Journal of Business & Economic Statistics

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ubes20

Estimating Trends in Male Earnings Volatility with the Panel Study of Income Dynamics

Robert Moffitt & Sisi Zhang

To cite this article: Robert Moffitt & Sisi Zhang (2022): Estimating Trends in Male Earnings Volatility with the Panel Study of Income Dynamics, Journal of Business & Economic Statistics, DOI: 10.1080/07350015.2022.2102024

To link to this article: https://doi.org/10.1080/07350015.2022.2102024

JOURNAL OF BUSINESS & ECONOMIC STATISTICS
2022, VOL. 00, NO. 0, 1–6
https://doi.org/10.1080/07350015.2022.2102024

Taylor & Francis
Taylor & Francis Group

Estimating Trends in Male Earnings Volatility with the Panel Study of Income Dynamics

Robert Moffitt[a] and Sisi Zhang[b]

[a]Johns Hopkins University, Baltimore, MD; [b]Jinan University, Guangzhou, China

ABSTRACT
The Panel Study of Income Dynamics (PSID) has been the workhorse dataset used to estimate trends in U.S. earnings volatility at the individual level. We provide updated estimates for male earnings volatility using additional years of data. The analysis confirms prior work showing upward trends in the 1970s and 1980s, with a near doubling of the level of volatility over that period. The results also confirm prior work showing a resumption of an upward trend starting in the 2000s, but the new years of data available show volatility to be falling in recent years. By 2018, volatility had grown by a modest amount relative to the 1990s, with a growth rate only one-fifth the magnitude of that in the 1970s and 1980s. We show that neither attrition or item nonresponse bias, nor other issues with the PSID, affect these conclusions.

ARTICLE HISTORY
Received August 2020
Accepted July 2022

KEYWORDS
Earnings; Panel data; Volatility

This article is part of a group project aimed at reconciling the disparate results on trends in male earnings volatility using four different datasets and six different data series, both survey and administrative. The Overview paper in this volume discusses the background literature in detail, describes all datasets and data series, including their differences and their comparability, and presents the results of the reconciliation exercise. This specific article provides more detailed results for one of the data series, the Panel Study of Income Dynamics (PSID). While some comparisons are made to the other datasets in the project and their results, most of the cross-data-series comparisons appear in the Overview paper.

The PSID is considered to be the workhorse dataset for estimating trends in individual earnings volatility in the United States. It is a longitudinal survey that has been ongoing since 1968 (and hence the longest-running general-purpose socioeconomic panel in the world), which has attempted to maintain reasonable population representativeness and which asks extensive questions on labor market activity. The use of the PSID for the study of male earnings volatility began with Gottschalk and Moffitt (1994), who found male earnings volatility to have increased from 1970 to 1987, with the largest increases occurring among the less educated. About a dozen PSID studies subsequent to the Gottschalk–Moffitt study have been conducted, almost all of which have also found increases in male earnings volatility over time. A full listing of these studies can be found in Moffitt and Zhang (2018).

This article reanalyzes the data used in past studies but provides updated results extending to more recent years of data not available in past work. In addition, this article conducts a series of examinations of the PSID data aimed at gauging the importance of attrition, of nonreporting and imputation of earnings in the survey data, and of a number of other threats to its population representativeness. We also conduct an examination of the role of trimming of the lower tail of the earnings distribution on estimated trends in earnings volatility.

The results of the analysis are 4-fold. First, we confirm prior work showing two phases of upward trends in male earnings volatility, one running from the 1970s to the 1980s or early 1990s, and one starting in the early or mid-2000s (with a phase of stable volatility in between, from the early 1990s to the early or mid-2000s). But the new results using recent years of data show volatility to be recently declining and to have fallen to a level only slightly above its level in the 1990s, with a net increase only one-fifth the size of the growth rate from the 1970s to the early 1990s. Consequently, the PSID is roughly consistent with results from other datasets showing little increase in volatility since 1990, as discussed in the Overview paper. Second, we find that the pattern of volatility trends is similar across all levels of the cross-sectional earnings distribution, but that the distribution of earnings changes shows trends to have been most pronounced in the tails of that distribution. Third, we find that neither attrition, imputation, nor other threats to the representativeness of the PSID are likely to be responsible for the patterns of volatility trends found in the data. Fourth, we find that differences in the volatility trends in the PSID and some past work with administrative data may be a result of differences in the size of the left tail of earnings and of possible differences in trimming methods at the bottom.

