

**Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.**

The Evolution of Household Income Volatility

Karen E. Dynan, Douglas W. Elmendorf, and Daniel E. Sichel

2007-61

NOTE: Staff working papers in the Finance and Economics Discussion Series (FEDS) are preliminary materials circulated to stimulate discussion and critical comment. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors. References in publications to the Finance and Economics Discussion Series (other than acknowledgement) should be cleared with the author(s) to protect the tentative character of these papers.

The Evolution of Household Income Volatility

Karen E. Dynan
Federal Reserve Board

Douglas W. Elmendorf
Brookings Institution

Daniel E. Sichel
Federal Reserve Board

October 2007

The views expressed in the paper are our own and not necessarily those of the Federal Reserve Board, the Brookings Institution, or other members of their staffs. We have benefited from comments by Alan Auerbach, Chris Carroll, Molly Dahl, Tom DeLeire, Jason Furman, Bill Gale, Peter Gosselin, Bob Hall, Jeff Kling, Dean Lillard, Annamaria Lusardi, Peter Orszag, Jim Poterba, Jon Schwabisch, Paul Willen, and Seth Zimmerman. Any errors are our own. This draft is preliminary; comments are welcome.

Abstract

Using data from the PSID, we find that household income has become noticeably more volatile during the past thirty years. We estimate that the standard deviation of percent changes in household income rose one-fourth between the early 1970s and early 2000s. This widening in the distribution of percent changes is concentrated in the tails of the distribution, and especially in the lower tail: Changes between the 25th and 75th percentiles are almost the same size now as thirty years ago, but changes at the 10th percentile look substantially more negative. The boost in volatility occurred throughout the 1970s, 1980s, and 1990s, albeit not at a steady pace. Households' labor earnings and transfer payments have both become more volatile over time.

1. Introduction

The aggregate U.S. economy has been markedly more stable since the mid-1980s than it had been in the preceding couple of decades. The reduction in volatility is widespread, showing up in real GDP and most of its components as well as other measures of economic activity. The source of this “Great Moderation,” as it has been labeled by some writers, has been the subject of considerable debate, with various papers arguing that volatility fell principally because of milder economic shocks, better monetary policy, improved inventory management, or financial innovation.

In contrast with this aggregate stabilization, many researchers have found that individual households and firms have faced *more* volatile economic circumstances over time. On the household side, Gottschalk and Moffitt (1994, 1995, 2002, and 2006), Dynarski and Gruber (1997), Haider (2001), and others estimated that individual labor earnings have become more volatile for certain subsets of earners during the past few decades, and Hacker (2006) argued the same for family income. On the firm side, Comin and Philippon (2005) and Comin and Mulani (2006) estimated that business performance has become more volatile as well. These estimates are consistent with the view of many commentators that the economy has become more “dynamic”—that globalization, deregulation, and technological change have increased the amount of creative destruction and thus the competitive pressures and risks faced by workers and firms.

However, the robustness of these empirical findings of increased microeconomic volatility is not clear. Contrary to other analyses, Davis, Haltiwanger, Jarmin, and Miranda (2006) estimated that the volatility of sales and employment growth at the firm

level has declined over time. In addition, the Congressional Budget Office (2007) estimated that individual earnings volatility has changed little since 1980.

In this paper, we build on our earlier work (Dynan, Elmendorf, and Sichel, 2006b) to examine the evolution of household income volatility using data from the Panel Study of Income Dynamics (PSID). To make the analysis as transparent as possible, we do not estimate a formal model of income. Instead, we simply document changes over time in the cross-sectional distribution of income changes. We focus on the appropriate treatment of extreme values in the data and on methodological changes in the PSID that can distort comparisons over time. We also explore the interplay between various sources of household income.

To summarize our results, we find that household income has become noticeably more volatile during the past thirty years. We estimate that the standard deviation of percent changes in household income rose one-fourth between the early 1970s and early 2000s. This finding contrasts sharply with a roughly one-half *decline* in aggregate economic volatility (as measured by the standard deviation of the quarterly growth rate of real GDP during the 1960-1983 and 1984-2004 periods). The widening in the distribution of percent changes in household income is concentrated in the tails of the distribution, and especially in the lower tail: Changes between the 25th and 75th percentiles are almost the same size now as thirty years ago, but changes at the 10th percentile look substantially more negative. The boost in volatility occurred throughout the 1970s, 1980s, and 1990s, albeit not at a steady pace. Households' labor earnings and transfer payments have both become more volatile over time.

The limitations of our analysis bear emphasis. First, an increase in the *volatility* of household income does not imply a corresponding increase in *risk*. Our calculations do not distinguish between voluntary and involuntary changes in income, they do not include shocks to desired spending, and they do not account for adjustments to saving and borrowing that can buffer income shifts. We are currently investigating all of these issues in related research on the evolving volatility of household spending. Second, our analysis includes only households whose heads are 25 years or older and who are not yet retired. Therefore, we have nothing to say about developments among the youngest and oldest households. Third, our findings are based on a particular methodology applied to a single dataset. Further research using different techniques and datasets will be needed before economists can have a high degree of confidence in the facts about household income volatility. Still more work will be needed to provide economic explanations for those facts and to consider the possible implications for economic policy.

2. Measuring Volatility in the PSID

In this section we briefly discuss how we measure volatility using data from the PSID. In subsequent sections we present our results regarding the evolution of volatility of households' labor earnings, capital income, transfer income, and total income.

PSID Data

The PSID contains longitudinal information about the income of a large set of households between 1967 and 2004, as well as information on those households' spending, employment, and demographic characteristics. Linking this information correctly is very labor-intensive, however, as the definition and construction of variables

have changed considerably over time. Households (dubbed “family units” by the PSID) are composed of people living together who are related by blood, marriage, or adoption—or living together permanently and sharing income and expenses. When households are headed by both a man and a woman, the PSID arbitrarily labels the man as the household head and the woman as his spouse; when households are headed by a woman alone, she is the head. The PSID collected data annually through 1996 and biannually thereafter, so we focus on two-year changes in earnings and income. We deflate nominal dollars into real 2002 dollars using the CPI for urban consumers.

We drop observations where the head is under age 25, to avoid most transitions between school and work. We also drop observations where the head is retired, to stay clear of transitions between work and retirement. We exclude income changes of people moving between households. For example, when a two-earner couple divorces, we account for the loss of the wife’s income to the husband (because he remains head of his household) but do not capture the change in the wife’s economic circumstances (because her new household, while typically added to the sample, did not exist previously). In addition, we drop observations where the head has changed from the previous observation and where farm income is positive (because such income is not reported in a comparable manner over time).

Our analysis focuses on the sub-sample of the PSID that was chosen in 1967 to be representative of the U.S. population. This sample has evolved over time to include both the original households and new households spun off the original households. The PSID also includes special samples of low-income households (in the survey since 1967) and

immigrants (in the survey since 1997).¹ Including these households in our analysis would be ideal because of the greater breadth of coverage, larger sample size at the bottom of the distribution, and the existence of population weights.² However, some of the variables that we need are not reported for this broader sample in some years or are available only in lower-quality form. We show later that focusing on the original representative sample does not appear to make a significant difference in the results.

Top-coding in the PSID can paint a misleading picture of volatility: Variables top-coded at the same level in consecutive years will appear more stable than they really are, and changes in the level of top-coding over time can make the distribution of income appear to change in ways it does not. To address these problems, we identify the year with the maximum share of the sample that is top-coded for each variable—for example, roughly ½ percent for total income—and drop that same share of observations from the upper end of that variable’s distribution in all years. In addition, some variables have been bottom-coded at \$1 in some years. For consistency over time and across variables, we replace any value of \$0 or below with \$1.

Data from the PSID, like all data, include a significant amount of measurement error. Accordingly, one should not take our estimates of the *level* of volatility literally. However, the crucial question for evaluating *changes* in volatility is whether measurement error has changed over time. A key concern, then, is that the PSID implemented two major methodological changes in the early 1990s, as described by Kim,

¹ The PSID also included a special sample of Latino households between 1990 and 1995, but this sample was surveyed for too short a period to be useful for our analysis.

² Population weights help to ensure that estimates reflect the experience of the country as a whole. For example, they keep the age distribution of the sample consistent with that of the population, and they diminish the consequences if attrition from the sample is greater for households in more volatile economic circumstances.

Loup, Lupton, and Stafford (2000) and Kim and Stafford (2000). Income data for 1992 and later were collected using Computer Assisted Telephone Interviewing (CATI) rather than traditional paper questionnaires, and income data for 1993 and later were processed using different software. Kim et al warned that these shifts create a “potential double seam” in the data of which researchers should be aware.³ We return to this issue below.

Measuring Volatility

Gottschalk and Moffitt’s seminal papers on labor earnings measured volatility as the magnitude of transitory earnings, which they calculated using two techniques. In some cases, they measured transitory earnings as actual earnings in a given year less a moving average of earnings; in others, they measured transitory earnings using decompositions based on time-series models. These studies, and research by other analysts that built on these techniques, have yielded interesting results.

However, this paper follows the simpler approach of measuring volatility using the cross-sectional distribution of changes in income. This approach does not allow us to say anything about the persistence of changes in income. However, at this early point in research on this topic, we think it is more important just to document the facts about household income changes in the least processed and filtered manner. Thus, we count it a virtue that our results do not depend on a particular model of income dynamics or a particular estimate of underlying income from which deviations represent volatility.

