

Trends in Income Volatility and Risk, 1970–2004

Peter Gosselin

The *Los Angeles Times* and the Urban Institute

Seth Zimmerman

The Urban Institute

Urban Institute Working Paper, May 2008

The Urban Institute

2100 M St. NW

Washington, DC 20037

Copyright © 2008. The Urban Institute. Permission is granted for reproduction of this document, with attribution to the Urban Institute. The views expressed are those of the authors and do not necessarily reflect those of the Urban Institute, its trustees, or its funders.

This paper was prepared for the Association for Public Policy Analysis and Management Fall Research Conference, November 8–10, 2007. The authors would like to thank Sheila Zedlewski and Austin Nichols for their help throughout the research process, and Greg Acs, Pamela Loprest and Dean Lillard for their comments. Please direct comments to Seth Zimmerman at szimmerm@ui.urban.org or Peter Gosselin at Peter.Gosselin@latimes.com.

The work presented in this paper was developed in tandem with the quantitative exercises presented in Peter Gosselin's forthcoming book, *High Wire: The Precarious Financial Lives of American Families*. The paper and the book follow nearly identical quantitative strategies and contain many of the same results. Differences between the paper and the book are the result of decisions made to make the book more accessible to the general reader. These decisions do not affect the conclusions drawn in either the book or the paper.

Introduction

Compared with previous decades, the past 20 years have been characterized by remarkable stability in the economy as a whole, a trend some authors have gone so far as to christen “the Great Moderation.” Between 1985 and 2006, real per capita GDP increased nearly 50 percent, while the variability of GDP growth declined precipitously. These signs of macro-level strength and stability, however, have not translated into increased perceptions of economic security.

In fact, in many ways, Americans are more worried about the prospects that the economy presents them with today than they were before the Great Moderation. In 1982, despite a major recession, only 12 percent of Americans told pollsters working for the International Survey Research Corporation that they worried frequently about being laid off. By 1998, at the top of an economic cycle, 46 percent reported concerns about layoffs in response to the same question. Although worries eased somewhat in the early 2000s, the equivalent figure in 2005 was still 35 percent, nearly three times the 1982 level.¹

Interpreting public concern, of course, is a difficult enterprise. Although increased risk is one potential explanation for these poll results, it is also possible that the phenomenon is primarily psychological. Nevertheless, such polls have helped elevate the discussion of trends in economic risk to a prominent place in media discussions of public policy. In 2004, Peter Gosselin (one of this paper’s coauthors) published a series of articles in the *Los Angeles Times* that focused on the life stories of individuals experiencing large and sudden drops in economic well-being. Two years later, in 2006,

¹ Towers Perrin-ISR annual survey, summer 2007. Patrick Kulesa, global research director for ISR, told the authors that the survey is conducted across more than 400,000 employees and asks for a yes/no response to the statement “I am frequently concerned about being laid off” (Patrick Kulesa, personal communication, October 2007).

Jacob Hacker published a book on economic risk that received wide attention in the popular press. President Bush joined the fray in October 2007, advancing his own explanation for the increasing worry over volatility in a television interview:²

A couple of factors, I think, trouble Americans. One is that there's a lot of churning in the job market. In other words, if you're under 30, you're likely to have had seven jobs by the time you're 30. And older people like me take a look at that kind of volatility or some would call it excitement in the job market, and they wonder whether or not this job turnover is going to affect them.

Popular accounts have recognized the need for empirical verification of a trend toward increased economic risk. Gosselin and Hacker each choose a general measure of income volatility, defined in terms of income variability over time, as a starting point. Gosselin adapts his measure from a 1994 paper by Gottschalk and Moffitt; Hacker relies on a measure developed in a 1995 paper by the same authors.³ When applied to data from the Panel Study of Income Dynamics (PSID), a publicly available data set that tracks individuals over long periods, such methods show fairly large increases in the volatility of family income between the early 1970s (the earliest panel years for which consistent income information is available) and the early 2000s.

But interpreting such findings as direct evidence of increasing economic insecurity is problematic. For one thing, it is unclear how well the volatility of family income reflects true economic risk. One plausible definition of economic risk is the amount that a family's realized economic well-being in a given year tends to diverge from its expected well-being. To believe that family income volatility provides a reasonable proxy for this concept requires two assumptions. First, one must assume that income and well-being are interchangeable. This is clearly not the case: choices that

² Interview with Maria Bartiromo, CNBC, October 11 2007.

³ Gottschalk and Moffitt collaborated on many papers on the measurement of economic volatility, with authors listed in various orders. In the text, we will abbreviate each such paper (GM).

make a family no worse or better off may produce changes in income, while circumstances that have no effect on income may alter a family's well-being. For example, consider a family with a working mother whose entire income is paid toward childcare. If the mother quits her job but is able to reduce childcare expenses to zero, the accompanying change in family income need have no effect on family resources, but will affect measured volatility. Since women in two-earner families earn more money relative to men today than they did thirty years ago, such choices could produce increases in measured volatility without necessarily increasing real risk (though labor market entry and exit *could* be induced by increased risk in other factors). Second, one must believe that income changes over time reflect variation in realized income at a point in time. This assertion cannot be easily evaluated, since we only observe one income for each family within each time unit (i.e., the income the family actually receives).

Even if we accept family income volatility as a valid risk measure, there remains the question of whether to trust the observed volatility trends: in studies like Hacker (2006), peak volatility levels in the early 1990s coincide with changes in methods used by PSID administrators to collect and record income data that have raised questions about data reliability. Further, it is often unclear whether increases in such volatility measures reflect a broad trend in the size of income swings, or merely the effect of a relatively small number of very large income swings. That is, different types of risk with the same consequences for measured volatility may have very different welfare consequences.