The outline of the article is as follows. Our first section briefly clarifies different definitions of volatility. The second section reviews the PSID dataset and our sample. The third section

CONTACT Robert Moffitt moffitt@jhu.edu Johns Hopkins University, Baltimore, MD.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/UBES.

reports our main findings on trends in male earnings volatility, while the fourth section conducts sensitivity tests related to attrition, imputation, immigration, and specification issues as well as an examination of the impact on estimated trends of methods of trimming at the bottom. A short summary concludes.

1. Measuring Volatility

Volatility is defined in different ways in the literature. Intuitively, volatility is just some measure of dispersion in the rate of change over time for some variable y. With a panel dataset consisting of individual observations on y_{it} for observations $i = 1,\ldots,N$ and $t = 1,\ldots,T$, the degree of volatility is often measured simply by the cross-sectional variance (or some other measure of dispersion) of the change in y between two time periods t and $t+1$, for example, the variance of $y_{i,t+1} - y_{it}$.

We shall use this definition of volatility in our work but emphasize that it is different—except in a special case—from the variance of the transitory component in a traditional permanent-transitory error components model. The special case in which they coincide is when $y_{it} = \mu_i + \nu_{it}$, with μ_i a traditional time-invariant permanent component and ν_{it} a traditional transitory component, and with the two components distributed independently. With ν_{it} iid, one-half the variance of $y_{i,t+1}$-y_{it} equals the variance of ν_{it}. However, while this simple model is still the standard in textbooks, the earnings dynamics literature has long moved beyond it, most importantly by allowing the permanent component to change over time. Most often the change is represented by a random walk, but often alternatively by a random growth factor or some other evolutionary process (Tables 1 and 2 in Moffitt and Zhang (2018) list the various error component models of income and earnings dynamics used in the literature). In the random walk case, with $y_{it} = \mu_{it} + \nu_{it}$ and $\mu_{i,t+1} = \mu_{it} + \omega_{it}$, and with ω_{it} distributed independently of μ_{it} and of ν_{it} at all t, the variance of $y_{i,t+1} - y_{it}$ contains the random walk variance in the permanent component as well as the transitory variance. We use the term "volatility" in our paper to mean *gross* volatility, composed of volatility in both the permanent and transitory components of an underlying error components model.

2. The Panel Study of Income Dynamics

The PSID is a longitudinal dataset based on a representative sample of household units in the 1968 U.S. population. The members of the households and their descendants have been followed over time, and so-called "splitoff" families—mainly children who leave the family and form new households—are also followed, allowing the survey to stay reasonably representative of the U.S. population, unlike most cohort studies. The 1968 sample also included a low income oversample (the so-called "SEO" sample) but we exclude this oversample in our analysis. We also exclude later Latino samples which were added to the survey in an attempt to address the exclusion of immigrants since 1968, an issue we examine below. We further exclude PSID "nonsample" members.

Families were interviewed annually until 1996 and have been interviewed biennually since that year. For this reason, we will look at 2-year volatility for all periods. Respondents are asked questions about the most recent calendar year's income, both for total family income and its components. The earnings components are separately identified only for household heads and their spouses and not for others in the household. The restriction to heads and spouses also opens up a possible difference with other datasets, which often include non-heads (and, in some administrative datasets, headship is not even identified). However, as discussed in the Overview paper, the Current Population Survey (CPS) and Survey of Income and Program Participation (SIPP) papers in this project compared volatility trends for heads and non-heads and found them to be the same, although having differences in levels (see the Overview paper). We also use only wage and salary income and exclude self-employment earnings (we include men who have positive wage and salary earnings even if they have some self-employment earnings as well). Our earnings measure also does not contain tips, overtime, commissions, or bonuses, as there is no consistent measure of these quantities over time in the PSID. As discussed at length in prior PSID work (e.g., Shin and Solon 2011; Dynan, Elmendorf, and Sichel 2012), the PSID asks self-employment earnings but the way it is defined has changed markedly over time and hence no consistently defined variable is available. Again, however, the Overview paper reports work from the CPS and SIPP showing estimated trends to be the same for wage and salary and self-employment earnings but with, again, differences in levels.