Moreover, people facing income changes experience the full changes, and they do not

³ Data for some recent years are available only in “early release” form, implying that some of the data have undergone limited processing and that some variables are not included. Other variables are released through the supplementary “Income Plus” files, which we use.

know how persistent those changes will be. Our results accurately characterize the full changes experienced by households.

We summarize the cross-sectional distribution of income changes using the standard deviation. In contrast, most research on the volatility of individuals' earnings, summarized later in the paper, has focused on variances. Because that research was designed to parse the widening of the earnings distribution into its permanent and transitory components, the additive nature of variances ($\text{var}[x+y] = \text{var}[x] + \text{var}[y] + 2\text{cov}[x,y]$) was crucial. However, this adding-up property is not needed to characterize the change in variability over time, and the squaring of terms used to calculate variances can be misleading. For example, consider an economy with three households in which one income rises 20 percent, one falls 20 percent, and one is unchanged. The standard deviation of income changes is 16 percentage points, measured in the same units as income growth and comparable to it. If the economy becomes more turbulent, and the corresponding changes become +30, -30, and 0, the standard deviation rises to 24 percentage points. The 50-percent increase in volatility is a sensible characterization of the evolution of this economy. However, the variance of income changes in this example rises from 266 to 600 percentage points squared; neither these levels nor the 125 percent increase between them provides useful information about this economy.

Our estimates do not control for changes in households' age or food needs (as estimated by the PSID staff based on demographic characteristics). Because we care about the evolution of income volatility for the working-age population as a whole, we do not want to exclude changes caused by compositional shifts among age groups or among groups based on other demographic characteristics. However, to gain greater insight into

that evolution, we do present results for various demographic groups. Changes in food needs are just one type of shock to desired consumption. Measuring the evolution of all risks on the spending side of household budgets is an important extension of the present paper that we are currently pursuing.

Measuring income changes in percentage form has various advantages and disadvantages. On the upside, percent changes are straightforward and easily understood. In addition, under the common assumption that utility displays constant relative risk aversion (CRRA), a given percentage decline in consumption causes the same percentage reduction in utility regardless of the absolute magnitude of the consumption decline. But simple percent changes also have an important disadvantage: Large decreases and increases of the same absolute size do not generate symmetric percent changes. For example, a drop in income from \$1000 to \$100 is a 90 percent decline, while a return to \$1000 represents a 900 percent increase. Moreover, such large values are especially influential when estimating the dispersion of changes.

To mitigate the problems with simple percent changes, we measure percent changes as $100*(Y_t - Y_{t-2})/Y_{\text{average}}$, where $Y_{\text{average}} = (Y_{t-2} + Y_{t-3} + Y_{t-4})/3$; missing values of Y_{t-3} and Y_{t-4} are excluded from the calculation.⁴ We calculate annual rates and then replace any increase greater than 100 percent with 100 percent.

⁴ Results based on simple percent changes are similar to those presented in the paper.

One alternative to our approach is to analyze deviations relative to a moving-average level (as in Gottschalk and Moffitt's work) rather than analyzing changes. A return of income to its previous long-run level represents stability in that calculation, but volatility in ours—which seems more appropriate to us.

Another alternative to our approach is to scale income changes by the level of income predicted by a household's demographic characteristics. We do not use this alternative because it is less transparent and because households presumably care about income movements relative to their previous income, not relative to an uninformed prediction of their income.

A third alternative is to replace percent changes with logarithms, which avoid the asymmetry between increases and decreases. For example, the income plunge and jump noted here represent log

Neither this paper nor previous ones in this literature distinguish between voluntary and involuntary changes in earnings and income. For example, we do not separate people whose earnings decline because of pay cuts or layoffs by their employers from people whose earnings decline because they choose to reduce their work hours or leave the labor force. We are pursuing this issue as well in ongoing research.

3. Volatility of Individual Labor Earnings

We begin this section by exploring the earnings of household heads. Then we turn to the earnings of spouses and to comparisons of our results with those of previous researchers.

Volatility of Earnings of Household Heads

Figure 1 documents the evolving volatility of household heads' labor earnings as recorded in the PSID. PSID labor earnings include wages and salaries, overtime pay, bonuses, commissions, and a portion of self-employment income determined by the PSID staff. The upper left panel plots, for each year, the standard deviation across heads of their annualized percent changes in earnings (as defined in the previous section). For years after 1996, when the PSID collected data biannually, we fill the gaps in the estimated line using linear interpolation. This measure of volatility rose modestly in the

changes of -2.3 and 2.3. However, logarithms are less transparent and would still give heavy weight to income movements involving very low income levels.

A final alternative is to use a more complex transformation that effectively gives weight to the absolute change as well as the percent change. For example, Carroll, Dynan, and Krane (2003) note that "effects [of risk on wealth] estimated using logs could give undue weight to responses at the lower end of the wealth distribution" (page 592), and they transform wealth using the inverse hyperbolic sine function instead. However, this approach would lose the clarity and simplicity of percent changes. Moreover, some people who argue that "a decline in income from \$1000 to \$1 is not as bad as a decline from \$100,000 to \$10,000" appear to be thinking that public and private safety nets provide more offset to the decline in consumption in the latter case; this argument is not relevant for our purpose, because we are explicitly studying income rather than consumption and because the PSID data incorporate public transfer payments.

1970s and 1980s and increased sharply in the early 1990s. The vertical line marks 1991, the last year of income data preceding the PSID methodological changes noted above.

The upper right panel displays three variants of this calculation. The solid line shows rolling estimates for the preceding three years, with the left endpoint covering the years 1971 through 1973 and the right endpoint covering the years 2002 and 2004. We use this smoothing throughout the paper, unless otherwise noted, to reduce the choppiness that can be induced by small sample sizes. The dashed line replaces the representative sample with the weighted full sample during the continuous period for which all of the data for the latter are available. This alteration confirms that the representative and weighted full samples generate similar results; we do not pursue results with the weighted full sample further. The dotted line Winsorizes the data by dropping the top tenth and bottom tenth of percent changes in each year. This modification substantially reduces estimated volatility in all years and nearly eliminates the jump in the early 1990s, implying that much of the action is occurring in the tails of the percent-change distribution.

To investigate these tails more closely, the middle panels focus on one source of very large increases and decreases in earnings: observations with reported earnings of precisely \$0. Such observations can arise for two broad economic reasons: People can be looking for work but unable to find a job, thereby having zero earnings involuntarily, or they can choose to be out of the labor force, thereby having zero earnings voluntarily. Because this study does not attempt to distinguish between voluntary and involuntary variation in income, we would like to include all of these observations in our estimates.

On the left, the solid line shows the share of all observations with zero earnings reported in the current year and two years earlier, the dashed line shows the share with zero earnings preceded by earnings between \$1 and \$10,000, and the dotted line shows the share with zero earnings preceded by earnings above \$10,000. On the right, the corresponding breakdown is shown for transitions out of zero earnings. Two aspects of these panels stand out: First, the frequency of consecutive (biannual) reported earnings of \$0 climbed in the late 1970s, retraced that gain in the late 1980s, and has been flat since the early 1990s. Because we re-code \$0 to \$1 in our percent change calculations, consecutive reports of \$0 appear as perfectly stable earnings. With such observations as common in 2004 as thirty years earlier, they have little net effect on comparisons of volatility between the beginning and end of the sample. Second, the frequencies with which zero reported earnings followed and preceded substantial reported earnings both stepped up markedly in the early 1990s. These combinations generate very large earnings declines and increases, so the boost in their occurrence helps to explain the thickening of the tails of the percent-change distribution.

Does the increase in the probability of zero earnings sandwiched between substantial earnings reflect changes in the economic environment or in measurement? We do not know of specific changes in PSID questions or procedures that had this effect. Nonetheless, we think the measurement-based explanation is more likely. First, the coincidence of timing is striking (and the PSID staff has noted that later data have a higher variance). Second, identifying changes in economic conditions that would have

had such a large and sudden effect is difficult.⁵ Third, we see no evidence of other outcomes that would be expected if economic conditions had become much more turbulent at that time: There is no reported change in the frequency of zero earnings following or preceding *low* earnings (as shown in the middle panels) or in the frequency of zero earnings sandwiched between substantial earnings for *spouses* (as shown later). Moreover, we think the measurement change had a permanent, rather than transitory, effect on the interpretation of zero reported earnings. In particular, the probability of zero earnings in-between substantial earnings has remained high through the past decade rather than receding.

To assess the evolution of true economic volatility, we need to place the data before and after the early 1990s on a comparable basis. Because we cannot adjust reports of zero earnings to make them comparable over time, we choose to exclude those observations altogether.⁶ This strategy has the disadvantage of ignoring some true, large drops in earnings—and thereby understating the probability of large changes in earnings—but we see little alternative. The bottom left panel shows that the standard deviation excluding observations with zero reported earnings (the solid line) lies everywhere below the standard deviation including those observations (the dashed line).

⁵ The long-term increase in female labor force participation has made it easier for men to take time off work to gain additional education or spend more time with their families. However, this logic would not easily explain a sudden jump in the frequency of zero earnings.

⁶ Whether the earlier or later data provide a more accurate picture of earnings volatility is unclear, but our main goal is comparability. One alternative to our approach is to replace any level of reported earnings below a threshold value with the threshold value itself (as Winship, 2007, does for household income). This strategy would be appropriate if, say, low but non-zero earnings were reported as such in earlier data but were reported as zero earnings in later data. However, making this adjustment using a threshold of either \$1000 or \$5000 leaves a marked step-up in measured volatility in the early 1990s: Because the additional recent observations of zero earnings are generally bracketed by earnings over \$10,000, even a \$5000 threshold allows for large gains and losses. Another alternative is to use data on hours worked to assess whether particular observations of zero earnings are accurate. This approach might also provide insight into whether the PSID measurement shift had a permanent or transitory effect. We plan to explore this strategy in future work.