This paper presents the results of two empirical exercises aimed at documenting trends in family income volatility, evaluating the robustness of these trends to measurement error, and exploring the relationship between changes in family income and

changes in well-being.⁴ In the first exercise, we measure family income volatility at the individual level using techniques outlined in GM (1994), and discuss the evolution of the volatility distribution over time, across income, age, and education subgroups and in the population as a whole. We consider the relationship between income volatility and patterns of labor market entry and exit within families, and also assess the extent to which various components of family income covary with family income itself—that is, the extent to which different income components insulate against or increase the size of income shocks. To address the question of data integrity, we compare volatility levels in the PSID to those observed in a separate data source, the Survey of Income and Program Participation (SIPP).

Like previous authors, we find that the volatility of family income has risen over time; this increase occurs across income, age, and education subgroups and cannot be written off as the product of outliers in the volatility distribution. Further, although individuals in two-earner families have family incomes that are substantially less volatile than individuals in single-earner or varying-earner families, income volatility increases over time within each earner group. These increases coincide with a gradual reduction in the stabilizing role played by the household head's labor income and transfer income. The trends in the volatility of family income and individual earnings measured in the PSID approximate those observed in the SIPP.

The second exercise, adapted from Burkhauser and Duncan (1989), is aimed exclusively at fleshing out the relationship between volatility and real risk. Here, we consider the relationship between seven potentially destabilizing events—divorce, the

⁴ We raise the issue of potential differences between income variation over time and the variation of realized income at a point in time only as a caveat.

death of a spouse, the birth of a child, loss of work due to retirement or disability, major unemployment of the household head, the loss of work hours due to illness, and a decline in the wife's work hours—and large income drops. Compared with general measures of income volatility, which are doubtless influenced by voluntary tradeoffs between income and other goods, income responses to destabilizing events seem more likely to be the outcome of real risk. To assess changes in the levels of income risk over time, we chart the changing relationship between destabilizing events and income loss over three decades, the most recent ending in 2003. We find that, while the probability of destabilizing events has decreased over the past 30 years, the chances that such an event will be accompanied by major income loss have generally increased. In combination with results from the first exercise, these findings suggest that increasing income volatility cannot be solely ascribed to changes in voluntary behavior.

Literature Review

Income Volatility

The literature on short-term income volatility has expanded rapidly over the past several years. Most findings point to general increases in the volatility of men's earnings and family income between the 1970s and the 2000s, although researchers remain divided about the magnitude of these changes. Much recent work draws heavily upon a series of articles by Gottschalk and Moffitt. In a 1994 paper, Gottschalk and Moffitt quantify volatility as the short-term variance of age-adjusted income, averaged over the population. Applying this definition to men's earnings in the Panel Study of Income Dynamics, they find that volatility between 1979 and 1987 was 42 percent higher than volatility between 1970 and 1978. Later Gottschalk and Moffitt papers like GM (1995)

and GM (2007) refine and update these findings, concluding that, while the volatility of men's earnings tends to vary cyclically, it underwent a secular increase in the early 1980s and has not returned to its previous level.

Hacker (2006) applies techniques from GM (1995) to family income, again using the PSID, and finds striking increases in volatility: roughly 200 percent between the early 1970s and the early 2000s, with volatility in the latter year representing a decline from still-higher levels in the early 1990s. Although many believe that this increase is unrealistically large, Dynan, Elmendorf, and Sichel (2007), hereafter DES, confirm the presence of modest but consistent increases in the volatility of PSID family income between 1970 and 2004.

Complicating these results, a CBO report by Dahl, DeLeire, and Schwabish (2007) finds that the volatility of the earnings of male household heads and the volatility of individual earnings in the Continuous Work History Sample (CWHS, a subset of tax data provided by the Social Security Administration) remained flat between 1980 and 2003. Upon examination, however, these findings do not necessarily contradict results from the PSID. CBO methodology differs in several respects from the methods used in the PSID studies; for instance, the CBO does not include self-employment income in its earnings measure, while DES (2007) do. After eliminating observations in which household heads have a stake in a business, PSID estimates of the volatility of the earnings of male household heads follow CWHS results more closely. Further, increases in the volatility of individual earnings are not a prerequisite for increases in the volatility of family income; in fact, DES (2007) produce both results using the PSID.

Still, the difficult questions remain about volatility measurement in the PSID, especially given the fact that dramatic volatility increases that studies like Hacker (2006) report in the early 1990s coincide with changes in the way the PSID collected and processed income data. Winship and Jacobs⁵ and DES (2007) note that these changes appear to have increased the frequency of erroneous observations of very low incomes. This phenomenon has the potential to produce large spikes in measures of income volatility that quantify income changes in percentage terms (or, equivalently, measure the volatility of log income). If, for instance, a person's observed income increases from an (erroneously measured) \$10 in one year to \$50,000 the next, the 499,900 percent change can dominate measures of volatility that rely on moments of the distribution of income changes or short-term income variances. DES (2007) argue that appropriate data trimming can produce volatility measures that are consistent through time. Thus, while PSID volatility measurements are not necessarily incompatible with those produced using other data sets, it is incumbent on PSID users to test the sensitivity of their results to different treatments of low-income observations.

