, this may not imply that it is unable to smooth consumption, but rather that it has already drawn down its wealth to finance consumption smoothing. Second, while there is information on wealth holding for two points in time (1984 and 1989) in the PSID, there is no wealth information in the CEX.[59]

57. Hubbard, Skinner, and Zeldes (1995).

58. Indeed, Hubbard, Skinner, and Zeldes (1995) suggest the redistributive social insurance structure as an explanation for the skewed nature of asset holding. Gruber and Yelowitz (1997) offer empirical support for this contention, finding that means-tested health insurance under the medicaid program has a large crowd-out effect on wealth holding.

59. In theory, the final CEX interview collects data on wealth; in practice, these data are missing for most households.

We therefore consider heterogeneity by the educational attainment of the head. Education is strongly correlated with wealth holding, as noted above, but it is exogenous to the underlying income variation over our sample period (since we exclude students). We divide our data into three categories for this analysis: high school dropouts, high school graduates who have not completed college, and college graduates.

The correlation of education with wealth holding is also documented in table 11, which shows the ratio of asset holdings to income loss from unemployment by education category. There is clearly a skewed distribution of gross liquid asset holdings, which might be associated with differential ability to smooth consumption. The median household headed by a high school dropout who separates from his job has gross liquid assets amounting to only 5 percent of the resulting income loss; only 38 percent of these households have liquid assets greater than 25 percent of the income loss. By contrast, the median household headed by a college graduate who separates from his job has gross liquid assets 1.2 times the income loss, and over three-quarters of these households have assets greater than 25 percent of the income loss. The distribution of net worth, relative to income loss, is somewhat less skewed.

The relationship between unemployment-induced earnings variation and consumption, by education, is shown in table 12. The upper panel considers the effect of unemployment-induced earnings variation on consumption by educational group. In the PSID, we find that the effects are much larger among households headed by high school dropouts and graduates than among those headed by college graduates; there is essentially no effect of unemployment-induced earnings variation on the consumption of the highly educated. In the CEX, the pattern for food and housing consumption is flatter. For total nondurables, there is a larger effect on households headed by high school dropouts than among those headed by college graduates, but the difference is not large. There is an enormous difference, however, for durables expenditures: there is no effect of unemployment-induced earnings variation on the durables purchases of households headed by college graduates, whereas there is a 53 cent drop in the durables purchases of households headed by high school dropouts for each dollar reduction in earnings due to unemployment. Thus for total consumption expenditures, there is large heterogeneity by educational group.

**Table 12. Estimating Responses to Unemployment-Induced Earnings Changes, by Educational Attainment[a]**

| | *Coefficient on head's earnings* | | | | | |
|---|---|---|---|---|---|---|
| | *PSID data* | | | *CEX data* | | |
| *Dependent variable* | *High school dropouts* | *High schoool graduates* | *College graduates* | *High school dropouts* | *High school graduates* | *College graduates* |
| *Expenditure* | | | | | | |
| Food | 0.075 | 0.098 | 0.019 | 0.096 | 0.049 | 0.055 |
| | (0.045) | (0.024) | (0.032) | (0.031) | (0.016) | (0.026) |
| Housing | 0.077 | 0.024 | −0.007 | −0.008 | 0.007 | 0.017 |
| | (0.043) | (0.023) | (0.028) | (0.018) | (0.013) | (0.020) |
| Food plus housing | 0.091 | 0.126 | −0.009 | 0.089 | 0.062 | 0.074 |
| | (0.078) | (0.037) | (0.041) | (0.038) | (0.023) | (0.037) |
| Nondurables | . . . | . . . | . . . | 0.197 | 0.079 | 0.129 |
| | | | | (0.103) | (0.062) | (0.097) |
| Durables | . . . | . . . | . . . | 0.525 | 0.142 | −0.079 |
| | | | | (0.197) | (0.132) | (0.180) |
| Total consumption | . . . | . . . | . . . | 0.697 | 0.200 | 0.072 |
| | | | | (0.239) | (0.152) | (0.213) |
| *Source of income smoothing* | | | | | | |
| Wife's earnings | −0.014 | −0.006 | −0.072 | 0.014 | −0.094 | −0.294 |
| | (0.081) | (0.045) | (0.048) | (0.061) | (0.042) | (0.075) |
| Transfer payments | | | | | | |
| Government | | | | | | |
| Unemployment insurance | −0.301 | −0.223 | −0.138 | −0.106 | −0.079 | −0.040 |
| | (0.052) | (0.024) | (0.034) | (0.017) | (0.007) | (0.006) |
| Total government | −0.317 | −0.254 | −0.151 | −0.166 | −0.111 | −0.056 |
| | (0.062) | (0.031) | (0.042) | (0.030) | (0.011) | (0.009) |
| Total transfers | −0.338 | −0.275 | −0.148 | −0.198 | −0.184 | −0.099 |
| | (0.066) | (0.034) | (0.041) | (0.048) | (0.023) | (0.027) |
| Taxes paid | 0.264 | 0.313 | 0.232 | 0.278 | 0.181 | 0.206 |
| | (0.040) | (0.031) | (0.053) | (0.066) | (0.041) | (0.077) |

Source: Authors' calculations based on data from the PSID and the CEX.

a. Coefficient is that on decrease in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument earnings loss with an unemployment dummy, as described in the text. The sample is confined to heads who experience an unanticipated period of unemployment in a given year. The PSID sample period is described in table 2, note a. The CEX sample period is 1980–93.

One potential problem with these findings is that the unemployment indicator may be capturing different types of shocks in the different educational groups. For example, the effect of unemployment-induced income variation on consumption may be larger for the lowest educational group because unemployment spells are more severe for this group, and consequently are perceived as more permanent in making consumption decisions. However, there is no evidence that unemployment spells are differentially severe across these groups. The average durations of unemployment spells within the three educational groups in the PSID are 15.3 weeks for high school dropouts, 14.3 weeks for high school graduates, and 13.8 weeks for college graduates. Similarly, examination of the effect of unemployment on earnings (the first stage of the two-stage least squares results in table 9) shows that the coefficients rise with education in the PSID (indicating that spells are actually more severe for the most educated) and are roughly equal in the CEX. Thus the differential consumption response to unemployment-induced income variation across educational groups shown in table 12 appears genuinely to reflect differential ability to smooth transitory income variation.

The fact that the effect of unemployment-induced earnings variation is larger for the lowest educational groups suggests that redistributive transfers are not offsetting underlying differences in wealth holding. Evidence on this proposition is shown in the lower panel of table 12, which documents the sources of consumption smoothing by educational group.

We find that transfers in response to unemployment-induced earnings variation are largest for the lowest educational group. This difference is driven by government transfers; nongovernment transfers appear to respond roughly equally across these groups. Off-setting these redistributive transfers, however, is a regressive response of spousal labor supply, particularly in the CEX. This may reflect the fact that the wives of the most educated have the highest potential spousal wages and commensurately greater ability to smooth earnings variation.[60] The tax

60. This finding is broadly consistent with the cross-sectional evidence of Juhn and Murphy (1997), who find that the wives of low-earning men who saw large declines in wages in the 1970s and 1980s did not earn more to off-set these declines. By contrast, they find a relative rise in earnings among the wives of high-earning men whose wages rose over this period.

**Table 13. Comparing Responses to Unemployment, by Wealth Group[a]**

| | *Coefficient on head's earnings* | |
|---|---|---|
| *Dependent variable* | *Low wealth*[b] | *High wealth*[b] |
| Food | 0.171 | 0.059 |
| | (0.077) | (0.036) |
| Housing | 0.061 | 0.026 |
| | (0.050) | (0.036) |
| Food plus housing | 0.255 | 0.084 |
| | (0.109) | (0.052) |

Source: Authors' calculations based on data from the PSID.

a. Coefficient is that on decrease in head's earnings from regression specifications like that described in table 2 and its notes. Standard errors are shown in parentheses. Estimates are by IV and instrument earnings loss with an unemployment dummy, as described in the text. The sample is confined to heads who experience an unanticipated period of unemployment in a given year. Sample period is 1985–92, excluding 1988–89.

b. Those in the bottom 75 percent of the wealth distribution in 1984.

c. Those in the top 25 percent of the wealth distribution in 1984.

effect is roughly equal across groups in both data sets, which is somewhat surprising, since marginal tax rates are, on average, higher for the more educated.

