Online Appendix for
"Earnings Risk in the Household:
Evidence from Millions of U.S. Tax Returns"
American Economic Review: Insights
by Seth Pruitt and Nicholas Turner

OA.1 Marriage

In order to abstract from earnings growth that results from changes in marital status, our baseline analysis conditions on families that remain married. 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 and show that our baseline results are qualitatively similar if we relax this restriction.

Figure OA1 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 dotted lines, the 90th in dashed lines and the 10th percentile in solid 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, the bottom panel of Figure OA1 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 earnings percentiles. There is generally an upward slope for each earnings change percentile moving across the long run earnings percentiles in both figures. This suggests that households with higher long run earnings are less likely to get divorced compared to households with lower long-run earnings 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

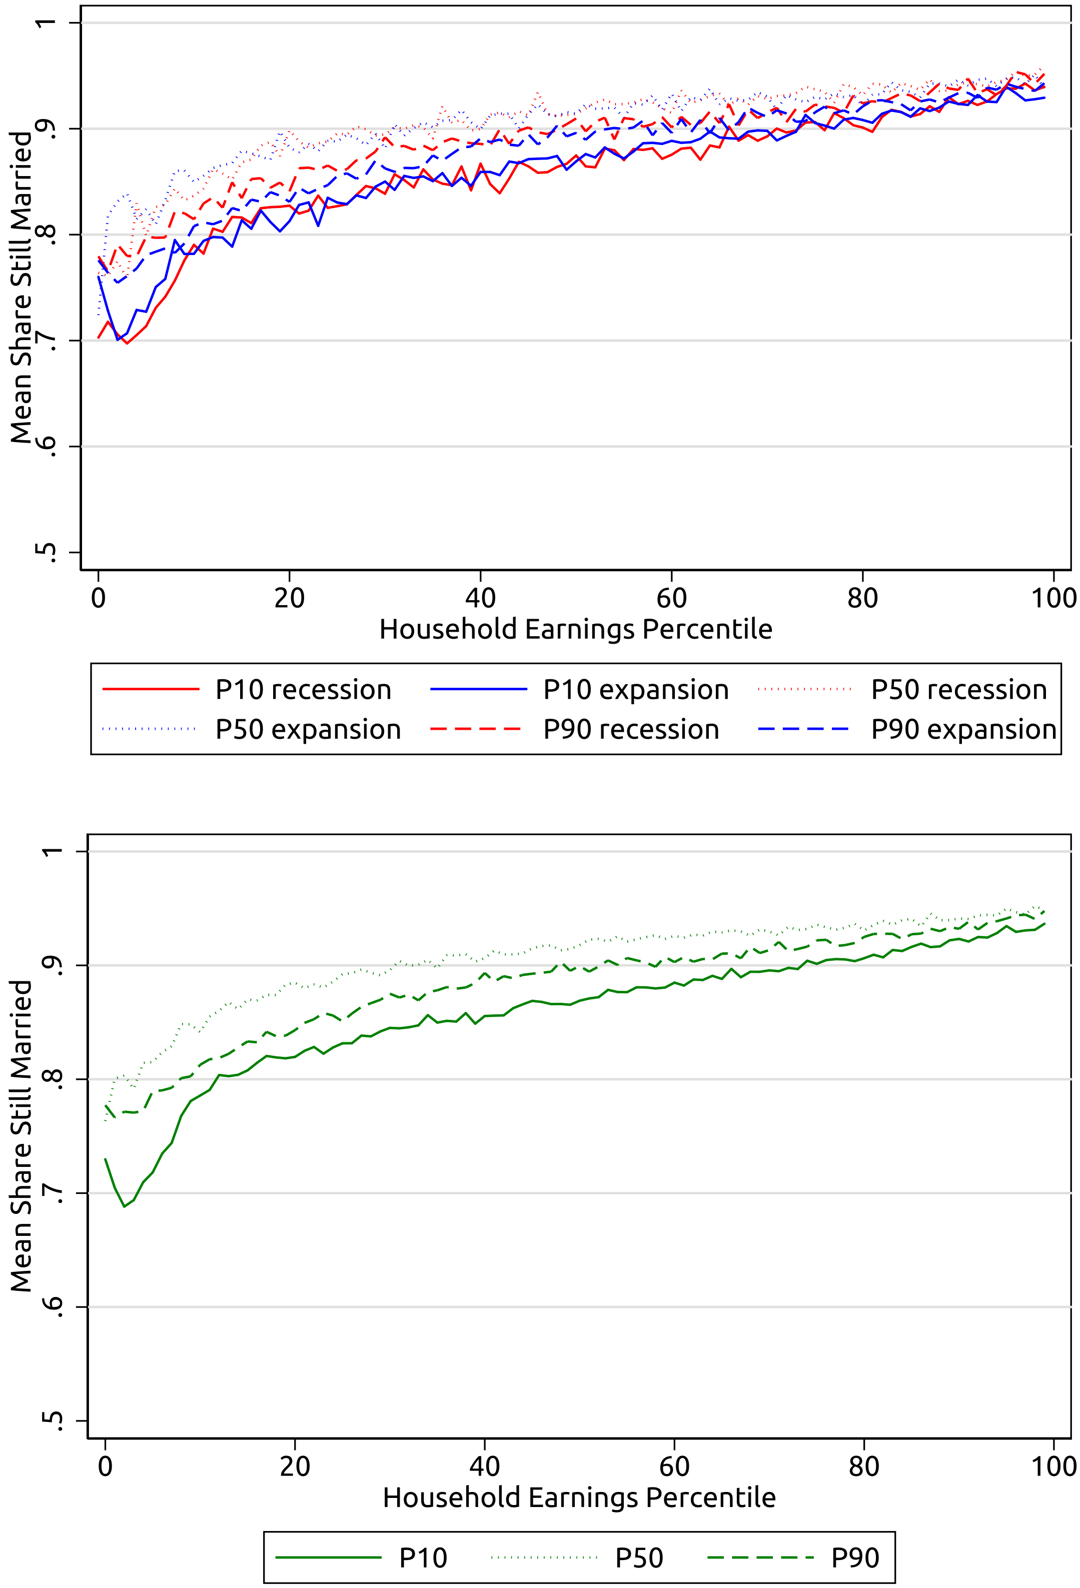

Figure OA1: Mean Share of Households Still Married After Four Years

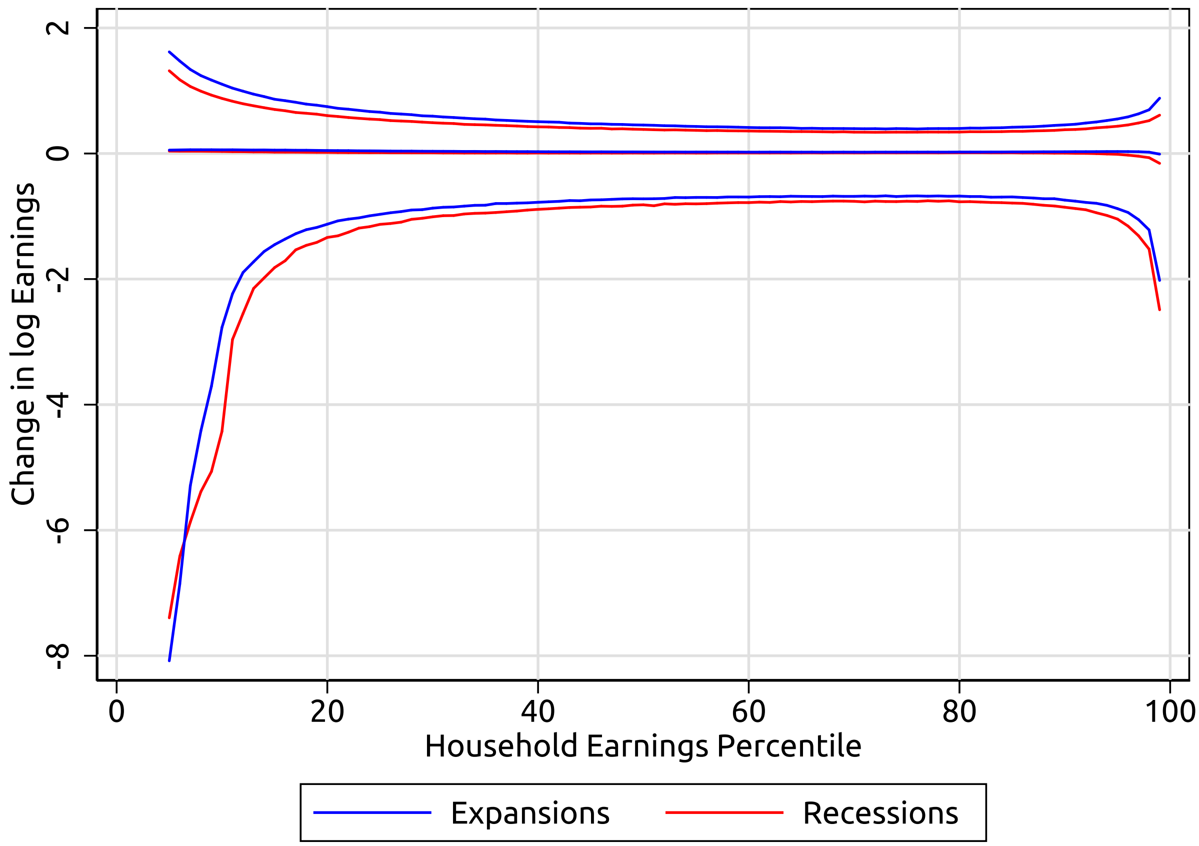

Figure OA2: Percentiles of the Household Persistent Earnings Growth Distribution Unconditional on Remaining Married

Notes: This is an analog to Figure 1 of the main text.

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 at the 90th percentile of male earnings change may reflect divorce induced male earnings shocks, although it is not possible to determine 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 earnings 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 (earnings shock induced divorce) and that there is less need to increase earnings among males following divorce (divorce induced earnings 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 $\bar{Y}(H)$.