The trends in the two series are similar, although the gap between them widens, unsurprisingly, in the early 1990s.⁷ Between the early 1970s and early 2000s, the standard deviation of percent changes in labor earnings of household heads increased 34 percent including all observations and 30 percent excluding observations with zero reported earnings. These figures are shown in the top lines of table 1, which also presents comparable figures for other categories of earnings and income that we discuss shortly.

It may be surprising that this measure of volatility has no evident correlation with aggregate business cycles. However, the bottom right panel reveals that the probability of large declines and increases in earnings is indeed related to business-cycle conditions. After allowing for the lag induced by the three-year rolling window, large declines (the solid line) were more common in the mid-1970s, early 1980s, early 1990s, and early 2000s than at other times—lining up well with recessions. Large increases (the dashed line) displayed a complementary pattern.

Volatility of Household Heads' Earnings by Age, Education, and Gender

The top panels of figure 2 report the evolution of earnings volatility for household heads under age 35, heads between ages 35 and 54, and heads age 55 and older (but not yet retired). Once again, we favor estimates of volatility that drop observations with zero reported earnings. Given that adjustment, the standard deviation of percent changes in earnings for all three age groups increased about 30 percent during the past thirty years.

The middle panels present comparable information for household heads with different levels of education. Volatility is higher throughout the past thirty years for heads without a high school degree than for heads with such a degree. Volatility

⁷ Taking the additional step of dropping observations with small positive earnings reduces the measured level of volatility further but has little effect on the contour.

increased during this period for all education levels, and by comparable amounts.

However, the run-up in volatility for heads with a college degree was concentrated in the latter part of the sample, while the rise in volatility for other heads occurred throughout the 1970s, 1980s, and 1990s.

The bottom panels divide household heads by gender. Because the PSID arbitrarily labels a male adult (if present) as the head, more than four-fifths of the heads are men. Unsurprisingly, then, male heads saw a rise in earnings volatility comparable to that for heads overall. By contrast, female heads experienced little net change in earnings volatility during the past thirty years.

Volatility of Spouses' Earnings

Figure 3 turns to the labor earnings of spouses as recorded in the PSID. The top left panel shows that the standard deviation of percent changes in spouses' earnings has fallen during the past three decades. This decline is a marked contrast with the increase in volatility for male household heads and a smaller contrast with the lack of change in volatility for female household heads. Spouses' earnings display no jump in volatility in the early 1990s and (in the middle panels) no rise in the frequency of observations with zero earnings that follow or precede substantial earnings. These findings reinforce our interpretation that the sharp increases in these phenomena for household heads likely stem from changes in PSID methodology specific to heads rather than changes in the economic environment. Spouses do show a considerable long-term decline in the frequency of observations with successive zero earnings, as women's labor force participation increased. Given a lack of evidence that reports of zero earnings for spouses are not comparable over time, we include these observations in our analysis. As

shown in the upper right panel and in table 1, the standard deviation of percent changes in earnings of spouses declined 10 percent between the early 1970s and early 2000s.

Comparison with Results in Recent CBO Report

A recent report by CBO (2007) analyzed the evolution of labor earnings volatility using the Continuous Work History Sample (CWHS) provided by the Social Security Administration. These data are based on administrative records rather than a survey; they cover individual workers only, not families, and include no information except for earnings, age, and gender. CBO reported that the standard deviation of percent changes in earnings declined a little between 1980 and 2003. CBO also extended the analysis back to 1960 for workers in the bottom two quintiles of the earnings distribution; earnings in higher quintiles sometimes exceeded the maximum level taxed by Social Security and therefore were not included in this analysis. For these individuals, earnings volatility changed little, on balance, between 1960 and 1980.

For comparison, the bottom left panel of figure 3 plots the evolution of earnings volatility for the pooled sample of household heads and spouses in the PSID. To focus on people who are workers, we exclude observations of successive zero earnings for spouses.⁸ We also exclude all observations of zero earnings for heads in light of the apparent non-comparability of these observations over time. The rough constancy of volatility in this picture seems not too different from the CBO results. However, the bottom left panel of figure 2 showed that earnings volatility for male household heads

⁸ Note that we include spouses' observations of zero earnings preceded or followed by positive earnings, which differs from the treatment in the preceding panels. When analyzing spouses as a group, we want to include women who are out of the labor force for an extended period, because the stability of their earnings (albeit at \$0) accurately captures the effect on household income. Yet, when analyzing workers as a group, we want to exclude women who are out of the labor force in order to get a better picture of labor market developments. Of course, some observations of non-successive zero earnings also reflect withdrawals from the labor force, but others reflect bad labor market outcomes, so we err on the side of including them.

rose noticeably after 1980, which contrasts with a slight decline reported by CBO.

Because this rise has an important effect on the volatility of household income, as we show shortly, the contrast deserves attention.

One difference between the two studies is that we include self-employment earnings, while CBO does not. Because we are ultimately interested in total household income, we use the most inclusive measure of earnings available. Whether excluding self-employment earnings raises or lowers the average level of volatility is unclear: Self-employment earnings are presumably more volatile than wages and salaries, but they also tend to be negatively correlated with wages and salaries. For example, leaving out self-employment earnings means that individuals moving from employment to self-employment will mistakenly appear to have suffered a large drop in labor earnings (and those moving the other way will mistakenly appear to have enjoyed a large rise); Fairlie (2005) reported that such transitions occur frequently. More important for our present purpose, excluding self-employment earnings may damp the change in volatility over time. In the bottom right panel, the solid line repeats the smoothed standard deviation of earnings for household heads (again dropping observations with reported zero earnings), and the dashed line excludes heads who report having a financial interest in a business. This adjustment—which removes an average of 15 percent of the sample—significantly flattens the uptrend in estimated volatility since 1980.⁹ In addition, because more men

⁹ In principle, this flattening might occur because people with business interests reported greater volatility over time in wages and salaries rather than self-employment earnings. However, it is not obvious why this would be so. Moreover, this alternative interpretation cannot be tested effectively. Research on self-employment generally distrusts self-employed workers' survey responses separating self-employment earnings from other earnings. For example, Fairlie (2005) reports that "more than half of the self-employed with positive earnings in the [National Longitudinal Survey of Youth] report wage and salary income, but do not report business income," and he explains why that is a natural mistake given the nature of the questions. Similarly, Holtz-Eakin, Rosen, and Weathers (2000) appear to use total labor earnings when studying the self-employed in the PSID. A further complication is that the PSID staff divides survey

than women are self-employed, self-employment earnings likely have a bigger effect on earnings volatility for men, for whom our results differ from CBO's, than for women, for whom our results are fairly similar.

A second difference between the two analyses is that CBO includes all workers who are 22 years or older, while we include only household heads who are 25 years or older (and their spouses). When we examine total household income, we include the labor earnings of all household members, but we focus on heads' (and spouses') earnings because they are usually larger. Also, we deliberately leave aside households where transitions between school and work are common, as with people in their early 20s. As a result, we exclude the 15 percent of the work force represented by people under age 25 as well as those older workers who are not a household head (or the spouse of one).

Yet a third difference between the two studies is that we cap all increases consistently at 100 percent, while CBO handles very large increases in various ways. In particular, CBO also caps all increases from zero earnings at 100 percent (say, from \$0 to \$1000), but allows other increases up to 1000 percent to show through (say, from \$100 to \$1000) and then drops increases larger than 1000 percent (say, from \$100 to \$1100). In calculating standard deviations, the CBO method puts greater weight than our method on some transitions from low earnings, but it also drops some transitions that we include.

A final prominent difference between the two analyses is the datasets used. Limitations of the PSID lead us to analyze two-year changes, while CBO analyzes one-year changes, and to drop transitions to and from zero earnings for household heads, while CBO does not. In addition, most economists would probably view the CWS

respondents' self-employment income between labor earnings and capital income. Any misallocation will distort measured labor earnings, but not total household income.

administrative data as more accurate for wages and salaries than the PSID survey data. However, it is not clear that this view is correct. Bound and Krueger (1991) viewed administrative data as the benchmark for examining measurement error in surveys, but Bound, Brown, Duncan, and Rogers (1994, page 357) were more skeptical: “Part of the reason that Bound and Krueger find larger errors than we do has to do with errors in the social security record and part has to do with [Current Population Survey, CPS] recording errors. ... We expected that the problems would be mostly with CPS reports; in fact in 15 of 26 cases [with the largest discrepancies] it was one of the SSA reports that seemed ‘out of line.’” Moreover, for self-employment earnings, where tax evasion may be substantial, it is not at all clear whether the CWS or PSID provides more accurate information.

Besides the different trends in volatility of men’s earnings, our findings also differ markedly from CBO’s in the frequency of very large declines in earnings. We estimate the probability of a 50 percent drop in earnings to be 5 percent or less throughout the sample period (not shown), while CBO puts the probability around 15 percent. The higher CBO figure could result from their exclusion of self-employment earnings (which might raise volatility, as discussed above), their inclusion of young household heads and workers who are not heads (who presumably have more volatile earnings), our exclusion of household heads’ transitions to and from zero earnings (which lowers volatility), or our use of two-year rather than one-year changes (which may lower volatility as well).

