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

The Nature of Household Labor Income Risk

Seth Pruitt and Nicholas Turner

2018-034

Please cite this paper as:
Pruitt, Seth, and Nicholas Turner (2018). "The Nature of Household Labor Income Risk," Finance and Economics Discussion Series 2018-034. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2018.034.

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 Nature of Household Labor Income Risk*

Seth Pruitt
Arizona State University

Nicholas Turner
Federal Reserve Board

May 9, 2018

Abstract

What is the nature of labor income risk facing households? We answer this question using detailed administrative data on household earnings from the U.S. Internal Revenue Service. By analyzing total household labor earnings as well as each member's earnings, we offer several new findings. One, households face substantially less risk than males in isolation. Second, households face roughly half the countercyclical increase in risk that males face. Third, spousal labor income ameliorates household earnings risk through both extensive and intensive margins.

*Pruitt: W.P. Carey School of Business, Arizona State University, seth.pruitt@asu.edu. Turner: Federal Reserve Board, nick.turner@frb.gov. **The data for the work were accessed when Nick Turner was employed at the U.S. Treasury Department. The views and opinions are solely those of the authors and do not necessarily reflect those of the Treasury Department or the Federal Reserve Board.**

1 Introduction

What is the nature of labor income risk facing households? How does it change over the business cycle? We answer these questions using administrative data on earnings from the Internal Revenue Service (IRS).[1] Our sample is a 1-in-5 random draw from the population of U.S. households with prime working age adults between 2000 and 2014, years that span both expansions and recessions including the Great Recession. This period exhibits historically high levels of female labor force participation, allowing us to analyze how spousal earnings decisions alter the earnings risk facing households relative to the risk inferred from studying male earners in isolation.[2]

Analyzing a large sample of administrative data on earnings for both adult members in a household offers important advantages relative to earlier work. In a seminal contribution, Guvenen, Ozkan and Song (2014, henceforth GOS) analyze a large sample of males using Social Security Administration data. This work focuses on male earners in isolation because household composition is not included in their data. Busch et al. (2015), Arellano (2017), and DeNardi, Fella and Pardo (2018) address this limitation by analyzing household earnings using the Panel Study of Income Dynamics (PSID).[3] The PSID includes information on labor earnings at both the individual and household levels over a relatively long period of time. However, the PSID contains only several thousand households each year so that it is not possible to implement the robust nonparametric approach used by GOS. This research design, which we implement using a sample of over 10 million households each year, allows for a flexible and accurate characterization of higher moments of the earnings growth distribution while imposing relatively little structure on the data. Additional drawbacks to the PSID include the fact that income is measured from survey responses and is likely to contain measurement error that may be substantial and may not be classical in nature (Bound,

[1] We refer to tax units as households, which differs from the Census definition that is based on location of residence. It is difficult to cleanly identify cohabiting individuals who do not file jointly in the tax data. As a result, our definition will not include these households, some of which are likely to act as a single economic unit to mitigate risk. However, unlike the Census definition, our approach will not include living arrangements where risk sharing would be unlikely such as non-related roommates.

[2] According to the Bureau of Labor Statistics, female participation peaked just over 60% around 2000 and was mostly flat through 2014, up from 51% in 1980. Following previous literature, we focus on the earnings risk of households and male workers, and then characterize how spousal earnings differentiate the two. Admittedly, this is a male-centric focus. However, male labor force participation has hovered around 12-15 percentage points higher than female labor force participation during our sample. Moreover, male earnings account for around two-thirds of household earnings on average. These simple averages, in addition to our aim to speak directly to previous literature, support the context in which we analyze female, male, and household earnings.

[3] Moffitt and Zhang (2018) provide a detailed discussion of additional prior work that uses the PSID to examine income volatility.

Brown and Mathiowetz 2001; Kreiner, Lassen and Leth-Petersen 2013), that earnings are top-coded, and that missing data may not reflect zero earnings. Our administrative data contains relatively less measurement error, are not top-coded, and missing data is far more likely to reflect a true zero value in earnings compared to survey data such as the PSID.[4]

Our results show that household earnings risk is countercyclical due to increasing negative-skewness during recessions, but not due to increasing dispersion. Intuitively, our findings suggest that large upward movements in the earnings distribution are less likely, and large downward movements more likely, during recessions. While these findings are consistent with the analysis of male labor earnings in GOS, we present new findings that build on this work in important ways. Using either skewness or dispersion to quantify earnings risk, we show that households face substantially less risk than the risk inferred from studying males in isolation.[5] Additionally, we show that households experience a smaller countercyclical increase in earnings risk than males. The presence of a second earner is the key reason why households face this different risk profile. Building on the intuition of GOS, we implement a new method of analyzing spousal earnings growth using nonparametric regressions. These results show that the second earner provides insurance against male earnings risk in two key ways. The first form of insurance comes from working spouses, who insure that male earnings changes are not equal to household earnings changes. As a result, male earnings losses are not amplified at the household level. The second key form of insurance comes from non-working spouses, who can enter the labor market to offset male earnings losses. We find evidence of this channel for almost all household income levels. Among median income households ($76,000), the probability that a non-working spouse enters the labor market is roughly 20% higher when the male experiences a 20% earnings loss instead of a 20% earnings gain.[6] Our findings that spousal labor supply provides meaningful insurance against permanent male earnings shocks is consistent with Blundell, Pistaferri, and Saporta-Eksten (2016) who find similar results between consumption and household earnings.

Overall, our findings emphasize the importance of analyzing earnings for both working-age adults in order to characterize the labor earnings risk facing households. Observing only one worker in isolation overstates the level of household earnings risk and overstates

[4]While most previous papers have been unable to use administrative data at the household level, an exception is the recent work by Panousi and Ramnath (2015). Using a panel of household data over the years 1987–2009 representing a 1-in-5000 sample of households, they find that male earnings growth has strong countercyclical volatility while total household income has weak countercyclical volatility and a-cyclical negative-skewness.

[5]These results contrast with Busch et al. (2015), who conclude that U.S. households face greater countercyclical skewness risk.

[6]In contrast, the probability of spousal exit increases for large male earnings gains only for relatively high-income households.

the countercyclical increase in household earnings risk. Using a simple model to quantify risk in certainty equivalent terms, our results suggest that households face half as much risk as males during expansions and experience half as much countercyclical risk over the business cycle. Given that many key economic decisions including consumption and labor supply are made at the household level (Mincer 1962, Becker 1965 and 1974, Kydland and Prescott 1982, Sargent 1987, Benhabib, Rogerson, and Wright 1991, Smets and Wouters 2003, Christiano, Eichenbaum, and Evans 2005, Ramey and Francis 2009, Chen, Michaux, and Roussanov 2015, Busch, Domeji, Guevenen and Madera 2015) and that household labor income risk is a key input into models of economic behavior including asset pricing (Mankiw 1986, Storesletten, Telmer and Yaron 2015, Constantinides and Ghosh 2017) and the costs of business cycles (Lucas 2003, Krebs 2007), the ability to directly characterize earnings risk for *households* is important contribution of our paper.

The paper proceeds as follows. In Section 2 we describe our data sample and variable definitions. We document statistical characteristics of household earnings growth and male earnings growth in Section 3. In Section 4, we focus on female earnings and their relationship to male earnings. We then conclude. The appendix contains additional discussion and extensive additional figures further documenting our main results.

2 Data

In this section, we describe the data and the sampling choices we implement.

2.1 Data on Earnings and Marriage

To construct our sample we draw on population-level administrative tax data on earnings and marriage combined with information from the Social Security Administration. The tax data we use is from the W2, that reports individual wage earnings, and the 1040, which is the individual income tax return. We also use date of birth and date of death from the Social Security Administration. To construct our sample, we first calculate the age of each male in the population of Social Security Administration data during the years 1999-2014 and retain household observations when the individual is between the ages of 25 and 60.[7] From these records, we draw a one in five random sample of households based on the last

[7]As fore mentioned, our data construction is admittedly male-earner-centric. Our reason for this is to better compare to existing work like GOS and in the labor literature broadly that has focused on labor market outcomes for prime-aged males.

digit of the Social Security number of the prime-aged male worker. For each individual in the households in our sample, we pull all W2s and all 1040s for the years 1999-2014. The use of W2s means that we are able to calculate earnings for nearly all individuals, including those who do not file taxes. Compared to survey data, earnings reported on the W2 form are less likely to suffer from measurement error, are not top-coded, and missing data is far more likely to represent a true $0 earnings. Using information on marriage from the 1040 tax form, we identify the spouse when present and merge in spousal W2 data. We use the Consumer Price Index to put all dollar amounts in $2014.

To focus on males with a strong labor-force attachment, we impose several restrictions. First, in cases where the male dies in the sample period we drop observations from the last two years of life. Chetty et al. (2016) show that on average there does not appear to be health-induced negative earnings shocks two years prior to death or earlier. Second, we remove any individuals with nonzero self-employment income. Such individuals may have weaker attachment to the labor market, and because self-employment income is not third-party verified so observed changes in income may not represent changes in real economic behavior but rather reporting behavior (Chetty, Friedman and Saez, 2013).[8] Third, we remove any individuals who receive Social Security disability payments. These individuals are likely to be disabled and therefore do not have a strong labor market attachment. Lastly, we require the male to file jointly in the years we measure income changes in order to abstract from earnings changes induced by divorce (we discuss the implications of this restriction in the appendix).[9] Overall, these restrictions allow us to characterize the behavior of most households in the U.S. For example, in 2010 our sample accounts for over 90% of all joint filing tax returns.[10]

2.2 Variable Definitions

Long-run earnings We construct a long-run earnings measure denoted $\bar{Y}_t(H)$ using household earnings. To construct this long-run measure, we track the male over the five years prior

[8] In a typical year, roughly 10% of observations that would otherwise be in the sample are dropped as a result of this restriction. Of these, about 2% report losses.

[9] There is a well-known decline in marriage rates over time in the U.S. (Isen and Stevenson 2010), which we observe in our sample period. Marriage rates among observations that meet all other sample selection criteria decline from 56% in 2000 to roughly 50% in 2013.

[10] Our sample includes 9.8 million observations (joint filing tax units) in 2010, representing about 49 million households given that our sample is a 1-in-5 random draw. By comparison, in 2010 the IRS reports that 54 million joint tax returns were filed.

to t, requiring that individuals be in this five-year period for at least one year.[11] While we condition on households that remain married when measuring income changes, we do not impose this restriction when measuring baseline earnings. For observations when the male is married, $\bar{Y}_t(H)$ is the sum of female and male wages and for non-married observations $\bar{Y}_t(H)$ is equal to male earnings. To account for differences in marital status when constructing the long-run measure we adjust for family size and household returns to scale by dividing earnings in married households by the square root of two, a common approach used in related work.[12]

Figure 1 shows the cumulative distribution function of long-run earnings $\bar{Y}(H)$ for each percentile $p_q(\bar{Y}(H))$, averaged across years 2000-2014 (we drop the t subscript when results are averages across multiple years). Similar to prior work that finds a relatively large concentration of annual earnings (Pikety and Saez 2003) and wealth (Saez and Zuchman 2016, Bricker et. al (2016)) in the upper-end of the distribution, we find considerable inequality in long-run earnings. As shown in the figure, the bottom five percent (the bottom ventile) of households account for roughly one-third of one percent of all long-run earnings. In contrast, the top five percent (the top ventile) accounts for nearly 20 percent of long-run earnings.[13] Values of household income increase dramatically moving up the $\bar{Y}(H)$ distribution: 5th percentile \$13,000, 10th percentile \$24,000, 25th percentile \$47,000, 50th percentile \$76,000, 75th percentile \$111,000, 90th percentile \$160,000, 95th percentile \$200,000. To be in the top 1%, households need more than \$770,000 of earnings.

Persistent earnings changes We define persistent earnings changes as the four-year change in log earnings.[14] To keep growth rates finite but enable us to consider extensive

[11]Excluding years at the beginning of the sample period when there are not five years of data available, virtually all (in excess of 98%) of individuals are in the data for the full five years. This means that for the early years in our sample, fewer than five years are used to construct the long-run income measure. For example, to define $\bar{Y}_t$ in the year 2000, we use only household earnings from the year 1999, for the year 2001 we use earnings from 1999 and 2000 etc. If we look only at the recession (2008) and expansion (2010) for which our $\bar{Y}_t(H)$ calculations have all five past years available, our main skewness and dispersion results are unchanged.

[12]Auten, Gee and Turner (2015) use this approach when measuring income mobility. This method is also commonly used by government agencies such as the U.S. Treasury, the Congressional Budget Office and the Department of Housing and Urban Development. This adjustment does not appear to have a large effect on our results.

[13]Note that this level of income concentration, while relatively high, is lower compared to that found by Pickety and Saez (2003) who use a broader measure of income compared to our measure based on household labor earnings. In forthcoming work, we analyze the income risk of total household income, including non-labor income.

[14]As shown in GOS, if log earnings (net of lifetime earnings effects) is modeled as a random walk and a purely transitory term then the within group variance reflects primarily permanent shocks as the time difference increases.

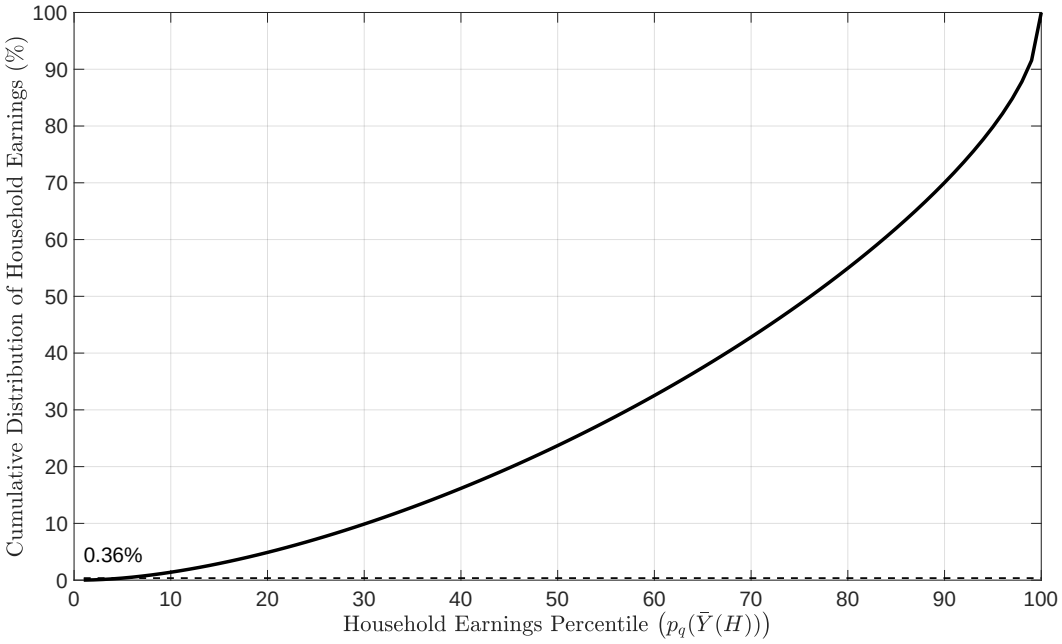

Figure 1: Aggregate Household Earnings Cumulative Distribution

Notes: The solid line shows the percentage of aggregate household earnings cumulatively represented by households in and below each household earnings percentile, on average over our whole sample. It uses the averaged past household earnings in year t, $\bar{Y}_t$. The 0.36% noted is the cumulative amount of aggregate household earnings represented by the bottom ventile of the household earnings distribution.

labor margins, we first recode \$0 earnings to \$1 earnings. This four-year change is the same as Busch et. al (2015) use, though it differs slightly from GOS who use a five-year change.[15] We use the four-year difference so that we can measure persistent earnings growth in the wake of the Great Recession. In our main results, we track the male earner over time and condition on marital status.[16] Similarly, Busch et. al (2015) use the sample of households whose marital status is unchanging. Persistent earnings is defined as $\Delta_4 y_{t+4} = y_{t+4} - y_t$, for log earnings y. We focus our analysis on these persistent earnings changes to allow for a reasonable adjustment time for the spousal responses we report in Section 4.[17]

[15]We found very small differences between statistics calculated for five-year changes and four-year changes, in years when both were available.

[16]However, our results are qualitatively similar if we do not condition on households that remain married. We report a subset of our main results using this restriction in the appendix.

[17]GOS also analyze transitory changes in earnings, defined as the one-year change. Results for transitory income changes $\Delta_1 y_{t+1}$ appear in the appendix.

2.3 Business Cycle Dates

We rely on data from the NBER Business Cycle Dating Committee to define recession and expansion years.[18] Therefore we define the following: recessionary years are 2001, 2002, 2008, 2009, and 2010; expansionary years are 2003, 2004, 2005, 2006, 2007, 2011, 2012, 2013 and 2014. For persistent earnings changes $\Delta_4 y_{t+4}$ we use the following years t to characterize earnings shocks over the business cycle.

- Recession: 2000, 2007
- Expansion: 2002, 2003, 2010

As a result, we focus attention on the roughly 50 million household year observations from these five years. Our reason for using $\Delta_4 y_{t+4}$ (instead of GOS's $\Delta_5 y_{t+5}$) is to be able to capture persistent earnings changes during the expansion following the Great Recession, beginning in year 2010.[19] Since we deal exclusively with persistent earnings growth in the main paper, we sometimes drop subscripts in the text for the sake of exposition.

It is not surprising that unemployment increases during recessions. However, as shown in Figure 2, there are differences in male and female unemployment rates over the business cycle. Figure 2 plots publicly-available monthly unemployment rates separately for males and for females between 2000 and 2014, with periods of recession shown in the shaded areas.[20] For much of our sample period, male and female unemployment rates are very similar. But, there are also times when the unemployment rates differ, such as during the Great Recession and its aftermath. This divergence is our *prima facie* evidence that the presence of a second earner within the household may serve as a form of insurance against male job loss. Intuitively, when the unemployment rates by gender differ, households with a second earner may have effectively diversified themselves against the earnings risk from job losses. Does this happen in our earnings data? And if so, how? These are the points of our analysis in Sections 3 and 4.