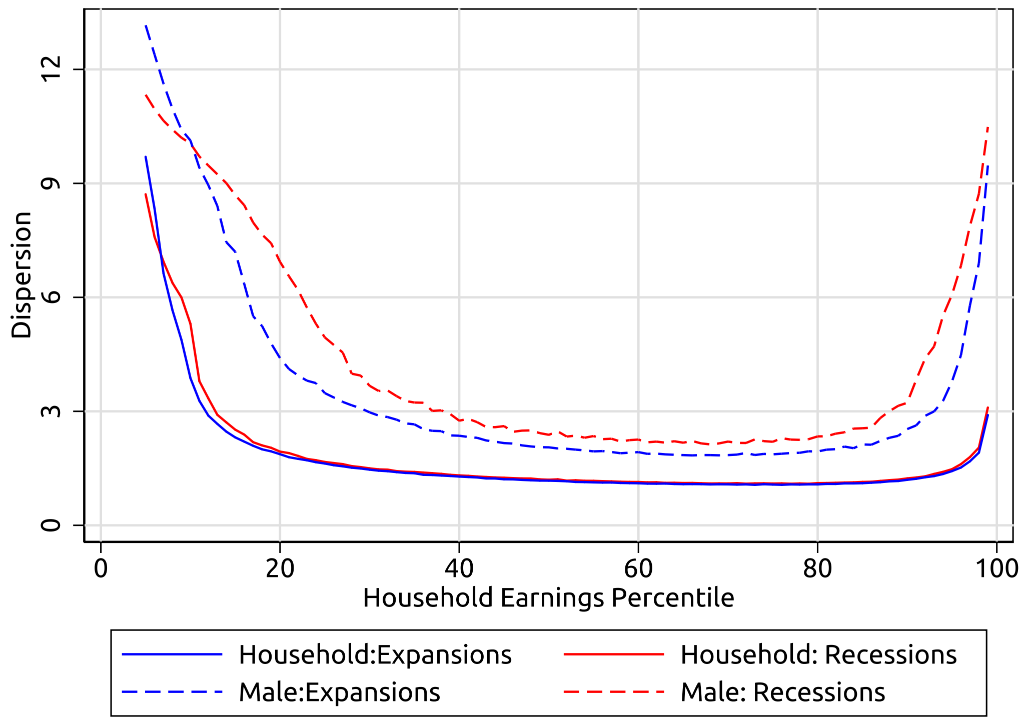

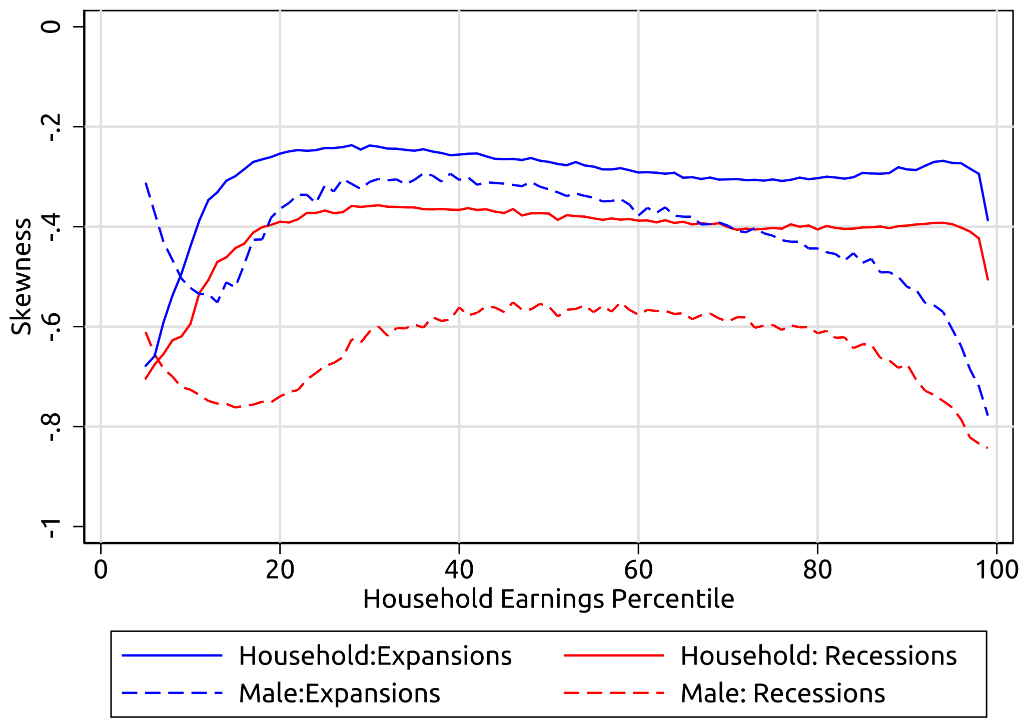

Figure OA3: Dispersion and Skewness of Household and Male Persistent Earnings Growth Unconditional on Remaining Married

Notes: These are analogs to the top panels of Figures 3 and 4 of the main text.

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 from the exit of a working spouse. Figure OA2 shows the change in household log earnings over the long run earnings 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 earnings 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 OA3 plots dispersion and skewness for comparison with the figures in the main text.

OA.2 Data

Source data We draw on population level data from U.S. tax records. Our sample includes households with prime age males who come into contact with the tax system as either tax filers, or non-filers who receive W2 forms. In order to focus on males in households with a strong attachment to the labor force we make several sample selections, including dropping person-year observations starting two years prior to the year of death and individuals who receive disability income. We require the individual to be present in at least one prior year in order to define the long run earnings measure. For individuals who meet these criteria, we draw a 1 in 5 random sample of all males (and their spouses if present), resulting in a sample of 235,552,817 observations at the person year level that span the period from 2000-2014. The key figures in the main analysis are based on the following aggregated data sets.

Aggregated data The main datasets are `DELTA1.xlsx` and `DELTA4.xlsx`, where 1 and 4 reflect the number of years in between earnings measurements: `DELTA4.xlsx` contains the data on persistent earnings growth that we use in the manuscript, and `DELTA1.xlsx` contains the data on transitory earnings growth that we use in the appendix. `DELTA1.xlsx` spans the base years (t) 2000-2013 (delta years ($t+1$) 2001-2014) and include one-year log-changes for persons aged 26-59 in the base year. `DELTA4.xlsx` spans the base years (t) 2000-2010 (delta years ($t+4$) 2004-2014) and include four-year log-changes for persons aged 26-56 in the base year. The datasets are aggregated at the percentile and year level. Percentiles range 0-99. The minimum number of counts in each cell is 25. Cells with fewer than 25 observations are dropped in order to meet disclosure rules and insure confidentiality.[1] In addition, the percentiles we report are averages across the actual percentile value and the four other closest values so that we do not report an actual value from a tax return. The difference between this averaged percentile value and the actual percentile value is small in all cases.

Percentile measures are defined in the following ways:

[1] For most datasets, this restriction is not binding. In some cases the percentiles are missing/not defined for the change in log earnings because there are many observations with no change in earnings, resulting in a relatively large mass at that point in the distribution.

ybar_pct Percentiles of long-run earnings are defined by age and year
deltam_pct Percentiles of male earnings changes are defined by year, *within* each ybar_pct[2]

1. DELTA4.xlsx: Observations are at the base_year and ybar_pct level. The data includes only observations that file joint returns in the base year as well as in the delta year. This dataset contains a variety of earnings change statistics ("deltas") with the following naming conventions: delta4_X_S, where X includes m for males, f for females, hh for households, and cfhh is for counterfactual households. The suffix S denotes the statistic, where _p is for percentile (followed by the percentile value including 10, 25, 50, 75, or 90), _std is for standard deviation, _mean is for mean. The dataset also includes a count variable and the mean of the household-size adjusted long-run earnings ybar.