Comparison with Results in Gottschalk-Moffitt Papers

Gottschalk and Moffitt (GM) launched the recent literature on earnings volatility (1994) and returned to the topic several times (1995, 2002, and 2006). In all of these

papers, they used the PSID to study wages and salaries received by male (or white male) household heads between the ages of 20 and 59. Their 1994 paper decomposed the cross-sectional variance of earnings into the variance of permanent earnings—defined as individuals’ average earnings during some time period (after allowing for the typical age-earnings profile)—and the variance of transitory earnings—defined as the difference between actual earnings and average earnings for that individual. Using this method, they estimated that the variance of transitory earnings rose sharply between the 1970-78 and 1979-87 periods. GM’s 1995 paper extended this analysis by developing a formal model of earnings dynamics. Using this technique, they estimated that the variance of transitory earnings followed no trend during the 1970s but rose strongly during the 1980s, which is consistent with their previous result.

In their 2002 paper, GM added data through 1996 and employed both their formal model and a less formal decomposition into permanent and transitory variances. Estimates from the model showed that transitory variance changed little in the 1970s and rose in the 1980s (confirmed their earlier result) before nearly doubling in the early 1990s and then quickly receding to its level in 1980. In their 2006 paper, GM extended the sample further through 2002. Their model estimates again showed an increase in transitory variance in the early 1980s, a further increase in the early 1990s, and a sharp retrenchment. The estimated magnitudes of the 1990s run-up and retrenchment were much smaller than in the 2002 paper, but a large estimated rise in the early 2000s left transitory variance much higher at the end of the sample than at the start.

All told, GM’s latest results imply that the standard deviation of transitory earnings nearly doubled between the early 1970s and early 2000s—a much more

dramatic rise than we find. The sources of this difference are unclear, as the GM analysis differs from ours in several important ways. They include younger household heads than we do, and they employ an upper cutoff based on age rather than dropping household heads who say they are retired. They also utilize a narrower definition of labor earnings, and they make no adjustment for the measurement changes whose effects we documented. More importantly, perhaps, GM focus on estimates from time series models of earnings dynamics rather than the cross-sectional distributions of earnings changes that we use. Although formal modeling clearly provides some benefits, the complexity of model calculations and the dependence of results on particular modeling assumptions can make interpreting such estimates difficult. For example, GM's (2006) informal decomposition implied that the variance of transitory earnings rose in the early 1980s—as in the model-based estimates—but moved sideways, on balance, thereafter—rather than increasing further as in the model-based estimates.

Comparison with Other Results on the Volatility of Individual Earnings

Table 2 summarizes selected studies on the evolution of earnings and income volatility, including the recent CBO report and the set of papers by Gottschalk and Moffitt.¹⁰ Here we briefly touch on some of the other papers that appear in the table.

¹⁰ A related strand of research has investigated the evolution of labor markets and institutions. Using data from the Retirement History Survey, the National Longitudinal Study of Older Men, and the Health and Retirement Study, Stevens (2005) concluded that the prevalence of long-term employment relationships for men was stable between 1969 and 2002. However, she studied only men born before 1944, so the broader relevance of this finding is unclear. Farber (2005) used data from the Displaced Workers Survey with adjustments for methodological changes. He found that the rate of job loss was essentially unchanged, on balance, between the early 1980s and the early 2000s, but that the gap between the job loss rate and the unemployment rate widened a good deal. An increase in worker displacement would tend to boost the volatility of earnings, all else equal. Separately, Lemieux, Macleod, and Parent (2006) argued that changes in labor market institutions during the past several decades (including the decline in unionization, drop in the real value of the minimum wage, and reduction in the cost of collecting and processing data on worker performance) have boosted the share of workers whose pay is linked to their performance. They showed that wages are more dispersed in performance-pay jobs than in other jobs, so that the shift has contributed

A number of papers have used data from the PSID. Daly and Duncan (1997) adopted the Gottschalk-Moffitt decomposition and estimated that the variance of transitory earnings for men aged 25 to 44 rose between the 1969-79 and 1979-89 periods but leveled off thereafter. By comparison, we find an ongoing increase in volatility for all men. Dynarski and Gruber (1997) estimated that the variance of transitory earnings for male household heads (defined as the deviation of earnings from individual-specific growth paths) rose in both the 1970s and 1980s, for a total gain of about 75 percent. This increase corresponds to an increase in standard deviation of roughly 30 percent, fairly close to our estimate for male heads during the same period. Haider (2001) developed a formal model of earnings different from that of Gottschalk and Moffitt. He estimated that the variance of transitory earnings rose during the 1970s but changed little, on balance, during the 1980s. In contrast, we find that all male heads experienced increasing earnings volatility in both the 1970s and 1980s; whether this discrepancy arises from differences in the sample or differences in the method of measuring volatility is unclear. Hyslop (2001) focused on the covariance of earnings within families during the short period from 1979 to 1985. He estimated that earnings volatility increased substantially for both men and women. This result is consistent with ours for men but not for women; again, determining the source of this difference is difficult. Comin, Groshen, and Rabin

to the rise in compensation inequality. Presumably, this shift would also make compensation more volatile from year-to-year.

Another related set of papers has examined the volatility experienced by individual businesses. Using Compustat data on publicly traded firms, Comin and Philippon (2005) and Comin and Mulani (2006) argued that firm performance (measured by the profits-to-sales ratio or the growth rates of sales, employment, and sales per worker) has become much more volatile during the past 35 years. Indeed, Comin, Groshen, and Rabin aimed to link the rise in earnings instability for household heads to this rise in firm-level volatility. However, Davis, Haltiwanger, Jarmin, and Miranda (2006) argued that Compustat data are not appropriate for this purpose because the selection of firms that are publicly traded has changed over time. Using instead the Longitudinal Business Database, which contains annual observations on employment and payroll for all U.S. firms, Davis et al found a significant decline during the past few decades in the volatility of sales and employment growth at the firm level.

(2006) followed Gottschalk and Moffitt's original approach. They estimated that the variance of transitory earnings for male household heads increased at a fairly steady pace between the early 1970s and early 1990s. This finding is consistent with ours.

Using data from the Current Population Survey, Cameron and Tracy (1998) estimated that the variance of transitory earnings for men increased in the 1970s and early 1980s but then retraced part of that run-up in the late 1980s and early 1990s. Altogether, Cameron and Tracy calculated that transitory variance rose about 65 percent between 1968 and 1997, with much of the rise stemming from an increase in the frequency of large earnings changes. Our analysis revealed a different time pattern of volatility, but a similar overall increase and a similar role for large changes.

To summarize, most previous papers on earnings volatility have concluded that volatility increased considerably during the past several decades. However, not all papers reach that conclusion, and even those that do disagree to some extent about the timing and magnitude of the increase. Determining the sources of these differences is difficult owing to the wide range of empirical techniques used and the variety of choices made about which individuals and types of income to analyze.

4. Volatility of Household Income

This section explores the volatility of household income. We examine, in turn, household labor earnings, capital income, transfer income, and total income.

Heads' and Spouses' Combined Earnings

The top panels of figure 4 document the evolving volatility of the combined labor earnings of household heads and their spouses. Because we think that observations with

zero reported earnings by heads are not fully comparable over time, we favor estimates that drop such observations. Dropping only observations in which *combined* earnings are reported to be zero does not solve the problem, because reports of positive earnings by a spouse can mask misleading reports of zero earnings by a head. With the correct adjustment, the standard deviation of percent changes in combined labor earnings increased 23 percent between the early 1970s and early 2000s (as reported in table 1).

This finding that volatility increased for combined head and spouse earnings might seem surprising in light of our earlier finding that volatility stayed roughly constant for individual earnings in the pooled sample of heads and spouses. The combination thus warrants further attention.

The increase in women's labor force participation does not provide an explanation. Consider a household with a husband in the labor force and his wife out of the labor force. Their household earnings equal his earnings and have the same volatility.¹¹ If his wife enters the labor force with the same earnings distribution as her husband, then the average volatility of individual earnings is unchanged. The volatility of household earnings will be the same in percentage terms if her earnings are perfectly correlated with his, but will fall to zero if her earnings are perfectly negatively correlated with his. Because actual correlations lie somewhere between these extremes, the volatility of household earnings declines.

Our two results also cannot be explained by changes in the correlation of earnings of household heads and their spouses. It might be expected that an individual would try to adjust his or her earnings to buffer changes in a partner's earnings—for example, by

¹¹ This statement assumes that the calculation of individual earnings volatility excludes the observations of successive zero earnings accruing to the wife—as our earlier calculation did.

taking a more demanding job if a partner lost a job, or by shifting toward home production if a partner's earnings rose significantly. At the same time, adults in the same household face some of the same earnings shocks—for example, changes in regional economic conditions, or even changes in economic conditions for workers in certain industries, occupations, or education levels. We use two simple techniques to estimate roughly the net effect of these competing forces over time.

The middle panels of figure 4 depict the frequencies with which decreases and increases in heads' earnings have been combined with decreases, increases, and no changes in their spouses' earnings. To focus on significant shifts, we deem any change smaller than 5 percent (for either the head or spouse) to be no change at all. We again exclude reports of zero earnings by heads. On the left, we show that declines in heads' earnings in the early 2000s are accompanied by increases in spouses' earnings about 35 percent of the time and by declines about 25 percent of the time. On the right, we show that increases in heads' earnings during the same period are accompanied by increases in spouses' earnings 40 percent of the time and by declines 25 percent of the time. The tendency for spouses' earnings to rise more than fall is not surprising given the secular increase in real wages and the rise in women's labor force participation. However, the tendency for spouses' earnings to rise a bit more often when heads' earnings rise than when they fall may be surprising. Most important for our analysis, this pattern has not varied much over time.