We form a sample of male heads from interview year 1971 through interview year 2019, with our earnings measure therefore, covering years 1970–2018. Our baseline sample includes all male heads of family between the ages 25 and 59 in the interview year. We also keep only those who have positive wage income and positive weeks worked in each year (at least for our baseline sample; we will examine nonworkers in an alternative sample) and those who are not full-time students, and we necessarily exclude those who are not interviewed in a year, which means excluding individuals who have attrited. We work with residuals from regressions of the 2-year change in log earnings on a polynomial in age, all estimated separately by calendar year (regression residuals have been used since the first paper by Gottschalk and Moffitt (1994) since it allows a separation of calendar time and life cycle volatility, but the specific age-specification we use is drawn from that of Shin and Solon (2011), which differs slightly from that of Gottschalk and Moffitt). We also trim the top and bottom 1% of the log earnings distribution each year (prior to the regression) to eliminate outliers which could distort our volatility measures, following the procedure initiated by Gottschalk and Moffitt (1994) and followed in many subsequent papers. We will conduct sensitivity tests to many of these choices. This gives us an unbalanced panel with 4285 men and 48,436 person-year observations, for an average of 11.3 year-observations per person. Earnings are in 2010 dollars, deflated by the Personal Consumption Expenditure (PCE) index.

We show the summary statistics in Appendix Tables 1 and 2, supplementary materials. Appendix Table 1, supplementary materials show the mean and dispersion of real earnings by year in our sample, showing that mean earnings increased over the entire sample period but experienced temporary declines in the

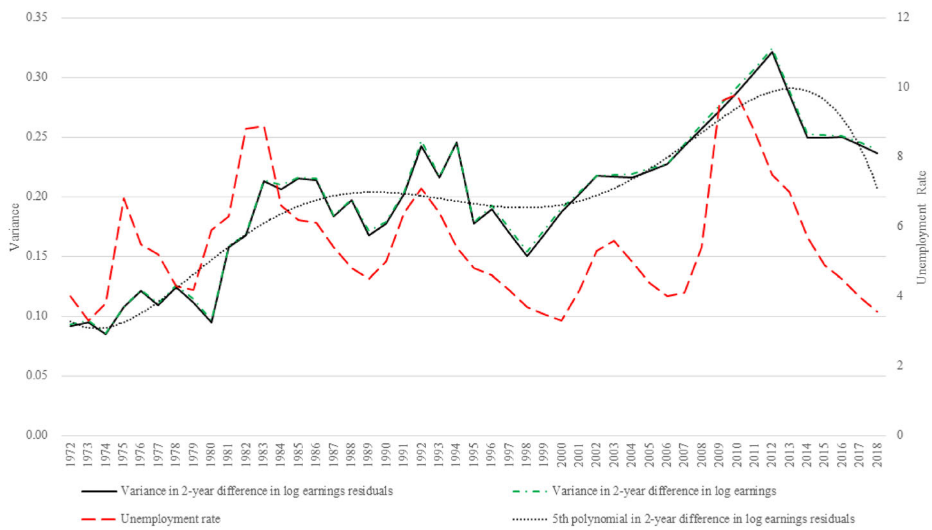

Figure 1. Variance in 2-year difference in log earnings and log earnings residuals.

NOTE: The first value, in 1972 represents the variance of the difference in log earnings residuals between 1970 and 1972 and the last value, in 2018, represents the variance of the difference in log earnings residuals between 2016 and 2018. We interpolate the volatility estimates between pairs of years after 1996 to be able to draw a continuous line.

1980s and during the Great Recession, particularly in the lower portion of the distribution. Cross-sectional inequality increased on average, especially in those same time periods. Appendix Table 2, supplementary materials show the mean and dispersion of earnings changes by year, and suggests an increasing variance because the dollar gap between the percentiles of differences is gradually increasing over time.

All survey data have measurement error in earnings reports. Appendix A, supplementary materials reports what is known about measurement error in the PSID and, most importantly, whether it might affect estimated trends in volatility. With one possible exception noted below, the evidence does not suggest that volatility trends should be expected to be affected.