[18] The NBER defines months March 2001 through November 2001 and December 2007 through June 2009 as recession periods. We follow GOS in categorizing 2007 as expansion because only its last month was in recession, and in adding 2002 and 2010 as recession years because unemployment stayed high after the NBER business cycle trough.

[19] Recession/expansion year definitions for transitory $\Delta_1 y_{t+1}$ income changes appear in the Appendix.

[20] Data for seasonally adjusted unemployment come from the Bureau of Labor Statistics and include men and women ages 20 and older. Data for recessions comes from the NBER.

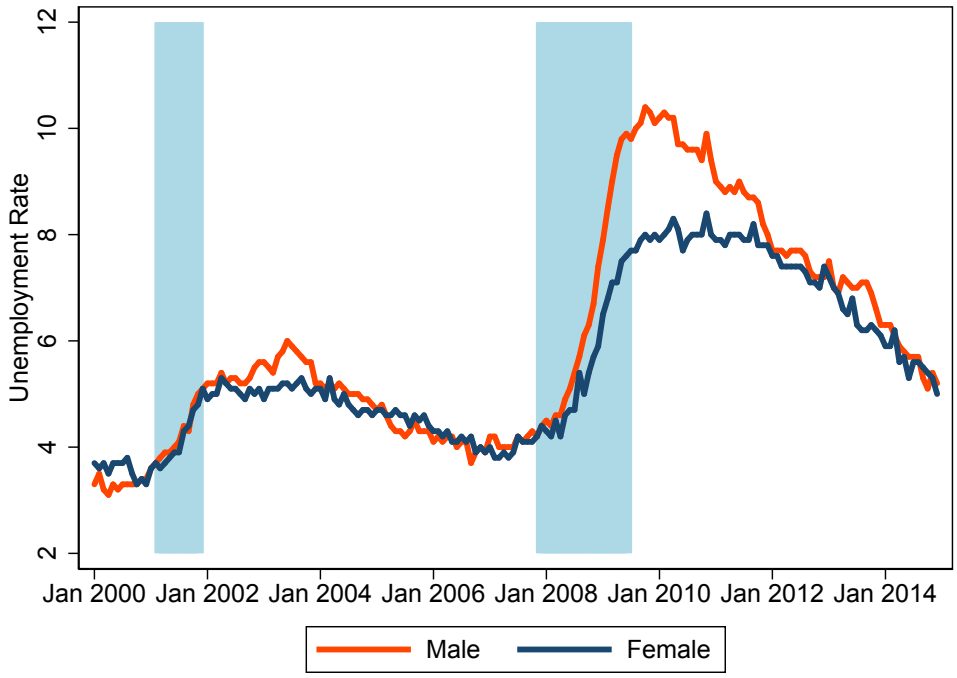

Figure 2: Unemployment Rates by Gender

Notes: National unemployment rates, for males (blue line) and females (orange line), ages 20 years and older. NBER recessions are shaded.

2.4 Isolating Idiosyncratic Earnings Shocks

To focus as closely as possible on idiosyncratic earnings risk used in earlier work (Constantinides Duffie, 1996; Storesletten Telmer Yaron, 2004; GOS, Busch et. al, 2015) we net out year effects and age effects in a flexible manner. We define the age of the household based on the male, and for each age 25-60 and each year 2000-2014 we calculate the percentiles of the long-run earnings measure.[21] This results in income groups that have very similar long-run earnings within a given percentile and that have identical year and age distributions across income percentiles. To be concrete, we denote the long-run earnings percentile q for year t and age a as $p_{q,t,a}(\bar{Y}(H))$. To construct the group of households at long-run earnings percentile q for year t, we put all ages a into the bin $p_{q,t}(\bar{Y}(H))$ – this yields bins containing about 105,000 households on average for each percentile q and year t. We next calculate for every household in this bin the persistent earnings changes for the male worker and the household as a whole. This approach removes both year effects and age effects while

[21] Our data run 1999–2014, but we calculate earnings changes only for the years 2000–2014 because the latter are the years where we can calculate a $\bar{Y}_t$.

imposing relatively little structure on the data.[22] We analyze the shock distribution following the method of GOS:[23] we calculate robust statistics (described below) that characterize the distribution of earnings growth $\Delta_4 y_{t+4}(M)$ and $\Delta_4 y_{t+4}(H)$ for each $p_{q,t}(\bar{Y}(H))$, and then average these statistics across recessionary t to derive recessionary statistics and across expansionary t to derive expansionary statistics.[24] Hence every point in the figures in Section 3 is a robust statistic from hundreds of thousands of observations, and since the associated confidence intervals around them are miniscule we abstract from plotting them.

In Section 4 we introduce a new method of analyzing spousal earnings that builds on the intuition of the approach we use for males and households.[25] The key insight of this approach is that we condition on both long-run earnings and *male earnings growth* when analyzing spousal earnings responses. To do this, we use kernel regressions, which allows us to mimic the approach used for male earnings while both smoothing the data (cell sizes are reduced to roughly 1,000) and to derive standard errors for the spousal response outcomes, which is now important because the statistical procedure is based on a much smaller number of (pre-aggregated) observations and now confidence intervals are not miniscule. Our approach allows us to determine if spousal responses vary significantly by long-run income, by male earnings growth and/or by business cycle phase.

[22]This approach of netting out life-cycle effects and year effects is similar to GOS who use a regression framework with dummy variables but there are three noteworthy differences. One, we define age groups by single years, whereas GOS uses 5-year age bins. Two, our approach can accommodate arbitrary shifts in the distribution of earnings by age. Results from Sabelhaus and Song (2010) based on an earlier period suggest that these distributions may differ by age. Three, our percentiles are guaranteed to have an identical age composition.

[23]Following prior work (Storesletten Telmer Yaron, 2004; GOS, Busch et. al, 2015), we refer to these idiosyncratic earnings changes as "shocks.' 'When analyzing spousal responses in Section 4, we refer to spousal earnings changes "responding" to male earnings shocks. We employ this imperfect terminology for the following two reasons. First, previous literature refers to the earnings changes of prime-aged males as shocks given the high rate of male labor force participation. We essentially extend this maintained assumption to the household level. Second, we find some evidence that various channels of spousal earnings changes (e.g. entrance into employment, and the earnings growth of spouses who work) are correlated with male earnings growth in intuitive ways. However, we recognize both that the earnings shocks of the male and the responses of the spouse may not be exogenous. In addition, we recognize that there are households where males may respond to spousal labor income shocks, rather than our assumption that spouses respond to male income shocks.

[24]The results are virtually identical if we instead grouped all the households in $\bar{Y}(H)$ percentile q for, say, recessionary years t_1 and t_2 and then calculated statistics from these roughly quarter-million household observations. The reason is: population changes only modestly over our sample period.

[25]When we aggregate the data this way, we calculate spousal earnings changes $\Delta_4 y_{t+4}(S)$ separately for the households in which the spouse enters employment, leaves employment, or remains employed. This allows us to distinguish the two extensive earnings margins (representing spouses who either enter or leave employment) from the intensive earnings margin (representing spouses who remain employed).

3 Household and Male Earnings Growth

In this section we report the earnings growth distributions for households and draw comparisons to the distributions of males in those households. We use robust measures of central tendency, skewness, and dispersion.

In our graphical analysis, we use $\bar{Y}(H)$ to denote the household earnings level, whose percentiles $p_q(\bar{Y}(H))$ define the "Household Earnings Percentile" on the horizontal axes. We use $\Delta y(H), \Delta y(M)$ to denote household earnings growth and male earnings growth, respectively, for which statistics are reported on the vertical axes.

3.1 Measures of Central Tendency, Skewness, and Dispersion

Negative-skewness is a natural component of economic risk because income losses have a larger impact than income gains of the same magnitude given concave utility. The negative-skewness of earnings growth means that large earnings losses are more likely than large earnings gains. In practice, we observe the mitigation of income losses with insurance and prior work characterizes nonzero probability on "bad" outcomes as a risk (Rietz 1988, Barro 2009, Gourio 2012, Hansen and Sargent 2008).

To document the skewness of earnings growth Δy, we follow GOS in using the measure of Kelley skewness:

$$\frac{(p_{90} - p_{50}) - (p_{50} - p_{10})}{p_{90} - p_{10}}.$$

This measure lies in $[-1, 1]$ and is unit-free, similar to correlation, so that we can meaningfully compare distributions with differently-scaled realizations (this is a difference between Kelley skewness and the third moment of the distribution). When we say "skewness" below, we mean Kelley skewness.[26]

Another component of risk is the dispersion of the distribution. We report the inter-decile range:

$$p_{90} - p_{10}.$$

This measure *is* scale-dependent, which is of course the point.

[26] GOS refer to this measure as Kelley skewness. We have found other sources referring to it as Kelly skewness.

Finally, to characterize the central tendency of earnings growth, we report the median:

$$p_{50}.$$

We also include the mean as a measure of central tendency, which provides a convenient comparison to the median, though one that is relatively more impacted by outliers.

We use these robust measures of central tendency, dispersion, and skewness for two main reasons. First, we want to avoid the effects of outliers, which can significantly affect first, second, and third moment estimates (see Kim and White 2004). Second, these robust measures show the possibility of changing skewness without changing dispersion, and vice versa, with a relationship to the central tendency.[27] If p_{50} shifts while p_{10}, p_{90} remain constant, skewness shifts but dispersion remains constant. If p_{50} remains constant while p_{10} and p_{90} shift away from p_{50} by particular proportional amounts, dispersion increases while skewness remains constant. Hence, robust dispersion and skewness are different dimensions of the nature of earning growth Δy risk.

3.2 Distribution over the Business Cycle

Figure 3 plots the 10th, 50th, and 90th percentiles of the $\Delta y(H)$ distribution. In the top panel we show these percentiles over the full income distribution and in the bottom panel we omit the bottom ventile of households based on the long-run income $\bar{Y}(H)$ distribution. In both cases, we separate the percentiles of the shock distribution by business cycle phase, with expansions shown in blue and recessions shown in red, a color convention we use throughout the paper. In the top panel, the shock distributions widely flare out for lowest and highest $\bar{Y}(H)$. While there appears to be a downward shift in the $\Delta y(H)$ distribution in recessions relative to expansions in the top panel, it is difficult to tell if this has to with a location shift, a greater negative-skewness during recessions, or both. This difficulty is due to the fact that the $\Delta y(H)$ distribution for households in the lowest $\bar{Y}(H)$ ventile have much more dispersion than other households, requiring an expanded range.

In order focus more clearly on the shock distributions all subsequent figures follow the approach in the bottom panel and omit observations for the bottom ventile of the long-run earnings $\bar{Y}(H)$ distribution. While this step omits 5% of households be definition, Figure 1 reports that it omits only 0.36% of aggregate long-run earnings. Moreover, households in

[27] Mechanics by which second and third moments can independently adjust are much more complicated to describe generally.

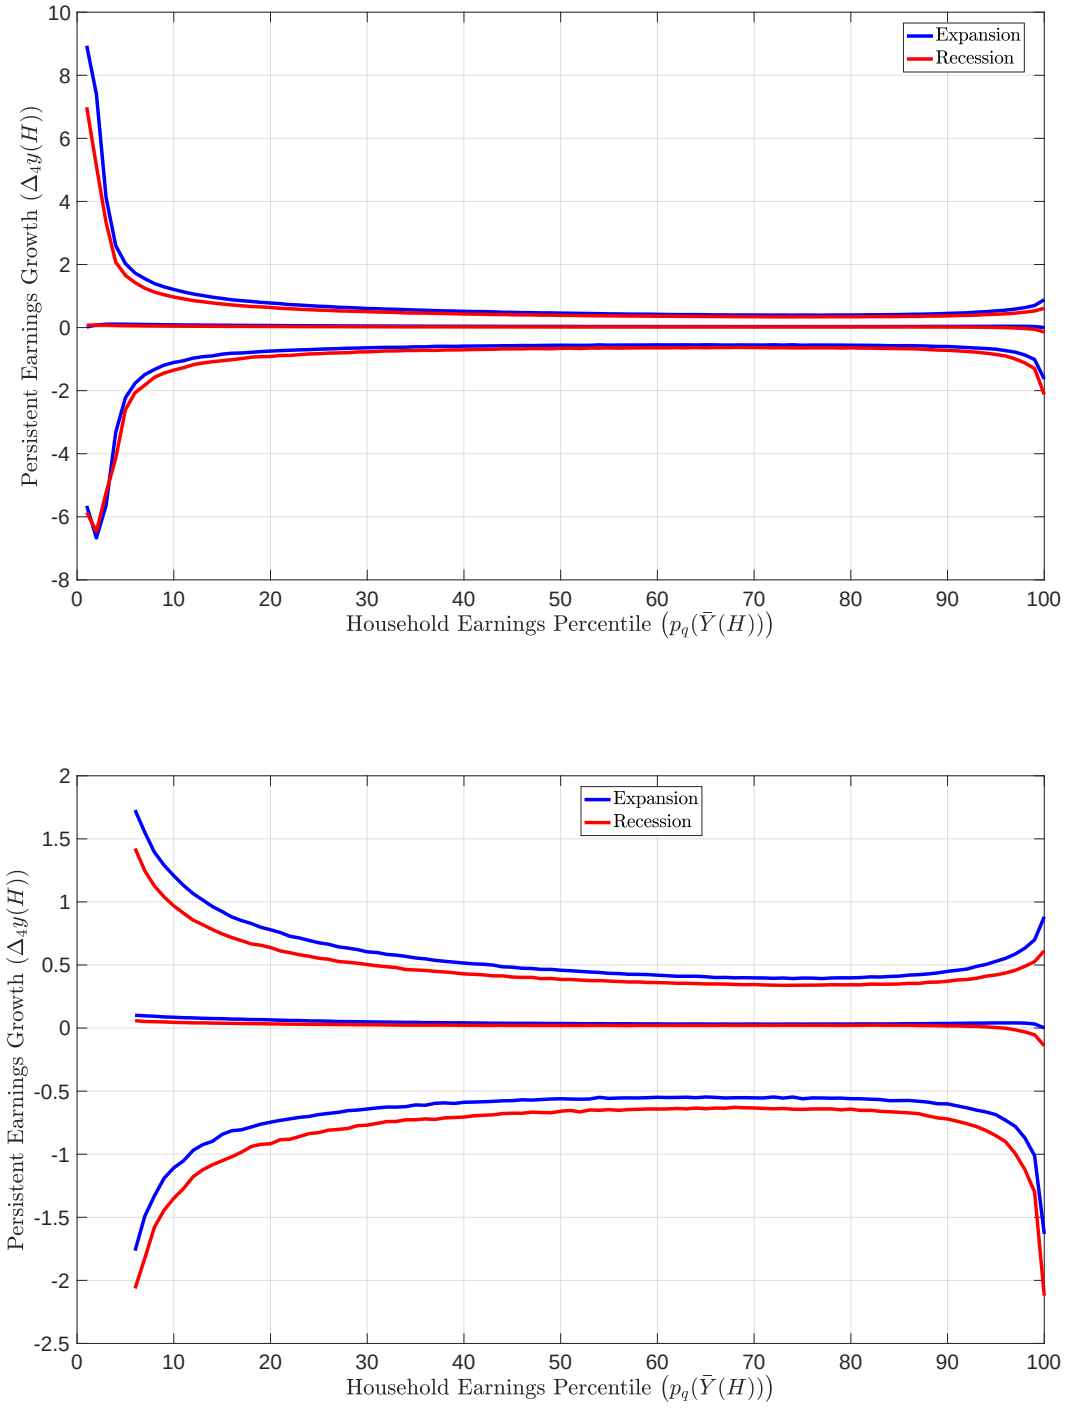

Figure 3: Percentiles of the Household Earnings Growth Distribution

Notes: The 10th, 50th, and 90th percentiles of the persistent household earnings growth $\Delta_4 y(H)$ distribution. For expansions (blue lines) and recessions (red lines) separately. The bottom panel omits observations for the bottom ventile of the $\bar{Y}(H)$ distribution.

the 5th percentile have long-run earnings of less than $13,000 (on average across years), less than what a single full-time minimum wage worker would earn annually. We assume these households have tenuous attachment to the labor force.[28]

Focusing on the bottom panel, the cyclical variation is clearer. For good outcomes measured as 90th percentile, shocks are smaller in recessions compared to expansions. The same holds true for bad outcomes as measured by the 10th percentile, as these losses are larger in recessions compared to expansions. The figure also suggests that in general, the distance from the median to the 90th percentile is less than the distance from the median to the 10th percentile, suggesting that shocks are generally negatively skewed and that negative-skewness increases during recessions. The shape of the earnings growth distributions also suggest meaningful differences in earnings risk over the income distribution. Both the 90th and 10th percentile shocks are relatively larger at the tails of the distribution and relatively smaller towards median long-run income, a pattern consistent with what GOS find in male earnings.

3.3 Central Tendency of Earnings Growth

We plot the median (solid) and mean (dashed) of household earnings shocks $\Delta y(H)$ over the business cycle in Figure 4. The magnitudes of median and mean shocks are noteworthy. The median shock suggests that over a four-year period half of all households will have weak earnings growth or worse, conditional on income and nearly independent of the business cycle. Consider a family at the median level of income ($76,000). Over a four-year period spanning recessions half of all of these families will have earnings growth of 2 log points or lower. Spanning expansions, the outlook is not much better, with half of all these families having growth of 3.5 log points or lower. This modest cyclicality is a characteristic for most of the household earnings percentiles.

The results for the mean suggest an even worse outlook over a four-year horizon. For the entire income distribution, the average growth rate implies a loss of 23 log points during recessions and 13 log points during expansions. The average effect is substantially worse for higher income families. During recessions, households at the 90th percentile of $\bar{Y}(H)$ see average earnings growth of about -25 log points. But, for the top 1% of households, the average growth is -79 log points. Even during expansions, we see largely negative average

[28] GOS and Busch et. al (2015) impose a similar sample restriction based on minimum earnings, though they use different thresholds. Sabelhaus and Song (2009) show omitting the bottom 10 percent of earners results in meaningfully smaller standard deviations of earnings growth.