2. DELTA1.xlsx: This data looks identical to DELTA4.xlsx, except that the the "delta" variables are calculated from 1-year growth rates in the underlying microdata. The composition of the percentile bins is different from DELTA4.xlsx because we condition on married status in a different delta year ($t + 1$ instead of $t + 4$).

3. ENTRYEXIT4.xlsx: Observations are at the base_year, ybar_pct, and delta4_m_pct level. The data includes only observations that file joint returns and remain married in the base year as well as in the delta year. Key variables in this dataset include: delta4_m the average 4-year earnings growth for males in the bin and work_x_y the share of the bin where the female transitioned from employment state x to employment state y, where 1 means employed and 0 means non-employed

4. INTENSIVE4.xlsx: Observations are at the base_year, ybar_pct, and delta4_m_pct level, as in ENTRYEXIT4.xlsx, but conditional on the household having a female who is employed in the base year and delta year. The key variable is delta4_f, the average 4-year earnings growth for females in the bin.

5. ENTRY4.xlsx: This data looks identical to INTENSIVE4.xlsx, except conditioning on the household having a female who is not employed in the base year and employed in the delta year.

6. EXIT4.xlsx: This data looks identical to INTENSIVE4.xlsx, except conditioning on the household having a female who is employed in the base year and not employed in the delta year.

7. ENTRYEXIT1.xlsx looks identical to the corresponding "4" version above, but report 1-year growth variables instead. The same compositional statement applies as in the description of DELTA1.xlsx.

[2] Our finding are very similar if instead we define these percentiles by age and year within each long-run income percentile and then average across ages. We use our baseline specification to reduce computation time.

8. `INTENSIVE1.xlsx` looks identical to the corresponding "4" version above, but report 1-year growth variables instead. The same compositional statement applies as in the description of `DELTA1.xlsx`.

9. `ENTRY1.xlsx` looks identical to the corresponding "4" version above, but report 1-year growth variables instead. The same compositional statement applies as in the description of `DELTA1.xlsx`.

10. `EXIT1.xlsx` looks identical to the corresponding "4" version above, but report 1-year growth variables instead. The same compositional statement applies as in the description of `DELTA1.xlsx`.

OA.3 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 OA4, which plots the 90th, 50th and 10th percentile earnings shocks over the long run earnings distribution. This figure retains many of the key findings from the comparable analysis of persistent earnings shocks in the main text, including: relatively small 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 OA5 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, 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.

Figure OA6 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, the figure suggests that transitory skewness is a widespread phenomena occurring at virtually all earnings levels in both recessions and expansions for households and for males. At the low end of the earnings distribution this does not hold, as some lower earnings households have positive skewness. Figure OA6 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.