3. Main Results

We begin with our baseline model. Gross volatility is measured by the variance of residuals in 2-year differences in log earnings. Results using log earnings itself show no difference in the two methods. Figure 1 shows that the trend in gross volatility follows the same three-phase pattern found in much of the PSID literature (Moffitt and Zhang 2018), rising from the 1970s to the mid-1980s, exhibiting a stable trend around significant fluctuations from the mid-1980s to the mid-2000s, and rising thereafter. The upward spike in the mid-1990s and the downward spike in the late-1990s and early 2000s are of unknown origin. We speculate that it may have been the result of a one-time change in interviewing and data handling procedures. See Appendix A, supplementary materials for a discussion. The rise in volatility starting in 2008 is no doubt related to the Great Recession and, in fact, the extra years that are now available from the PSID show that volatility falls afterwards, including using the most recent data point 2016–2018. By 2016–2018, it had fallen to its value just prior to the Recession.

The dotted line in Figure 1 is a fit of the data to a fifth order polynomial, which shows the three phases of the change more visibly. When we compare the earnings volatility with unemployment rate (also shown in the figure), we find that volatility exhibits strong countercyclicality, although, on average, volatility only falls in a recovery period after a lag. But despite the countercyclicality of volatility, it did not return to its pre-downturn value in the late 1980s and it also rose in the mid-2000s when the unemployment rate was dropping prior to the Great Recession.

An important feature of Figure 1 is that there is a large difference in the net trend prior to the early 1990s and afterwards. It is easiest to see this visually by normalizing volatility trends in the first half of the period to the average of 1972–1976 values, and normalizing them to the average of 1992–2002 values in the second half. Appendix Figure 1(a), supplementary materials show that volatility doubled from the early 1970s to 1990 and Appendix Figure 1(b), supplementary materials show that volatility ended up in 2018 only about 20% higher than it had been in the early 1990s—only one-fifth as large a percent increase as in the first period. Thus, the evidence that the upward trend in volatility was much greater from the 1970s to the 1980s than it has been anytime since then is quite strong.

We also measure volatility using the arc percent change (APC) as reported in the Overview paper—the APC is just the change in earnings divided by the mean of the two values–and compare it to the log earnings difference measure in Figure 1. As shown in Appendix Figure 2, supplementary materials, the APC volatility measure is lower than that shown by log earnings differences and has a somewhat flatter trend, but the same pattern of faster growth in the first two decades compared to the second two still appears. The difference is a result of nonlinearities in the log transformation. We also find the APC trend using residuals to be very similar to the APC trend without residuals. The APC also allows us to include men who are nonworkers in one of the two years. When zero earnings are included, as shown in the upper line in that Figure, the level of volatility is much higher in magnitude and cyclicality is greater, which is not surprising since a movement from work to nonwork or vice-versa is likely to generate a larger change in earnings than within-work changes and since employment is likely to be more cyclical than conditional earnings. As for trends, the trend including zeroes is somewhat steeper than that without zeroes but again slows down after the 1990s. In addition, volatility falls much faster after the Great Recession than before and reaches an ending level in 2018 that is about the same as in 1990, eliminating any net trend over that period (see Appendix Figures 3(a) and (b), supplementary materials for renormed trends).

We further explore the source of the increases in volatility within the distribution of earnings changes in two ways. First, we examine the percentile points of the distribution of the 2-year difference in log earnings residuals following Shin and Solon (2011), which shows whether the average trend is stronger in some parts of the distribution than in others. As shown in Appendix Figure 4, supplementary materials, volatility widens out at all percentile points but with the largest widening occurring at the top and bottom of the change distribution. Second, we examine whether volatility trended differently in different quartiles of the cross-sectional distribution. Appendix Figures 5(a)–(d), supplementary materials show the level of volatility to be higher in the bottom quartile than in the rest of the distribution, but Appendix Figure 6, supplementary materials show that volatility trends for all four quartiles are approximately the same, but with more instability in volatility in the bottom quartile.

4. Issues with the PSID and with Comparisons to Other Datasets

4.1. Trimming

We conduct sensitivity tests to our percentile point trimming but also examine the impact of real-dollar trimming used in past work. We find that almost all of the trimming methods we use in our basic results have no effect on our conclusions regarding volatility trends, as long as we trim symmetrically using percentile points (see below for results when real dollar values are used instead). Appendix Figure 7, supplementary materials show that whether we conduct our trim at the 1%/99% on log earnings by year (as we do in the basic results), not do any trimming at all, or trim at the 5%/95% level instead, has no effect on the main trend patterns; those alternatives just introduce more or less noise in the year-by-year movements. This necessarily implies that earnings volatility is not trending differently in the tails or at least not enough to show up in our average results.