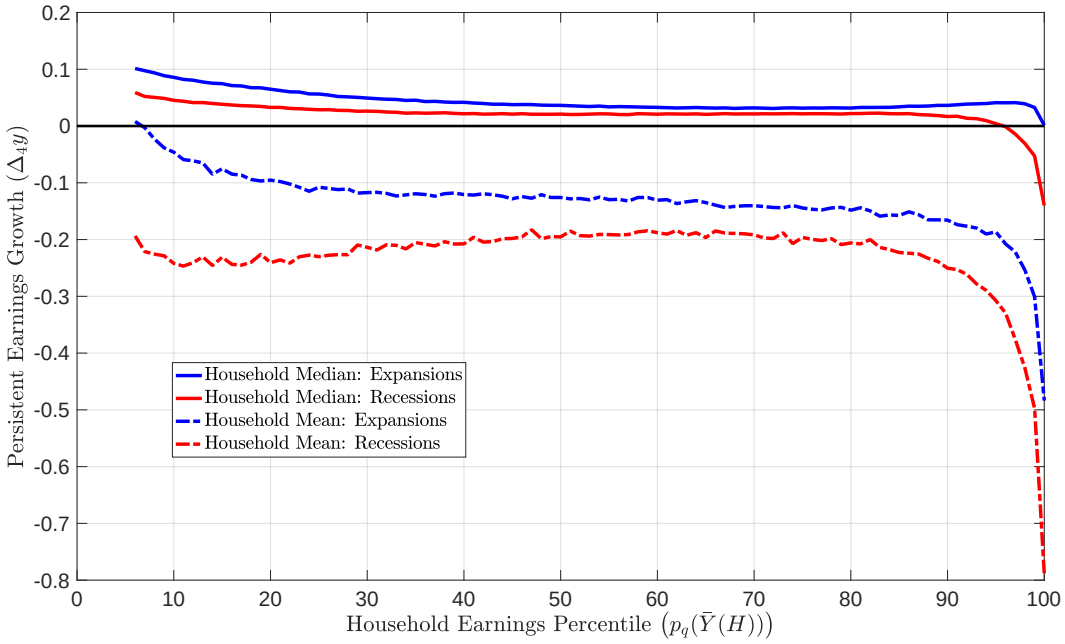

Figure 4: Median and Mean Household Earnings Growth

Notes: The median and mean of household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$. Blue lines are for expansions and red lines are for recessions. Household earnings $\Delta y(H)$ median shown as solid lines, and male earnings $\Delta y(M)$ median shown as dashed lines.

growth for higher income households. At the 90th percentile, average growth is -17 log points, and for the top 1% average growth is -48 log points. The contrast between median and mean shocks is attributable to the sensitivity of mean values to outlying observations. Given the negative-skewness clearly reported below, it is not surprising that the mean is so far below the median. Likewise, the sensitivity of the mean to this negative-skewness creates much larger business cycle effects relative to the median.

3.4 Skewness of Earnings Growth

We show the skewness of household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$ over the business cycle in Figure 5. As suggested by the patterns in Figure 3, household earnings shocks are negatively skewed at all income levels in both recessions and expansions. This implies that negative-skewness is a pervasive characteristic of earnings risk for households in the U.S., which we refer to as "skewness risk." The skewness risk of household earnings growth generally increases in household earnings, with the highest income

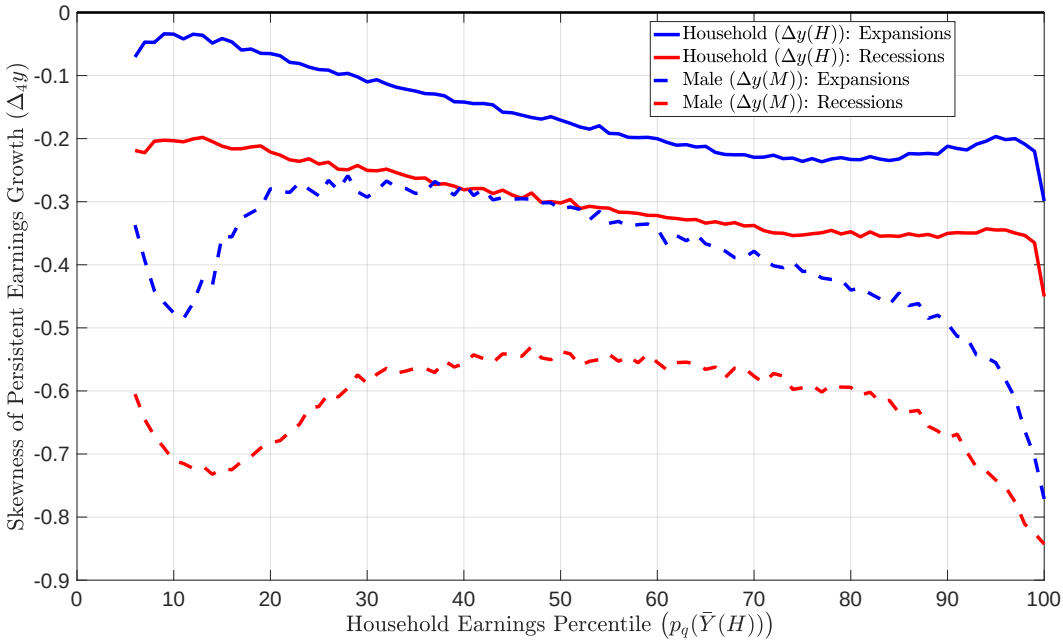

Figure 5: Skewness of Household and Male Earnings Growth

Notes: Measures of the skewness of household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$. Blue lines are for expansions and red lines are for recessions. Household earnings $\Delta y(H)$ skewness shown as solid lines, and male earnings $\Delta y(M)$ skewness shown as dashed lines.

households experiencing the most skewness risk.

The pattern for males is similar, though in this case there is a region where negative-skewness decreases (from the 15th to the 30th percentile in recessions and the 10th to 20th percentile in expansions). We find that higher income households experience the greatest male skewness risk. Male skewness risk is far greater than the household's, at each phase of the business cycle and at each household earnings percentile. For example, at the 20th percentile of income ($42,000) during recessions, household skewness is -0.22 compared to -0.68 for males. On average, household skewness risk is less than half of male skewness risk.

Negative-skewness increases in recessions, a pattern we refer to as "countercyclical skewness risk." To show countercyclical skewness risk clearly, in Figure 6 we plot the difference in the skewness of earnings growth for households and males as the solid and dashed purple lines, respectively. This figure shows how much more negatively skewed are earnings shocks in recessions relative to expansions, shedding light on how earnings risk changes over the business cycle.

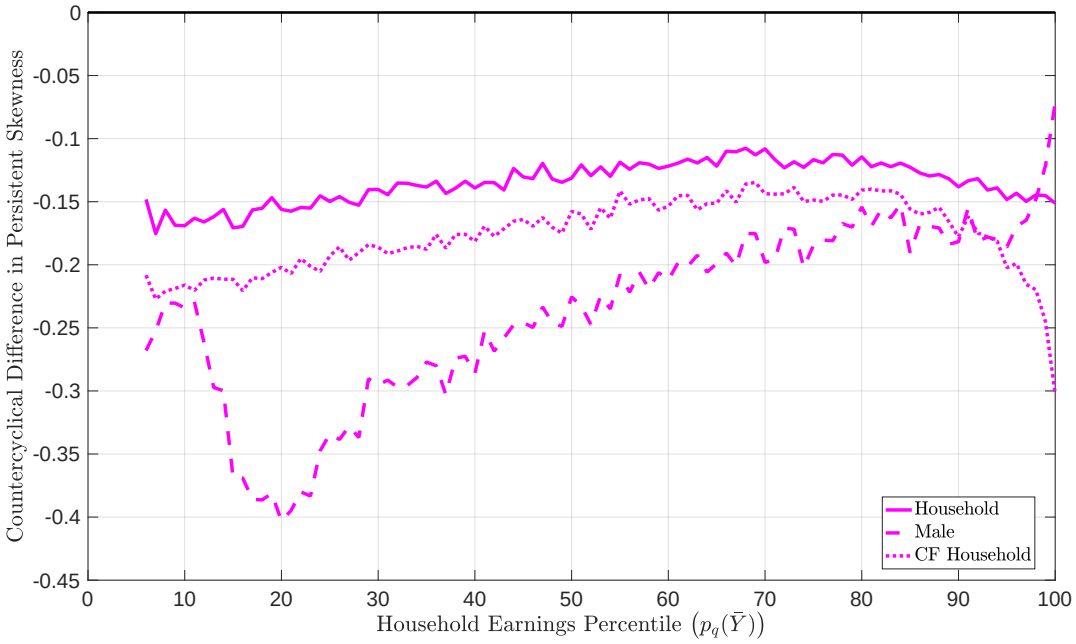

Figure 6: Countercyclical Change in Skewness

Notes: The difference between earnings growth skewness in recessions and in expansions. Defined as the recessionary skewness minus the expansionary skewness. Shown for household earnings $\Delta y(H)$ (solid line), male earnings $\Delta y(M)$ (dashed line), and counter-factual household earnings (dotted line, labeled "CF Household"). Counter-factual household earnings are constructed by assuming zero earnings growth for the spouse.

To help decompose the gap between household and male countercyclical skewness risk, we construct counterfactual household earnings growth. The counterfactual we consider holds fixed spousal income, isolating the change in household income coming only from the male earner. The skewness risk for this counterfactual is shown as the dotted line. Generally, counterfactual skewness risk lies between the series for males and for households.[29] The figure suggests that a nonzero second earner reduces the change in household earnings risk. At the household level, this risk is partially offset by the second earner. Graphically, this channel is the difference between the male skewness change and the counterfactual skewness

[29]There are two reasons why this relationship does not hold at higher income levels (specifically in the top 2 income percentiles) for two reasons. First, among these households male skewness change is smaller in magnitude than the household skewness change. This occurs because males in these percentiles have relatively high levels of skewness risk generally with small cyclical effects, as shown in Figure 5. Second, the counterfactual skewness change increases in magnitude at the top of the income distribution. Among these households, the rate at which spouses exit from the labor market increases in expansions and working spouses in this income range have high earnings. As a result, the counterfactual level of skewness is dampened relatively more for expansions than recessions creating a relatively large skewness change over the business cycle.

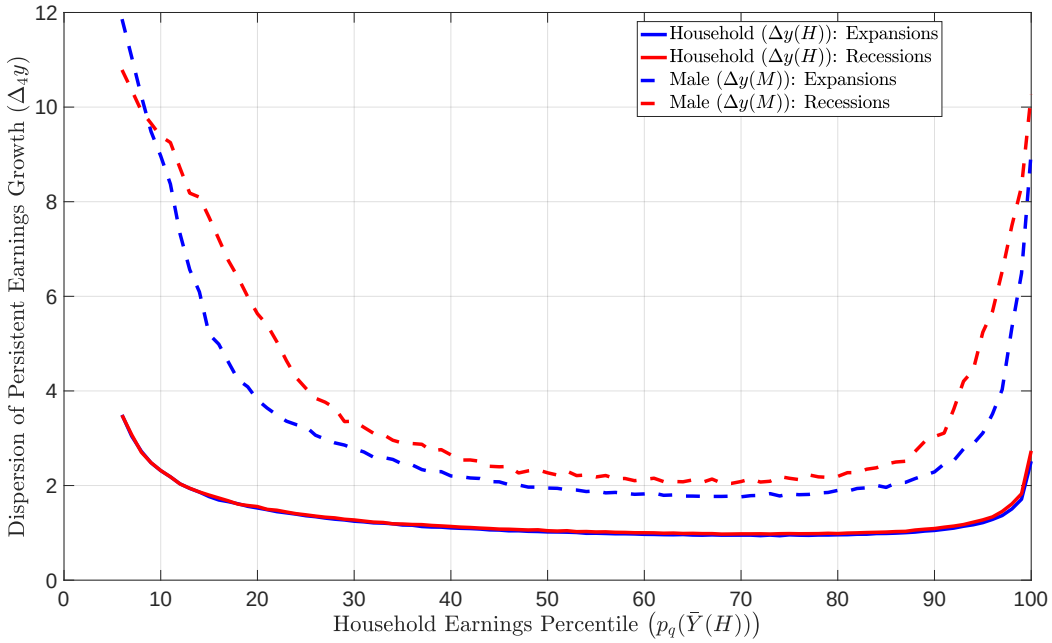

Figure 7: Dispersion of Household and Male Earnings Growth

Notes: Measures of the dispersion of household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$. Blue lines are for expansions and red lines are for recessions. Household earnings $\Delta y(H)$ dispersion shown as solid lines, and male earnings $\Delta y(M)$ dispersion shown as dashed lines.

change for the household. Moving from the counterfactual to the actual household skewness change reflects active labor market decisions of households, such as spousal entry into the labor market or an increase in labor supply among working spouses. The fact that the actual skewness change is less than the counterfactual skewness change suggests that the active decisions made by spouses further mitigate the business cycle effects of skewness risk. We explore this channel further in Section 4.

3.5 Dispersion of Earnings Growth

Figure 7 shows the dispersion of earnings shocks for households (solid lines) and for males (dashed lines) with recessions in red and expansions in blue. This figure suggests that households experience less dispersion than males in isolation. The difference is largest at the bottom and the top of the household earnings distribution, where households dispersion is only 20-25% of male earnings dispersion. For both males and households, dispersion is lowest in the middle of the income distribution. This differs from the pattern for skewness,

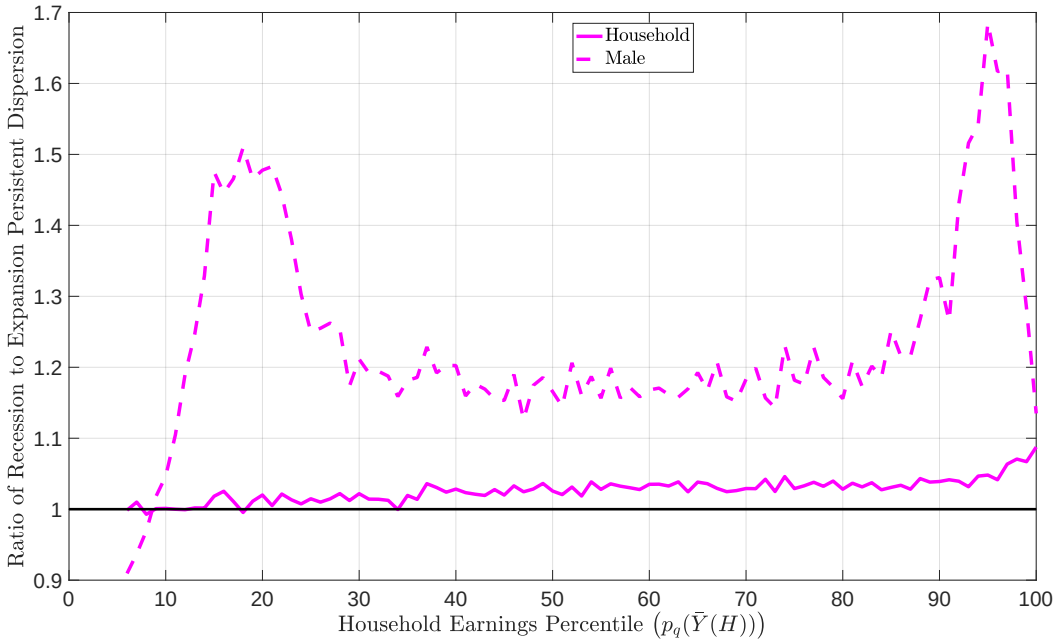

Figure 8: Countercyclical Ratio of Dispersion

Notes: The ratio of earnings growth dispersion in recessions to earnings growth dispersion in expansions. For households and males, separately.

where lower income households had the lowest level of skewness risk.

Figure 7 suggests that households experience very little increase in dispersion over the business cycle, which we refer to as "countercyclical dispersion." Figure 8 make this point clearer by plotting the ratio of recessionary dispersion to expansionary dispersion. This figure suggests that there is very little countercyclical dispersion in household earnings, because the ratio lies near 1 (the solid line). However, we find that that male earnings experience non-negligible countercyclical dispersion, as the ratio lies near 1.2 and hits 1.5 and 1.6 for low- and high-earning households, respectively (the dashed line). Our findings on household dispersion match Busch et. al's (2015), who also find no evidence of meaningful cyclicality. In contrast, our results for male dispersion differ from both GOS and Busch et. al (2015), which report no cyclicality in their longer sample periods.

3.6 Modeling Idiosyncratic Household Earnings Growth

How do earnings shocks evolve over the business cycle? We present several new insights. One, cyclical changes to the median of household earnings shocks are consistently modest over much of the household earnings $\bar{Y}(H)$ distribution. Two, household earnings shocks exhibit negative-skewness generally and recessions exacerbate skewness risk. Three, the dispersion of household earnings shocks does not change meaningfully over the business cycle. Together, these findings are consistent with the idea that household earnings shocks experience a median preserving shift around relatively low levels of growth over the business cycle.

These results provide important stylized facts shaping how macroeconomic theory should model the household side of the economy. While much prior work uses a different characterization of household earnings shocks over the business cycle, modeling it as a change in variance (Storesletten, Telmer, Yaron 2004, Favilukis, Ludvigson and Van Nieuwerburgh 2017), more recent work builds on the results from GOS to characterize household labor income. Schmidt (2016) adopts results from GOS to quantify labor income risk and shows that idiosyncratic earnings risk is a key driver of asset prices. Constantinides and Ghosh (2017) report that household consumption shocks are negatively skewed and meaningful drivers of asset prices. Their estimates also fit the moments of the equity premium, building on earlier work (Mankiw 1986, Brav et. al 2002). Our findings offer empirical support for these newer approaches.