Our results therefore suggest important heterogeneity in the response of consumption to unemployment-induced earnings variation. For low-education heads, there is an enormous response of durables expenditures to earnings reductions through unemployment, and a large response of net government transfers as well; as a result, there is no implied smoothing through saving. This is consistent with the low asset holdings of this group. For college graduate heads, there is little consumption response to unemployment-induced earnings reductions, and a somewhat smaller role for government transfers than for the low-education group (although a larger role for spousal labor supply); for this group, dissaving offset roughly one-half of the earnings loss due to unemployment.

While cutting by education solves the problems noted above, it is a somewhat indirect approach to describing the effect of liquidity constraints. Therefore in table 13 we confirm our results for the educational groups by directly dividing our PSID data in terms of ex ante wealth holdings. We divide the post-1984 observations on the basis of their gross liquid asset levels in 1984; by using ex ante wealth, we hope to mitigate the dependence of wealth holding on unemployment.[61] We

61. However, this does not fully solve this problem if unemployment is serially correlated.

compare the top quarter of the wealth distribution to the bottom three quarters. The results are consistent with our findings by educational group: there is a much larger effect of income changes on consumption within the low-wealth group than within the high-wealth group.

## Consumption Instability in the United States

Our findings from the microdata analysis of the previous sections suggest that families are fairly well able to smooth consumption against earnings variation. For nondurables consumption, we find that each dollar of earnings change leads to a change in consumption of less than 10 cents. Even when one adds durables expenditures, the total consumption response is only 25 cents for each dollar change in earnings.

In this section, we extend the aggregate analysis discussed at the start of this paper to describe time patterns in the variation of consumption. We follow the same approach as that described above, replacing head's earnings with food or food plus housing expenditures as the dependent variable. Figure 3 plots the time trends for variation in labor income, food consumption, and food plus housing consumption. This figure suggests two findings.

First, there is little countercyclicality in the variation of food consumption. There is more countercyclicality in the food plus housing series, but this may reflect cyclical asset pricing effects in the housing market that are not appropriately captured by the housing deflator.

Second, there is a secular rise in the variation of consumption. The proportional rise in the variation of food consumption is of almost exactly the same magnitude as that of head's earnings; over the entire period, food variation rises by 74 percent, while earnings variation rises by 76 percent. For food plus housing, the rise is smaller (43 percent), but still quite sizable. If we decompose the food series into spending on food at home and spending on food away from home, we find that the rise is driven by an increase in the variation of food at home; the variation of food away from home actually declines. Figure 4 shows the results using our fixed effects model. They are fairly similar; the rise in food plus housing variation over this period is roughly one-half as large as the rise in labor earnings variation.

The first finding, that consumption variation does not follow the

**Figure 3. Variation of Earnings and Consumption Expenditures, 1970–91**[a]

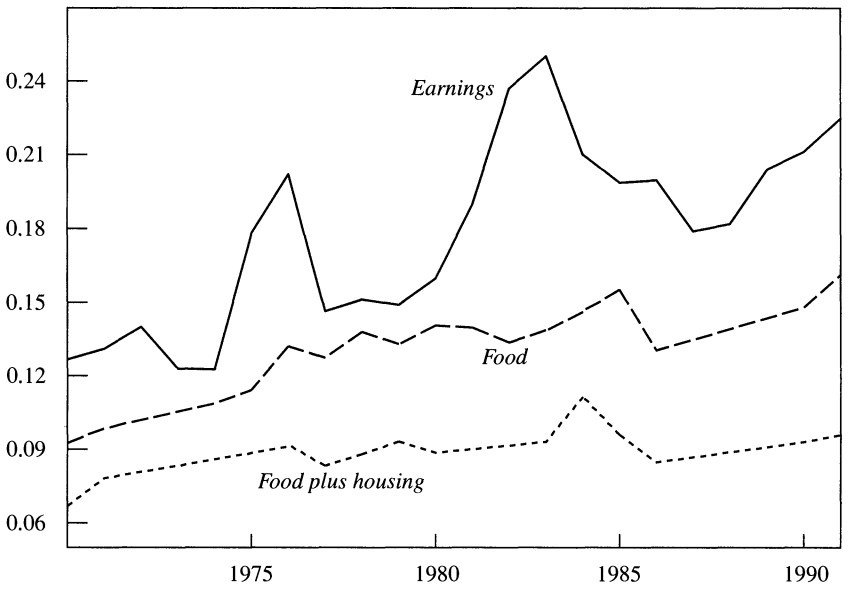


Source: Authors' calculations based on data from the PSID.

a. Figure plots, by year, the mean of the squared residuals produced in regressing changes in log earnings or consumption expenditure on the family characteristic controls and year dummies described in the text and used throughout the paper.

countercyclical pattern of earnings variation, is consistent with our micro-level evidence, which shows that individuals are largely able to smooth year-to-year variation in earnings. This point is made even more starkly in figures 5 and 6, which decompose the aggregate trend in earnings instability into its two components: variation in hours and variation in earnings per hour.[62] It is clear from these figures that the countercyclical pattern in earnings instability is driven by variation in hours of work, not in the wage rate. But in the microdata regressions we show that variation in hours of work is readily smoothed by house-

62. This is a literal decomposition of the earnings last year variable (using reported earnings divided by reported hours), so that earnings per hour is not our hourly wage instrument, which is an independent measure of wages. When we instead use our hourly wage instrument to compute this time series, the pattern is similar.

**Figure 4. Variation of Earnings and Consumption Expenditures, Individual Fixed Effects Included, 1970–91[a]**

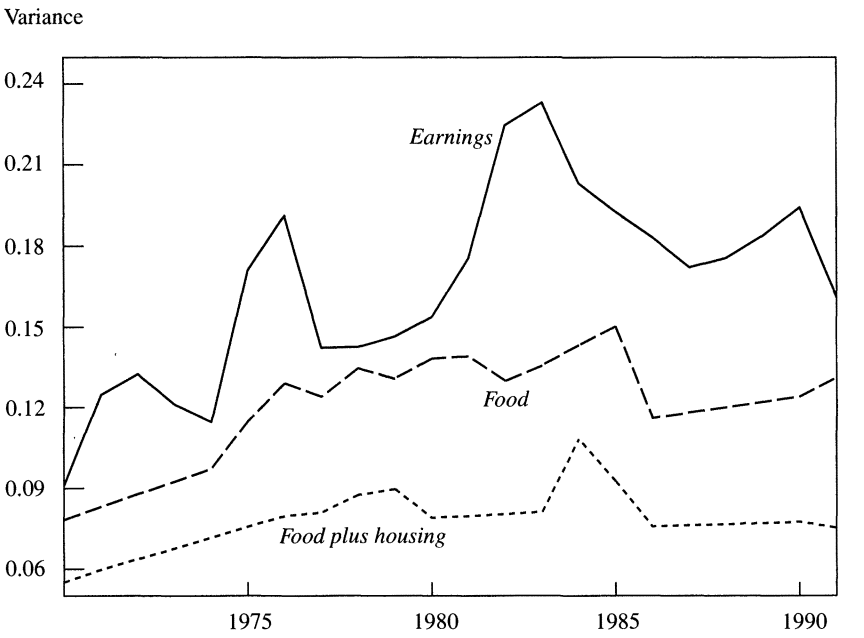


Source: Authors' calculations based on data from the PSID.
a. Produced by the same method as figure 3, except that this equation adds a person-specific fixed effect.

holds.[63] Thus it is not surprising that the variation of consumption does not display the countercyclicality seen in the variation of earnings.