We build upon existing volatility literature in several ways. First, we confirm trends that previous authors have observed in the year-by-year evolution of family income volatility. Second, we examine trends in income volatility for families exhibiting various patterns of labor force involvement—two-earner families, families switching from one to two earners, and so on—to provide an exploratory starting-point for separating changes in decisions about labor force participation from changes in real income risk. Third, we present results from an alternative publicly available data set, the SIPP, to address concerns about PSID data error. Fourth, and finally, we use a concept

⁵ Scott Winship and Elisabeth Jacobs, personal communication, February and March 2007.

analogous to the beta of a stock to describe trends in the buffering and exacerbating effects that changes in various income components have on changes in family income as a whole.

Income Drops and Destabilizing Events

Many volatility studies compute the risks of income drops of various sizes. Dahl and colleagues (2007) report that, between the late 1980s and 2003, approximately 12–15 percent of individuals in the CWHS saw their earnings drop by half from one year to the next. Using SIPP data, Orszag (2007) finds that about 10 percent of individuals lost half their income between 2001 and 2002. Similarly, using PSID data, DES (2007) find that roughly 10 percent of individuals experienced 44 percent income drops over two years (44 percent declines being the result of consecutive years of 25 percent drops).

The literature on the connection between large income drops and destabilizing life events, however, is not as developed. Burkhauser and Duncan (1989) consider the responses of needs-adjusted family income to four family composition events (divorce/separation, death of spouse, birth of a child, and the departure of a family member), four labor market events (loss of hours due to retirement or disability, major unemployment of the household head, work loss due to illness, a major decline of the wife's work hours), and one asset-related event (a large decrease in asset income). Using PSID data, they find that, between 1968 and 1983, individuals had between a 16 and a 22 percent probability of experiencing one of these events in a given year, depending on age and gender. The probability that an event would be associated with an income drop ranged widely over events and age/gender groups: only 2 percent of men between 26 and 35 who experienced a large drop in asset income also experienced a large drop in overall

income, compared with 23 percent of women in that age group who experienced a divorce. We extend this work through the most recent PSID panel (although we eliminate the asset income and departure of family member event categories) and compute separate statistics for three different decades to document changes in the relationship between destabilizing events and income drops over time.

Methods and Data

Volatility

At the individual level, we define volatility as the short-term variance of age-adjusted income, following GM (1994). More formally, consider a data set containing longitudinal income and age observations for a sample of individuals over years ranging from one through n . If y_{it} is the residual from a regression of log income on a quartic in age for individual i in year t , then volatility for individual i is equal to the estimated variance of the set $\{y_{i1}, y_{i2}, \dots, y_{in}\}$. The intuition behind this measure, which Gottschalk and Moffitt term the transitory variance, is fairly straightforward: an individual's volatility is equal to the variance of his log income over time, after accounting for the expected evolution of income over age.⁶ Gottschalk and Moffitt compute volatility for the population as a whole by taking the average of the individual volatility levels and, in general, we follow their example. However, it is also informative to consider various percentiles of the volatility distribution, and we will do so as the need arises.

We apply this measure to PSID data from panel years 1970 through 2005, which provide annual income data for the years 1970–96 and even-year income data thereafter.

⁶ It is possible to further simplify this method by skipping the regression and computing individual volatility as the short-term variance of income or log income. Results from Nichols and Zimmerman (forthcoming) indicate that this method produces results similar to those presented here.

Begun in 1968 with a sample of approximately 5,000 families, the PSID follows individuals and their descendants over time, tracking changes in incomes, behaviors, and living situations. The PSID was administered annually between 1968 and 1997, and biannually thereafter. The most recent wave, administered in 2005, includes weighted data for more than 8,000 families and 16,000 individuals. The main PSID sample is divided into several subsamples: a core sample, a low-income oversample, and a 1997 refresher sample containing families that immigrated after 1968 as well as their adult children. We consider individuals from each of these groups.

The PSID collects data on many types of annual pretax income, including earned income, asset income, and cash transfer income. We focus on family income—the sum of earned, asset, and transfer income over all family members, less out-transfers like alimony and child support—on the grounds that an inclusive measure best captures overall well-being.⁷ With an eye toward general accessibility, we have not adjusted income for needs; however, scaling income to the PSID-provided USDA needs standard⁸ does not substantively alter the results presented here (see figure A1).

To compute volatility in year t for a given individual, we consider that individual's income as observed in years t , $t + 2$, $t + 4$, and $t + 6$. Including only alternate years in our calculation allows us to measure volatility consistently across the switch to biannual data collection in the late 1990s.⁹ We limit our sample by deleting observations

⁷ Following the PSID, we define the family unit broadly, as a group of individuals permanently cohabiting and sharing income and expenses. Note that our income measure excludes the value of in-kind transfers like food stamps.

⁸ This is an Orshanky-type needs standard, based on weekly food needs for individuals of different ages and genders and converted to account for economies of scale.

⁹ Gosselin (2008) measures volatility in year t using income in years t , $t + 1$, $t + 2$, $t + 3$, and $t + 4$. Choosing a five-consecutive-year analysis window does not substantively alter the results presented here, and makes findings more accessible to a general audience.

in which the respondent either (1) is under 25 or over 64 years old, or (2) reports living in a family with income of less than \$10 (in 2007 dollars) before out-transfers.