Beyond asset pricing, updating the characterization of household labor income shocks to match our results also has implications for the literature analyzing household savings (Leland 1968, Sandmo 1970, Miller 1976, Sibley 1975, Guiso et. al 1992, Aiyagari 1994, Carroll and Samwick 1998, Kazarosian 1995, Fulford 2015). While a better understanding of how our characterization of household risk affects asset prices and precautionary savings is beyond the scope of this paper, below we offer a simple analysis that shows how our risk characterization impacts household welfare over the business cycle.

3.7 Welfare Analysis

To give an economic context for the statistical quantities we just presented, now we put our characterization of earnings risk into an simple framework. We consider workers facing the following choice: either receive random earnings growth α following distribution P_j, or receive a certain amount $1-\rho$ to avoid the gamble. For the utility function U this certainty

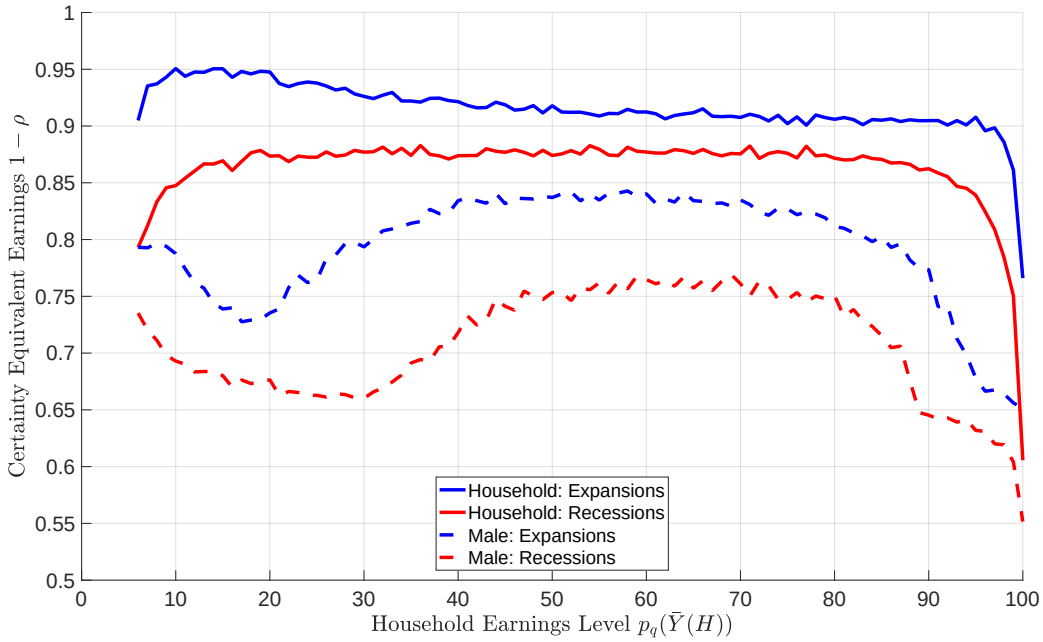

Figure 9: Certainty Equivalent Earnings

Notes: Certainty equivalent earnings (in proportion of base year earnings) implied by CRRA utility with parameter 1.25 for each household earnings level $\bar{Y}(H)$.

equivalent calculation is given by

$$U(1-\rho) = E_j[U(1+\alpha)]$$

where E_j is the expectation operator using distribution P_j. We refer to ρ as the risk premium and $1-\rho$ as the certainty equivalent earnings (a share of baseline earnings $\bar{Y}(H)$). We assume that $U(\cdot)$ is CRRA utility with a risk aversion parameter of 1.25. Of course, CRRA-type utility is widely used in macroeconomic models (e.g. Ljungdvist Sargent 2004), and the value of 1.25 implies that agents have only moderate risk aversion, as a parameter value of 1 gives log utility (for instance, Mehra and Prescott's 1985 equity premium analysis uses 10, and cites important early work that placed the values between 1–2). To derive a distribution j, we fit a piecewise uniform distribution to the data (further details in the appendix).[30]

Figure 9 plots the certainty equivalent earnings $1-\rho$. Solid lines plot the premia for households, and the dashed lines show premia for male workers alone. This figure suggests that

[30]We specify the distribution as piecewise uniform for two reasons. One, it can exactly match the distribution's percentiles to those measured in the data. Two, a piecewise uniform distribution can approximate any distribution arbitrarily well as the set of specified percentiles becomes dense in $[0, 1]$.

households would be willing to trade a meaningful fraction of their earnings to avoid the risk we characterized in the previous sections. The figure also underscores the importance of analyzing household earnings, as the tradeoff households are willing to make is meaningfully lower than what would be inferred from studying male earners in isolation. Averaged across $\bar{Y}(H)$ and business cycle phases, households would give up 11% of their baseline income to avoid their earnings risk while males would give up 25%. The certainty equivalents in Figure 9 vary meaningfully over the $\bar{Y}(H)$ distribution, mirroring the patterns we found above for dispersion and skewness. Of particular note, the risk premia for the highest income households is strikingly large – almost 25% in expansions and 40% in recessions – indicating a tremendous degree of earnings risk over both phases of the business cycle.

The certainty equivalents are also meaningfully different across households and males at each point of the business cycle. Among median-income households, the certainty equivalent is 92% during expansions for the household overall, but 83% for male earners in these households. A similar pattern holds in recessions, with a certainty equivalent of 83% for median-income households and 75% for males from these households.

These estimates hold interesting implications for the costs of business cycles. On average, male countercyclical risk is worth 10% of baseline earnings, or around 13% of their certainty equivalent earnings during expansions. By comparison, household countercyclical risk is worth 6% of baseline earnings, or just under 7% of their certainty equivalent earnings during expansions. These results imply that households face a countercyclical increase in risk that is roughly half as large as what males face in isolation. This implies that intra-household insurance reduces the cost of business cycles, relative to the cost facing males in isolation, by an economically significant amount.

Existing literature, using more sophisticated and dynamic methods, also suggests that increased risk over the business cycle results in meaningful welfare loses. Krebs (2007) finds that job losses lead to increased negative-skewness of earnings shocks, quantifying the idea that job losses permanently harm workers, resulting in relatively larger costs to households compared to prior work (Lucas 1987, 2003). Berger et. al (2016) build on the intuition from Krebs (2007) to show that household welfare increases by roughly 4% of lifetime income in response to optimal policy (designed to reduce the effects of recessions), not too far removed from our simple calculation for households.

Our results illustrate the economically meaningful attenuation of earnings risk that occurs when multiple working age adults form a household. The public seems generally aware that marriage offers a form of insurance for earnings risk. In a recent poll, 76% of respon-

dents list financial stability as an important reason for marriage.[31] Yet, trends in marriage rates suggest that this form of insurance is decreasing over time generally, with relatively larger declines among less educated individuals.[32] Instead of marriage, individuals are increasingly cohabiting.[33] Understanding how cohabiting households respond to male labor income risk is outside the scope of this paper, though remains an important avenue for future research.

4 Spousal Earnings

Our results so far suggest that households mitigate male labor income risk. But what are the important channels they use to do so? Do households insure against male labor earnings risk by the active behavior of spouses who alter labor supply? Or, does this insurance result mostly from a passive component reflecting the fact that males do not account for all household earnings, and that spouses' earning growth only imperfectly correlated to male earnings growth? In this section, we attempt to shed light on these important questions by quantifying spousal labor supply responses. This line of inquiry is a novel component of our analysis, as it allows us to simultaneously consider male and household earnings risk in a detailed way while highlighting the channels households use to insure their labor income risk.

4.1 Decomposing Spousal Labor Supply Responses

The spouse in the household is either employed or not employed, creating four categories: those who enter the labor market, those who exit the labor market, those who stay in the labor market, and those who stay out of the labor market.[34] In each of these cases, we

[31]This total reflects survey respondents who selected financial stability as "very important" and "somewhat important." For details see "A Survey of LGBT Americans: Attitudes, Experiences and Values in Changing Times." June 13, 2013. Pew Research Center.

[32]According to the Pew Research Center, 69% of adults age 25 or older with at least a bachelor's degree were married in 1990, compared to 65% in 2015. For individuals with only a high school education or less, the marriage rate was 63% in 1990 compared to 50% in 2015. Lundberg, Pollak and Stearns (2016) report a similar pattern of marriage rates by education over time among individuals with at least a high school degree or equivalent.

[33]The cohabitation rate for women increased 14 percentage points (42%) from 1995 to 2010 among women age 15-44 for their first "union" Casey et. al (2013).

[34]We refer to changes in earnings of working spouses as changes in the intensive margin, and both exit for working spouses and entry for nonworking spoues a changes in the extensive margin. However, all we observe are annual earnings. As a result, it is possible that some responses that we classify as intensive are actually movements along the extensive margin, for example by leaving one job and entering another.

characterize both the share of spouses in that category and the earnings of the spouse.[35] We denote the shares in the following way: $\theta(S, E)$ is the share of non-working spouses who *enter* employment, $\theta(S, L)$ is the share of spouses who *leave* employment, $\theta(S, R)$ of spouses who *remain* employed. The share of spouses staying out of employment $\theta(S, O)$ is pinned down by $\theta(S, O) = 1 - \theta(S, L) - \theta(S, E) - \theta(S, R)$.

The $\theta(S, E)$, $\theta(S, L)$, and $\theta(S, R)$ are *unconditional* probabilities of being in one of the three states E, L, R. In our main results, we restrict attention to *conditional probabilities* that take the current worker or non-worker status of the spouse as given. We take this step to focus on the the particular extensive and intensive margin behaviors that households use to mitigate male labor income risk. For example, when studying the likelihood that a spouse enters the labor market, we restrict the analysis to households where the spouse was not working so that this margin is a possible outcome.[36] We define the following conditional probabilities

$$\theta(S, E|N) \equiv \frac{\theta(S, E)}{\theta(S, E) + \theta(S, O)} \quad \text{and} \quad \theta(S, L|W) \equiv \frac{\theta(S, L)}{\theta(S, L) + \theta(S, R)}.$$

Note that $\theta(S, E|N)$ is the entry probability conditional on currently being a non-worker, and $\theta(S, L|W)$ is the leave probability conditional on currently being a worker.

We furthermore can measure the earnings change associated with the E, L, R states, denoted $\Delta y(S, E), \Delta y(S, L), \Delta y(S, R)$. The quantity $\Delta y(S, E)$ is the average log earnings gained by newly employed spouses. The quantity $\Delta y(S, L)$ is the average log earnings left behind by spouses leaving employment. And the quantity $\Delta y(S, R)$ is the average earnings growth seen by spouses who remain employed in years t and $t + 4$. We focus attention in the main text on $\Delta y(S, R)$.[37]

4.2 Empirical Method

We characterize spousal responses conditional on the joint distribution of $\bar{Y}_t(H)$ and $\Delta_4 y(M)$. For example, we can make comparisons across households who differ in the male labor earnings shock, but that have very similar long-run income. Likewise, we can look at how the effect of a given male earnings shock varies over the long-run income distribution. In order to clarify these relationships, we estimate a surface using nonparametric kernel regressions. There are some important advantages to using this nonparametric method. First, although

[35] Recall that we have recoded all \$0 earnings to \$1 earnings. This allows our entry and exit earnings growth rates to remain finite.

[36] In the appendix we show results unconditional on spousal employment.

[37] In the appendix we report results for $\Delta y(S, E), \Delta y(S, L)$.

less flexible than the approach used so far, we make weak assumptions on the nature of the conditional mean relationship. Second, it allows us to draw useful statistical inferences in a way that was not possible in our earlier approach. A third advantage is that the method lends itself to graphical analysis used earlier. Specifically, we plot slices of the estimated surface. Fourth, as described below, the regressions help smooth the data and allow us to make meaningful comparisons even when analyzing relatively disaggregated data.

Formally, in the main paper we estimate the conditional expectation function

$$E(x(S)|\Delta y(M), \bar{y}(H))$$

where $x(S)$ is a spousal variable from the set $\{\theta(S, E|N), \theta(S, L|W), \Delta y(S, R)\}$. Let $z_i \equiv (x(S)_i, \Delta y(M)_i, \bar{y}(H)_i)^\top$ with generic index i indexing the data from $y(H) \times \Delta y(M)$ bins included in estimation. We assume the variables have a common probability density function (p.d.f.) $f(\zeta) \equiv f(\zeta_f, \zeta_m, \zeta_{hh})$. The first step begins by using a product kernel function $K(z_i - \zeta, \kappa)$ to estimate this vector's probability density function.[38] Li and Racine (2007) show that this $\hat{f}(\zeta)$ is consistent and asymptotically normal in the presence of dependence across the i and between the elements of z_i. The second step is to find the conditional mean by integration. The estimated conditional mean function is given by

$$\hat{g}(\zeta_m, \zeta_{hh}) = \frac{\sum_{i=1}^{n} x_{f,i} K(z_i - \zeta, \kappa)}{\sum_{i=1}^{n} K(z_i - \zeta, \kappa)}.$$

This implies that for each point (ζ_m, ζ_{hh}) the fitted value is a weighted average of $x(S)_i$, with weights coming from that ratio of product kernels. This is why the procedure is referred to as "local constant" and shows how this technique is a more sophisticated version of taking averages. Li Racine (2007) show that $\hat{g}$ is asymptotically normal for the population conditional mean function. To give statistical inference in the statements we make, we use asymptotic standard errors for each conditional mean estimate and we provide these standard errors as shaded regions in all of our figures.

[38] The $k(\cdot)$ are univariate kernel functions satisfying $\int k(v)dv = 1$, $k(v) = k(-v)$, $\infty > \int v^2 k(v)dv > 0$. The product kernel function is $K(z_i - \zeta, \kappa) \equiv k\left(\frac{x(S)_i - \zeta_f}{\kappa_f}\right) \times k\left(\frac{\Delta y_4(M)_i - \zeta_m}{\kappa_m}\right) \times k\left(\frac{\bar{y}(H)_i - \zeta_{hh}}{\kappa_{hh}}\right)$. The pdf is given by $\hat{f}(\zeta) = \frac{1}{n\kappa_f\kappa_m\kappa_{hh}} \sum_{i=1}^{n} K(z_i - \zeta, \kappa)$. We use as our $k(\cdot)$ a second-order Gaussian kernel, using data-dependent bandwidth parameters κ consistently chosen by following Li Racine (2004, 2007) and using least squares cross-validation. Ultimately, the population object of interest is $E(\zeta_f|\zeta_m, \zeta_{hh}) \equiv \int \zeta_f f_{\zeta_f,(\zeta_m,\zeta_{hh})}((\zeta_m, \zeta_{hh}), \zeta_f)\, d\zeta_f$. We use $\hat{f}$ to estimate the conditional p.d.f. All of these choices follow Li Racine (2007).

4.2.1 Sample Restriction

We estimate the function separately in recessions and expansions for the following types of responses: entry, exit, remaining, and earnings growth associated with each. We restrict our sample in two ways. First, we focus on households between the 15th and 95th $\bar{y}(H)$ percentiles (roughly \$34,000 to \$200,000). Second, we focus on households with $\Delta_4 y(M)$ between -0.7 and 0.7 (a loss of 50% to a gain of 100%). These filters reduce the impact of extreme $\bar{y}(H)$ and $\Delta_4 y(M)$ observations that represent a very small number of households in our data. With this $\Delta_4 y(M)$ filter in place, our results focus on 75% of the households with long-run earnings between \$34,000 and \$200,000. The bottom line of this is to say that our regression results for, say, recessions come from about 15,000 observations, as opposed to the "full" 20,000 observations.[39]

4.3 Results

We focus primarily on results that can help address the following thought experiment. Among families that have similar long run income, how does the spousal response vary over the male earnings shock distribution? Do households with large male income loses respond differently compared to households with large male earnings gains? For insight into these questions, we plot the estimated spousal labor supply response over the $\Delta y(M)$ distribution for various slices of $\bar{Y}(H)$. These slices and confidence intervals come from estimates of the surface defined by $x(S))$, $\bar{Y}(H)$ and $\Delta y(M)$ and estimated confidence surfaces. In the appendix, we show the results for other slices of $\bar{Y}_(H)$, and also show results that plot spousal labor supply over the $\bar{Y}(H)$ for various slices of $\Delta y(M)$.

Exit and Entry Probabilities by Household Earnings Before exploring how spousal entry and exit vary among families with similar $\bar{Y}(H)$ but different $\Delta y(M)$, we first characterize these transitions over the $\bar{Y}(H)$ distribution among families where the male earnings shock is zero. Figure 10 plots the conditional entry rate in the top panel and the conditional

[39] More detail on what we mean. For each year t, we have 10,000 pre-aggregated observations of the vector $(\bar{y}(H), \Delta_4 y_{t+4}(M), x(S))^\top$. It may be easier to explain each observation's construction by indexing it by (q, r) for $q, r = 1, 2, \ldots, 100$. Observation (q, r) has as its first element the average value of long-run log earnings in $p_q(\bar{y}(H))$, as its second element the average value of male earnings growth in male earnings growth percentile r of the distribution of male earnings growth in $p_q(\bar{y}(H))$, and as its third element the average value of a spousal response variable $x(S)$ in that particular bin of households. Put another way: we perform a sequential percentile sort for each year t, first by $\bar{y}(H)$ and then by $\Delta_4 y(M)$, and define the observation vector to be the average values of long-run log earnings, male earnings growth, and the spousal response variable within each of those $100 \times 100 = 10,000$ groups.