The second finding, that variation in consumption is trending upward along with variation in earnings, is harder to reconcile with the micro-level evidence. The PSID data do indicate that there is a significant relationship between wage variation and consumption variation, and wage variation shows a upward trend in figures 5 and 6. But the magnitude of this relationship cannot explain the close parallel between consumption variation and earnings variation. In particular, for each 10 percent rise in the wage, there is a 1.13 percent rise in food consumption. This implies that the increased variation of the change in wages can explain less than 2 percent of the increased variation of the

63. Even variation through unemployment, which may be more meaningful when considering the cyclical pattern, is mostly smoothed.

**Figure 6. Variation of Earnings, Hours Worked, and Wages, Individual Fixed Effects Included, 1970–91[a]**

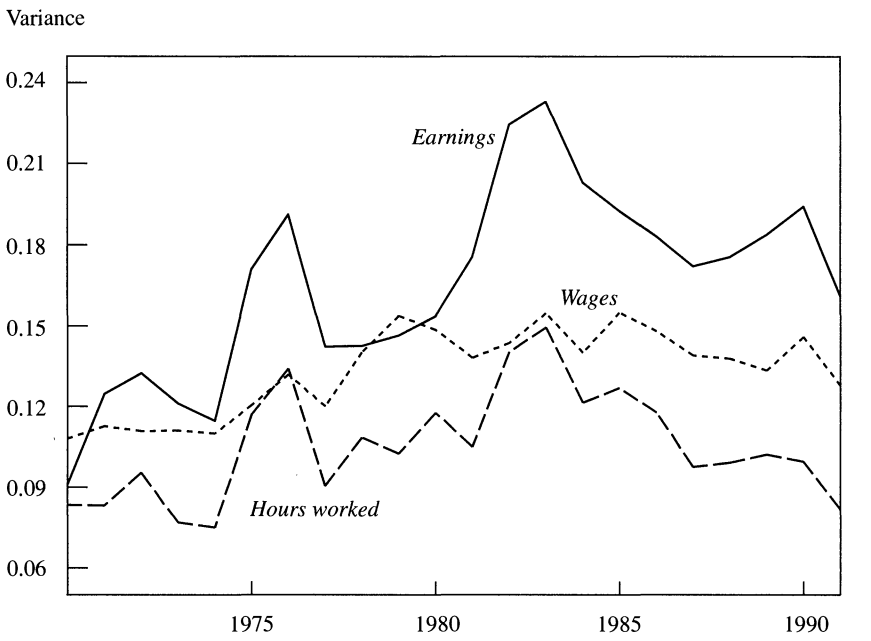


Source: Authors' calculations based on data from the PSID.
a. Produced by the same method as figure 5, except that this equation adds a person-specific fixed effect.

change in food consumption. As confirmation of this point, we plotted (not shown) the yearly squared residuals from an instrumental variables regression of change in consumption on change in earnings—the basic IV model from tables 3 and 4. Even after conditioning out the effect of earnings in this way, the variation of consumption rises virtually as fast as the absolute variation in consumption shown in figures 3 and 4.

Two additional pieces of evidence suggest that the time trends in earnings and consumption instability are not causally related. First, the time patterns of these series do not match. Three-quarters of the secular rise in earnings instability in the PSID occurs after 1980. But only 40 percent of the rise in food consumption instability and 10 percent of the rise in food plus housing consumption instability occur after 1980. Second, the patterns of these series by educational group do not match. If we disaggregate the earnings instability trends by educational group, we find a much smaller effect on college graduates than on lower

**Figure 5. Variation of Earnings, Hours Worked, and Wages, 1970–91**[a]

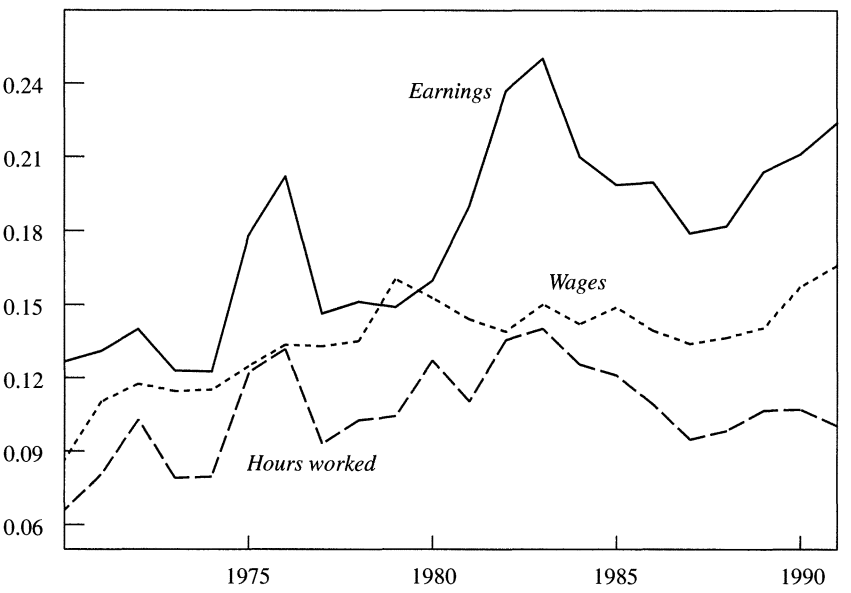


Source: Authors' calculations based on data from the PSID.

a. Figure plots, by year, the mean of the squared residuals produced in regressing changes in log earnings, log wages, or log hours worked on the family characteristic controls and year dummies described in the text and used throughout the paper.

educational groups. But if we disaggregate the consumption instability trends, we find a much larger effect on college graduates than on lower educational groups. Thus it does not appear that rising instability in earnings is driving the time-series pattern of consumption instability.

### *Are These Trends Spurious?*

Given the lack of a causal link between rising variation in labor earnings and in consumption, a concern arises that one or both of these trends might be spurious. One possible source of a spurious change in the variation of consumption over time is change in family composition. But we include a detailed set of compositional controls in our regression and also limit our sample to families with male heads. If we further restrict the sample to heads with no change in marital status over the

sample period, the results are similar: the magnitude of the rise in variation is smaller for both earnings and consumption (roughly, a 50 percent rise in both earnings and food consumption instability, rather than the 75 percent rise shown above), but the increases over time are of the same relative size. Thus change in family composition does not appear to be driving these trends.

Another potential concern is that our findings arise from increased early retirement among the old or weakening attachment to the labor force among the young, either of which could produce increasingly correlated changes in consumption and income. However, the time patterns in consumption and earnings are basically unchanged if we restrict our sample to males aged twenty-five to fifty-four, rather than twenty to fifty-nine.

Another concern is that the trends in both earnings and consumption might be driven by increasing measurement error in the PSID. The approach that we use in our aggregate time-series analysis is unable to distinguish true transitory movements from movements due to measurement error. As Gottschalk and Moffitt point out, there is no reason to think that measurement error in the PSID grew systematically worse over this period.[64] The fraction of interviews that is carried out by phone, for example, has remained constant (roughly 90 percent) since 1975. There has been a shift to computer-aided interviewing, but this did not begin until 1991. Nor should attrition from the PSID explain this finding, since we are using the sample weights.

The PSID does impute income and consumption for some households, but the imputation rate is relatively low. Over the sample period, on average, fewer than 3 percent of observations are imputed for any of our key variables. If the imputation rate is changing over time it could bias our time trends, since the distribution of imputed observations is presumably more compressed than the true distribution. However, there is little noticeable time trend in the imputation rate, with the exception of the first three years of the food consumption data, when it was relatively high. And if we reestimate our time trends excluding imputed values, the results are substantively unchanged.

Furthermore, if the quality of PSID interview data were simply deteriorating over time, one would expect that all dollar quantities would

64. Gottschalk and Moffitt (1994).

similarly show a pattern of rising instability. But this is not the case. If we plot variation in the earnings of other family members, for example, we see a secular decline in instability.

Thus we conclude that the time-series increase in the variation of both earnings and consumption is real, but that these two upward trends are not causally related. This raises the question of why the variation in consumption has increased over time. One possibility is that we are obtaining a misleading picture from our subcomponents of consumption. Unfortunately, we cannot test this hypothesis in the CEX, since the majority of the rise in consumption variation occurred in the 1970s, before our sample begins. Over the 1980s, the rise in variation of total consumption in the CEX is similar to that of food plus housing consumption in the PSID, but this is not a strong test, since both series are basically flat in these years.