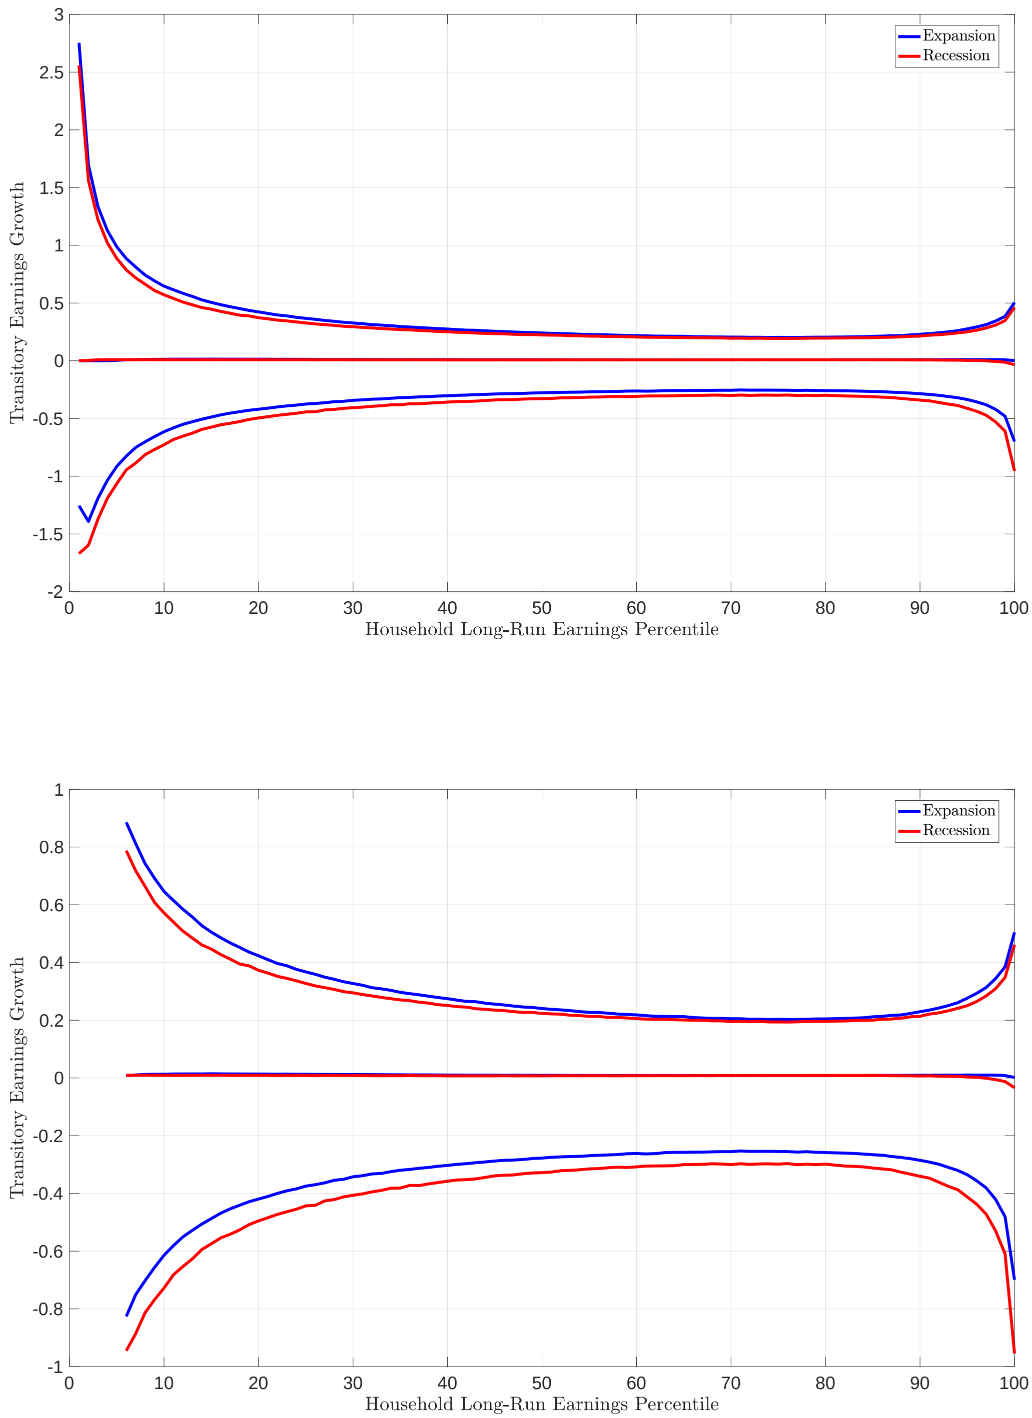

Figure OA4: Percentiles of the Household Transitory Earnings Growth Distribution

Notes: The 10th, 50th, and 90th percentiles of the transitory household earnings growth $\Delta_1 y(H)$ 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.