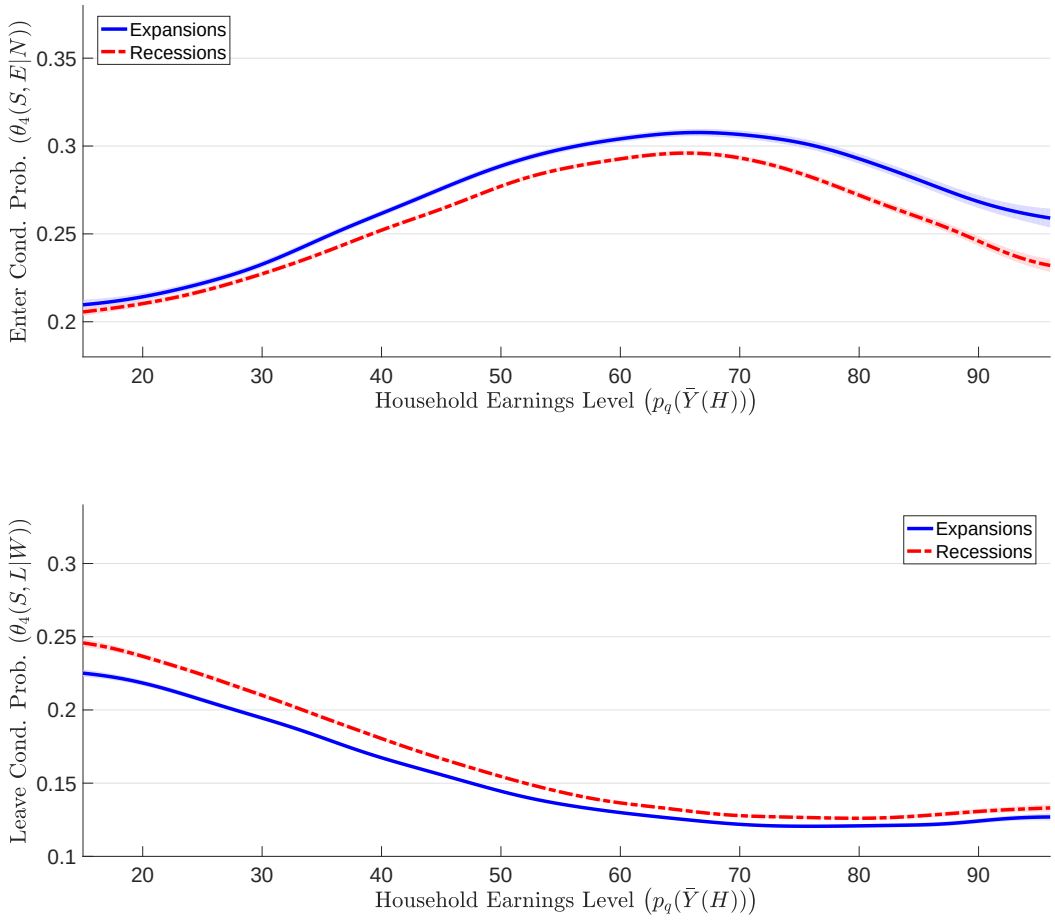

Figure 10: Conditional Entry and Exit Probabilities, at Zero Growth

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. At $\Delta_4 y(M) = 0$.

exit rate in the bottom panel, over the $\bar{Y}(H)$ distribution.[40]

The top panel of Figure 10 shows that there is substantial variation in the probability of entry during expansions over the income distribution, increasing from 0.2 at the low-end, peaking around the 65th percentile at 0.3, and decreasing to 0.26 for the highest income families. There is also substantial variation in the probability of exit during expansions over the income distribution, decreasing from 0.23 at the low-end, bottoming out around the 80th percentile at 0.12, and barely increasing to 0.13 for the highest income families. Lower-income households are less likely to enter employment and more likely to leave employment, and higher-income households are more likely to enter employment and less likely to leave employment. These appear to be characteristics, perhaps endogenously, of households at

[40] $\Delta y(M)$ of zero occurs about at the 45th percentile.

different earnings levels.

The same broad patterns hold during recessions, but meaningful differences exist over the business cycle. The probability of entry is lower during recessions relative to expansions, likely reflecting the fact that during recessions the job finding rate decreases. For lower-income households, the cyclical gap is negligible. But thereafter the gap increases as we move up the $\bar{Y}(H)$ distribution.

As shown in the bottom panel of Figure 10, there is also meaningful cyclical variation in the likelihood of spousal exit. The probability of exit during a recession is higher than during expansions, likely reflecting the fact that during recessions the job separation rate increases. But the countercyclical increase in the likelihood of spousal exit varies over the household income distribution. For lower-income households, the difference is meaningful and statistically significant. But this gap decreases as we move up the $\bar{Y}(H)$ distribution. For 60th percentile $\bar{Y}(H)$ households and above, the gap is substantively smaller, though still statistically significant.

Spousal Entry Non-working spouses are a potentially important way that households mitigate male labor income risk. If the spouse is not employed and the male worker experiences a negative earnings shock, the household can attenuate the earnings loss when the spouse gains employment. Consistent with this idea, entry rates for spouses are higher among families where the male experiences a negative earnings change than among families where the male experiences a similar sized positive earnings change, conditional on income. Figure 11plots the estimated conditional entry rate over $\Delta y(M)$ values that range from -0.7 to 0.7 (roughly a 50% earnings loss to a 100% earnings gain). The three panels show this relationship at three percentiles of $\bar{Y}(H)$: p_{25} (roughly \$47,000), p_{50} (roughly \$76,000), and p_{75} (roughly \$111,000).

Among lower-income families in the top panel, the probability that a spouse enters the labor market during expansions is 4 percentage points higher (18% higher) when the $\Delta y(M)$ is a 60% decrease (roughly 50 log point decrease) compared to when the $\Delta y(M) = 0$. This pattern also holds at higher income levels. For families at the median $\bar{Y}(H)$, the probability of entry is 4 percentage points (15%) higher and for families at the 75th percentile $\bar{Y}(H)$, the entry probability is 5 percentage points (17%) higher. For median and higher-income households, the entry rate decreases as we move from $\Delta y(M) = 0$ to positive male earnings growth. For families at the median $\bar{Y}(H)$, the entry rate is 1 percentage point lower (4%) lower when $\Delta y(M)$ is a 40% increase compared to when $\Delta y(M) = 0$. For families at the

$p_{25}\left(\bar{Y}(H)\right)$

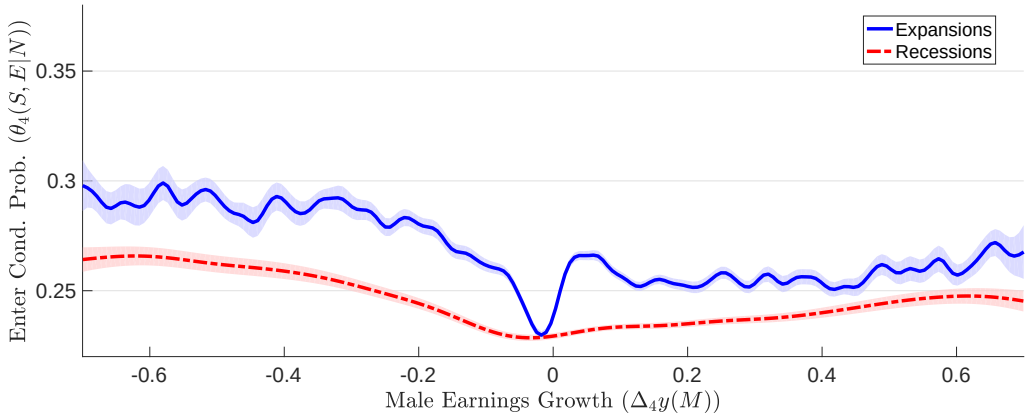

$p_{50}\left(\bar{Y}(H)\right)$

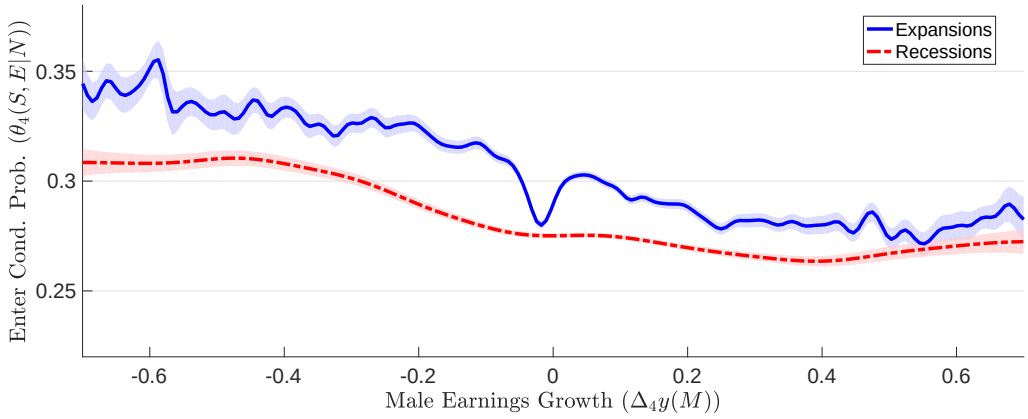

$p_{75}\left(\bar{Y}(H)\right)$

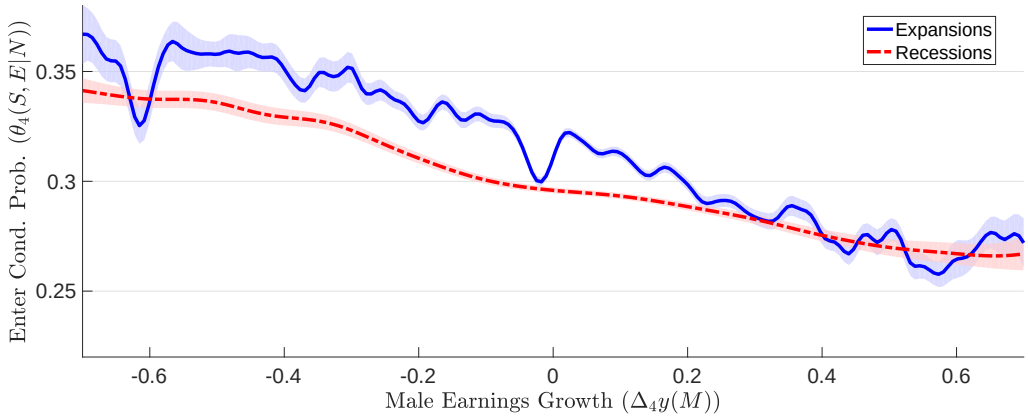

Figure 11: Conditional Entry Probabilities

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. The horizontal axis is in units of log growth.

75th percentile $\bar{Y}(H)$, the comparable difference is 2.5 percentage points (9%) lower.

Recessions generally attenuate spousal entry, but interesting patterns are visible.[41] First, across income levels in Figure 11 there is negligible cyclical difference when $\Delta y(M)$ is slightly negative. However, when $\Delta y(M)$ becomes more negative or becomes positive, there are meaningful differences. The cyclical gap is as big as 3 percentage points for negative $\Delta y(M)$ in lower- and median-income households. For large and positive $\Delta y(M)$, the cyclical gap is statistically insignificant, in median- and higher-income households.

Spousal Earnings Growth Working spouses also mitigate male labor income risk. Figure 12 shows $\Delta_4 y(S)$ values over the same $\Delta y(M)$ range and for the same $\bar{Y}(H)$ slices as Figure 11. Spousal earnings changes are not strongly related to negative male earnings shocks, evident by the the blue and red lines being roughly horizontal over negative $\Delta y(M)$ values in Figure 12. Given that many households (on average 65% over the $\bar{Y}(H)$ distribution) have a working spouse who continues to work, the zero correlation between spousal earnings changes and male earnings losses reduces household labor income risk. Intuitively, this pattern insures that male earnings losses are not amplified by comparable losses in spousal earnings, effectively maintaining a form of earnings diversification.

For positive male earnings shocks, spousal earnings growth decreases, and by more during recessions. Among lower-income households, the decline in $\Delta y(S, R)$ is 8 log points during expansions and 10 during recessions, moving from small positive male shocks to large (60 log point or roughly 80%) shocks. The lower panels show that this patten is not isolated to lower-income families. Therefore, we find widespread evidence of negative correlation between male and spousal earnings changes *only for positive male earnings growth*. This pattern could arise if spouses lower their earnings growth as the male experiences large earnings gains. For median-income households, the contrast during recessions is large. When male earnings growth is negative or only modestly positive, then spousal earnings growth is nearly zero, but as male earnings growth becomes strongly positive, then spousal earnings tend to *decline*. For higher-income households the pattern is similar. This negative correlation for male earnings losses serves to reduce household earnings dispersion but could amplify negative-skewness. We next consider an extreme reduction in spousal earnings, labor market exit.

[41] Note that the recessionary kernel regression estimates are much smoother than the expansionary estimates, when viewed as $\bar{Y}(H)$ cross-sections. We have estimated these surfaces separately, hence each has separately estimated bandwidths from Li and Racine's (2004) cross-validation procedure. It turns out that the recessionary $\Delta y(M)$ bandwidths are estimated to be larger, which helps lead to a smoother surface. However, the cross-validation routine chooses as large a bandwidth as leads to small changes in the estimated surface. Hence the recessionary surface in the data appears to be smoother than the expansionary surface.

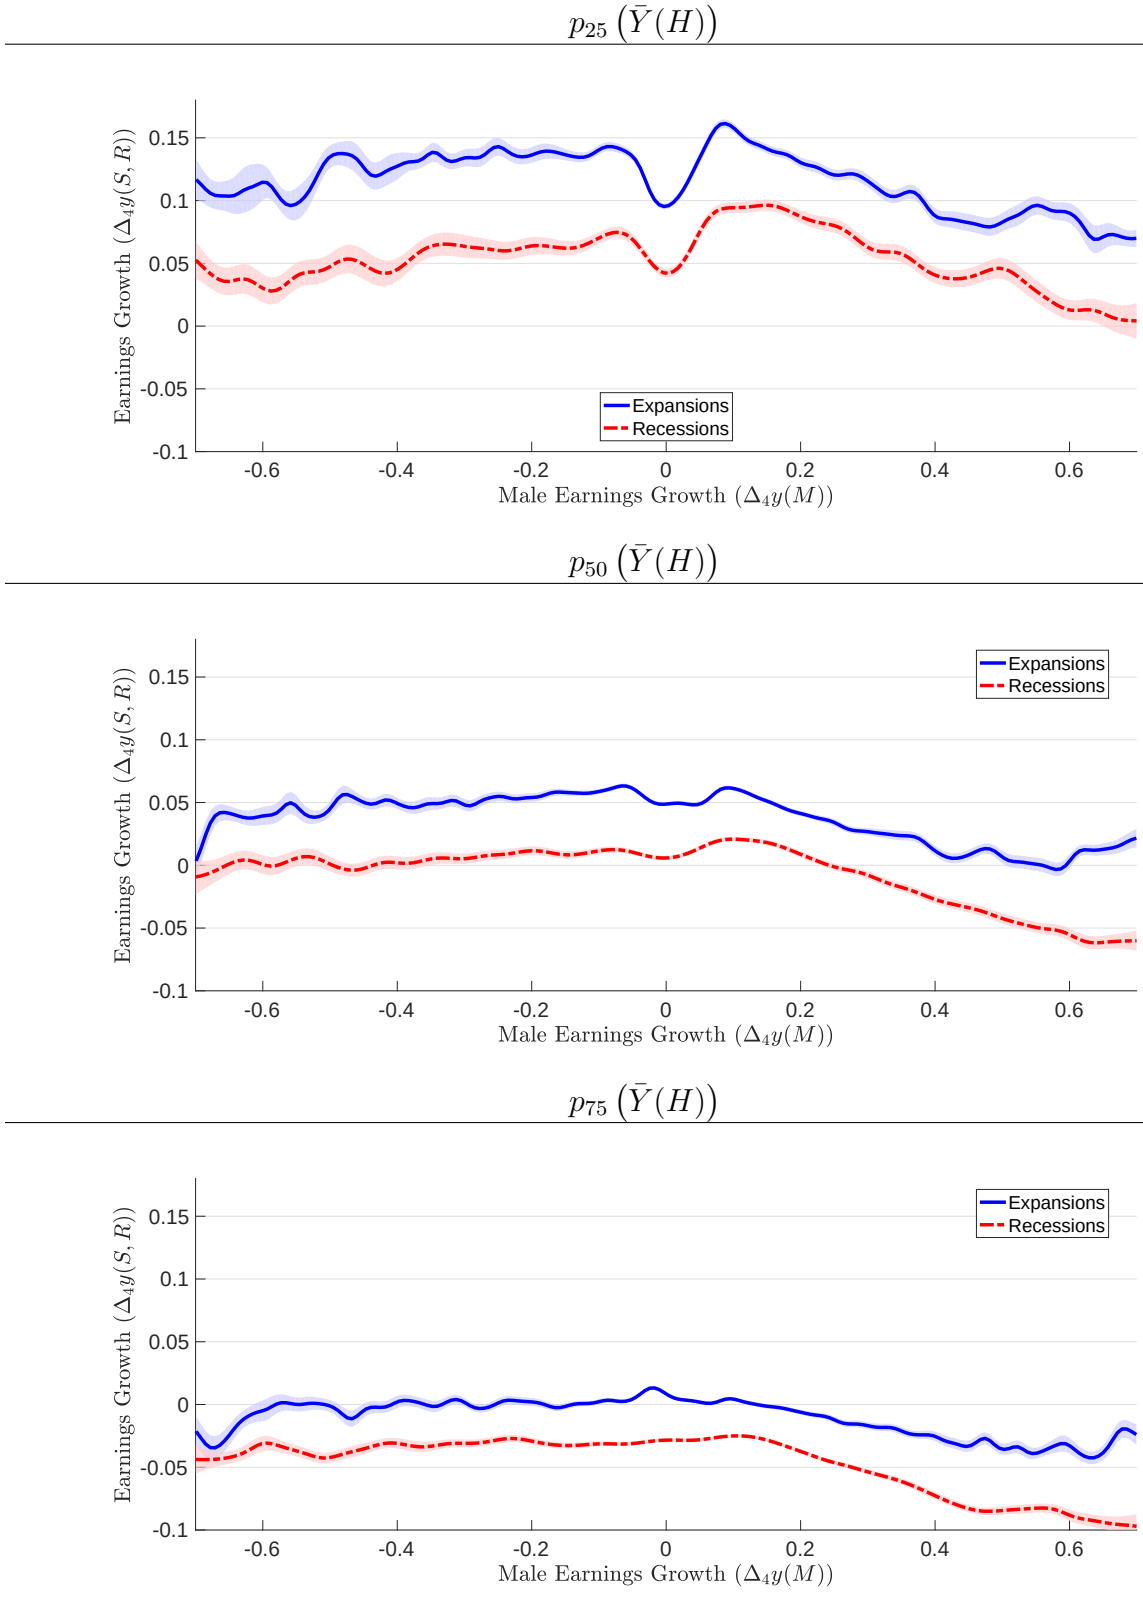

Figure 12: Intensive Earnings Change

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. The horizontal axis is in units of log growth.