Another possibility is that the dynamic of consumption smoothing has changed over time. Perhaps because of some change in the underlying process that determines earnings variation (for example, more severe unemployment spells), families may be increasingly unable to smooth earnings variation. But if we allow our microdata estimates to vary with time, we find no evidence of a linear time trend in the ability of families to smooth consumption.[65]

## Conclusions

Earnings variation is a persistent and growing feature of the U.S. labor market. The past twenty years have seen a 76 percent rise in the earnings variation of male household heads. Key to interpreting the welfare implications of this growing earnings instability is understanding the extent to which earnings variation is translated into household consumption. We present a variety of evidence on the ability of families to smooth earnings variation in order to address this issue.

Our primary conclusion is that families are well able to smooth variation in the earnings of household heads. Our instrumental variables

65. In the PSID, there is a larger effect of unemployment-induced variation in earnings in later years, as is apparent from comparison of the post-1984 coefficients in table 13 with those in table 9. But the difference in effects over time is not large enough to explain our consumption time trend. Nor is this finding robust to the CEX.

estimates suggest that only roughly 10 percent of the variation in the head's earnings is translated into nondurables consumption. There is a large effect on durables expenditures, but it is difficult to draw strong welfare inferences from this finding without a clean measure of the consumption flow from durables. Our results are robust across both year-to-year variation in earnings and large exogenous movements in earnings due to unemployment.

We also find that the government is an important source of insulation against earnings variation for family consumption, through both transfers and—more important—taxes. For year-to-year variation in earnings, roughly half of the smoothing that is done by families is through government taxes and transfers and roughly half is through saving and dissaving. For earnings loss due to unemployment, transfers from the government are somewhat larger, so that saving plays a smaller role.

But our findings for unemployment point out important heterogeneity in the ability of families to smooth earnings losses. Consumption expenditures, particularly on durables, are much more responsive to unemployment-induced earnings reductions for low-education or low-wealth groups than for high-education or high-wealth groups. These differentials emerge despite a redistributive government transfer system that replaces a higher share of the earnings loss due to unemployment for lower income groups. These findings are consistent with the skewed nature of wealth holding across educational groups: the median high school dropout head of household who becomes unemployed has gross liquid assets of only 5 percent of the income loss from unemployment, while the median college graduate head has liquid assets that are 124 percent of the income loss.

We also present aggregate evidence on the variation of earnings and of consumption and report four key findings. First, we confirm the conclusion of Gottschalk and Moffitt that there has been a secular rise in earnings instability.[66] Second, we highlight the countercyclical nature of this instability, which is entirely driven by instability in hours of work. Third, we find that this countercyclical earnings instability is not reflected in consumption, which is consistent with our contention that individuals can smooth variation in their hours of work.

But we conclude with a puzzle. Despite the finding that earnings

66. Gottschalk and Moffitt (1994).

variation is largely smoothed, we find that there are parallel secular rises in the instability of earnings and of consumption. This suggests that some time-series mechanism is increasing the instability of consumption independent of earnings variation. An important priority for future work is to confirm and sort out the source of this rising instability in consumption.

Our findings raise two important issues for policy design. First, the government plays an important consumption smoothing role through both the tax and transfer systems. In particular, our results suggest that discussions of tax policy should not ignore the role of taxation as a consumption smoothing mechanism. In addition, while the wealth holdings of the unemployed appear adequate to finance one-quarter of their earnings reduction due to unemployment, families would be hard pressed to finance a much larger share. This suggests that the role of smoothing through taxes and transfers is particularly important for the unemployed. Yet this is a static comparison, which does not account for the fact that the earnings loss from unemployment might respond to the underlying tax and transfer scheme. Future work incorporating the dynamics of unemployment into calculations such as these could usefully inform policymaking in this area.

Second, the redistribution inherent in government transfers for unemployment is not sufficient to overcome the differential ability of households to smooth consumption over unemployment spells. The dramatic skewness in asset holding suggests that programs such as unemployment insurance could increase their value as total consumption insurance through further progressivity in the determination of benefits. Such increased progressivity would, once again, have implications for unemployment behavior that must be counterbalanced against the gains in consumption smoothing. But unless the elasticity of unemployment with respect to benefits is larger at low income levels, more redistribution would increase insurance without reducing incentives.[67]

67. A more radical alternative would be the asset testing of unemployment insurance benefits, but this might have important negative implications for asset accumulation. Powers (1996) and Gruber and Yelowitz (1997) document that asset holding is very responsive to asset testing through the aid to families with dependent children and medicaid programs.

# *Comments and Discussion*

**Robert A. Moffitt:** This paper by Susan Dynarski and Jonathan Gruber is a follow-up to, and extension of, a series of papers that Gruber has produced over the past few years addressing afresh the consumption smoothing and insurance effects of several public transfer programs.[1] In his paper on the unemployment insurance program, Gruber, improving on earlier work by Daniel Hamermesh, empirically examines whether the program smoothes consumption between periods of employment and unemployment and finds considerable evidence for smoothing.[2] In a greater departure from prior work, he studies whether the aid to families with dependent children (AFDC) program smoothes the consumption of women who experience the "event" of becoming a single mother (unmarried or divorced).[3] While a social insurance view of the AFDC program is not, by itself, a new idea, Gruber is the first to conduct a serious empirical examination of the program from this perspective. Janet Currie and Gruber examine the effect of the medicaid program on the health of babies in low-income families and find that it significantly reduces adverse outcomes.[4] Although the medicaid program, like AFDC, is not a traditional social insurance program, and although this study does not address consumption smoothing, it does reflect a more general interest in measuring the potential benefits of public programs rather than their disincentive effects, which have dominated the empirical literature on those programs.

1. The author would like to thank Christopher Carroll for a helpful discussion.
2. Gruber (1997); Hamermesh (1982).
3. Gruber (1996a).
4. Currie and Gruber (1996).

Dynarski and Gruber go far beyond a concern with the insurance effects of narrowly defined public transfer programs to ask whether there are general sources of insurance against shocks to income. Taking as their departure point a series of empirical studies that test whether risk pooling effectively eliminates cross-sectional variation in consumption changes in response to idiosyncratic income shocks—in particular, a paper by Cochrane—they seek to quantify the degree to which such insurance is present, by estimating the fraction of individual income variation that is smoothed away.[5] The most notable finding of the paper is that almost all income variation is smoothed. Their maximum estimate of the effect of a $1 change in head's earnings on family consumption, including spending on durables, is only 24 cents, implying that 76 cents is smoothed away. This estimate seems very high, as compared with the conventional wisdom that there is relatively little insurance against income shocks in the United States. Dynarski and Gruber conclude that familes are "fairly well able to smooth consumption in the face of variation in the heads' earnings." Dynarski and Gruber also find that smoothing occurs more due to an insurance effect arising from tax and government transfer payments than due to private transfers and the earnings of other family members, although there is a significant role for saving as well.

The authors' study is, in part, motivated by the findings of Gottschalk and myself that the transitory variation of earnings in the United States increased during the 1980s. They are to be commended for taking seriously the comment of Dickens concerning the welfare implications of such an increase.[6] More work needs to be done on both the source of that increase in variation and its effects on the consumption and labor market behavior of individuals and families. Dynarski and Gruber infer from their own results that the welfare implications of the increase in variation are not as unfavorable as might be expected, because consumption smoothing and insurance seem to be so effective in response to such transitory shocks.

I address three issues. The first concerns the prima facie plausibility of the amount of smoothing that Dynarski and Gruber find; the second concerns the interpretation of the estimates from the viewpoint of the

5. Cochrane (1991).
6. Gottschalk and Moffitt (1994); Moffitt and Gottschalk (1995); Dickens (1994).

permanent income hypothesis literature and the relationship of the authors' approach to the approaches in that literature; and the third concerns the puzzle in the time-series trends in income and consumption variation that they uncover.