The motivation for the first trim is straightforward: we do not want the fluctuations in income that accompany the transition from education to full-time work or full-time work to retirement to affect our measure of volatility. The second trim requires a fuller explanation. Other authors have proposed a variety of bottom trims aimed at minimizing the impact of changes in the levels of measurement error in the mid-1990s. We considered several of their suggestions, including (1) bottom-coding family income at \$4,000,¹⁰ and (2) eliminating observations in which the head reports zero labor income (DES 2007). Figure A2 shows mean and median volatility levels produced using each of these techniques. In general, the more extensive the trim, the lower the volatility level, as would be expected, given that we view income in log terms and people with low incomes tend to experience larger percent changes in income. However, an upward trend in volatility is evident in each case.

In the knowledge that trim choice would not affect the way we described our results, we ultimately chose to eliminate only trivially low levels of income. Similarly, after determining that 1–2 percent top trims did not substantively alter our results, we left the upper end of the income distribution entirely untrimmed. Though our method may overstate the magnitude of the increase in volatility, it largely avoids ascribing real volatility to measurement error, which is arguably an equal hazard. After eliminating

¹⁰ Scott Winship and Elisabeth Jacobs, personal communication, February and March 2007.

observations that fail to meet these two specifications, we drop all individuals who are observed less than three times in the analysis window.¹¹

When considering income volatility within various subgroups, we introduce several other trims. For age subgroups, we drop observations outside the given age range. Similarly, for education subgroups, we drop observations in which individuals report education levels outside the appropriate group. For income subgroups, we compute income quintiles using average family income (i.e., permanent family income) over each year in the analysis window, then drop individuals with average incomes outside the quintile in question. Grouping by earner pattern presents somewhat more of a challenge. We assign earner groups based on changes in the number of earners in a given family—one group for people who are consistently in two-earner families, one group for people who are in one-earner families in one year and two-earner families in a later year, and so forth. Families in which neither the head nor the wife had any earned income are excluded. The problem here is differentiating between income shifts that result from changes in the number of earners within a household and income shifts that result from an individual moving to a household with a different number of earners. Because the aim of this subgroup split is to explore the volatility response to voluntary changes in labor force participation on the household level, we restrict individuals included in the earner subgroups to those who have the same household head in each window year.

¹¹ Requiring the presence of individuals in three out of seven years may affect the representativeness of our sample. Each year, a small but non-negligible group of PSID families, typically one to three percent (Hill 1991, chapter 5), either decline to be interviewed or cannot be located by PSID administrators. There is no reason to expect that attrition occurs randomly. In fact, it seems likely that individuals who administrators cannot locate live less stable economic lives than people who are consistently surveyed. If this is the case, our sample restrictions may tend to lower observed volatility. See Nichols and Zimmerman (forthcoming) for a discussion of the effects of sample restrictions on measured volatility.

To investigate the relationship between the swings in family income and swings in its components, we compute a measure of the contribution of each component of family income to the variability of family income. This measure, described in Jenkins (2000), is analogous to the beta of a stock, and it can be obtained by conducting individual-level regressions of the income component in question, labeled income component k , on family income (using real income, not residualized log income as above), then taking the average of the slope coefficients (β_k s) over all individuals.¹² We then compare the variability contribution of each component to that component's contribution to total income. Note that, for a given individual, if $\beta_k/(u_k/u) = 1$, where u_k is the longitudinal mean of income component k and u is the longitudinal mean of family income, then component k contributes equally to the level and variability of family income. If $\beta_k/(u_k/u) < 1$, component k contributes more to the longitudinal average of family income than to the longitudinal variability. Conversely, $\beta_k/(u_k/u) > 1$ implies that component k contributes disproportionately to variability. To limit the impact of changes in family composition on our findings, we consider only individual in families with stable head/wife structures over the relevant analysis window.

Income Drops

We consider seven potentially destabilizing events—divorce, death of a spouse, birth of a child, the loss of work due to retirement or disability, major unemployment of the household head, loss of work due to illness, and a decline in the wife's work hours—and the relationship between these events and 50 percent income drops. As in our volatility analysis, we use data from the PSID.

¹²Stated rigorously, $\beta_k = \rho_k \sigma_k / \sigma$, where ρ_k is the correlation between income component k and total income, σ_k is the standard deviation of component k , and σ is the standard deviation of total income.

In a given panel year, the PSID collects data on current living situations (e.g., marital status and family structure) and prior-year income and work hours. To accommodate this disjoint, the procedures we use to identify life events and match them to income drops differ depending on whether the life event in question involves a change in living situation or a loss of income or hours.¹³ To allow for consistency across the switch to biannual data collection, we assess living-arrangement events such as divorce or death of a spouse over two-year intervals and code the event to the later year: if a respondent reports being divorced in 1988 but not in 1986, we mark 1988 as the divorce year. We match events recorded in year y to income drops measured between years $y - 3$ and $y - 1$ or $y - 1$ and $y + 1$. Continuing the example, if a respondent who became divorced between 1986 and 1988 lost 50 percent of his income between 1985 and 1987 or 1987 and 1989, we code that divorce as coinciding with an income drop.

For income- and hours-related events that can be identified using a single year of data, like major unemployment or loss of work hours due to illness, we record events in year y based on work losses experienced in year $y - 1$ (and reported in the year y panel). We match these events to income drops that occur between years $y - 3$ and $y - 1$ or $y - 1$ and $y + 1$. So, if a respondent in the 1988 panel reports major unemployment in 1987, we record the event in 1988 and note its association with an income drop if that respondent lost 50 percent of his income between either 1985 and 1987 or 1987 and 1989. Assessing income drops over a pair of two-year intervals allows for consistency across the 1997 switch to biannual data collection and yields estimates of income drop/event associations that are at least roughly comparable to those for living-arrangement events.