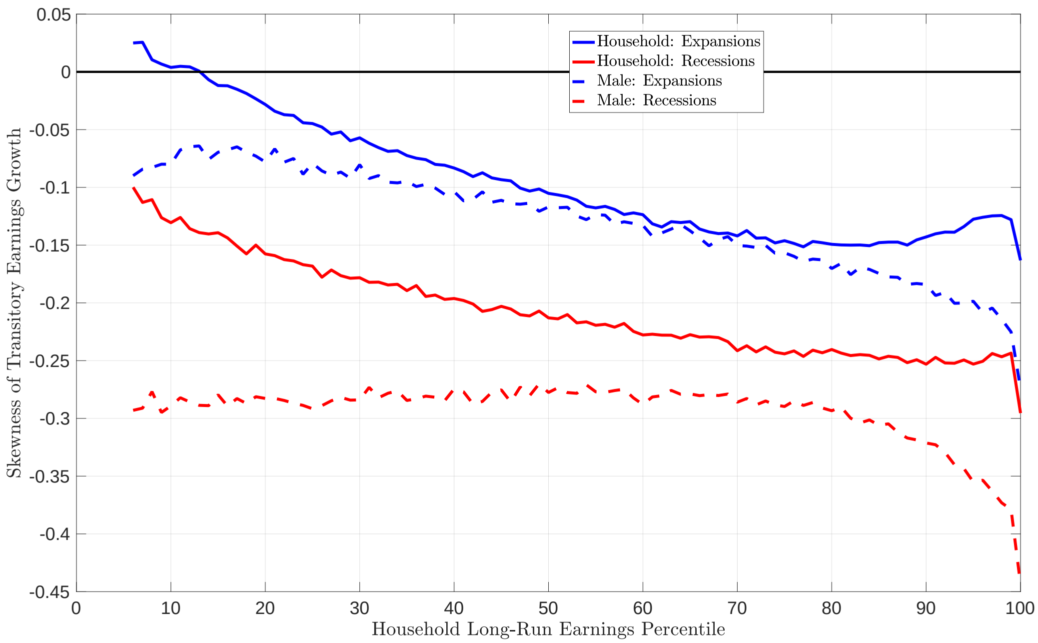

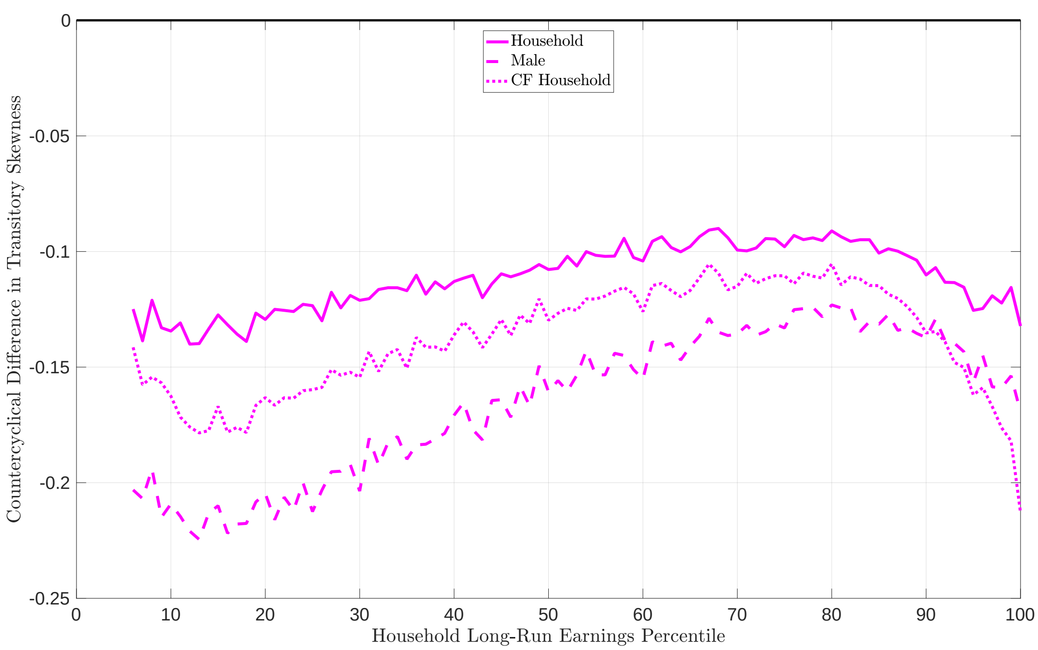

Figure OA5: Dispersion of Household and Male Transitory Earnings Growth

Notes: Measures of the dispersion of household earnings growth (solid lines), male earnings growth (dashed lines), and counterfactual (CF) household earnings growth (dotted lines). Counterfactual household earnings are constructed by assuming zero earnings growth for the spouse. *Panel A:* Blue lines are for expansions and red lines are for recessions. *Panel B:* The ratio of earnings growth dispersion in recessions and in expansions.

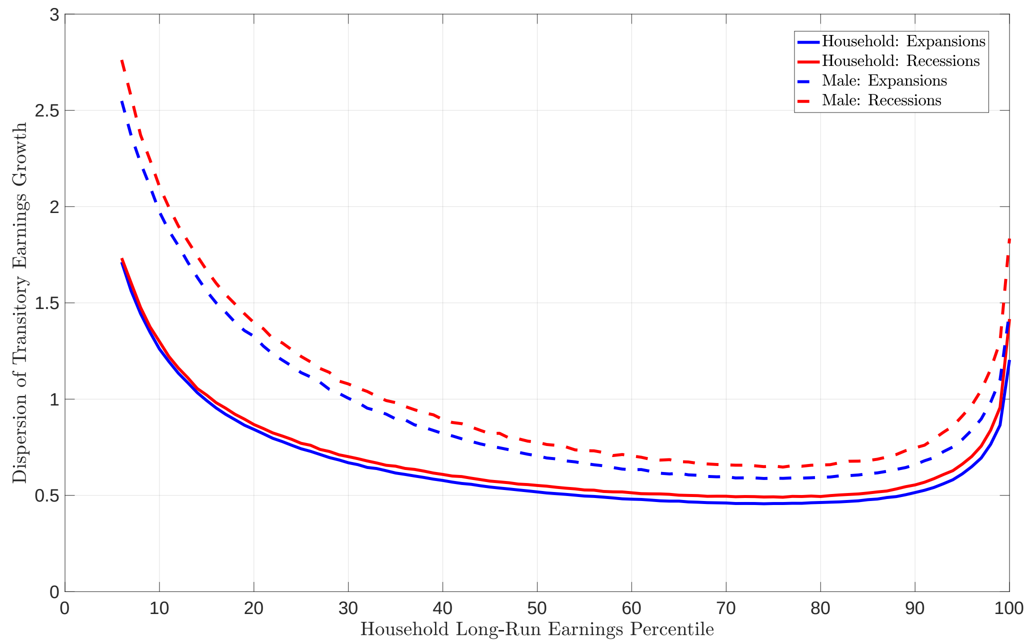

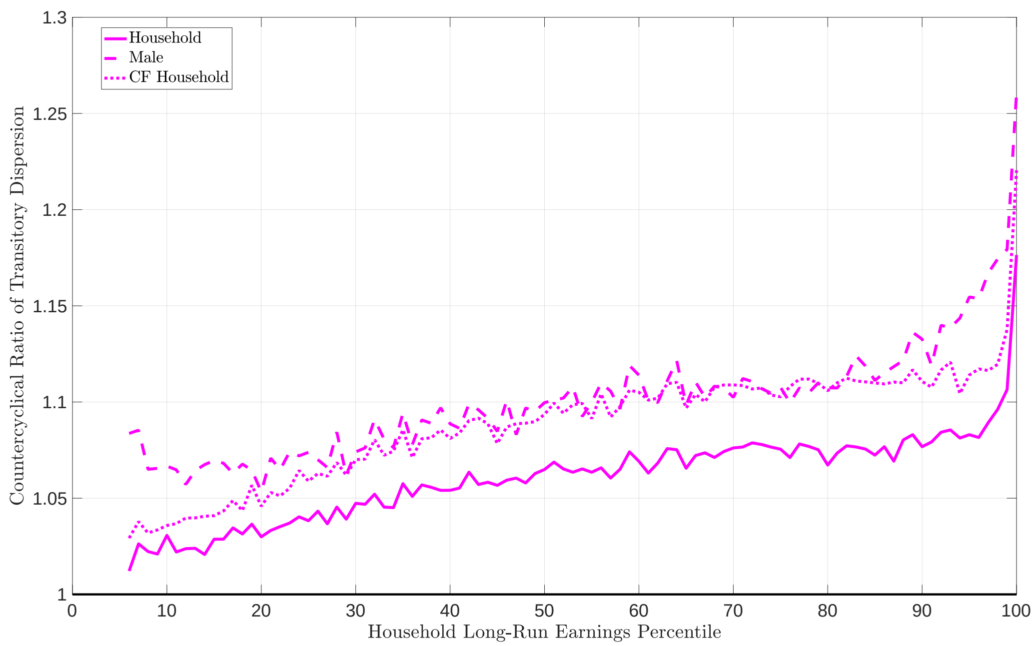

Figure OA6: Skewness of Household and Male Transitory Earnings Growth

Notes: Measures of the skewness of earnings growth (solid lines), male earnings growth (dashed lines), and counterfactual (CF) household earnings growth (dotted lines). Counterfactual household earnings are constructed by assuming zero earnings growth for the spouse. *Panel A:* Blue lines are for expansions and red lines are for recessions. *Panel B:* The difference between earnings growth skewness in recessions and in expansions.