The plausibility of the authors' results on smoothing depends, in part, on whether they are interpreted as arising from intertemporal smoothing—self-insurance—or from consumption insurance across individuals. I am not so bold as to attempt any statement of consensus on what the large empirical literature on the permanent income hypothesis has already shown, given the wide dispersion of estimates illustrated by Browning and Annamaria Lusardi.[7] However, it is fair to say that the conventional view is that there are relatively few opportunities for income insurance in the United States. For example, most negative earnings shocks are not compensated for by the unemployment insurance program, because they do not result in unemployment but are instead declines in earnings due to a change in jobs or the bad fortunes of a worker's employer. On average, the percent of income loss compensated by unemployment insurance equals the product of the fraction of income variation due to unemployment variation (that is, due to inflows and outflows from unemployment) and the replacement rate of the unemployment insurance program. My guess is that, at most, 10 percent of the variation of the year-to-year change in individual earnings is due to inflows and outflows from unemployment; and, if the replacement rate of the program is 0.50, as it is conventionally taken to be, this implies that smoothing of only 5 percent can arise from this source. Food stamps may be an important source of compensating income, but the participation rate in the program has never been more than 10 percent and has averaged closer to 8 percent over its history.[8] Yet Dynarski and Gruber find that as much as 12 percent of earnings variation, averaged over the entire population, is compensated by changes in government transfers (see table 8).

The authors find the tax system to be of greater importance for consumption smoothing; they find that as much as 35 percent of earnings variation is compensated by changes in tax burdens, implying that the average marginal tax rate in the United States is 0.35.[9] While such

7. See Browning and Lusardi (1996, table 5.1).
8. U.S. House of Representatives, Committee on Ways and Means (1996, p. 874).
9. In the PSID, the source of this estimate, taxes are calculated rather than actually

a rate could conceivably have occurred over the 1970s, and is consistent with some other calculations over that period, it is difficult to believe that the average marginal federal income tax rate today—averaged over the entire population—is so high, given the expansion of the Earned Income Tax Credit, the increase in the zero-tax threshold, and reductions in marginal rates embodied in the Tax Reform Act of 1986. It would have been helpful if Dynarski and Gruber had estimated separate replacement rates over time, for the increase in the transitory variation that Gottschalk and I find occurred primarily in the 1980s, when tax rates were lower. In addition, current policy should be based on current replacement rates, not historic ones.

Turning to the consistency of the authors' approach with that in the permanent income hypothesis literature, a major difference of this study is the emphasis on insurance and risk pooling. The authors are aware that considerable care is needed to separate insurance and risk-pooling effects from intertemporal mechanisms for consumption smoothing. A particularly clear discussion of the distinction is provided by Cochrane, who emphasizes that insurance is entirely a static—or cross-sectional—concept concerning interpersonal resource flows, holding aggregate output constant, whereas intertemporal smoothing, in theory, is an individual-specific mechanism for using saving and borrowing to smooth consumption flows over time (over a long enough time period, one could conceivably examine intertemporal smoothing on an individual-by-individual basis, without conducting any cross-sectional comparisons).[10] Dynarski and Gruber concede that it is difficult to separate consumption smoothing effects arising from insurance from the effects arising from intertemporal considerations, such as those suggested by the permanent income hypothesis literature, a point which has considerable importance. Ultimately, they decide to examine the effects of consumption smoothing from any source, and to take the simple approach of regressing contemporaneous consumption on contemporaneous earnings—though in first differences—as a type of reduced-form analysis that yields a coefficient reflecting smoothing in general.

A major difficulty with this position is that the specification that they assume is not a meaningful reduced form, in the usual sense. The

observed, which probably biases this coefficient upward. The estimate of 26 percent from the CEX may be more accurate.

10. Cochrane (1991).

authors mean to imply that their estimated coefficient is some average of smoothing from insurance and from intertemporal mechanisms, but their specification is inconsistent with those used in the permanent income hypothesis literature for measuring the latter; it is not measuring the same thing. In fact, the authors' equation 1 is most closely akin not to the permanent income hypothesis, but to the traditional Keynesian consumption function relating current income to current consumption.

The permanent income hypothesis literature is large.[11] In my reading, that literature provides at least five lessons for the Dynarski-Gruber study. The first is that a single, period-to-period change in income may contain a permanent shock arising from the presence of a random walk component in income. In the micro-level literature on earnings dynamics, Thomas MaCurdy, Abowd and Card, and Gottschalk and I all find evidence—also from the PSID—for a random walk in individual earnings in the United States.[12] To some extent, Dynarski and Gruber might welcome this source of bias because it would lead to a coefficient on earnings ($\Delta Y$) that is too large, whereas their study finds the coefficient to be smaller than expected. However, they instead seem to wish to interpret their coefficient as reflecting the effects of changes in permanent as well as transitory earnings, for they devote one section of their study to instrumenting actual earnings with hours worked (assumed to be transitory) and the hourly wage rate (assumed to be permanent). Leaving aside the issues of whether fluctuations in the wage rate can be taken as permanent—a position for which there is no obvious evidence—or fluctuations in hours of work can be taken to be transitory (what about permanent changes in health and work capacity?)—this section suggests that they take their estimates as partly reflecting the effects of permanent income on consumption. If so, that effect is not what is referred to as smoothing in the permanent income hypothesis literature, which is very explicit in interpreting smoothing as a response to unanticipated and transitory changes in income.

A second lesson relates to the possible presence of serial correlation in individual income or earnings. The major studies of the permanent income hypothesis—for example, Marjorie Flavin's work using aggre-

11. See Deaton (1992b) and Browning and Lusardi (1996).
12. MaCurdy (1982); Abowd and Card (1989); Moffitt and Gottschalk (1995).

gate data and Hall and Mishkin's using micro-level data—have gone to some pains to model the earnings or income process, assuming either autoregressive structures (as does Flavin), a combination of autoregressive and moving-average specifications (as do Hall and Mishkin), or these and a random walk as well.[13] Also, the three micro-level studies already referred to, which examine the earnings process but not consumption, find autoregressive structures of either order one or order two plus a moving-average component of order one in individual earnings. For this reason, most studies of the permanent income hypothesis include some form of lagged $\Delta Y$ to guarantee that the contemporaneous $\Delta Y$ is measuring unexpected shocks. Because Dynarski and Gruber do not include such lags, it is possible that their coefficient is, in part, picking up the effects of lagged shocks. A leading possibility, for example, is that consumption has been adjusted at some point in the past and that current consumption appears not to change in response to current changes in income because the adjustment has already taken place. More generally, current consumption should not respond to expected changes in income, but only to unexpected shocks; this is one of the key points of the permanent income hypothesis literature.

A third lesson is the more general point that it is, ultimately, difficult to avoid the need to model the dynamics of the earnings process in one way or another in order to isolate the components of income change that are permanent or transitory, or those that are expected or unexpected. Although making that distinction can be quite difficult and can be sensitive to specification, as the authors argue, in its absence the coefficient on the contemporaneous change in income is virtually impossible to interpret and, as already noted, is very far from corresponding to the effect of smoothing as that term is ordinarily used.

A fourth lesson from the permanent income hypothesis literature concerns the importance of measurement error in earnings, an issue examined at some length by Dynarski and Gruber. However, the instruments that they use for last year's earnings—their key regressor—are last year's hours of work and the current hourly wage rate at the

13. Flavin (1981); Hall and Mishkin (1982). An additional issue is that Dynarski and Gruber examine the effect of the earnings of the head—rather than family earnings—on family consumption, so that if the earnings of other family members are experiencing movements correlated with those of the head, the authors' estimated coefficient proxies some unknown change in family earnings.

time of the survey. The authors argue that because these three variables come from different questions on the survey, they are independent and free of correlated measurement error. Unfortunately, this is not sufficient to guarantee the absence of such correlation. For example, if the respondents in the survey are trying to be even halfway internally consistent—and let us hope that they are, in general!—then last year's earnings should be arithmetically a function of last year's hours of work, and measurement error in the latter should generate measurement error in the former. The contemporaneous wage rate is a stronger variable, because it is separated by a period of time from that covered by the earnings question, but how much of the measurement error is eliminated by this separation depends largely on the source of that error. For example, a respondent's omission of casual or informal earnings or of income from a second job (which should be reported but may not be) is likely to be made both last year and currently.