¹³ Certain events, such as the loss of work hours due to retirement or disability, require variations on the basic identification strategies described here. See table A1 for detailed event definitions and matching procedures.

One hours-related event—a fall in the wife’s work hours—requires two years of data and is identified and matched with income drops using a different method. Specifically, we record a decline in the wife’s work hours in year y based on changes the wife’s hours between year $y - 1$ and year $y - 3$. We then match the event with income drops occurring over the same set of years.

Note that, as stated, these definitions will not produce comparable estimates of yearly rates of work-related and non-work-related events. For instance, the divorce event captures divorces that took place over a two-year period, while the major unemployment event reflects only one year of unemployment outcomes. To allow for easier side-by-side comparisons of event rates, we approximate the annual rates of events calculated over a two-year interval by dividing computed rates by two. We account for this same problem in joint probabilities (e.g., in the process of assembling the “any event” category”) by randomly assigning two-year events to one year or the other.

Despite this adjustment, some cross-event inconsistencies remain. For instance, we match declines in wives’ hours to a narrower (though more precisely targeted) set of income drops than we use when measuring the associations between drops and other events. This will produce measured association rates between declines in wives’ hours and income drops that are lower relative to the true association rate than are measured association rates for other events. We therefore focus our discussion on changes in event likelihoods and event-drop associations over time, rather across events at a point in time.

We consider three decades of event years: 1974–83, 1984–93, and 1994–2003. Because our matching techniques involve data collected two years before and two years after the each event year, these three periods include data from panel years 1972–85,

1982–95, and 1992–2005, respectively. As in the volatility exercise, we restrict our sample in several ways. First, we only include individuals who are present in the survey and between the ages of 25 and 64 in each relevant year. We also eliminate individuals in the top or bottom 2 percent of the distribution of income changes in any relevant year. After making these trims, we compute year-by-year rates of income drops, life events, and the association between the two, and define the decade-by-decade rates as the average of the yearly rates.

Results

Calculating Income Volatility

Evolution of the Volatility Distribution

In the population as a whole, average volatility rose significantly between 1970 and 1998. The median, the 25th percentile, and the 75th percentile of the volatility distribution increased as well, though not as dramatically (figure 1). Mean volatility increased more or less monotonically from 0.116 to 0.252, or 118 percent, while median volatility rose 63 percent, from 0.044 to 0.072. The 25th percentile of the volatility distribution also rose, increasing 32 percent, from 0.016 to 0.021. Mean volatility tracks the 75th percentile of the volatility distribution very closely until about 1986, at which point it begins to increase more quickly. The explanation for this separation can be found in the rapid increase in the upper percentiles of volatility, particularly the 99th percentile (figure 2). The general picture here is of a distribution slowly rising in the median and increasing in positive skew.

Volatility by Income, Education, and Age Group

Increases in mean volatility persist throughout income quintiles, education levels, and age groups. Although volatility levels tend to be much higher for people with low income or little education, percentage increases in volatility in high-income and high-education groups are often similar to or greater than percentage increases in the low-income and low-education groups, respectively. For instance, between 1970 and 1998, the income volatility of those in the bottom income quintile increased 156 percent, from 0.23 to 0.59. Over that same period, volatility of individuals in the top income quintile increased 162 percent, from 0.068 to 0.178 (figure 3). Mean volatility also rose for individuals in the middle three quintiles, by about 80 percent in each case.

Similarly, percentage increases in mean volatility levels were the highest for people with college degrees (figure 4). Volatility for individuals with college degrees increased from 0.077 to 0.226, or about 193 percent, while volatility for people with less than a high school education increased 158 percent, from 0.142 to 0.365. Individuals with high school–educated heads saw percentage increases along the lines of the high school dropout group.

Differences in volatility levels across age groups were less dramatic (figure 5). Volatility for people between 25 and 34 years old increased from 0.110 to 0.234, compared with shifts from 0.103 to 0.250 and from 0.132 to 0.244 for the 35–54 and 55–64 categories, respectively.

Volatility by Earner Group

Trends in volatility were similar across workforce participation groups, with one notable exception: people in families with two earners in each year of the analysis window (figure 6). For this group, volatility increased 39 percent between 1970 and 1998, from

0.050 to 0.070. For the other groups, percentage increases ranged from 62 percent for people in families that switched from two earners to one earner to 135 percent for people in families with a single earner across all analysis years. The latter finding, the result of a steady increase over time, is worth highlighting, as it indicates that increases in volatility are not solely the result of individual decisions to shift in and out of the labor force.

Figure 7 charts the changes in the proportion of the population in each earner group over time, with the most noticeable trend being a gradual increase in the proportion of two-earner families peaking in the mid-1980s and declining slowly thereafter.

Decomposing Income Volatility

As mentioned above, one of the more paradoxical findings in the volatility literature is that the volatility of individual earnings seems to remain flat over time, even as the volatility of family income and earnings have risen. We reproduce those findings in both the PSID and the SIPP. In the PSID, the volatility of individual earnings rises only 10 percent between 1970 and 1998, from 0.286 to 0.315 (figure 8). In the SIPP, volatility computed over a window of three consecutive years rises 10 percent between 1983 and 2001. In contrast, the volatility of family income in the PSID rises 118 percent, while the volatility of family income in the SIPP rises 48 percent (figure 9).