We provide an extended discussion of using the dollar-denominated trimming at the bottom of the distribution, given the demonstration in the Overview paper of the importance in the left tail of earnings in explaining differences in estimated volatility across different data series. Several studies in the literature using administrative data on Social Security earnings—all of which found, to varying degrees, declines in earnings volatility—used some version of dollar-denominated trims to trim the bottom of their earnings distributions. Using dollar-denominated trims is hazardous if earnings inequality is increasing, as it has been for several decades in the United States and as it has in our PSID dataset (Appendix Table 1, supplementary materials). With earnings inequality growing, a constant dollar trim will systematically exclude an increasing fraction of the lower tail of the earnings distribution (unlike a percentile point trim). If volatility levels are higher in the lower tail (as they are in the PSID), then deleting an increasing fraction of that tail will bias the trend in average volatility in a downward direction.

We follow Carr and Wiemers (2021) to test the effect of three different dollar-denominated trimming methods employed by several studies using Social Security earnings data. One excludes observations with real annual earnings below one quarter of full-time full-year work at the 2011 federal minimum wage, a method used by Kopczuk, Saez, and Song (2010) (this is $3685 in our 2010 dollars). A second excludes observations with real annual earnings below a quarter of a year of full-time full-year work at half the federal minimum wage, but using the actual minimum in each year, a method used by Guvenen, Ozkan, and Song (2014) and Bloom et al. (2017). A third excludes observations below the annual earnings need to qualify for the Social Security threshold for credit, a method used by Sabelhaus and Song (2009, 2010).

The results are shown in Figure 2, and include our baseline results for comparison. While the second method yields approximately the same trend as that using our percentile point trim (ending up at about the same place, although differing in some periods), the first and third methods yield considerably flatter profiles of volatility growth over the whole period from 1972 to 2018. Those methods yield higher volatility in the 1970s and lower volatility in the 2000s and even decline at the end, reaching levels approximately those in 2000. Those two methods increasingly trim out greater fractions of the left tail over time, biasing their volatility trends downward. This could explain part of the reason for the difference in volatility trends in administrative data studies that use those methods and in the PSID studies that use percentile point trims.

4.2. Attrition

Appendix A, supplementary materials discusses what is known about attrition and attrition bias in the PSID. While cumulative attrition reached 50% by the 1990s, studies have shown that the cross-sectional distribution of income has been little affected. However, the one study examining its effect on earnings instability (Fitzgerald, Gottschalk, and Moffitt 1998) suggested that higher volatility individuals are more likely to attrite, which would bias trends in volatility downward, not upward. Nevertheless, we conduct an examination of possible attrition bias using traditional inverse probability weighting, first estimating a model of attrition on observables and then using the predicted probabilities to reweight the data on the nonattriters (of course, this does not address attrition on unobservables). Our attrition model relates current attrition probabilities to three lagged observables for each individual: their earnings in the previous period, their mean earnings over the past six years, and the standard deviation of their earnings over the past six years. Appendix Figures 8 and 9, supplementary materials show, respectively, year-by-year attrition rates and unit nonresponse rates in our sample, and the reweighted volatility trends. Appendix Figure 9, supplementary materials show that attrition-adjusted volatility is higher in level than the unad-