A fifth lesson concerns the importance of precautionary—buffer-stock—saving, which has been the focus of much of the recent research on the permanent income hypothesis. For example, Christopher Carroll finds in a simulation study that a simple model of buffer-stock saving can generate marginal propensities to consume out of transitory income that are in line with the empirical literature.[14] Taking Dynarski and Gruber's estimates at face value, the portion of income smoothed by saving might arise from this source. Yet the interpretation is very different if this is a major source of the smoothing that Dynarski and Gruber find, because it has significant welfare implications. For example, if precautionary saving rises with the variation of transitory income, there is a welfare loss associated with that change that is missed by the examination of consumption smoothing.

This brings me to my final remark, concerning the puzzle posed by the data showing an increase in consumption variation over time, which, if the cross-sectional smoothing result is true, should not have occurred in response to the increase in transitory income variation. The authors' figures are intriguing, and they are to be credited for being the first to show time trends in consumption variation matched up with those for income. Dynarski and Gruber argue that the two trends may not be causally related, but an alternative explanation is that the cross-

14. Carroll (1997, p. 28).

sectional smoothing coefficient found in their microdata analysis is a mixture of permanent and transitory consumption responses of various kinds, and that the mixture has changed over time. Put differently, it suggests that the structure of the relationship that the authors estimate may have changed over time, a hypothesis which could readily be examined with their model. In addition, their finding that the growth in consumption variation was greater in the 1970s than in the 1980s, exactly the reverse of the timing of the growth of variation in earnings, might be traceable to an increase in precautionary saving in the 1980s. There is much additional work to be done to explore this interesting set of issues.

**Gary Burtless:** The proposition examined in this paper would astonish most noneconomists. Susan Dynarski and Jonathan Gruber take seriously the idea that families can insure completely against variability in the earnings of their principal breadwinners. According to the full insurance hypothesis, individual consumption should not vary in response to idiosyncratic shocks in that individual's wealth or earnings. The existence of a variety of risk-sharing institutions and arrangements permits individuals and their families to smooth consumption fully in the face of individual-specific fluctuations in earnings. These arrangements allow consumption to remain constant, even when the breadwinner's wages take a nose dive.

While most noneconomists will be skeptical of this theory, many economists find it attractive. At least a few find the evidence for it persuasive.[1] The basic idea is similar to, although not quite the same as, that behind the permanent income—or life-cycle consumption—hypothesis. According to that theory, far-sighted workers rationally plan consumption over a full lifetime. In doing so, they take account of the likely path of their labor earnings as they age and prudently accumulate savings in anticipation of their retirement. Any transitory deviations in earned income will be smoothed by additions to or subtractions from household savings. Changes in the flow of earnings that are expected to be permanent will cause breadwinners to recalibratetheir lifetime consumption plans in order to stay within their lifetime budget constraints.

1. See, for example, Mace (1991) and Cochrane (1991).

The full insurance hypothesis goes beyond the permanent income model in one important respect. It assumes that workers and their families smooth consumption in the face of *all* idiosyncratic fluctuations in income, even those that are expected to be permanent. The permanent income model makes a clear distinction between unanticipated changes in flows of income that can be expected to last and changes that are only temporary. An unexpected income improvement that is permanent, such as an earnings gain associated with a promotion, will have a much bigger impact on the worker's consumption than an improvement that is only temporary, such as a one-time bonus for outstanding job performance. According to the full insurance hypothesis, however, neither of these kinds of earnings changes should affect the flow of consumption, so long as they are idiosyncratic to the individual earner.[2] Individual consumption should only be affected if the income fluctuation reflects an economywide change.

In estimating the permanent income model, the trick is to distinguish between changes that are thought to be temporary and those that are expected to be permanent. Making this distinction is not easy for the typical consumer; making it accurately is impossible for the econometrician. Economists have invested great ingenuity in plausibly separating out transitory and permanent income changes in order to estimate their different effects. In the full insurance model, it is not necessary to make this distinction, but the statistician must instead distinguish between income changes that are idiosyncratic to the individual and those that reflect permanent economywide movements.

Dynarski and Gruber emphasize another distinction between the full insurance and permanent income models. They suggest that the latter relies on self-insurance against earnings fluctuations (through saving and borrowing), whereas the former also considers interpersonal transfers, for example, across extended families and through social insurance. This will come as a surprise to many economists who have worked within the permanent income—or life-cycle consumption—framework. Martin Feldstein argues strongly for a version of the life-cycle model in which anticipated social security retirement benefits fully or partially offset private retirement saving. Hubbard, Skinner, and Zeldes offer a version in which, for a sizable minority of households, asset-tested

2. This follows from the model as presented in Mace (1991) and Cochrane (1991).

transfer programs erode the incentive to save.[3] In these life-cycle models, choices about the level and timing of household saving are made in light of incentives created by the social insurance and public assistance systems. The observed pattern of wealth accumulation is a predictable consequence of the design of those systems. Workers accumulate too little private wealth to finance their own retirement because they anticipate receiving social security pensions. Workers with low lifetime earnings accumulate proportionately less precautionary savings than workers with high wages because unemployment insurance and means-tested transfer programs offer them better protection when their earnings decline than is available to high-income workers. Table 12 offers indirect support for this theory. Workers with low educational attainment (and low expected earnings) have less net worth or liquid assets in relation to typical earnings loss due to unemployment than do workers with greater educational attainment.

The authors take the full consumption insurance model seriously, but they do not take it literally. They play down the importance of statistically rejecting the implications of full consumption insurance. With a large enough and good enough data set, the hypothesis that consumption is invariant to idiosyncratic movements in the earnings of the principal breadwinner would certainly be rejected. The authors focus instead on the more interesting question of how far actual consumption deviates from its predicted path under full insurance, and they closely examine the mechanisms that permit consumption to remain much more stable than earnings.

The paper offers a good introduction to the subject of consumption smoothing. It treats several interesting aspects of the issue in ingenious ways. Its conclusions rely on evidence drawn from two data sets rather than one, as is usual. This difference is particularly important. Poor data is the Achilles's heel of research in this area. Few data sets offer reliable measures of consumer income, and almost none provides good information about consumption. Dynarski and Gruber note that between 15 and 30 percent of the cross-sectional variation in earnings in the PSID is due to measurement error. Between 20 and 25 percent of the variance in the first difference of earnings in the CPS is apparently due to measurement error. It seems inevitable that measurement error in

3. Feldstein (1974); Hubbard, Skinner, and Zeldes (1994).

income will bias most studies toward a finding of full consumption smoothing, unless respondents' errors in reporting consumption are correlated with their errors in reporting income. Using two data sets rather than one does not eliminate this source of bias, but it assures us that the findings are not due to idiosyncratic measurement problems in a single data source.

It is natural to ask whether the authors confirm or reject the hypothesis of full consumption insurance. If they reject the model, how far does the actual path of consumption deviate from its predicted path under full consumption smoothing? My interpretation of the paper is that they reject it, but do not think deviations from full smoothing are particularly large, except in special circumstances. Their basic results (in tables 2–4) imply that families do not succeed in smoothing consumption fully. In both their PSID and their CEX samples, and under both the OLS and IV specifications, the authors statistically reject the hypothesis of full consumption insurance.

Whether the practical difference between actual consumption smoothing and full smoothing is large or small depends on the statistical specification that one favors. The OLS estimates imply much lower responsiveness of consumption to earnings changes than do the IV estimates. Although the authors appear to favor the latter, there is no clear explanation for the large differences between the results under the two specifications.