The cause of the observed increase in family income volatility therefore remains somewhat of a mystery: it does not appear to be driven by an increase in the volatility of individual earnings, nor can it be fully explained by workers' decisions to enter or leave the labor force. A decomposition of longitudinal family income variability explores the issue in more depth. Figure 10 plots the population averages of the beta coefficients for head's labor income, wife's labor income, taxable income from other household

members, transfer income, and other income. The coefficients stay relatively flat over time, with one notable exception being the coefficient on transfer income, which hovers around and often below zero until the mid-1980s, when it increases into positive territory. Figure 11 shows the proportion of family income made up by each of the components, that is, (u_k/u) . The most notable trend here is the slow decrease in the contribution of head's labor income to total income, accompanied by a gradual rise in the role played by the wife's labor income.

Figure 12 plots $\beta_k/(u_k/u)$, the contribution of each income component to income variation as a fraction of its contribution to total income. In the 1970s, head's labor income made a disproportionately small contribution to overall variability, while wife's labor income made a disproportionately large contribution. By the 1990s, however, both types of labor income contributed almost equally to family income levels and income variation. But total labor income (labor income of the head and wife combined) contributed proportionally more to the variability of family income in the last period than in the first. In 1970, labor income accounted for 79 percent of total income and 65 percent of total variability; in 1998, labor income accounted for 77 percent of income and 72 percent of variability. Increasing variability in head's labor income thus outweighs decreasing variability in wife's labor income.

Another striking trend is the increase in the contribution of transfer income to overall income variation. In the 1970s and early 1980s, increases and decreases in transfer income contributed very little to the variation of family income, and sometimes even offset other changes. In the mid- to late 1980s and early 1990s, however, transfer income provided less of a buffer against variation in other income components.

Life Events and Income Drops

Between the ten-year period beginning in 1974 and the ten-year period beginning in 1994, the probability of 25 percent drops in income and income scaled to needs increased modestly (table 1). In the earlier decade, individuals between ages 25 and 64 had a 17.1 percent chance of experiencing a 25 percent drop in real income over a given two-year period, compared with a 19.3 percent chance in the later decade. Over the same interval, chances of a 50 percent income drop increased from 4.6 percent to 7.8 percent.

This increase in the risk of income drops coincided with a slight decrease in the risk of destabilizing events (table 2). While the probability of some events, like divorce and declines in wives' work hours, did increase slightly, the probability of experiencing any of the seven events in a given year declined from 21.0 percent to 19.2 percent. People who experienced these events, however, grew substantially more likely to experience 50 percent income drops (table 3). For instance, between 1974 and 1983, 17.3 percent of individuals in families in which the head experienced major unemployment also experienced a 50 percent income drop. Between 1994 and 2003, that figure was 25.5 percent. Similarly, the proportion of individuals in families experiencing a divorce that also experienced an income drop grew from 29.6 percent in the earlier decade to 36.2 percent in the later decade. Overall, the percentage of individuals experiencing income drops associated with destabilizing life events increased by almost half, from 14.3 percent to 20.2 percent. Destabilizing events also became increasingly associated with 50 percent drops in income scaled to needs (table A2).

Table A3 reports the frequency of each event in an average year in each analysis window. Though some of the events are relatively rare within the sample population in a

given year, each analysis window contains at least a few hundred observations of each event. For instance, only 39 sample individuals experienced the lowest-frequency event—the death of a spouse—in an average year in the 1994-2003 window. However, 39 occurrences per year times six years (1994, 1995, 1997, 1999, 2001, and 2003) implies that the reported event-income drop associations were based on 234 potential associations.¹⁴ Other events averaged, at minimum, roughly one hundred occurrences per year (table A3).

Discussion

A flurry of recent research has documented changes in the volatility of family income between the 1970s and the early 2000s. Most authors report substantial increases in volatility, within various subgroups as well as the population as a whole. The causes and consequences of this trend, however, remain relatively uninvestigated, leaving open the possibility that changes in volatility do not reflect true changes in income risk. Specifically, observed increases in overall volatility may be the product of data error in the PSID, the outsize contribution few high-volatility outliers, or voluntary, utility-neutral behavior.

This paper addresses the first two concerns directly and finds that increases in volatility cannot be fully explained by measurement error or the presence of a few individuals with very volatile incomes. Though large increases in PSID-measured average volatility in the early 1990s could be traced in part to concurrent jumps in the

¹⁴ Considering only odd-numbered years to avoid double counting events computed over two-year intervals yields a somewhat smaller total of 169 observations.

upper percentiles of the volatility distribution (i.e., a relatively small number of individuals with very volatile incomes), the median and quartiles of the volatility distribution also increased over time. Nor does PSID-specific measurement error appear to be the principal force behind the observed increase, as the volatility of annualized family income in the SIPP increased over the period as well. This result is consistent with previous studies, which report that the increase in the volatility of family income persists even after making various adjustments to reduce measurement error.

To address the third concern, we present two exercises exploring the relationship between income volatility and real risk. The first considers a straightforward attempt to ascribe changes in volatility to (potentially) voluntary transitions into and out of the labor force. We find that volatility has increased even for individuals in families that consistently have only one earner, which suggests that volatility has increased even disregarding voluntary labor force transitions. The second exercise assesses the contribution of various components of income to income variability. Here, we find that neither family labor income nor transfer income provided as much of a buffer against overall income variability in the late 1990s and early 2000s as in the 1970s. This result is consistent with the findings of DES (2007), who trace part of the increase in family income volatility to increases in the volatility of labor and transfer income. The impact of transfer income on overall volatility is particularly important to note, as fluctuations in transfer income are presumably somewhat less likely to be the result of voluntarily accepted tradeoffs than fluctuations in other types of income.

Results from our analysis of the association between destabilizing life events and large income drops provide more concrete evidence that real income risk has increased.