The bottom panels show the shares of decreases and increases in heads' earnings that have been offset by changes in their spouses' earnings. Declines in heads' earnings at the end of the sample were offset an average of -10 percent—that is, they were

accentuated by an average of 10 percent. Increases in heads' earnings during the same period were offset essentially not at all. These patterns have varied just a little over time, with heads' earnings declines offset to a small positive extent during most of the past three decades, and heads' earnings increases generally accentuated by changes in spouses' earnings.

These four panels together suggest that the correlation of movements in heads' and spouses' earnings has been close to zero throughout the past thirty years.¹² Warren (2005) argued that the rise in two-earner families has reduced people's scope for getting a job when their partners' earnings falter; others might speculate that the rise in two-earner families makes it easier for people to work more hours when their partners' earnings falter. Neither of these hypotheses receives support in our estimates.

Thus, the question persists: Why did the volatility of combined head and spouse earnings increase while the volatility of individual earnings did not? The answer is that household heads have higher average earnings than spouses, so the rise in volatility they experienced had a bigger effect on household earnings than the decline in volatility for spouses. Specifically, when calculating volatility for households, each person's dollar change in earnings is combined with his or her partner's dollar change to calculate the change for the household as a whole. Among two-earner couples in our sample, spouses earn less than half what heads earn, on average, so they get less weight in household

¹² To explore possible non-linearities in the relationship between partners' earnings, we also produced versions of these charts in which any change smaller than 25 percent was treated as no change. The results were not affected qualitatively by this modification.

At least two previous authors used PSID data to conduct careful investigations of this relationship. Focusing on the 1979-95 period, Hyslop (2001) estimated that wives' earnings were positively correlated with their husbands' earnings in both preceding and successive years. Because we analyze two-year changes in earnings, our calculations should capture this effect. In contrast, Shore (2006) concluded that innovations to husbands' and wives' permanent earnings were slightly negatively correlated, on balance, between 1968 and 2001. Because we do not decompose innovations into permanent and transitory components, we cannot compare our results directly to Shore's.

volatility. However, when calculating volatility for individuals, each person's percent change in earnings receives the same weight regardless of the dollar change in their earnings. The existence of one-earner couples strengthens this point. In a world with one two-earner couple and one one-earner couple, and all individuals earning the same average amount, the single head's earnings receive a one-third weight in individual volatility and a one-half weight in household volatility.

To verify the importance of this phenomenon in our estimates, we calculated the evolution of individual earnings volatility using percent changes weighted by earnings levels. The result (not shown) increases over time almost exactly in line with the volatility of combined head and spouse earnings.

Capital Income

Now we shift from labor earnings to capital income. Capital income in the PSID equals total market income less labor earnings; it does not include capital gains. The top left panel of figure 5 shows that the volatility of household heads' and spouses' combined capital income has changed little over time. Because the data display no break following the changes in PSID methodology and no unusual patterns involving the frequency of reported zero values, we make no adjustments to the data.

Transfer Income

The top right panel of figure 5 documents the evolving volatility of transfer income received by household heads and their spouses. The PSID collects data only on monetary transfers; the value of in-kind transfers is ignored. The standard deviation of percent changes in transfers rose in the 1970s, changed little in the 1980s, and moved further up in the 1990s. This latter increase does not seem to be related to measurement

changes in the PSID: There are no notable shifts in the tails of the distribution, no change in the frequency of very large increases or decreases, and no sudden change in the frequency of zero transfers. Therefore, we accept the data as reported. All told, the standard deviation of percent changes in transfer income rose 20 percent between the early 1970s and early 2000s.

One might expect that shifts in transfer income would be negatively correlated with shifts in market income—that is, the sum of labor earnings and capital income, which is labeled “taxable income” by the PSID. This correlation might arise because transfers act as a safety net when market incomes decline, because people earn more market income when public benefits decline, or both. The remaining panels of the exhibit examine this hypothesis.

The middle panels depict the frequencies with which decreases and increases in market income are combined with decreases, increases, and no changes in transfers. Once again, we ignore any change in market or transfer income smaller than 5 percent, and we exclude observations in which heads report zero earnings. Declines in market income in the early 2000s are accompanied by increases in transfers nearly 30 percent of the time and by declines only 10 percent of the time; increases in market income during the same period are accompanied by increases in transfers about 15 percent of the time and by declines only 10 percent of the time. Thus, shifts in market income and transfer income tend to provide some offset to each other, but the tendency is fairly weak. This pattern has not varied much during the past three decades.

The bottom panels confirm these conclusions. Declines in market income were offset an average of 0 percent to 15 percent in the past thirty years, and increases in

market income an average of 0 percent to 5 percent. Once again, these relationships do not seem to have changed to an important extent over time.

Total Household Income

Total household income, labeled “total money income” by the PSID, equals the combined labor earnings, capital income, and monetary transfer income of the head and spouse, as well as the income of other household members. After-tax income is not available consistently in the PSID, so we study pre-tax income. As a reminder, our sample includes only those households with heads age 25 or older and not retired.

The top left panel of figure 6 shows that the volatility of total household income rose sharply in the early 1990s following the important methodological changes in the PSID. Once again, because we think that observations with zero reported earnings by heads are not fully comparable over time, we favor estimates that exclude such observations. As before, excluding only observations in which total income is reported to be zero does not solve the problem, because misleading reports of zero earnings by a head can be hidden by positive income from other sources. With the appropriate adjustment, shown by the solid line in the top right panel, the standard deviation of percent changes in total household income rose 24 percent between the early 1970s and early 2000s. (Without this adjustment, estimated volatility climbed 32 percent.) The boost in volatility of household income appears to have occurred throughout the 1970s, 1980s, and 1990s, albeit not at a steady pace. It can be attributed principally to increases in the volatility of labor earnings and transfer income.

The remaining panels take a closer look at the distribution of percent changes in total household income. In interpreting these figures, recall that we analyze two-year

changes at annual rates, so a 25 percent annual-rate decline corresponds to a 44 percent drop over a two-year period ($.75 \times .75 = .56$), and a 50 percent annual-rate decline corresponds to a 75 percent drop. Moreover, because some of the excluded observations where household heads report zero earnings represent true zero earnings, we understate the probability of large changes in income (in order to ensure comparability over time). With that in mind, the middle panels show that the probabilities of large and very large income declines, as well as large and very large income increases, have trended up during the past thirty years. The bottom panels show that the magnitudes of income changes between the 25th and 90th percentiles have shifted little over time, but that the magnitude of changes at the 10th percentile has become much larger (in absolute value). All four panels suggest that household income volatility has increased over time primarily because of thickening in the tails of the percent-change distribution.

Comparison with Previous Results on the Volatility of Household Income

Relatively little research has explored the evolving volatility of household income. Table 2 summarizes this work, and we review it briefly here.

A few studies have used data from the PSID. Batchelder (2003) examined the variation in households' incomes around their average incomes, focusing on households with heads aged 44 to 49. She found that market income volatility increased significantly between 1968 and 1992, which is consistent with our results. Gosselin (2004) analyzed fluctuations in total income for households with heads aged 25 to 64. He reported that volatility rose in the 1970s, leveled out in the early 1980s, and climbed further in the late 1980s and early 1990s. The timing of these increases differs from the time pattern that we estimate. Hacker (2006) applied the Gottschalk-Moffitt decomposition and showed

that the transitory variance of total household income rose at a moderate pace in the 1970s and 1980s, surged in the early 1990s, and retraced part of that run-up by the early 2000s (page 2). All told, Hacker concluded that the volatility of pre-tax income tripled between the early 1970s and early 2000s (page 27); this increase in variance amounts roughly to a 75 percent increase in standard deviation, a much larger run-up than we find. In ongoing work refining his calculations, Hacker (2007) has estimated that transitory variance increased to a somewhat lesser extent, especially if outlying observations are trimmed. Winship (2007) argued that Hacker's (2006) estimates were unduly influenced by observations with reported family income below \$4000, and that such observations should be re-coded at \$4000. With this modification, Winship estimated that income volatility increased a moderate amount between 1975 and 2000.

Using CPS data, Hertz (2006) analyzed dollar (not percent) changes in households' incomes from one year to the next. He estimated that income volatility increased significantly between 1990-91 and 1997-98 and then rose further by 2003-04. Based also on the CPS, Bollinger and Zilliak (2007) showed that income volatility for households headed by women was stable in the 1980s and early 1990s, but rose 60 percent between 1995 and 2004. Using data from the Survey of Income and Program Participation (SIPP), Bania and Leete (2007) studied monthly deviations in households' incomes from their average incomes. Focusing on low-income households, they estimated that volatility increased substantially between 1992 and 2003.

5. Conclusion

Using data from the Panel Study of Income Dynamics, we estimate that the standard deviation of percent changes in household income increased one-fourth between the early 1970s and early 2000s. Because several features of the data suggest that methodological changes in the early 1990s created a spurious jump in the frequency of household heads reporting zero labor earnings, we focus on calculations that exclude these observations. If one includes those observations instead, household income volatility appears to have increased somewhat more. The boost in volatility occurred throughout the past three decades, and it stemmed primarily from a rise in the frequency of large income changes rather than shifts throughout the income-change distribution. Households' labor earnings and transfer payments have both fluctuated more over time.

While household incomes have become more volatile since the early 1970s, aggregate income has become much less volatile. Connecting these developments is a central goal of our ongoing research. In preliminary work (Dynan, Elmendorf, and Sichel, 2006b), we showed that aggregate income constructed from PSID data has become less volatile in line with aggregate income from the national accounts. We reconciled this finding with greater household-level volatility by showing various ways in which the covariance of income across households appears to have declined.