OA.4 Treatment of $0 observations

When quantifying earnings growth, measured as the change in log earnings, we include in the sample observations with $0 values by recoding these cases to $1 before taking the log transformation. This decision differs from Guvenen, Ozkan and Song (2014), who instead limit the sample to men with earnings that exceed a minimum threshold (roughly $1,300 in $2005) in each year used to calculate earnings growth.[3] Guvenen, Ozkan and Song (2014) make this decision because they focus on males with strong labor force attachment. As a result, their findings reflect patterns for men who remain employed for at least part of the year. In contrast, our focus is on the entire household, where movement into and out of employment are potentially important factors affecting household earnings risk. In order to include these factors, we allow observations with $0 to enter the sample, using the recoding to $1 to insure the growth rates are defined.

We are unable to directly analyze the effect of this recoding in our analysis sample because we no longer have access to the underlying taxpayer-level data. However, below we explore the effects of this decision using an alternate dataset that also draws from administrative earnings data. We use data from the 2006 Earnings Public Use File (EPUF) from the Social Security Administration.[4] In related work, Pruitt and Turner (2019), we use these data to explore both male and female labor earnings risk since the 1950s and analyze the effect of increased female employment on earnings risk.

The EPUF data are limited in several important ways relative to the tax data used in the main analysis. First, the EPUF data are top-coded at the Social Security earnings cap. We address this limitation by focusing on percentiles in the distribution where top-coding is not generally binding, analyzing the interquartile range.[5]Second, the EPUF data are at the individual level, and therefore cannot be used to construct households. As a result, in the analysis below we consider only males. Third, the data end in 2006. To align as closely as possible with the period considered in the main results, we analyze earnings growth starting in 1996 and we show results for transitory (one year changes) so that we can include data for year t through 2005. Despite these limitations, the EPUF has several important advantages. First, it is drawn from administrative sources, making it less susceptible to measurement error relative to survey data. Second, it is relatively large, including a random sample of 1 percent of all SSNs issued prior to 2007. The analysis below draws from 5.3 million person-year observations for prime age male earners between 1996 and 2005.

For this analysis, the only sample selection we impose (beyond the age restriction) is to exclude observations in year t that have zero values of long-run earnings. As shown in Figure

[3]This is not the only difference in the research designs. Relative to Guvenen, Ozkan and Song (2014), we also use a different method of controlling for life-cycle earnings effects, and we analyze a different set of years.

[4]Data are available here: `https://www.ssa.gov/policy/docs/microdata/epuf/index.html`

[5]Between 1996 and 2005, top-coding affects just over 4 percent of observations at the 75th percentile of long run earnings, less than 1 percent at the 50th percentile, and less than 0.5 percent at the 25th percentile. This level of top-coding will not affect our outlier robust statistics that rely on the 10th and 90th percentiles to characterize the tails of the distributions.

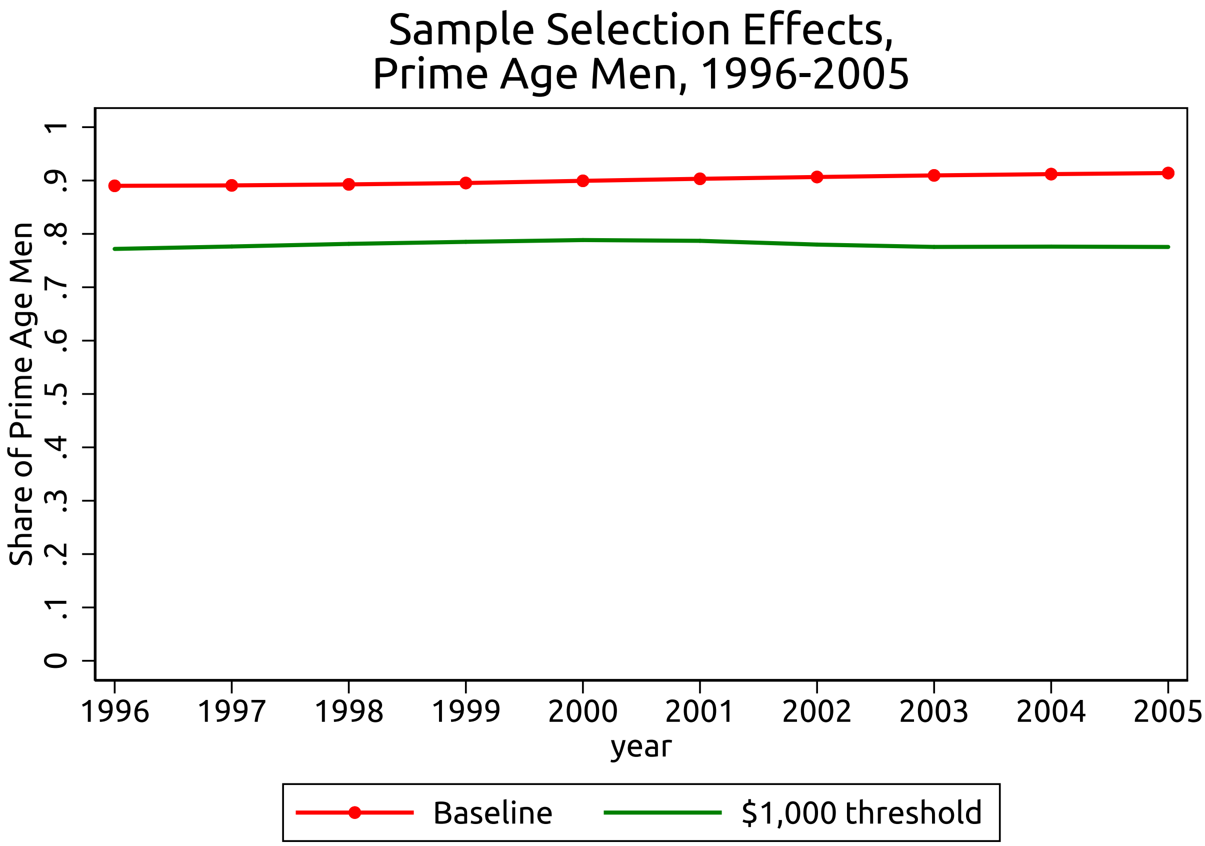

Figure OA7: Sample selection

Notes: The figure plots the share of observations retained under each sample selection approach, relative to all observations that meet the age restriction.