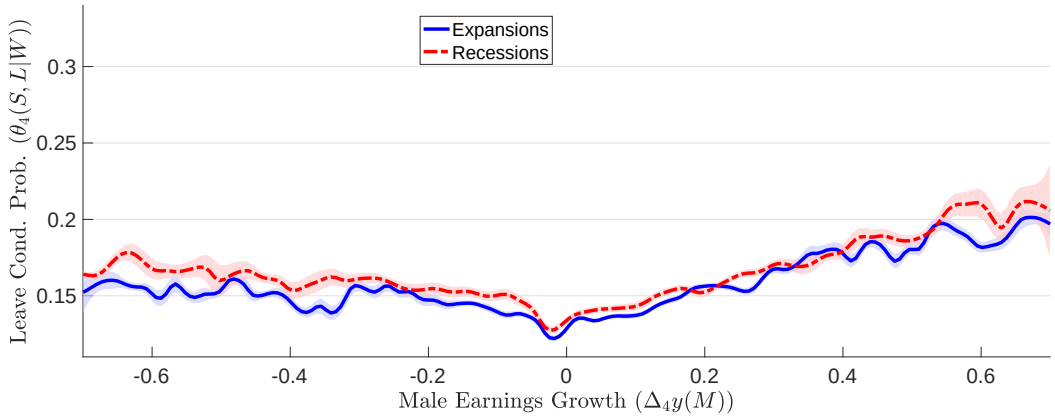

Figure 13: Conditional Exit Probabilities by Household Earnings, High Earnings

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. For 95th percentile of $\bar{Y}(H)$. The horizontal axis is in units of log growth.

Spousal Exit Somewhat surprisingly, spousal exit from the labor market in response to large positive male earnings shocks is not common: it is present only for relatively high-income families. Figure 13 shows conditional spousal exit over the same $\Delta y(M)$ values in Figures 11 and 12, but only for households at $p_{95}(\bar{Y}(H))$ (roughly \$200,000). Among these families, the probability of spousal exit increases from about 0.13 at $\Delta y(M) = 0$ to about 0.20 for $\Delta y(M) = 0.6$ (an increase of roughly 80%). This suggests that among high-income families, spousal exit is about 50% more likely when male earnings grow by 80% than when the male earnings do not grow at all. As shown in the appendix, we do not find a similar pattern for families at the 25th, 50th, or 75th percentiles of $\bar{Y}(H)$. Perhaps this channel reflects the spouse's substitution of market-based-production for home production, which might be more preferable the larger are the actual dollar gains experienced by male earnings. Note that spousal exit results in a reduction of spousal earnings that reduces dispersion in household earnings growth, all else equal. However, somewhat surprisingly, it does not feature as an important avenue for most long-run earnings levels.

4.4 Discussion

Section 4 points to at least three reasons why we found in Section 3 that household earnings risk was less than male earnings risk.

First, spousal entry responds to male earnings outcomes. Non-working spouses are more likely to enter employment when the male worker experiences earnings losses than when

the male worker experiences earnings gains. All else equal, this reduces household skewness risk. Second, spousal earnings are not positively correlated with male earnings loses. This reduces skewness risk and dispersion as male earnings losses are not amplified by spousal earnings losses. Third, spousal earnings are (weakly) negatively correlated with positive male earnings growth. This also reduces the dispersion of household earnings growth. Overall, these patterns suggest that households make active choices, including when to enter and when to leave the labor market, as well as which job to select, that effectively insure the household against earnings risk from the male earner.

These empirical results have noteworthy economic implications. Our data show that when male earnings decline, working spouses do not increase their earnings but that non-working spouses are more likely to enter employment. This empirical finding differs from recent work that expands commonly used models to capture household insurance mechanisms. For instance Wu and Krueger (2018) build on empirical work by Blundell, Pistaferri, and Saporta-Eksten (2016) finding that spousal labor supply responses are key to household insurance. However, the specific mechanism identified by Wu and Krueger(2018) is the ability of working spouses to increase earnings by increasing hours. In fact, Wu and Krueger (2018) shut off the extensive margin entirely, deeming it as secondary importance. The fact that we document a strong empirical relationship between spousal labor market entry and negative earnings shocks for males suggests that updated models may want to explicitly incorporate this mechanism in addition to the mechanism highlighted in Wu and Krueger (2018).

5 Conclusion

Analyzing the earnings of roughly one-fifth of all married households in the United States that contain prime age males for the years 2000 to 2014, we find evidence that households face increased earnings risk over the business cycle from increasing negative-skewness but not from greater dispersion in earnings growth. This pattern is consistent with the idea that large earnings losses are more likely and large earnings gains are less likely during recessions compared to expansions. Existing research documented this pattern in male earnings. We show that this pattern holds for households and we document important differences between male earnings growth and household earnings growth.

The presence of an actual or potential second earner affords households an opportunity to dampen earnings shocks of the primary earner. We find that the probability of the

spouse entering employment rises when the male experiences earnings losses. We also find that working spouses experience earnings growth that is uncorrelated with male earnings losses and weakly negatively correlated with male earnings gains. These responses insure the household against male earnings risk and result in households experiencing less risk compared to the risk inferred from studying males in isolation. And yet household earnings risk remains consequential during both phases of the business cycle. This prompts additional questions. Do earnings alone adequately represent household income? What are the roles for capital income and income sources that mix returns to labor and capital? What is the role of a progressive tax system? We plan to address these questions in future work.

References

Aiyagari, S. Rao. 1994. "Uninsured Idiosyncratic Risk and Aggregate Saving." *Quarterly Journal of Economics*, 109(3): 659–684.

Arellano, Manuel, Richard Blundell, and Stphane Bonhomme. 2017. "Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework." *Econometrica*, 85(3): 693–734.

Barro, Robert J. 2009. "Rare Disasters, Asset Prices, and Welfare Costs." *American Economic Review*, 99(1): 243–64.

Becker, Gary S. 1965. "A Theory of the Allocation of Time." *The Economic Journal*, 75(299): 493–517.

Becker, Gary S. 1974. "A Theory of Social Interactions." *Journal of Political Economy*, 82(6): 1963–93.

Benhabib, Jess, Richard Rogerson, and Randall Wright. 1991. "Homework in Macroeconomics: Household Production and Aggregate Fluctuations." *Journal of Political Economy*, 99(6): 1166–87.

Berger, David, Ian Dew-Becker, Konstantin Milbradt, Lawrence Schmidt, and Yuta Takahashi. 2016. "Layoff risk, the welfare cost of business cycles, and monetary policy." Northwestern University.

Blundell, Richard, Luigi Pistaferri, and Itay Saporta-Eksten. 2016. "Consumption Inequality and Family Labor Supply." *American Economic Review*, 106(2): 387–435.

Bound, John, Charles Brown, and Nancy Mathiowetz. 2001. "Measurement error in survey data." *Handbook of Econometrics*, Chapter 59, 3705–3843. Elsevier.

Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus. 2016. "Measuring Income and Wealth at the Top Using Administrative and Survey Data." *Brookings Papers on Economic Activity*, 261–312.

Busch, Christopher, David Domeij, Fatih Guvenen, and Rocio Madera. 2015. "Asymmetric Business Cycle Risk and Government Insurance." University of Minnesota.

Chetty, Raj, John Friedman, and Emmanuel Saez. 2013. "Using Differences in Knowledge across Neighborhoods to Uncover the Impacts of the EITC on Earnings." *American Economic Review*, 103(7): 2683–2721.

Chetty, Raj, Michael Stepner, Sarah Abraham, Shelby Lin, Benjamin Scuderi, Nicholas Turner, Augustin Bergeron, and David Cutler. 2016. "The Association Between Income and Life Expectancy in the United States, 2001-2014." *Journal of the American Medical Association*, 315(16): 1750–1766.

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Constantinides, George M, and Darrell Duffie. 1996. "Asset Pricing with Heterogeneous Consumers." *Journal of Political Economy*, 104(2): 219–40.

Friedman, Milton. 1962. *Capitalism and Freedom.* University of Chicago Press.

Fulford, Scott. 2015. "The Suprisingly low imnportance of income uncertainty for precaution." *European Economic Review*, 79(2): 151–171.

Gourio, Franois. 2012. "Disaster Risk and Business Cycles." *American Economic Review*, 102(6): 2734–66.

Guiso, Luigi, Tullio Jappelli, and Daniele Terlizzese. 1992. "Earnings uncertainty and precautionary saving." *Journal of Monetary Economics*, 30(2): 307–337.

Guvenen, Fatih, Serdar Ozkan, and Jae Song. 2014. "The Nature of Countercyclical Income Risk." *Journal of Political Economy*, 122(3): 621 – 660.

Hansen, Lars P, and Thomas J Sargent. 2008. *Robustness.* . 2 ed., Princeton University Press.

Hayfield, Tristen, and Jeffrey Racine. 2008. "Nonparametric Econometrics: The np Package." *Journal of Statistical Software*, 27(1): 1–32.

Kazarosian, Mark. 1997. "Precautionary savings: a panel study." *Review of Economic Statistics*, 79(2): 241–247.

Kim, Tae-Hwan, and Halbert White. 2004. "On more robust estimation of skewness and kurtosis." *Finance Research Letters*, 1(1): 56–73.

Krebs, Tom. 2007. "Job Displacement Risk and the Cost of Business Cycles." *American Economic Review*, 97(3): 664–686.

Kreiner, Claus Thustrup, David Dreyer Lassen, and Soren Leth-Petersen. 2014. "Measuring the Accuracy of Survey Responses Using Administrative Register Data: Evidence from Denmark." In *Improving the Measurement of Consumer Expenditures. NBER Chapters*, 289–307. National Bureau of Economic Research, Inc.

Kydland, Finn, and Edward Prescott. 1982. "Time to Build and Aggregate Fluctuations." *Econometrica*, 50(6): 1345–70.

Leland, Hayne E. 1968. "Saving and Uncertainty: The Precautionary Demand for Saving." *Quarterly Journal of Economics*, 82: 465–473.

Li, Qi, and Jeffrey S Racine. 2004. "Nonparametric estimation of regression functions with both categorical and continuous data." *Journal of Econometrics*, 119: 99–130.

Li, Qi, and Jeffrey S Racine. 2007. *Nonparametric Econometrics: Theory and Practice.* Princeton University Press.

Ljungqvist, Lars, and Thomas Sargent. 2004. *Recursive Macroeconomic Theory, 2nd Edition.* Vol. 1. 2 ed., The MIT Press.

Miller, Bruce L. 1976. "The Effect on Optimal Consumption of Increased Uncertainty in Labor Income in the Multiperiod Case." *Journal of Economic Theory*, 13: 154–167.

Mincer, Jacob. 1962. "Labor Force Participation of Married Women: A Study of Labor Supply." In *Aspects of Labor Economics*. 63–105. National Bureau of Economic Research, Inc.

Moffitt, Robert A., and Sisi Zhang. 2018. "Income Volatility and the PSID: Past Research and New Results." National Bureau of Economic Research Working Paper 24390.

Nardi, Mariacristina De, Giulio Fella, and Gonzalo Paz Pardo. 2018. "Nonlinear Household Earnings Dynamics, Self-insurance, and Welfare." National Bureau of Economic Research Working Paper 24326.

Panousi, Vasia, and Shanthi Ramnath. 2015. "The nature of income risk: Self-insurance and social insurance at the household level." Department of the Treasury.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913?1998." *The Quarterly Journal of Economics*, 118(1): 1–41.

Ramey, Valerie A., and Neville Francis. 2009. "A Century of Work and Leisure." *American Economic Journal: Macroeconomics*, 1(2): 189–224.

Reports, National Health Statistics. 2013. "First Premarital Cohabitation in the United States: 2006-2010 National Survey of Family Growth." National Health Statistics Reports 64.

Rietz, Thomas A. 1988. "The equity risk premium a solution." *Journal of Monetary Economics*, 22(1): 117–131.

Sabelhaus, John, and Jae Song. 2009. "Earnings Volatility Across Groups and Time." *National Tax Journal*, 62(2): 347–364.

Sabelhaus, John, and Jae Song. 2010. "The great moderation in micro labor earnings." *Journal of Monetary Economics*, 57(4): 391–403.

Saez, Emmanuel, and Gabriel Zucman. 2016. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data." *Quarterly Journal of Economics*, 131(2): 519–578.

Sandmo, Agnar. 1970. "The Effect of Uncertainty on Saving Decisions." *Review of Economic Studies*, 37: 353–360.

Sargent, Thomas J. 1987. *Dynamic Macroeconomic Theory.* Harvard University Press.

Sibley, David S. 1975. "Permanent and Transitory Income Effects in a Model of Optimal Consumption with Wage Income Uncertainty." *Journal of Economic Theory*, 11: 68–82.

Smets, Frank, and Raf Wouters. 2003. "An Estimated Dynamic Stochastic General Equilibrium Model of the Euro Area." *Journal of the European Economic Association*, 1(5): 1123–1175.

Storesletten, Kjetil, Chris I. Telmer, and Amir Yaron. 2004. "Cyclical Dynamics in Idiosyncratic Labor Market Risk." *Journal of Political Economy*, 112(3): 695–717.

Wu, Chunzan, and Dirk Krueger. 2018. "How Much Consumption Insurance in Bewley Models with Endogenous Family Labor Supply?" National Bureau of Economic Research Working Paper 24472.

A.1 Marriage

In order to abstract from income growth that results from changes in marital status, our baseline analysis conditions on families that remain married over the period used to measure income growth. To the extent that marital dissolution is correlated with financial stress that is brought about by the combination of male earnings changes and the inability of spouses to respond, our baseline results will overstate the spousal labor supply responses. As a result, our baseline results may be an underestimate of the total level of labor earnings risk that households face. In the remainder of this section, we briefly describe the effect of this decision on the composition of the sample.

Figure A.1 plots the share of joint families who are still married after four years over the long-run earnings distribution for households where the male experiences different earnings shocks, with the median in solid lines, the 90th in dotted lines and the 10th percentile in dashed lines. Overall, there does not appear to be a material difference in the mean rates of remaining married, conditional on long run earnings percentile and male earnings growth percentile, over the business cycle. To simplify the interpretation, Figure A.2 plots the mean (averaged over recessions and expansions) with a smaller range on the vertical axis. While these figures suggests that conditioning on joint filing status in both years results in differential selection in two key ways, the magnitudes are relatively small.

First, the figures suggests differences in the share remaining married across long-run income percentiles. There is generally an upward slope for each earnings change percentile moving across the long run income percentiles in both figures. This suggests that households with higher long run income are less likely to get divorced compared to households with lower long-run incomes holding fixed the percentile of male earnings changes. However, the magnitude of this upward slope is relatively small. For example, pooling all years and all male earnings percentiles, the mean of still being married after one year among the lowest $\bar{Y}(H)$ percentile is just over 0.90, increasing to 0.95 by the 20th percentile, to 0.96 by the 50th percentile and to 0.98 by the highest percentile. Four years later, the means are 0.77, increasing to 0.87 by the 20th percentile, to 0.90 by the 50th percentile and to 0.94 by the highest percentile

A second source of selection comes within each long run earnings percentile as there are differences in the likelihood of remaining married across the percentiles of male earnings changes. Interestingly, the likelihood of remaining married is generally highest at the median percentile for male earnings changes and lower at the 10th and 90th percentile of male earnings changes. The relatively lower share remaining married at the 10th percentile of male earnings changes may reflect male earnings shock-induced divorce, whereas the lower share

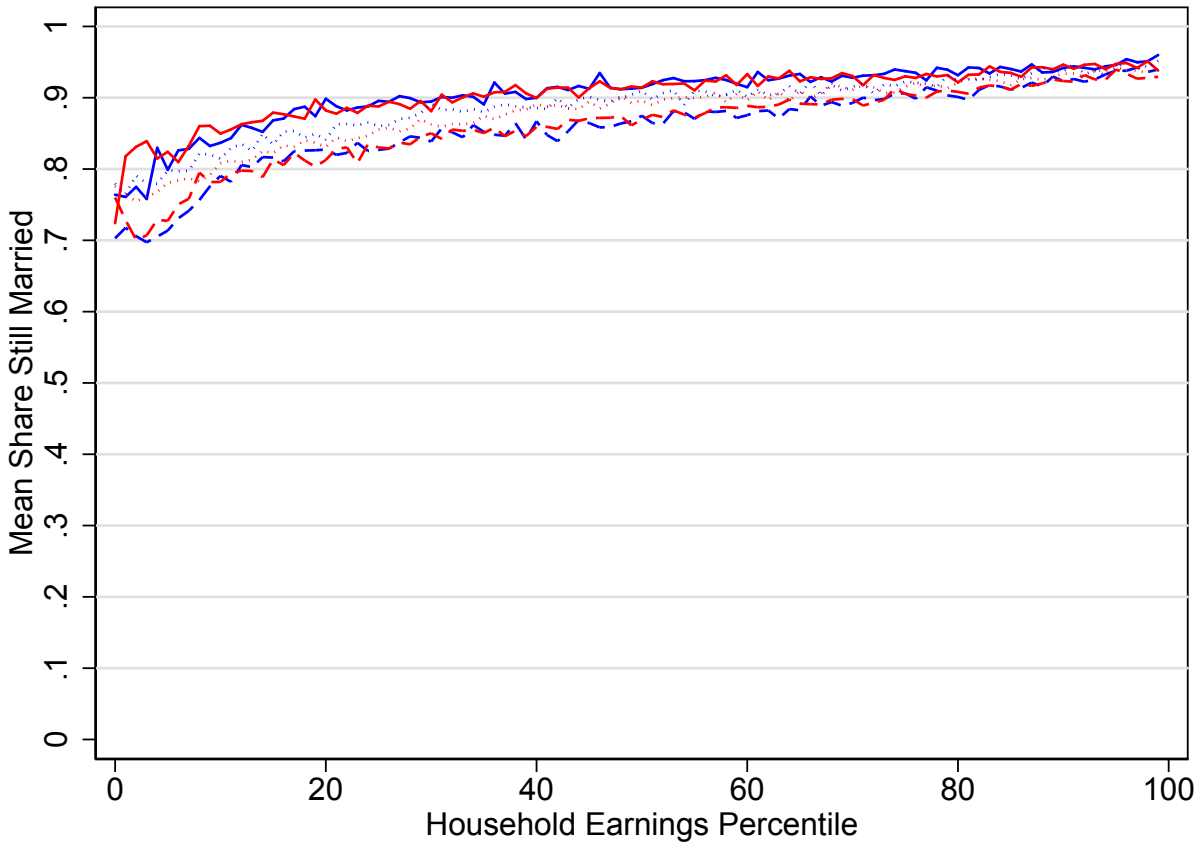

Figure A.1: Marriage: Mean Share

at the 90th percentile of male earnings change may reflect divorce-induced male earnings shocks, although it is not possible to differentiate the direction of causality. The relative difference in the likelihood of remaining married across male earnings change percentiles is smaller for households at the top of the long-run income distribution, compared to households lower down in the distribution. This may reflect the fact that bad earnings shocks are less consequential for households with higher levels of long-run earnings (income shock induced divorce) and that there is less need to increase earnings among males following divorce (divorce induced income shocks). Yet, these magnitudes are also relatively small. Pooling all years, the standard deviation in the share still married one year later across male earnings change percentiles is 0.037 at the lowest $\bar{Y}(H)$ percentile, decreasing to 0.014 by the 20th percentile of $\bar{Y}(H)$, to 0.007 at the 50th percentile of $\bar{Y}(H)$ and to 0.004 at the highest percentile of $\bar{Y}(H)$. Over four years these standard deviations are 0.069 at the lowest $\bar{Y}(H)$ percentile, decreasing to 0.034 by the 20th percentile of $\bar{Y}(H)$, to 0.022 at the 50th percentile of $\bar{Y}(H)$ and to 0.018 at the highest percentile of $\bar{Y}(H)$.