Assuming that the IV estimates are more accurate, how should one interpret the findings? The authors state in their introduction and again in their conclusion that families are "well able" to smooth consumption in the face of fluctuations in the earnings of the male breadwinner. They report that the average estimated elasticity of food consumption with respect to male earnings ranges between 0.142 (in the CEX) and 0.205 (in the PSID). The average elasticity of total consumption is 0.240, and that of nondurables consumption is 0.110. It is important to consider one's benchmark in assessing whether these deviations from full consumption smoothing are large or small. In comparison with the change in gross male earnings, these changes in consumption seem modest.

But male earnings represent only part of household income. If male earnings account for 70 percent of household income, a 1 percent reduction in male earnings will represent a loss of just 0.7 percent of family income (assuming other sources of family income remain un-

changed). This means that a 1 percent reduction in family income (caused by a 1.4 percent reduction in male earnings) reduces food consumption by between 0.2 and 0.3 percent, reduces total consumption by 0.34 percent, and reduces nondurables consumption by 0.16 percent. In addition, earned income is taxed under a progressive schedule. Thus a 1 percent change in gross earnings causes less than a 1 percent change in after-tax income from employment. For example, if the average tax on earned income is 15 percent, while the marginal tax is 30 percent, a 1 percent rise in gross earned income will increase net earnings by just 0.82 percent. Stated another way, gross earnings must increase by 1.21 percent to produce a 1 percent gain in net earned income. By implication, a 1 percent increase in *net* family income (produced by a 1.7 percent increase in *gross* male earnings) will boost food consumption by between 0.25 percent and 0.35 percent, will increase total consumption by 0.42 percent, and will raise nondurables consumption by 0.19 percent.

These estimates of the implied elasticity of consumption with respect to net family income are not intended to be exact. Exact calculations require more information than is provided in the paper. Rather, they show that the reported elasticities may understate the responsiveness of household consumption to changes in after-tax family income. In a naive model of household consumption, spending in each period is financed entirely out of income received in the period. Actual consumption does not come close to following this model; but neither does it come close to following the full consumption insurance model. This suggests to me that the welfare loss associated with increased earnings variability is sizable.

Results in the paper suggest that a large part of consumption smoothing is attributable to changes in government transfers and tax payments, and much of the remainder may be the result of changes in household saving. It is interesting to consider how consumption smoothing would be affected if government transfers or taxes were reduced. If consumers are rational and far-sighted, they will off-set the loss of government insurance by increasing their accumulation of savings. Whether they will boost saving enough to fully offset the loss of government insurance is an empirical question. Government insurance for prime-age men has declined over the past two decades. The fraction of new job losers who collect unemployment insurance benefits dropped by about 20

percent between 1980 and 1985. Also, the after-tax value of these benefits fell when compensation payments, which were once tax free, became fully taxable between 1978 and 1987. Marginal income tax rates were reduced, particularly for high-wage earners, as a result of changes to the tax law passed in 1981 and 1986. On the one hand, the reductions in unemployment compensation and marginal tax rates has meant that changes in gross male earnings are more fully reflected by equivalent changes in net family incomes. On the other hand, male earnings have become a less important component of family income as wives' earnings have become more important.

Near the end of the paper, Dynarski and Gruber present evidence that both variation in male earnings and variation in family consumption have increased over time. They find these parallel increases puzzling. The increase in consumption variation is much larger than can be accounted for by the increase in the variation of male earnings, in light of the fact that families are largely successful in smoothing the variation in the earnings of the male breadwinner. The parallel increases in consumption and in earnings variation suggest that families are less successful in smoothing consumption than was the case in the 1970s. Not only is earnings variation greater than it once was, but some source of consumption insurance that was available to families in the 1970s is weaker than it once was.

**General discussion:** Ben Bernanke reinforced a point made by both discussants, that the paper does not distinguish sharply between smoothing and insurance as explanations for the observed insensitivity of consumption to income changes. He suggested that it would have been useful to estimate directly the response of individuals' consumption to aggregate income and test whether that response is significantly greater than their response to idiosyncratic changes in their own income, as it should be with full insurance. Robert Hall added that the appropriate benchmark for testing the insurance model is the response implied by a simple permanent income model, rather than zero. A larger response than that implied by the permanent income model would suggest a failure such as liquidity constraints, while a smaller response would suggest a role for interfamily transfers. Hall also agreed with the discussants that it was difficult to interpret the results without making a better distinction between transitory and permanent shocks. Gruber responded that the paper did not fully test

the permanent income hypothesis, because such a test requires a full specification of transitory and permanent shocks, which depends crucially on what procedure is used to distinguish these shocks. He observed that while the equations being estimated should be interpreted as reduced-form equations, the procedures used were consistent with some common identification assumptions. One such assumption is that age and education affect permanent income; the variation remaining after conditioning on these is labeled transitory. Since the regressions in the paper include most of these variables on the right-hand side, these effects should be captured. Another common identifying assumptions is that all mean-reverting shocks are transitory. But transitory shocks may largely reflect measurement error, which the paper handles by the use of instrumental variables.

Robert Shiller questioned the implicit assumption that the relevant horizon for smoothing or insurance was one year. Conceptually, one can argue that the relevant horizon for insurance is an individual's lifetime, although, with just twenty-two years of PSID samples and thirteen from the CEX, the data for testing a model with such assumptions is not available. He also observed that the limited response of consumption to annual changes may reflect, in part, the difficulty of adjusting consumption rather than insurance. Having children in private schools or owning an expensive house are consumption decisions that cannot be quickly changed. The results may simply indicate that habits and consumption commitments are important.

John Abowd doubted that the two measures of earnings used in the paper actually have independent measurement errors. He suggested other ways to investigate the importance of measurement error in the PSID and CEX, such as looking at the within-quarter variation in the CEX and cross-year variation in the PSID. Lawrence Katz suggested it might be useful to concentrate on specific groups that would be expected to have quite different resources for smoothing the effects of income variation, such as college graduates and high school dropouts. Gruber replied that they had explored differences between some groups. They were baffled by the finding that the highest educational group, whose earnings variability had risen least, had the greatest increase in consumption variability. But other comparisons were more in line with expectations: those in the top quartile of the wealth distribution smoothed the effect of becoming unemployed much more than those in the lowest quartile.

George Borjas cautioned against drawing policy conclusions from the

estimated importance of government programs in smoothing shocks. In the absence of such programs, either the wage structure or the saving behavior of individuals might be quite different. Katz said that the paper had changed his view of which government programs are most important in helping people to smooth income shocks. While unemployment insurance and welfare help in smoothing transitory shocks, longer-term programs, such as the Earned Income Tax Credit, are needed to help the less educated workers who have taken the largest permanent earnings losses in the past twenty years. Hall speculated that the paper's finding of a change in the character of consumption and earnings dynamics could be explained by a trend away from short unemployment spells toward longer spells as a result of displacement. The recent Bureau of Labor Statistics survey shows that displacement has remained at very high levels following the 1990 recession. Since the consequences of permanent job loss have been estimated at 1.2 years of earnings, this may explain the change in the response of consumption.

## References

Abowd, John M., and David Card. 1989. "On the Covariance Structure of Earnings and Hours Changes." *Econometrica* 57(2): 411–45.

Altonji, Joseph G., and Aloysius Siow. 1987. "Testing the Response of Consumption to Income Changes with (Noisy) Panel Data." *Quarterly Journal of Economics* 102(2): 293–328.

Altug, Sumru, and Robert A. Miller. 1990. "Household Choices in Equilibrium." *Econometrica* 58(3): 543–70.

Attanasio, Orazio, and Steven J. Davis. 1996. "Relative Wage Movements and the Distribution of Consumption." *Journal of Political Economy* 104(6): 1227–62.

Barro, Robert J., and Chaipat Sahasakul. 1986. "Average Marginal Tax Rates from Social Security and the Individual Income Tax." *Journal of Business* 59(4, part 1): 555–66.

Blank, Rebecca M., and David E. Card. 1991. "Recent Trends in Insured and Uninsured Employment: Is There an Explanation?" *Quarterly Journal of Economics* 106(4): 1157–89.

Bound, John, and Alan B. Krueger. 1991. "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?" *Journal of Labor Economics* 9(1): 1–24.