OA7, this results in a sample that includes just over 90% of all prime-age observations in most years. Imposing a threshold of $1,000, similar to the approach used by Guvenen, Ozkan and Song (2014), would result in a smaller sample, including roughly 75% of observations in each year.[6]

[6]This level of attrition is consistent with Figure 2 in Guvenen, Ozkan and Song (2014).

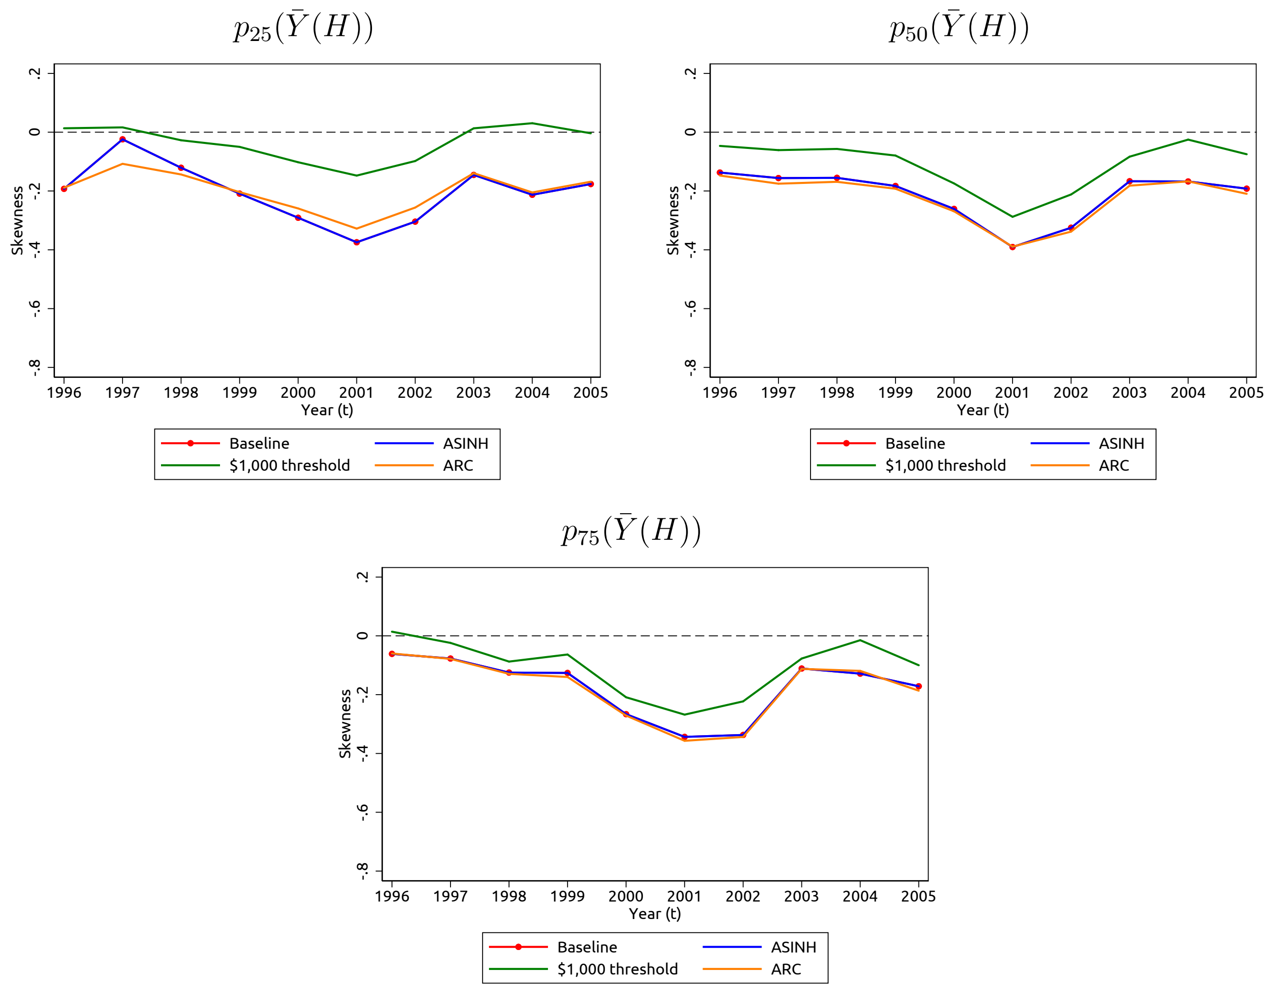

Figure OA8: Comparison of Skewness Measures, Transitory Earnings Growth

Notes: The figures show skewness in the EPUF data using different methods to measure earnings growth.

Figure OA8 shows skewness for three different long-run earnings percentiles, the 25th, 50th and 75th, using alternate methods to quantify earnings growth. In these figures, we hold fixed the long-run earnings percentiles to those defined for our baseline sample.[7] The red series show the results for our baseline method that retains all observations and recodes \$0 to \$1 before taking logs, while the other colors show results for different methods of handling low earnings observations. The green series show the results using a \$1,000 earnings threshold in t and $t+4$ for inclusion in the sample. The blue series uses the inverse hyperbolic sine transformation, which is defined for \$0 values (equal to 0) and the orange shows the arc percent growth rate, which allows for zero values in one period and is bound between -2 and $+2$. These transformations are used in a variety of contexts in the literature. For example, Johannesen et al. (2018) use the hyperbolic sine transformation when analyzing reported earnings to allow for zero and negative values, while the arc percent growth rate is commonly used to analyze labor flows Davis, Haltiwanger and Schuh (1996).

In each panel, skewness values are higher in the baseline approach (red lines) relative to imposing a \$1,000 threshold (green lines). Intuitively, this threshold compresses the earnings growth distributions by removing cases of both large earnings growth (from below the threshold to large positive values, including entry to employment) and large earnings losses (from large positive values to values below the threshold, including exit from employment). In general, our baseline approach gives very similar skewness values compared to using the inverse hyperbolic sine transformation to measure earnings growth. The baseline approach is also similar to the results using the arc percent growth rates at the 50th and 75th percentile. Slight differences between the baseline approach and the arc growth rate at the 25th percentile of the long run distribution occurs because there are relatively more \$0 earners at the 25th percentile and the arc percent and baseline approaches differ in the treatment of these cases. Specifically, the arc percent formula bounds extensive entry and exit to have values of 2 and -2, whereas the baseline approach does not bound these cases. The arc percent formula is also not defined if both values are zero, whereas the baseline approach calculates a growth rate of zero in these cases. At the 25th percentile of long-run earnings, the mean value of working in both periods is 79 percent, while 10 percent have extensive movements into or out of employment. By comparison, at the 50th (75th) percentile the comparable statistics are 93 percent (96 percent) and 4 percent (2 percent). Table OA1 provides the mean values of not working for each of the three percentiles across the years 1996-2005.