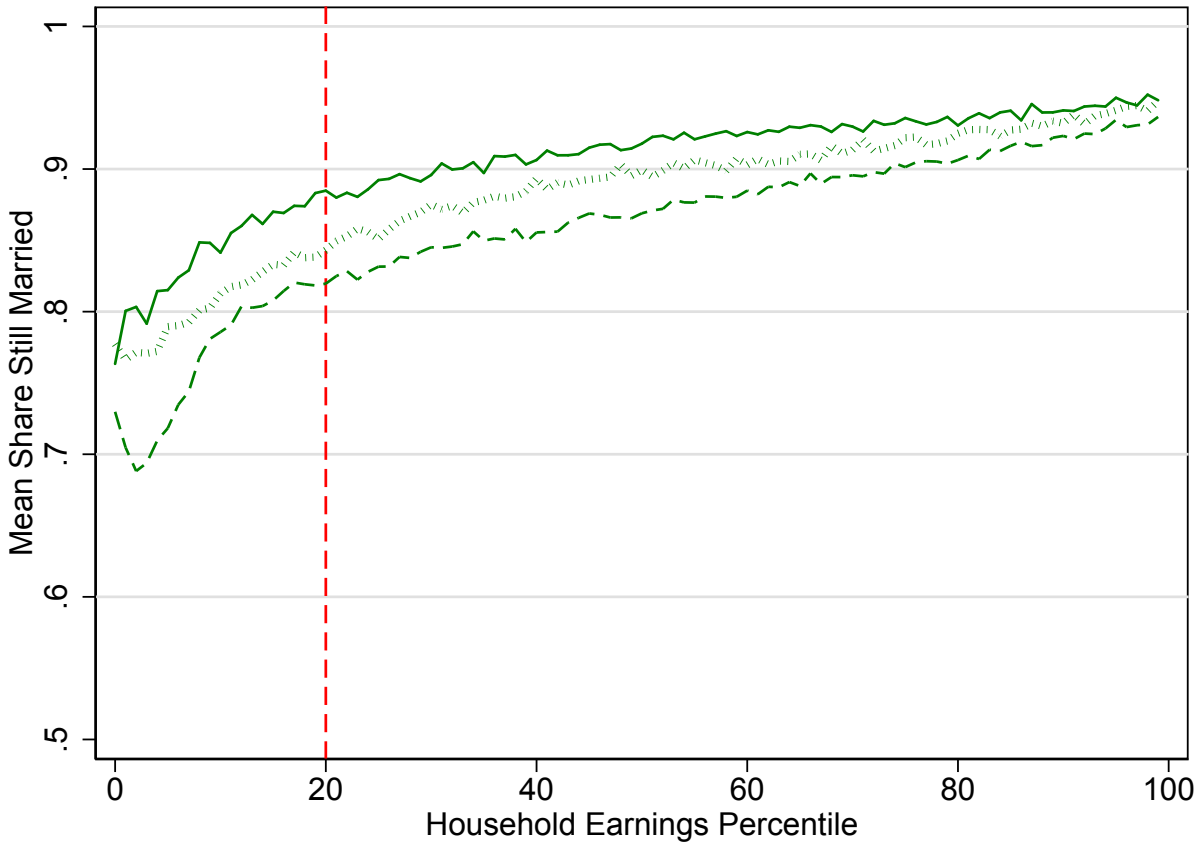

Figure A.2: Marriage: Mean Share, Focused

Not surprisingly, when we do not condition on the set of households that remain married, there is considerably larger negative earnings shocks. Intuitively, when we allow for divorce, there are cases where the household will experience large earnings losses simply from the exit of a working spouse. Figure A.3 shows the change in household log earnings over the long run income distribution for the 90th, 50th and 10th percentiles of household earnings shocks for the sample that does not condition on marriage in the second year when income is measured. Yet, even in this case when there are large swings in earnings that result from marital changes the patterns in skewness and dispersion still suggest that households face less risk relative to male earners in isolation. Figure A.4 plots skewness and Figure A.5 plots dispersion.

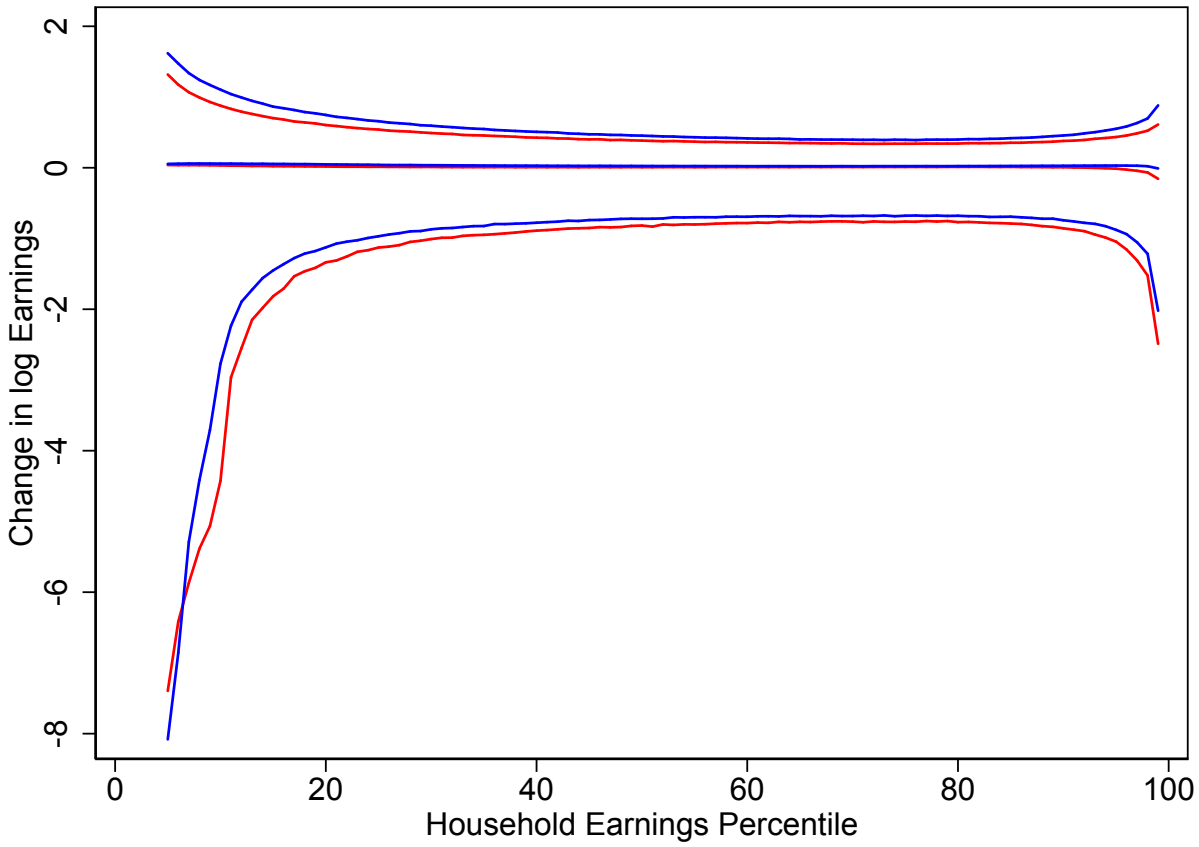

Figure A.3: Marriage: Distribution

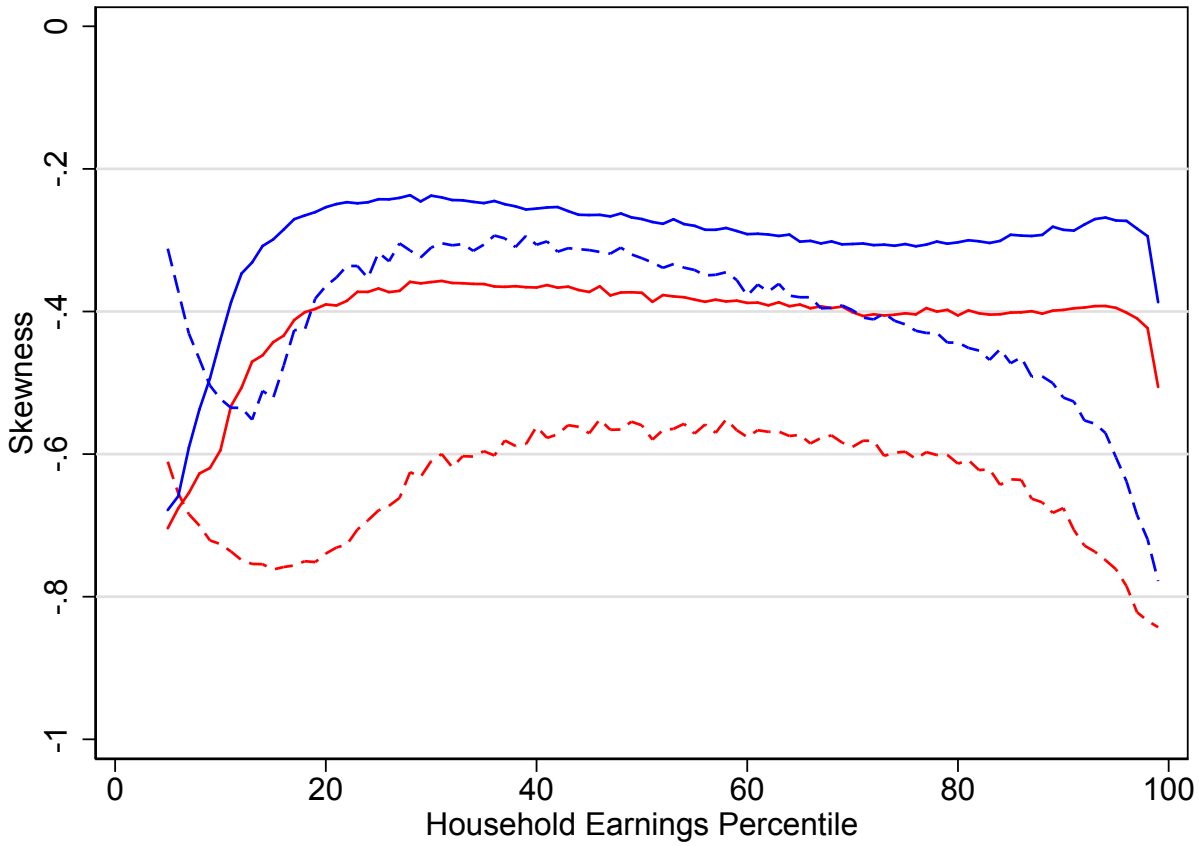

Figure A.4: Marriage: Skewness

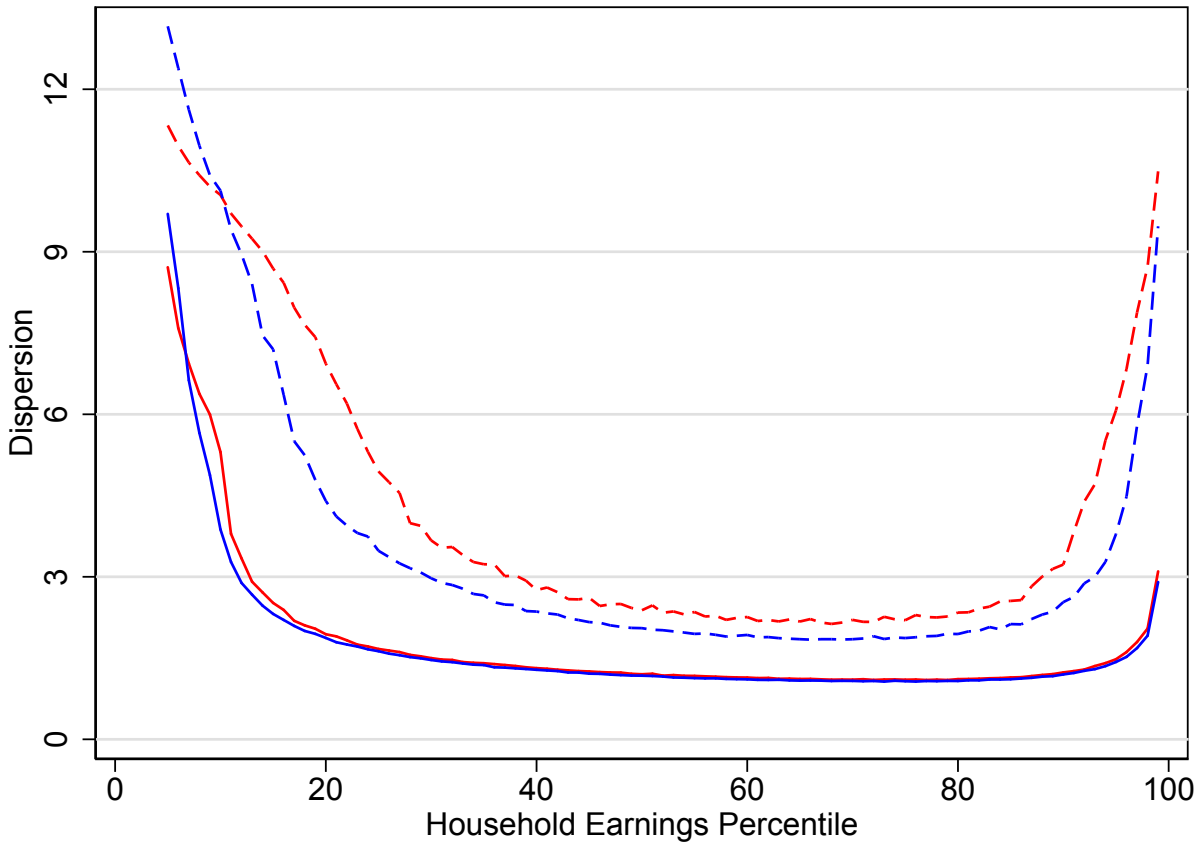

Figure A.5: Marriage: Dispersion

A.2 Welfare Analysis

We consider workers facing the following choice: either receive random earnings growth α following distribution P_j, or receive a certain amount $1-\rho$ to avoid the gamble. For the utility function U this certainty equivalent calculation is given by

$$U(1-\rho) = E_j[U(1+\alpha)]$$

where E_j is the expectation operator using distribution P_j. We refer to ρ as the risk premium and $1-\rho$ as the certainty equivalent earnings (a share of baseline earnings $\bar{Y}(H)$). Therefore there are two main ingredients for this calculation: the utility function and the earnings growth distribution. For simplicity we assume that $U(\cdot)$ is CRRA utility with a risk aversion parameter of 1.25, so that

$$U(c) = \frac{c^{1.25-1}-1}{1.25-1}.$$

To derive a distribution j, we fit a piecewise uniform distribution to the data as follows. For each type of economic unit (male worker or household), business-cycle phase (expansion or recession), and $\bar{Y}(H)$ percentile, we have calculated the 10th, 25th, 50th, 75th, and 90th percentiles of the Δy distribution from the data. For the purposes of this exercise, we choose to calibrate the minimum possible Δy value as -3 (a earnings loss of about 95%) and the maximum possible Δy value as 3 (earnings growth of about 2000%).[42]

We therefore have seven Δy values, $p_0, p_{0.1}, p_{0.25}, p_{0.5}, p_{0.75}, p_{0.9}, p_1$, that define six distinct regions, denoted $p_{0.1}-p_0$, $p_{0.25}-p_{0.1}$, etc. We specify that earnings growth is such that: (1) the probability of being in any region $p_r - p_q$ is $r-q>0$; and (2) the earnings growth values within each region are uniformly distributed. It is simple to draw from such a distribution using two independent uniform random variables $u_1, u_2 \sim \mathcal{U}[0,1]$. Suppose u_1 falls in region $p_r - p_q$ for $r > q$.[43] Then earnings growth is drawn as $p_q + (p_r - p_q)u_2$. We calculate $E_j[U(1+\alpha)]$ by approximating the expectation using 100,000 draws from the piecewise uniform distribution j. We then numerically solve for ρ using an equation solver. In practice, the entire procedure takes a fraction of a second for each j.