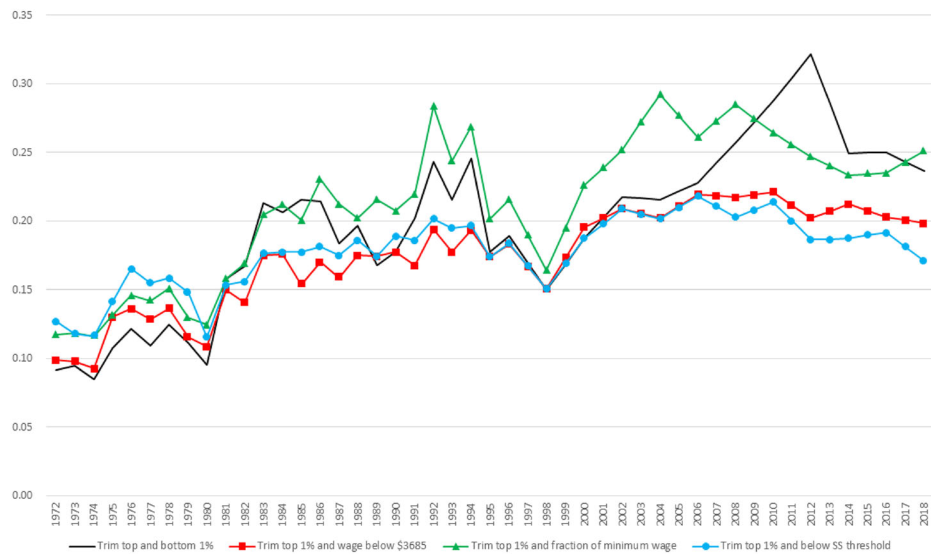

Figure 2. Variance in 2-year difference in log earnings residuals, different trimming at the bottom.

justed series, which is expected since those with high (lagged) volatility are disproportionately missing from the responding sample. The weighted trend is, as prior results predict, more positive, not less positive, than in Figure 1, at least in the second half of the period. In addition, much of the trend is a result of some very low predicted weights; when these observations are deleted, volatility levels and trends are close to those for unadjusted volatility, implying that attrition bias mainly arises from a small number of observations with high volatility. In any case, this analysis does not support attrition as a possible reason for the positive long-run trend in PSID volatility among male heads of household and, if anything, implies a stronger upward trend.

4.3. Imputation

Appendix A, supplementary materials describes what is known about earnings nonresponse and imputation in the PSID, with the available evidence suggesting quite low rates of imputation. Appendix Figure 10(a), supplementary materials show the rates of imputation of earnings because of missing values of earnings from 1970 to 2018 in our sample of male working heads, indicating that the percent of wage and salary income observations that are imputed ranges from a low of 0.30 to a high of 4.7, with the high value occurring in 1992, a period when the PSID changed its methodology and interviewing method (see Appendix A, supplementary materials). But while the low value of 0.30 is unlikely to change the results much, the higher value of 4.7 could if imputation is strongly correlated with volatility. But Appendix Figure 10(b), supplementary materials show that estimated volatility trends with and without imputed observations included are very close to one another. This is unlikely to occur unless nonresponse is mostly ignorable. We conclude from this simple exercise that item nonresponse and imputation for earnings in the PSID are unlikely to be a reason for the greater upward trend in volatility compared to that in other datasets.

4.4. Immigrants

One significant difference between the PSID and other datasets concerns the representation of immigrants. The core PSID sample was representative of the 1968 U.S. population and has been followed since that time, but necessarily does not include those immigrating to the United States since 1968, who now constitute about 10% of the U.S. population. The PSID has attempted three times to enroll immigrants into the sample to represent this population. In 1990, about two thousand Latino households were added to the PSID which, though not representing all post-1968 migrants, represented an important migrant group of interest. But because of a lack of sufficient funding, the households were dropped after 1995. In 1997, a sample of 441 immigrant families was added to the PSID and another 70 immigrant families were added in 1999, for a total of 511 families. And the 2017 New Immigrant Refresher Sample adds approximately 500 new immigrant families to the PSID. While the sample size is small, they have continued to be followed and their sample sizes have grown through childbearing and splitoffs.

We briefly analyze the volatility trends with and without this additional immigrant sample to determine whether there is any suggestion that the exclusion of immigrants might be contributing to upward PSID trends compared to those of other datasets which include immigrants. Since the immigrant sample was only begun in income years 1996 or 1998, we start our analysis in 2000. Appendix Figure 11, supplementary materials show the results. If anything, the volatility of the immigrant sample increases faster than that of the baseline sample, not slower. The volatility trend for the immigrant sample alone bounces around more, probably because of its smaller sample size. The gross volatility of the combined sample seems to increase slightly faster than the main sample after the late 2000s. Thus, while minimal, this evidence does not indicate a markedly slower growth of volatility for immigrants.

We should also note that the CPS has a question in immigrant status which makes it capable of examining volatility trend differences between immigrants and natives. That examination finds no difference in volatility trends for the two groups (see the CPS paper in this volume, Appendix Figure S.8, supplementary materials, as well the Overview paper).