We close by emphasizing that an increase in the *volatility* of household income does not imply a corresponding increase in the *risk* faced by households. First, only part of income variability reflects involuntary job loss and wage cuts, while part reflects voluntary choices to obtain more education or spend more time with one's family. The evolution over time of these two types of volatility is an important subject for future

research. In ongoing work, we are using the wealth of information collected about households in the PSID to distinguish between alternative sources of volatility.

Second, risk can arise from a variety of economic shocks besides the income variations we study in this paper. For example, leaving aside households whose heads are retired means that we did not learn anything about the volatility of pension benefits. By restricting our analysis to households with the same head over time, we ducked some of the risks stemming from divorce. Because of limitations in the PSID data, we cannot study changes in taxes or non-monetary transfer payments, which might be quite important. And we did not examine risks on the spending side, such as a loss of health insurance and an unexpected need to pay for additional health care.

Third, shocks to income can be buffered to some extent by adjustments to saving and borrowing, thereby reducing their impact on consumption. In an earlier paper (2006a), we argued that financial innovation has enhanced households' access to credit over time and thus strengthened their ability to smooth consumption in the face of income shocks. Using aggregate data, we showed that consumer spending has become less sensitive in the past few decades to movements in contemporaneous income. In preliminary work using household data (2006b), we confirmed that spending has become less responsive to income shifts. We also showed that the response of spending to negative income shifts is larger than the response to positive shifts—consistent with the existence of liquidity constraints—but has fallen more over time—consistent with the relaxation of those constraints through financial innovation. In terms of people's well-being, this improvement in the ability to smooth consumption relative to income provides a partial counterweight to the increase in income volatility documented in this paper.

References

- Bania, Neil and Laura Leete, "Income Volatility and Food Insufficiency in U.S. Low-Income Households, 1992-2003," Institute for Research on Poverty Discussion Paper #1325-07, April 2007.
- Batchelder, Lily L., "Taxing the Poor: Income Averaging Reconsidered," *Harvard Journal on Legislation*, 2003, pages 395-452.
- Bollinger, Christopher and James P. Ziliak, "Welfare Reform and the Level, Composition, and Volatility of Income," Mimeo, March 2007.
- Bound, John, Charles Brown, Greg J. Duncan, and Willard L. Rogers, "Evidence on the Validity of Cross-Sectional and Longitudinal Earnings Data," *Journal of Labor Economics*, July 1994, pp 345-368.
- Bound, John, and Alan B. Krueger, "The Extent of Measurement Error in Longitudinal Earnings Data: Do Two Wrongs Make a Right?," *Journal of Labor Economics*, January 1991.
- Cameron, Stephen and Joseph Tracy, "Earnings Variability in the United States: An Examination Using Matched-CPS Data," Mimeo, October 1998.
- Carroll, Christopher D., Karen E. Dynan, and Spencer D. Krane, "Unemployment Risk and Precautionary Wealth," *Review of Income and Statistics*, August 2003.
- Comin, Diego and Sunil Mulani, "Diverging Trends in Aggregate and Firm Volatility," *Review of Economics and Statistics*, 2006, pages 374-383.
- Comin, Diego and Thomas Philippon, "The Rise in Firm-Level Volatility: Causes and Consequences," *NBER Macroeconomics Annual*, 2005.
- Comin, Diego, Erica L. Groshen, and Bess Rabin, "Turbulent Firms, Turbulent Wages?," National Bureau of Economic Research Working Paper #12032, February 2006.
- Congressional Budget Office, "Trends in Earnings Variability Over the Past 20 Years," April 2007.
- Daly, Mary C. and Greg J. Duncan, "Earnings Mobility and Instability, 1969-1995," Mimeo, 1997.
- Davis, Steven J., John Haltiwanger, Ron Jarmin, and Javier Miranda, "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," National Bureau of Economic Research Working Paper #12354, June 2006.

- Dynarski, Susan and Jonathan Gruber, "Can Families Smooth Variable Earnings?," *Brookings Papers on Economic Activity*, 1997, pages 229-303.
- Dynan, Karen E., Douglas W. Elmendorf, and Daniel E. Sichel, "Can Financial Innovation Help to Explain the Reduced Volatility of Economic Activity?," *Journal of Monetary Economics*, 2006a, pages 124-150.
- Dynan, Karen E., Douglas W. Elmendorf, and Daniel E. Sichel, "Financial Innovation and the Great Moderation: What Do Household Data Say?," Mimeo, November 2006b.
- Fairlie, Robert W., "Self-Employment, Entrepreneurship, and the NLSY79," *Monthly Labor Review*, February 2005, pages 40-47.
- Farber, Henry, "What Do We Know about Job Loss in the United States? Evidence from the Displaced Workers Survey, 1984-2004," Princeton University Industrial Relations Section Working Paper #498, January 2005.
- Gosselin, Peter G., "The Poor Have More Things Today – Including Wild Income Swings," *Los Angeles Times*, December 12, 2004.
- Gottschalk, Peter and Robert Moffitt, "The Growth of Earnings Instability in the U.S. Labor Market," *Brookings Papers on Economic Activity*, 1994, pages 217-272.
- Gottschalk, Peter and Robert Moffitt, "Trends in Earnings Volatility in the U.S.: 1970-2002," Mimeo, December 2006.
- Hacker, Jacob S., *The Great Risk Shift*, Oxford University Press, 2006.
- Hacker, Jacob S., Personal conversation, 2007.
- Haider, Steven J., "Earnings Instability and Earnings Inequality of Males in the United States: 1967-1991," *Journal of Labor Economics*, October 2001, pages 799-836.
- Hertz, Tom, "Understanding Mobility in America," Center for American Progress Discussion Paper, April 2006.
- Holtz-Eakin, Douglas, Harvey S. Rosen, and Robert Weathers, "Horatio Alger Meets the Mobility Tables," National Bureau of Economic Research Working Paper #7619, March 2000.
- Hyslop, Dean R., "Rising U.S. Earnings Inequality and Family Labor Supply: The Covariance Structure of Intrafamily Earnings," *American Economic Review*, September 2001, pages 755-777.

- Kim, Yong-Seong, Tecla Loup, Joseph Lupton, and Frank P. Stafford, "Notes on the 'Income Plus' Files: 1994-1997 Family Income and Components Files," Mimeo, December 2000.
- Kim, Yong-Seong and Frank P. Stafford, "The Quality of PSID Income Data in the 1990s and Beyond," Mimeo, December 2000.
- Lemieux, Thomas, W. Bentley Macleod, and Daniel Parent, "Performance Pay and Wage Inequality," Mimeo, September 2006.
- Moffitt, Robert A. and Peter Gottschalk, "Trends in the Covariance Structure of Earnings in the U.S." Mimeo, July 1995.
- Moffitt, Robert A. and Peter Gottschalk, "Trends in the Transitory Variance of Earnings in the United States," *The Economic Journal*, March 2002, pages C68-C73.
- Shore, Stephen H., "The Co-Movement of Couple's Incomes," Mimeo, October 2006.
- Stevens, Ann Huff, "The More Things Change, the More They Stay the Same: Trends in Long-Term Employment in the United States, 1969-2002," National Bureau of Economic Research Working Paper #11878, December 2005.
- Warren, Elizabeth, "Rewriting the Rules: Families, Money, and Risk," Social Science Research Council Paper, 2005.
- Winship, Scott, "Income Volatility, *The Great Risk Shift*, and the Democratic Agenda," Mimeo, 2007.

Table 1
Three-Year Rolling Standard Deviations of Percent Changes

	1973	2004	Change	Percent Change
Earnings of household heads				
Baseline	22	30	8	34
Dropping values = \$0	20	26	6	30
Dropping values = \$0 and households with financial interest in businesses	20	24	4	23
Earnings of heads, dropping values = \$0				
Under age 35	22	29	7	32
Age 35 to 54	19	25	5	28
Age 55 and older	20	27	7	36
No high school degree	22	30	8	35
High school (but no college) degree	19	26	7	35
College degree	20	27	7	32
Male	19	26	8	42
Female	29	27	-2	-6
Earnings of spouses				
Baseline	39	35	-4	-10
Earnings of heads and spouses pooled				
Dropping head earnings = \$0 and successive spouse earnings = \$0	29	29	0	1
Combined earnings of heads and spouses				
Dropping head earnings = \$0	20	24	5	23
Capital income of heads and spouses	44	46	2	4
Transfer income of heads and spouses	37	44	7	20
Total household income				
Baseline	19	25	6	32
Dropping head earnings = \$0	18	23	4	24
<p>Note. Percent changes as described in text, at annual rates, with nominal values deflated by the CPI. Based on PSID representative sample excluding observations with household head under age 25, household head retired, household head changed from previous observation, or farm income positive. Based on consistent top-coding and bottom-coding of levels, with maximum percent change set at 100.</p>				