Browning, Martin, and Thomas Crossley. 1996. "Unemployment Insurance Benefit Levels and Consumption Changes." Unpublished paper. McMaster University.

Browning, Martin, and Annamaria Lusardi. 1996. "Household Saving: Micro Theories and Micro Facts." *Journal of Economic Literature* 34(4): 1797–855.

Burgess, Paul L., and others. 1981. "Changes in Spending Patterns Following Unemployment." Unemployment Insurance Occasional Paper 81–3. Washington: Unemployment Insurance Service.

Carroll, Christopher D. 1997. "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis." *Quarterly Journal of Economics* 112(1): 1–55.

Cochrane, John H. 1991. "A Simple Test of Consumption Insurance." *Journal of Political Economy* 99(5): 957–76.

Cox, Donald. 1987. "Motives for Private Income Transfers." *Journal of Political Economy* 95(3): 508–46.

Currie, Janet, and Jonathan Gruber. 1996. "Saving Babies: The Efficacy and Cost of Recent Changes in the Medicaid Eligibility of Pregnant Women." *Journal of Political Economy* 104(6): 1263–96.

Cutler, David M., and Lawrence F. Katz. 1991. "Macroeconomic Performance and the Disadvantaged." *BPEA, 2:1991*, 1–61.

Deaton, Angus. 1992a. "Saving and Income Smoothing in Côte d'Ivoire." *Journal of African Economics* 1(1): 1–24.

———. 1992b. *Understanding Consumption*. Oxford: Clarendon.

Dickens, William T. 1994. "Comment." *BPEA, 2:1994*, 262–67.

Duncan, Greg J., and Daniel H. Hill. 1985. "An Investigation of the Extent and Consequences of Measurement Error in Labor-economic Survey Data." *Journal of Labor Economics* 3(4): 508–32.

Dynarski, Mark, and Steven M. Sheffrin. 1987. "Consumption and Unemployment." *Quarterly Jourrnal of Economics* 102(2): 411–28.

Farber, Henry S., and Robert Gibbons. 1996. "Learning and Wage Dynamics." *Quarterly Journal of Economics* 111(4): 1007–47.

Feldstein, Martin S. 1974. "Social Security, Induced Retirement, and Aggregate Capital Accumulation." *Journal of Political Economy* 82(5): 905–26.

Flavin, Marjorie A. 1981. "The Adjustment of Consumption to Changing Expectations about Future Income." *Journal of Political Economy* 89(5): 974–1009.

Friedman, Milton. 1957. *A Theory of the Consumption Function*. Princeton University Press.

Garcia, René, Annamaria Lusardi, and Serena Ng. 1997. "Excess Sensitivity and Asymmetries in Consumption: An Empirical Investigation." *Journal of Money, Credit, and Banking* 29(2): 154–76.

Garner, Thesia I. 1993. "Consumer Expenditures and Inequality: An Analysis Based on Decomposition of the Gini Coefficient." *Review of Economics and Statistics* 75(1): 134–38.

Gertler, Paul, and Jonathan Gruber. 1997. "Insuring Consumption Against Illness." Unpublished paper. Massachusetts Institute of Technology (April).

Gottschalk, Peter, and Robert Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *BPEA, 2:1994*, 217–54.

Gruber, Jonathan. 1996a. "Cash Welfare as a Consumption Smoothing Mechanism for Single Mothers." Working Paper 5738. Cambridge, Mass.: National Bureau of Economic Research (September).

———. 1996b. "The Wealth of the Unemployed: Levels, Distribution, and Utilization." Unpublished paper. Massachusetts Institute of Technology (September).

———. 1997. "The Consumption Smoothing Benefits of Unemployment Insurance." *American Economic Review* 87(1): 192–205.

Gruber, Jonathan, and Julie Berry Cullen. 1996. "Spousal Labor Supply As Insurance: Does Unemployment Insurance Crowd Out the Added Worker Effect?" Working Paper 5608. Cambridge, Mass.: National Bureau of Economic Research (June).

Gruber, Jonathan, and Aaron Yelowitz. 1997. "Public Health Insurance and

Private Savings." Unpublished paper. Massachusetts Institute of Technology (April).

Hall, Robert E., and Frederic S. Mishkin. 1982. "The Sensitivity of Consumption to Transitory Income: Estimates from Panel Data on Households." *Econometrica* 50(2): 461–81.

Hamermesh, Daniel. 1982. "Social Insurance and Consumption: An Empirical Enquiry." *American Economic Review* 72 (1): 101–13.

Hayashi, Fumio. 1985. "The Effect of Liquidity Constraints on Consumption: A Cross-Sectional Analysis." *Quarterly Journal of Economics* 100(1): 183–206.

Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff. 1996. "Risk-Sharing Between and Within Families." *Econometrica* 64(2): 261–94.

Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1994. "Expanding the Life-Cycle Model: Precautionary Saving and Public Policy." *American Economic Review, Papers and Proceedings* 84(2): 174–79.

———. 1995. "Precautionary Saving and Social Insurance." *Journal of Political Economy* 103(2): 360–99.

Juhn, Chinhui, and Kevin M. Murphy. 1997. "Wage Inequality and Family Labor Supply." *Journal of Labor Economics* 15(1, part 1): 72–97.

Karoly, Lynn A., and Gary Burtless. 1995. "Demographic Change, Rising Earnings Inequality, and the Distribution of Personal Well-Being, 1959–1989." *Demography* 32(3): 379–405.

Lusardi, Annamaria. 1996. "Permanent Income, Current Income, and Consumption: Evidence from Two Panel Data Sets." *Journal of Business and Economic Statistics* 14(1): 81–90.

MaCurdy, Thomas. 1982. "The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis." *Journal of Econometrics* 18(1): 83–114.

Mace, Barbara J. 1991. "Full Insurance in the Presence of Aggregate Uncertainty." *Journal of Political Economy* 99(5): 928–56.

Mariger, Randall P., and Kathryn Shaw. 1993. "Unanticipated Aggregate Disturbances and Tests of the Life-Cycle Consumption Model Using Panel Data." *Review of Economics and Statistics* 75(1): 48–56.

Moffitt, Robert, and Peter Gottschalk. 1993. "Trends in Covariance Structure of Earnings in the United States: 1969–1987." Unpublished paper. University of Wisconsin (April).

———. 1995. "Trends in the Autocovariance Structure of Earnings, 1967–1987." Unpublished paper. Johns Hopkins University.

Morduch, Jonathan. 1995. "Income Smoothing and Consumption Smoothing." *Journal of Economic Perspectives* 9(3): 103–14.

Nelson, Julie A. 1994. "On Testing for Full Insurance Using Consumer Ex-

penditure Survey Data: Comment.'' *Journal of Political Economy* 102(2): 384–94.

Paxson, Christina H. 1992. ''Using Weather Variability to Estimate the Response of Savings to Transitory Income in Thailand.'' *American Economic Review* 82(1): 15–33.

Powers, Elizabeth. 1996. ''Does Means-Testing Welfare Discourage Saving? Evidence from the National Longitudinal Survey of Women.'' Unpublished paper. Federal Reserve Bank of Cleveland.

Slesnick, Daniel T. 1992. ''Gaining Ground: Poverty in the Postwar United States.'' *Journal of Political Economy* 101(1): 1–38

Townsend, Robert M. 1994. ''Risk and Insurance in Village India.'' *Econometrica* 62(3): 539–91.

———. 1995. ''Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies.'' *Journal of Economic Perspectives* 9(3): 83–102.

U.S. House of Representatives. Committee on Ways and Means. 1996. *Background Material and Data on Programs within the Jurisdiction of the Committee on Ways and Means: 1996 Green Book.* WMCP 104–14, 104 Cong., 2 sess. Government Printing Office.

Varian, Hal. 1980. ''Redistributive Taxation as Social Insurance.'' *Journal of Public Economics* 14(1): 49–68.

Zeldes, Stephen P. 1989. ''Consumption and Liquidity Constraints: An Empirical Investigation.'' *Journal of Political Economy* 97(2): 305–46.