5. Summary and Conclusions

In light of conflicting evidence from different datasets and research papers on how male earnings volatility has evolved in the United States, this article has provided a new study of male earnings volatility from the workhorse dataset in the literature, the Panel Study on Income Dynamics (PSID). We have four findings. First, contrary to reports that the PSID shows increasing volatility over the past five decades, we find that increases in volatility are primarily concentrated in the period from the 1970s to the 1980s. After approximately 1990, the rate of growth of volatility has been very slow. Second, we find that the pattern of volatility trends is similar across all levels of the cross-sectional earnings distribution, but that the volatility trends have been most pronounced in the tails of the distribution of earnings changes. Third, we find that neither attrition, imputation, nor other threats to the representativeness of the PSID are likely to be responsible for the patterns of volatility trends in the data. Fourth, we find that differences in the volatility trends in the PSID and in administrative datasets may be a result of differences in trimming of the left tail of the earnings distribution.

Going forward, more work on volatility among subgroups and decompositions would be warranted. More work on decompositions of volatility into its three components–permanent, persistent, and transitory—is warranted, as well as work on the relation of job and occupational mobility to earnings volatility. Datasets which match firms to workers would be valuable to ascertain the role of firms in worker earnings volatility, given much recent work on the importance of firms to the understanding of trends in labor market earnings (e.g., Song et al. 2019).

Supplementary Materials

The supplementary appendix to this article provides discussion on response error, attrition, and nonresponse and imputation, and additional results discussed in the article.

Acknowledgments

The authors would like to thank Joseph Altonji for comments on a version presented at the 2019 AEA meetings as well as the participants of a July 2019 conference sponsored by Equitable Growth in Cambridge, Massachusetts for their comments. Comments from James Ziliak and Michael Carr are appreciated as well as input from the other team members on this joint project: John Abowd, Christopher Bollinger, Charles Hokayem, Kevin Mckinney, and Emily Wiemers. Assistance from David Johnson was valuable and the comments from an Associate Editor and two referees were also helpful. The authors would like to dedicate this article to the memory of our long-time collaborator, mentor, and friend, Peter Gottschalk, who passed away on March 25, 2021.

Disclosure Statement

The authors have no competing interests to declare.

Funding

Zhang would like to thank the financial support from the National Natural Science Foundation of China (NSF72174074) and the National Social Science Fund of China (21FJYB022).

References

Bloom, N., Guvenen, F., Pistaferri, L., Sabelhaus, J., Selgado, S., and Song, J. (2017). "The Great Micro Moderation." Working Paper. [4]

Carr, M. D., and Wiemers, E. E. (2021), "The Role of Low Earnings in Differing Trends in Male Earnings Volatility," *Economics Letters*, 199, 109702. [4]

Dynan, K., Elmendorf, D., and Sichel, D. (2012), "The Evolution of Household Income Volatility," *The B.E. Journal of Economic Analysis & Policy,* 12, 1–42. [2]

Fitzgerald, J., Gottschalk, P., and Moffitt, R. A. (1998), "An Analysis of Sample Attrition in Panel Data: The Michigan Panel Study of Income Dynamics," *Journal of Human Resources,* 33, 251–299. [4]

Gottschalk, P., and Moffitt, R. (1994), "The Growth of Earnings Instability in the U. S. Labor Market," *Brookings Papers on Economic Activity,* 2, 217–254. [1,2]

Guvenen, F., Ozkan, S., and Song, J. (2014), "The Nature of Countercyclical Income Risk," *Journal of Political Economy,* 122, 621–660. [4]

Kopczuk, W., Saez, E., and Song, J. (2010), "Earnings Inequality and Mobility in the United States: Evidence from Social Security Data since 1937," *Quarterly Journal of Economics,* 125, 91–128.

Moffitt, R. A., and Zhang, S. (2018), "Income Volatility and the PSID: Past Research and New Results," *American Economic Association Papers and Proceedings*, 108, 277–280. [1,2,3]

Sabelhaus, J., and Song, J. (2009), "Earnings Volatility Across Groups and Time," *National Tax Journal,* 62, 347–364. [4]

Sabelhaus, J., and Song, J. (2010), "The Great Moderation in Micro Labor Earnings," *Journal of Monetary Economics,* 57, 391–403. [4]

Shin, D., and Solon, G. (2011), "Trends in Men's Earnings Volatility: What Does the Panel Study of Income Dynamics Show?" *Journal of Public Economics,* 95, 973–982. [2,4]

Song, J., Price, D. J., Guvenen, F., Bloom, N., and von Wachter, T. (2019), "Firming Up Inequality*," "*The Quarterly Journal of Economics,* 134, 1–50. [6]