Table 2
Selected Studies of Earnings and Income Volatility

Authors (Date)	Data	Measure of Volatility	Key Conclusions
Gottschalk and Moffitt (1994)	PSID; 1970 to 1987; white male household heads aged 20-59; wages and salaries	Variance of transitory earnings defined as gap between actual earnings and individual average earnings	Volatility of earnings rose between the 1970s and 1980s
Gottschalk and Moffitt (1995)	PSID; 1970 to 1987; white male household heads aged 20-59; wages and salaries	Variance of transitory earnings estimated using formal model of earnings dynamics	Volatility of earnings rose between the late 1960s and late 1980s
Daly and Duncan (1997)	PSID; 1969 to 1995; male household heads aged 25-44; labor earnings	Variance of transitory earnings and other measures	Volatility of earnings rose between the 1970s and 1980s
Dynarski and Gruber (1997)	PSID; 1970 to 1991; male household heads aged 20-59; labor earnings	Variance of transitory earnings defined as gap between actual earnings and individual earnings growth path	Volatility of earnings rose in the late 1970s and early 1980s
Cameron and Tracy (1998)	CPS; 1968 to 1997; men; wages and salaries	Variance of transitory earnings	Volatility of earnings rose in 1970s and early 1980s, and later retraced part of run-up
Haider (2001)	PSID; 1968 to 1992; white male household heads aged 25-60; labor earnings	Variance of transitory earnings estimated using formal model of earnings dynamics	Volatility of earnings rose between early 1970s and late 1980s
Hyslop (2001)	PSID; 1979 to 1985; men and women aged 18-60; labor earnings	Variance of transitory earnings estimated using formal model of earnings dynamics	Volatility of earnings rose in 1980s

Moffitt and Gottschalk (2002)	PSID; 1970 to 1996; male household heads aged 20-59; wages and salaries	Variance of transitory earnings defined using decomposition and estimated using model	Volatility of earnings rose in early 1980s and early 1990s, and later retraced run-up
Batchelder (2003)	PSID; 1968 to 1992; households; market income	Variation in income around average income	Volatility of household income rose between 1968 and 1992
Gosselin (2004)	PSID; 1970 to 2000; households; total income	Fluctuations in income	Volatility of household income rose in 1970s and jumped in early 1990s
Hacker (2006)	PSID; 1970 to 2002; households headed by people aged 25-61; total income	Variance of transitory income defined using decomposition	Volatility of household income rose in 1970s and 1980s, surged in 1990s, and later retraced part of run-up
Comin, Groshen, and Rabin (2006)	PSID; 1970 to 1993; household heads; labor earnings	Variance of transitory earnings defined as gap between actual earnings and individual average earnings	Volatility of earnings rose between early 1970s and early 1990s
Gottschalk and Moffitt (2006)	PSID; 1970 to 2002; male household heads aged 20-59; wages and salaries	Variance of transitory earnings defined using decomposition and estimated using model	Volatility of earnings rose in 1970s, 1980s, 1990s, and early 2000s
Hertz (2006)	CPS; 1990 to 2004; households; income	Median absolute value of dollar changes in income	Volatility of household income rose between early 1990s and early 2000s
Winship (2007)	PSID; 1974 to 2000; households; total income	Variance of transitory income defined using decomposition	Volatility of household income rose between mid-1970s and 2000
Bollinger and Ziliak (2007)	CPS; 1979 to 2004; households headed by women aged 16-54; total income	Variance of transitory income defined using decomposition	Volatility of household income was stable in 1980s and early 1990s, but rose between mid-1990s and early 2000s

Bania and Leete (2007)	SIPP; 1992 to 2003; low-income households; total income	Variation in monthly income around average income	Volatility of household income rose between early 1990s and early 2000s
Congressional Budget Office (2007)	CWHS; 1980 to 2003; men and women aged 22-59; labor earnings excl. self-employment	Percent changes in earnings	Volatility of earnings has changed little since 1980

Figure 1

Volatility of Household Head Labor Earnings

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample Excluding Those with Farm Income; 1973-2004 (Vertical Line at 1991)

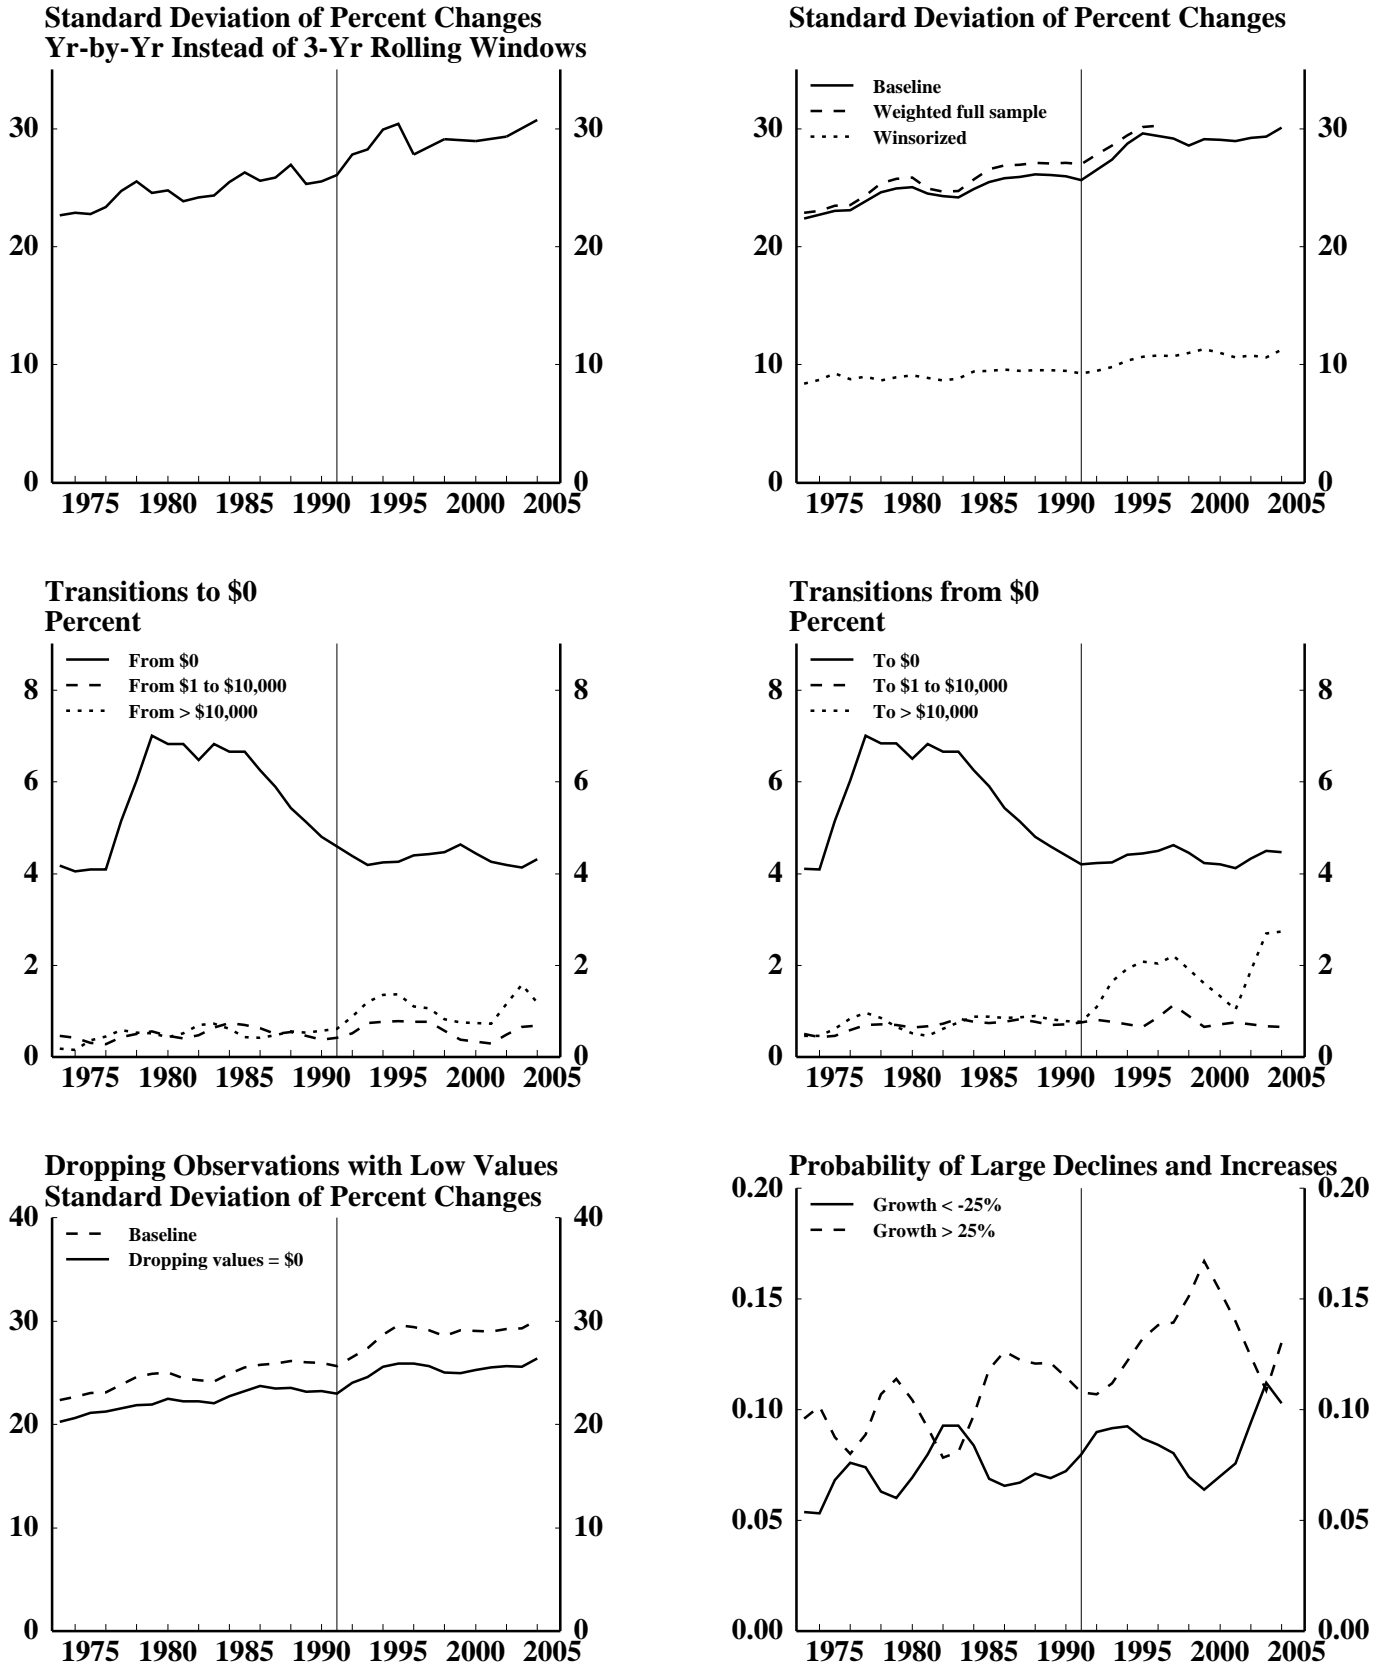



Figure 2


Volatility of Head Earnings by Demographic Characteristics

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample Excluding Those with Farm Income; 1973-2004 (Vertical Line at 1991)