[42] We furthermore winsorize the 10th, 25th, 75th, and 90th percentiles as follows: the 10th percentile is no smaller than -2.5 (loss of about 91%), the 25th percentile is no smaller than -2 (loss of about 86%), the 75th percentile is no bigger than 2 (growth of about 740%), and the 90th percentile is no bigger than 2.5 (growth of about 1210%). This only affects male earnings growth percentiles, meaning that our modeled male earnings growth distribution is *weakly less risky* than a straight reading of the data.

[43] Ie. if $u_1 \in [0, 0.1]$ then earnings growth is drawn from region $p_{0.1} - p_0$, if $u_1 \in (0.1, 0.25]$ then earnings growth is drawn from region $p_{0.25} - p_{0.1}$, etc.

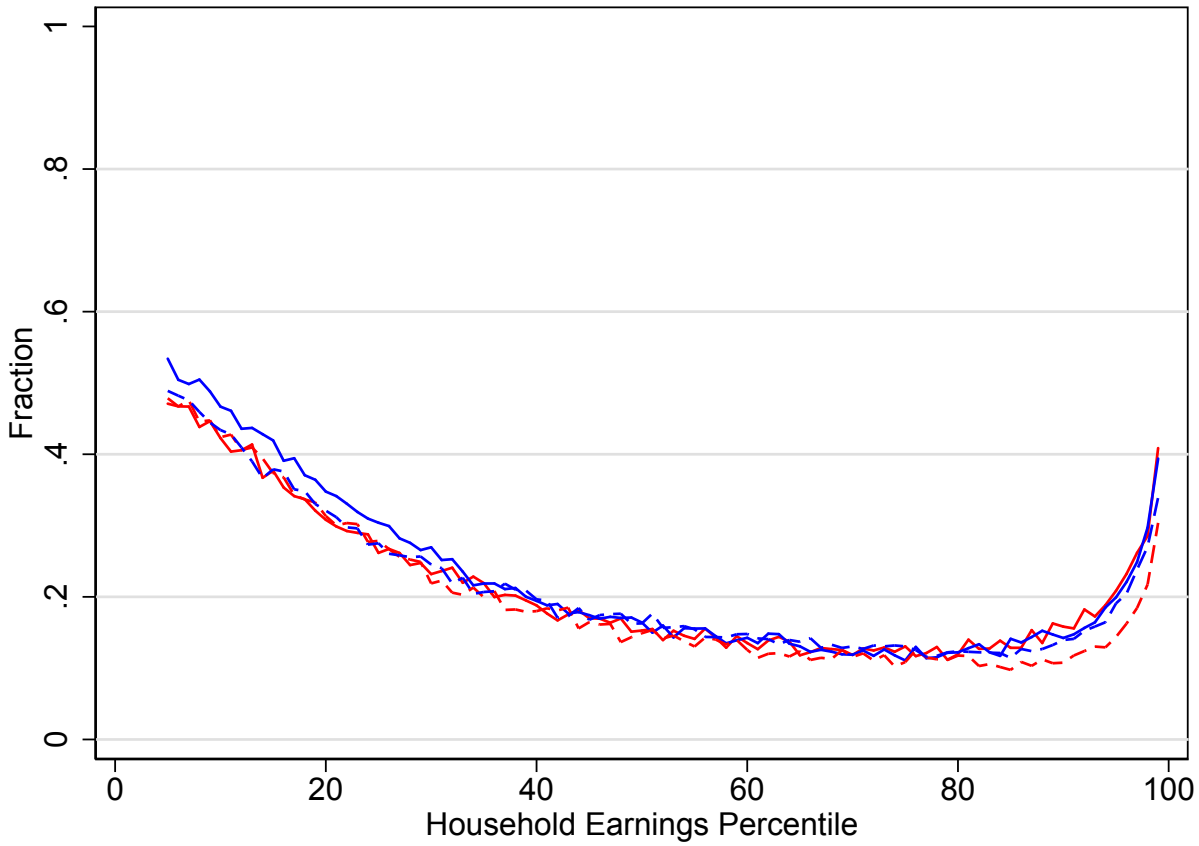

Figure A.6: Sample Selection

A.3 Sample Selection for Spousal Response

When analyzing the likelihood that a spouse enters the labor market, we condition on households with a non-working spouse. Likewise, when looking at the likelihood that a spouse stays employed or leaves employment, we restrict attention to households with a working spouse. Figure A.6 shows the likelihood of having a non-working spouse across the $\bar{Y}(H)$ distribution among households at the 25th percentile of make shocks (solid lines) and at the 75th percentile (dashed lines). The likelihood of having a non-working spouse in the initial year has no business cycle component, though it does vary over the $\bar{Y}(H)$ distribution. Both high and low-income households are more likely to have a non-working spouse relative to households in the middle of the income distribution. There is only modest variation across the male earnings shocks, conditional on income, as suggested by the figure at the 25th and 75th percentiles of the male earnings shocks distributions.

A.4 Transitory Shocks

Our baseline analysis examines persistent (four year) changes in earnings. In this section, we briefly discuss the results from transitory (one year) changes in earnings. These results suggest that households face less earnings risk and a smaller change in risk over the business cycle, relative to males, exactly the conclusions we reach in the main text when analyzing persistent earnings shocks.

The distribution of transitory earnings shocks are shown in Figure A.7, which plots the 90th, 50th and 10th percentile earnings shocks over the long run income distribution. This figure retains many of the key findings from the comparable analysis of persistent earnings shocks (Figure 3 in the main text), including: relatively anemic median shocks with limited business cycle effects, meaningful business cycle effects at the tails of the distribution with worse outcomes in recessions at both the 90th and 10th percentiles, and evidence of negative-skewness in both recessions and expansions.

Figure A.8 shows the skewness of the transitory household and transitory male earnings shocks in each phase of the business cycle. Mirroring our skewness results in the main text (Figure 5), the figure suggests that transitory skewness is a pervasive phenomena occurring at virtually all income levels in both recessions and expansions for households and for males. At the low end of the income distribution this does not hold, as some lower income households have positive skewness. Figure A.8 also suggests that the change in skewness is larger for males than for households, following the results for persistent shocks discussed in the main text. In general, the results for transitory skewness also imply that transitory shocks are less negatively skewed compared to persistent shocks.

Figure A.9 shows the dispersion of the transitory household and transitory male earnings shocks in each phase of the business cycle. Following our findings for persistent dispersion in the main text (Figure 7), the figure suggests that households face less transitory dispersion relative to males. The difference in dispersion between households and males is less for transitory changes compared to persistent changes in large part because male earnings shocks are far less disperse over a one year period. Once again, male dispersion appears countercyclical, as in persistent growth shown in Figures 7 and 8.

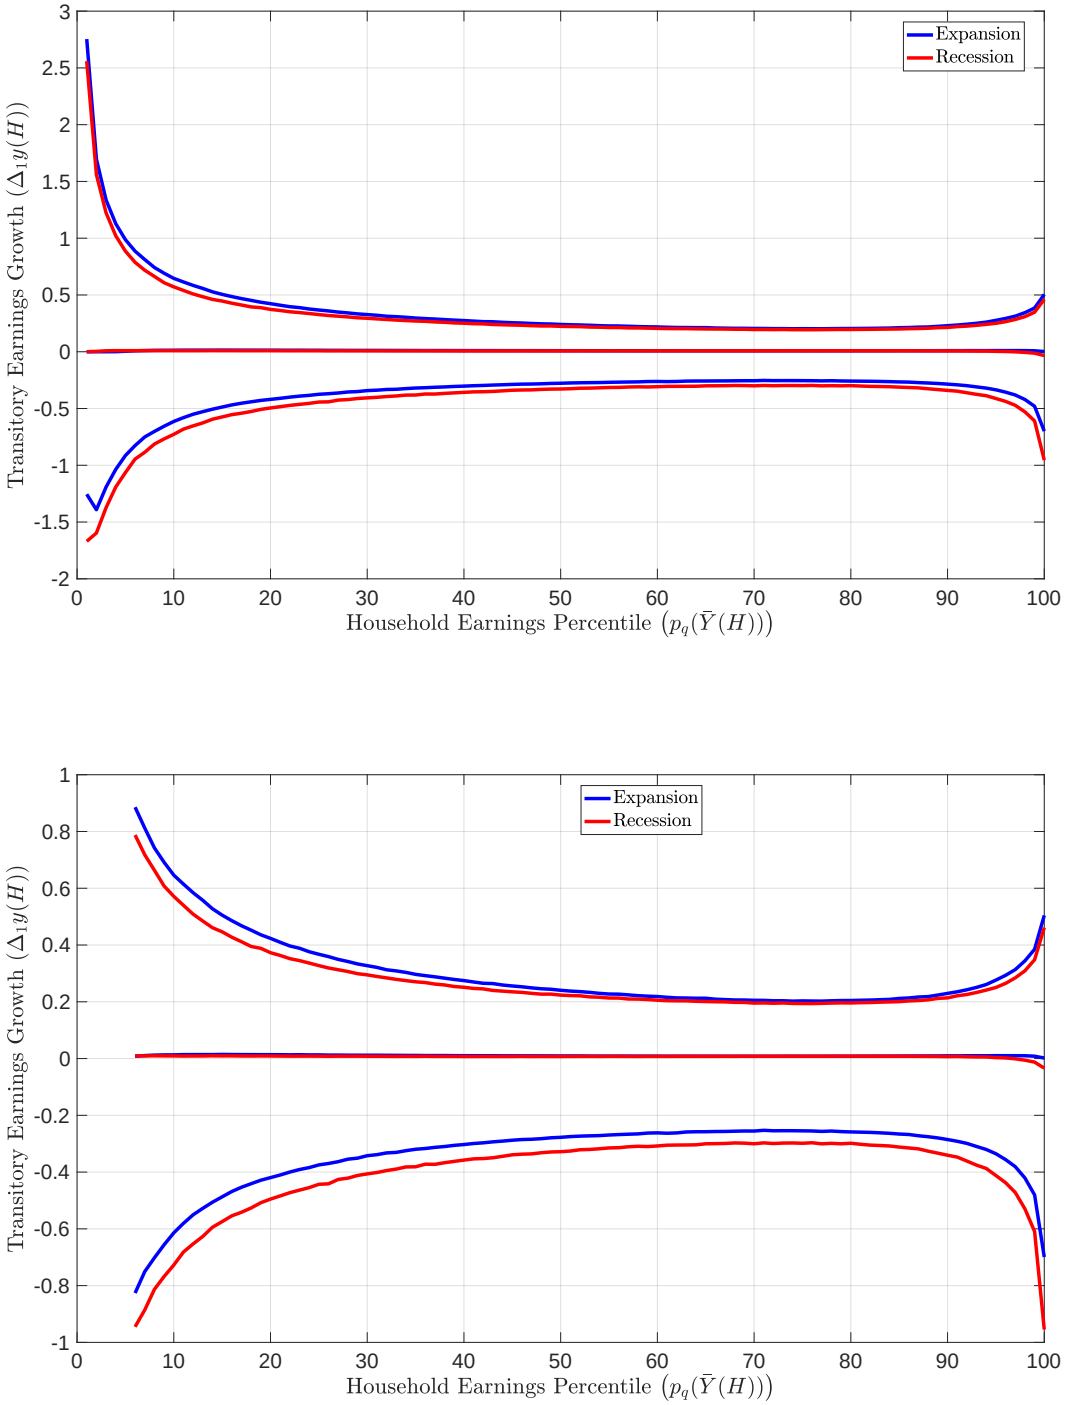

Figure A.7: Percentiles of the Household Earnings Growth Distribution, Transitory

Notes: The 10th, 50th, and 90th percentiles of the transitory household earnings growth $\Delta_4 y(H)$ distribution. For expansions (blue lines) and recessions (red lines) separately.

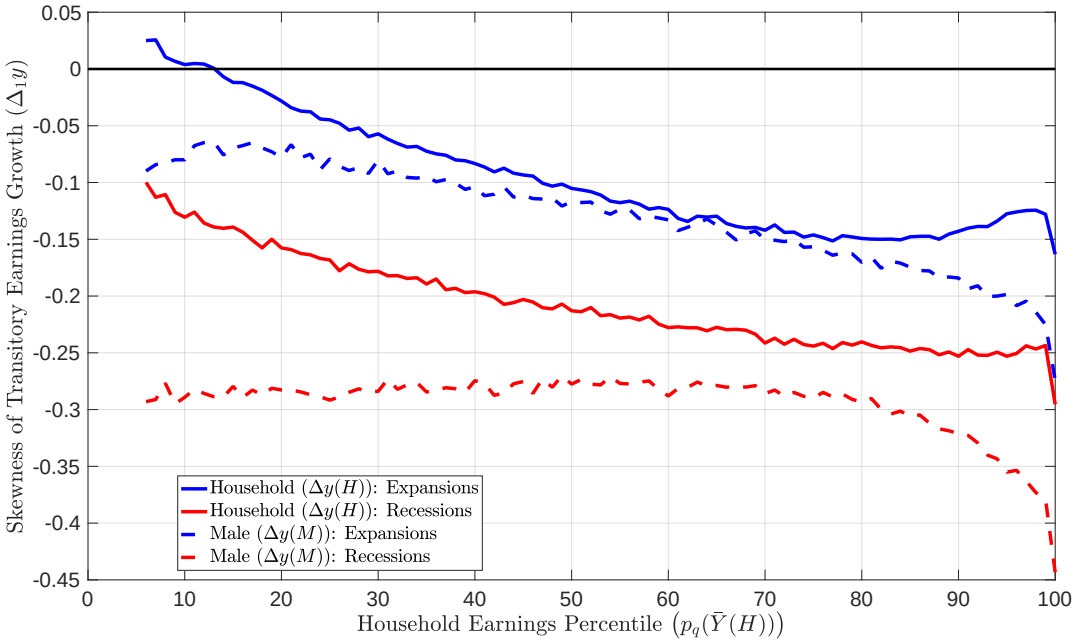

Figure A.8: Skewness of Household and Male Earnings Growth, Transitory

Notes: Measures of the skewness of transitory household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$. Blue lines are for expansions and red lines are for recessions. Household earnings $\Delta y(H)$ skewness shown as solid lines, and male earnings $\Delta y(M)$ skewness shown as dashed lines.

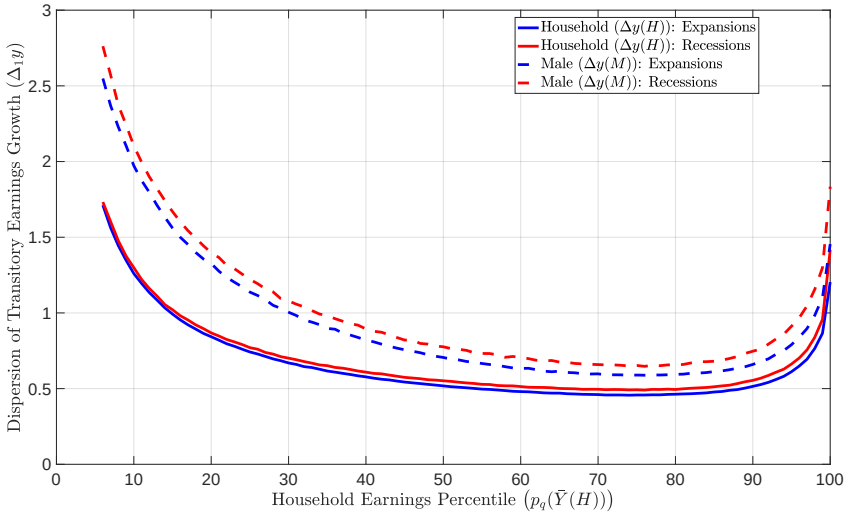

Figure A.9: Dispersion of Household and Male Earnings Growth, Transitory

Notes: Measures of the dispersion of transitory household earnings growth $\Delta y(H)$ and male earnings growth $\Delta y(M)$.

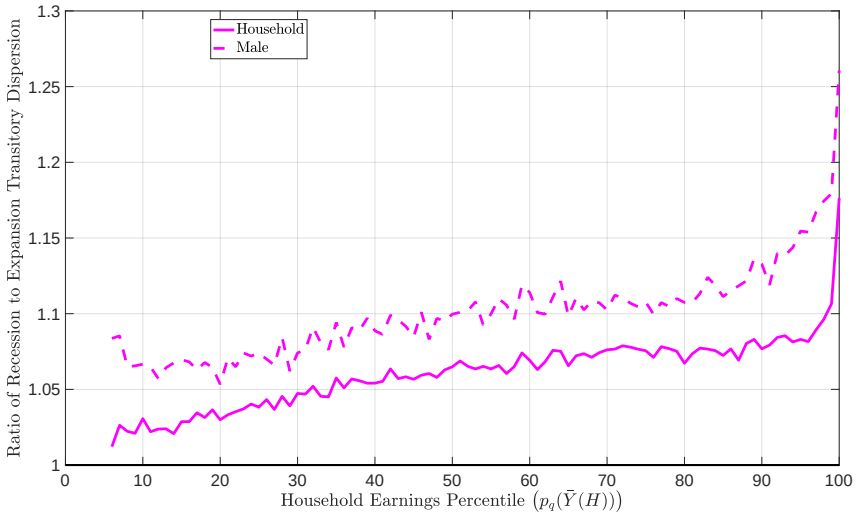

Figure A.10: Countercyclical Ratio of Dispersion, Transitory

Notes: The ratio of transitory earnings growth dispersion in recessions to earnings growth dispersion in expansions. For households and males, separately.