Like other researchers, we find that the frequency of 50 percent income drops has increased moderately over time. Further, we find that large income drops were more likely to accompany destabilizing events in the 1990s and early 2000s than in the 1970s and 1980s. While abstract volatility measures are vulnerable to criticism because they conflate real risk with changes in voluntary behavior, it is more difficult to argue that income drops in families in which the head reported major unemployment are the result of an increased preference for leisure.

It should of course be noted that our analysis considers only pretax income in an era over which many benefits programs were incorporated into the tax code. Further, we do not consider data on spending, a potentially problematic omission given research suggesting that decreasing liquidity constraints have allowed families to smooth consumption more effectively over changes in income. If families today are better able to use credit to maintain a given standard of well-being over low-income months or years than they were in the past, increases in income volatility may have few practical consequences. But, insofar as income and the ability to count on future income are important to Americans' perceptions of their own security, our work suggests that the increasing public concern about risk may well have some basis in reality.

Works Cited

- Burkhauser, Richard V., and Greg J. Duncan. 1989. "Economic Risks of Gender Roles: Income Loss and Life Events over the Life Course." *Social Science Quarterly* 70(1): 3–23.
- Dahl, Molly, Thomas DeLeire, and Jonathan Schwabish. 2007. "Trends in Earnings Variability over the Past 20 Years." Washington, DC: Congressional Budget Office.
- Dynan, Karen, Douglas Elmendorf, and Daniel Sichel. 2007. "The Evolution of Household Income Volatility." Draft, June 2007.
- Gosselin, Peter. 2008. *High Wire: The Precarious Financial Lives of American Families*. New York, NY: Basic Books.
- Gottschalk, Peter, and Robert Moffitt. 1994. "The Growth of Earnings Instability in the U.S. Labor Market." *Brookings Papers on Economic Activity* 1994(2): 217–72.
- Gottschalk, Peter, and Robert Moffitt. 2007. "Trends in Earnings Volatility in the U.S.: 1970–2002." Draft, January 2007.
- Hacker, Jacob S. 2006. *The Great Risk Shift*. New York: Oxford University Press.
- Hill, Martha S. 1991. *The Panel Study of Income Dynamics: A User's Guide*. Newbury Park, CA: Sage.
- Jenkins, Stephen P. 2000. "Modelling Household Income Dynamics." *Journal of Population Economics* 13(4): 529–67.
- Moffitt, Robert A., and Peter Gottschalk. 1995. "Trends in the Covariance Structure of Earnings in the U.S.: 1969–1987." Draft, July 1995.
- Nichols, Austin, and Seth Zimmerman. Forthcoming. "Measuring Trends in Income Variability." Washington, DC: The Urban Institute.
- Orszag, Peter. 2007. "CBO Testimony: Economic Volatility." Washington, DC: Congressional Budget Office.

Figure 1: Mean, Median, and Quartiles of the Family Income Volatility Distribution





Figure 2: Upper Percentiles of the Family Income Volatility Distribution



Note that the results presented on this page differ from those presented in Gosselin (2008); see the methods section for details.

Figure 3: Volatility of Family Income by Income Quintile





Figure 4: Volatility of Family Income by Education Level



Note that the results presented on this page differ from those presented in Gosselin (2008); see the methods section for details.

Figure 5: Volatility of Family Income by Age Group





Figure 6: Volatility of Family Income by Earner Group



Note that the results presented on this page differ from those presented in Gosselin (2008); see the methods section for details.

Figure 7: Population Proportion by Earner Pattern





Figure 8: Volatility of Individual Earnings in the PSID and the SIPP



Note that the results presented on this page differ from those presented in Gosselin (2008); see the methods section for details.

Figure 9: Volatility of Family Income in the PSID and the SIPP




Figure 10: Betas of Components of Family Income




Figure 11: Proportion of Family Income, by Component




Figure 12: Beta Coefficients Divided by Component Proportions




Table 1: Probability of Income Drops of Various Sizes over Two Years

Period	Age group	Income	Falling by	Income/Needs	Falling by
		falling by at least 25%	at least 50%	Falling by at least 25%	at least 50%
1974–83	25–65 years old	17.13%	4.60%	15.05%	3.63%
	35–55 years old	15.96%	3.95%	12.50%	2.79%
1984–93	25–65 years old	16.94%	5.71%	15.36%	4.73%
	35–55 years old	15.77%	5.12%	13.38%	4.05%
1994–2003	25–65 years old	19.31%	7.75%	18.49%	6.99%
	35–55 years old	18.07%	7.24%	17.06%	6.61%

Table 2: Annual Risk of Various Life Events, by Decade

	Age in year before drop in income	Divorce/ sep	Death of spouse	Birth of child	Reduction in head's hours due to retirement or disability	Unemployment of head	Work loss of head due to illness	Fall in work hours of wife	Any of the seven events
1974–83	25–65 years old	1.43%	0.46%	3.06%	1.74%	6.94%	3.11%	6.99%	21.00%
	35–55 years old	1.09%	0.40%	0.66%	1.09%	7.03%	3.29%	6.51%	18.22%
1984–93	25–65 years old	1.61%	0.43%	3.29%	1.82%	6.94%	2.39%	7.85%	21.58%
	35–55 years old	1.55%	0.34%	0.96%	1.05%	6.30%	2.69%	7.35%	18.12%
1994–2003	25–65 years old	1.55%	0.37%	3.06%	1.39%	5.06%	2.08%	8.20%	19.23%
	35–55 years old	1.51%	0.35%	1.16%	0.98%	5.12%	2.07%	7.45%	16.72%