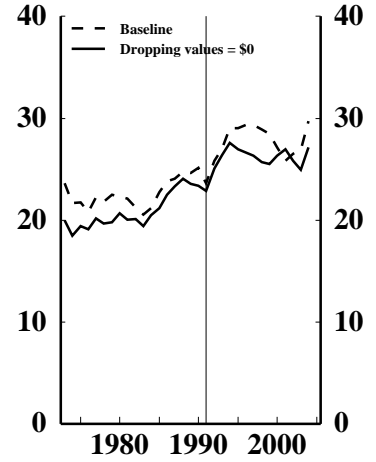
Heads Less than 35 Years Old
Standard Deviation of Percent Changes




Heads 35 to 54 Years Old
Standard Deviation of Percent Changes




Heads 55 Years Old and Older
Standard Deviation of Percent Changes




Heads with No High School Degree
Standard Deviation of Percent Changes




Heads with High School but No College Degree
Standard Deviation of Percent Changes



Heads with College Degree
Standard Deviation of Percent Changes



Male Heads
Standard Deviation of Percent Changes



Female Heads
Standard Deviation of Percent Changes




Figure 3

Volatility of Labor Earnings for Other Groups

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample Excluding Those with Farm Income; 1973-2004 (Vertical Line at 1991)

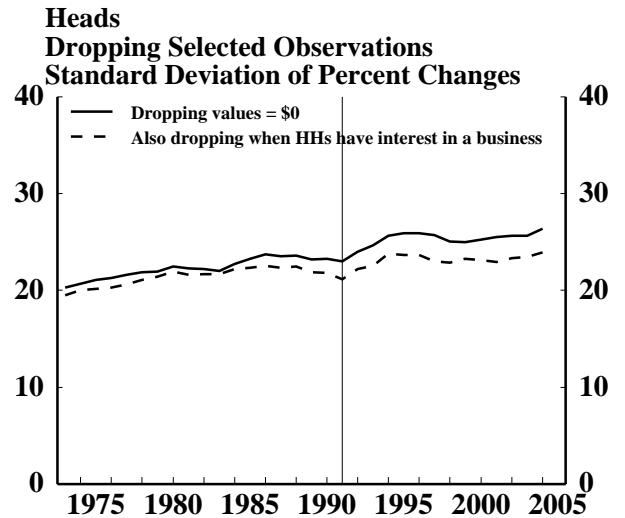








Figure 4

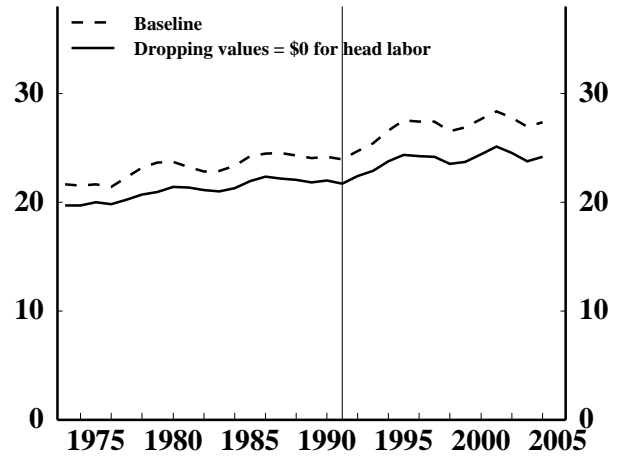
Volatility of Head and Spouse Combined Labor Earnings

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample Excluding Those with Farm Income; 1973-2004 (Vertical Line at 1991)


**Standard Deviation of Percent Changes
Yr-by-Yr Instead of 3-Yr Rolling Windows**



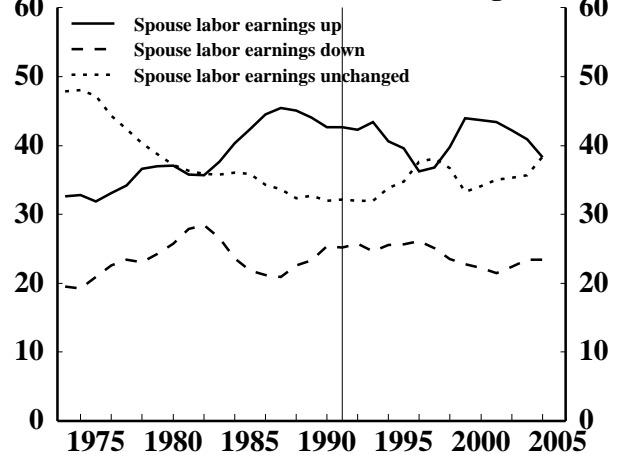
Standard Deviation of Percent Changes



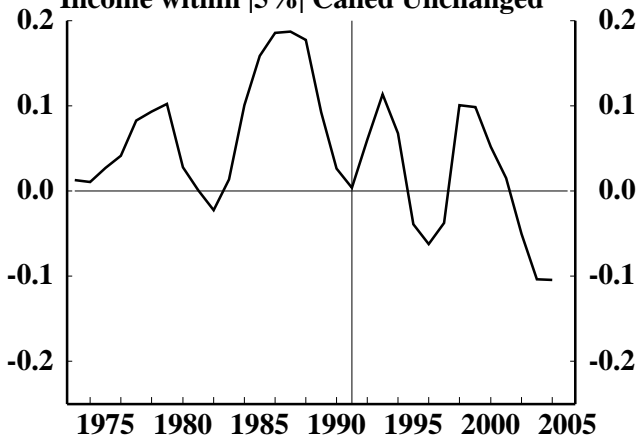
**Head Labor Earnings Down
Income within |5%| Called Unchanged**



**Head Labor Earnings Up
Income within |5%| Called Unchanged**



**Head Labor Earnings Down
Amount Replaced by Spouse Labor Earnings
Income within |5%| Called Unchanged**



**Head Labor Earnings Up
Amount Replaced by Spouse Labor Earnings
Income within |5%| Called Unchanged**





Figure 5


Volatility of Head and Spouse Capital and Transfer Income

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample; 1973-2004 (Vertical Line at 1991)

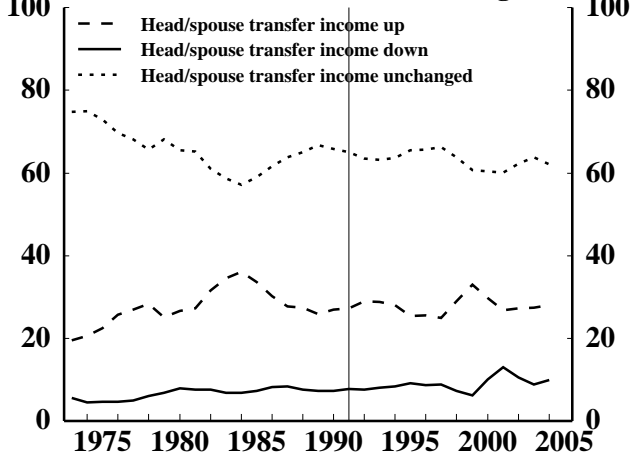
Head and Spouse Capital Income
Standard Deviation of Percent Changes




Head and Spouse Transfer Income
Standard Deviation of Percent Changes




Market Income Down
Income within |5%| Called Unchanged



Market Income Up
Income within |5%| Called Unchanged



Market Income Down
Amount Replaced by Transfer Income
Income within |5%| Called Unchanged



Market Income Up
Amount Replaced by Transfer Income
Income within |5%| Called Unchanged





Figure 6


Volatility of Total Household Income

3-Yr Rolling Windows; Annualized 2-Yr Change Scaled by 3-Yr Average Lagged Income, Topcoded at 100 Percent; Representative Sample; 1973-2004 (Vertical Line at 1991)


**Standard Deviation of Percent Changes
Yr-by-Yr Instead of 3-Yr Rolling Windows**




Standard Deviation of Percent Changes




**Probability of Large Declines
Dropping Values = \$0 for Head Labor**



**Probability of Large Increases
Dropping Values = \$0 for Head Labor**



**Percentiles in Lower Half of Distribution
Dropping Values = \$0 for Head Labor**



**Percentiles in Upper Half of Distribution
Dropping Values = \$0 for Head Labor**