Table 3: Percentage of Life Events Associated with 50% Income Drops

	Age in year before drop in income	Divorce/ sep	Death of spouse	Birth of child	Reduction in head's hours due to retirement or disability	Unemployment of head	Work loss of head due to illness	Fall in work hours of wife	Any of the seven events
1974–83	25–65 years old	29.57%	25.09%	7.24%	25.19%	17.27%	14.91%	8.47%	14.25%
	35–55 years old	31.72%	22.56%	6.44%	28.81%	16.04%	14.90%	7.13%	14.15%
1984–93	25–65 years old	29.32%	30.01%	8.07%	30.94%	21.79%	16.39%	9.30%	16.86%
	35–55 years old	32.59%	32.04%	10.10%	29.26%	21.90%	13.38%	9.53%	17.18%
1994–2003	25–65 years old	36.17%	39.56%	9.53%	34.91%	25.53%	19.29%	12.21%	20.15%
	35–55 years old	35.88%	36.54%	7.18%	33.19%	24.67%	19.86%	11.43%	20.23%

Figure A1: Volatility of Income Scaled to USDA Needs Threshold





Figure A2: Volatility of Family Income under Alternate Trims



*DES trim: Eliminating all observations in which the head reports zero labor income.
WJ trim: Coding all sub-\$4,000 observations at \$4,000.*

Note that the results presented on this page differ from those presented in Gosselin (2008); see the methods section for details.

Table A1: Event Definitions and Matches to Income Drops

Event Definitions

Divorce:

A head/wife experienced divorce in year y if he reported being divorced in year y but not in year $y - 2$.

Widowhood:

A head/wife experienced widowhood in year y if he reported being in widowed in year y but not in year $y - 2$.

New Child:

A family was tagged as having a new child if a child under age 2 was present in the household in year y and that same child was not present in the family in year $y - 2$.

Major Unemployment of Head:

Family units experienced major unemployment of the head in year y if the head reported losing at least 320 hours of work to unemployment, layoffs, or strikes in year $y - 1$.

Fall in work hours of Wife:

Family units experienced a fall in the wife's work hours in year y if the unit wife worked at least 320 hours less in year $y - 1$ than year $y - 3$. In contrast to our treatment of head's unemployment, we do not require that the wife attribute her work loss to unemployment, layoffs, or strikes. We make this decision because we wish to assess the effect of the increasing importance of women's earnings to family income using a broader sample of declines in work hours.

Loss of Work Hours to Illness:

Family units fell into this category in year y if the head reported losing at least eight weeks of work due to his own illness or the illnesses of others in year $y - 1$.

Loss of Work Hours to Retirement or Disability:

Families fell into this category if (a) the unit head reported being retired or disabled in year y but not in year $y - 2$ and (b) the unit either worked 320 less hours in year $y - 3$ than in $y - 1$, or worked 320 fewer hours in year $y + 1$ than in year $y - 1$.

Event Matching

Events other than a fall in the wife's work hours that were recorded in year y were matched to income drops over the two year intervals ending in years $y - 1$ and $y + 1$. Drops in the wife's work hours in year y were matched to income drops ending in year $y - 1$ only. Note that this last event receives special treatment because it is the only event that can be directly matched to changes in income. The matching process is fairly easy to understand when represented visually. The Xs represent the beginning of a recording window, and the text represents the end.

Fall in Work Hours of Wife

<i>Year</i>	Y - 3	Y - 2	Y - 1	Y
<i>Income</i>	X=====		→ Lost Income?	
<i>Event</i>	X=====		→ Lost Hours?	

If yes to both, association recorded

Other Events

<i>Year</i>	Y - 3	Y - 2	Y - 1	Y	Y + 1
<i>Income(1)</i>	X=====		→ Lost Income?		
<i>Income(2)</i>			X=====	→ Lost Income?	
<i>Event</i>		X=====	→ Became Divorced?		

If yes to both, association recorded

Table A2: Percentage of Life Events Associated with 50% Income Drops in Income/Needs

	Age in year before drop in income	Divorce/ sep	Death of spouse	Birth of child	Reduction in head's hours due to retirement or disability	Unemployment of head	Work loss of head due to illness	Fall in work hours of wife	Any of the seven events
1973–83	25–65 years old	16.68%	14.66%	8.33%	20.41%	15.16%	12.34%	5.36%	11.31%
	35–55 years old	16.66%	10.59%	7.33%	21.01%	12.04%	10.30%	3.95%	9.53%
1983–93	25–65 years old	17.23%	17.81%	9.55%	27.02%	18.74%	13.93%	6.85%	13.92%
	35–55 years old	18.28%	20.83%	14.34%	25.48%	17.60%	10.31%	6.39%	13.03%
1993–2003	25–65 years old	22.48%	24.56%	11.26%	30.51%	24.46%	14.36%	10.58%	17.39%
	35–55 years old	21.15%	21.84%	9.40%	25.33%	24.21%	14.44%	9.73%	16.80%

Table A3: Average Annual Number of Individuals Experiencing Destabilizing Events

	Divorce/ separation	Death of spouse	Birth of child	Reduction in head's hours due to retirement or disability	Unemployment of head	Work loss of head due to illness	Fall in work hours of wife
1974-83	167.1	58.7	284.5	189.7	416.2	185.9	680.3
1984-93	198.8	62.4	358.9	195.3	489.5	139.5	800.0
1994-2003	145.7	39.0	264.8	129.2	273.7	96.5	708.7