OA.5 Cumulative distribution of long-run household earnings

Figure OA9 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 (Piketty and Saez (2003)) and wealth (Saez and Zucman (2016),

[7] Given this restriction, the percentiles are of equal size only for the baseline and the ASINH results. In the other cases, the measures are not defined for all observations, so that the size of the percentiles differ.

Table OA1: Share of Observations with Zero Earnings

	Long-Run Earnings Percentile		
	25	50	75
1996	0.15	0.04	0.02
1997	0.15	0.04	0.02
1998	0.14	0.04	0.01
1999	0.13	0.04	0.02
2000	0.14	0.04	0.02
2001	0.15	0.04	0.02
2002	0.16	0.05	0.02
2003	0.17	0.05	0.03
2004	0.18	0.05	0.03
2005	0.18	0.06	0.03

Notes: The share of observations with zero earnings in the EPUF data.

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. Values of household earnings increase dramatically moving up the $\bar{Y}(H)$ distribution: 5th percentile $9,500; 10th percentile $17,500; 25th percentile $34,000; 50th percentile $54,500; 75th percentile $79,500; 90th percentile $112,500; 95th percentile $143,500. Households in the top 1% have on average $553,000 in long-run earnings.

OA.6 Spousal earnings: nonparametric regressions

Our results suggest that households mitigate male labor earnings risk. But what are the important channels that spouses use? We use nonparametric kernel regressions to understand spousal responses. In this section, we describe our empirical procedure.

OA.6.1 Decomposing spousal labor supply responses

In a given year t, a spouse in the household is either employed or not employed. Looking ahead to $t+4$ (or $t+1$), this creates four categories of spouses: those who enter the labor market, those who exit the labor market, those who stay in the labor market, and those who

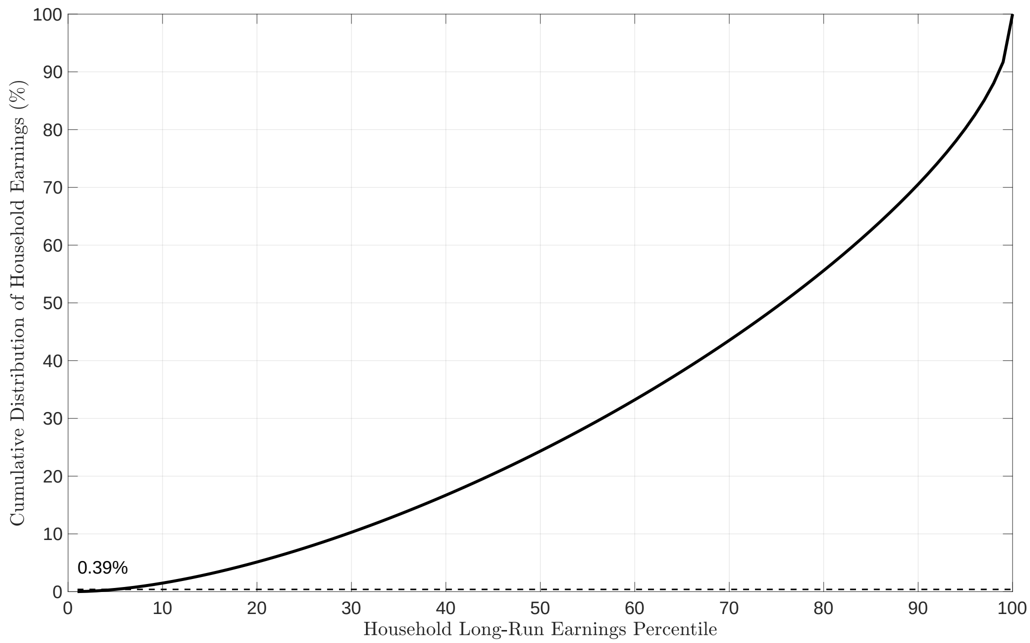

Figure OA9: Aggregate Household Earnings Cumulative Distribution

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 0.39% noted is the cumulative amount of aggregate household earnings represented by the bottom ventile of the household earnings distribution.

stay out of the labor market.[8] In each of these cases, we characterize both the share of spouses in that category and the earnings of the spouse.[9] We denote the shares in the following way: $\theta(S, E)$ is the share of non-working spouses who *enter* employment, $\theta(S, X)$ is the share of spouses who *exit* employment, $\theta(S, R)$ of spouses who *remain* employed. The share of spouses staying out of employment $\theta(S, O)$ is pinned down by $\theta(S, O) = 1 - \theta(S, X) - \theta(S, E) - \theta(S, R)$.

The $\theta(S, E)$, $\theta(S, X)$, and $\theta(S, R)$ are *unconditional* probabilities of being in one of the three states E, X, 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 earnings 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

[8] 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 spouses 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.

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

so that this margin is a possible outcome. We define the following conditional probabilities

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

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

We furthermore can measure the earnings change associated with the E, X, R states, denoted $\Delta y(S,E), \Delta y(S,X), \Delta y(S,R)$. The quantity $\Delta y(S,E)$ is the average log earnings gained by newly employed spouses. The quantity $\Delta y(S,X)$ is the average log earnings left behind by spouses exiting 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 $\Delta y(S,R)$.

OA.6.2 Empirical method

We characterize spousal responses conditional on the joint distribution of $\log(\bar{Y}(H)) \equiv \bar{y}_t(H)$ 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 earnings. Likewise, we can look at how the effect of a given male earnings shock varies over the long-run earnings 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 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

$$\{\theta(S,E|N), \theta(S,X|W), \theta(S,E), \theta(S,X), \Delta y(S,R), \Delta y(S,E), \Delta y(S,X)\}.$$

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.[10] Li and Racine

[10] 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

(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 and 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.

OA.6.2.1 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" approximately 20,000 observations.[11]

Here is 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.

---

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 and Racine (2007).

[11]Further subject to the rule that we can only release data for bins with at least 25 households.

OA.6.3 Additional results

In Figure OA10 we plot the manuscript's Figure 5, but translate male earnings growth $\Delta_4 y(M)$ to the percentiles of the male earnings growth distribution: this helps us visualize where the population mass is located in Figure 5. Because of the countercyclical skewness we document, notice that -0.7 and $+0.7$ are higher percentiles for recessions than for expansions, and therefore the red and blue lines do not cover the same horizontal values. In Figure OA11 we plot the results for conditional leaving (exit) probabilities. In Figure OA12 we plot the results for unconditional entry and exit probabilities, which are provided for completeness but do not have a natural interpretation as do the conditional probabilities. In Figure OA13 we plot the results for entry and leaving earnings (i.e. how much job entrants and job exiters obtain/leave-behind, respectively).

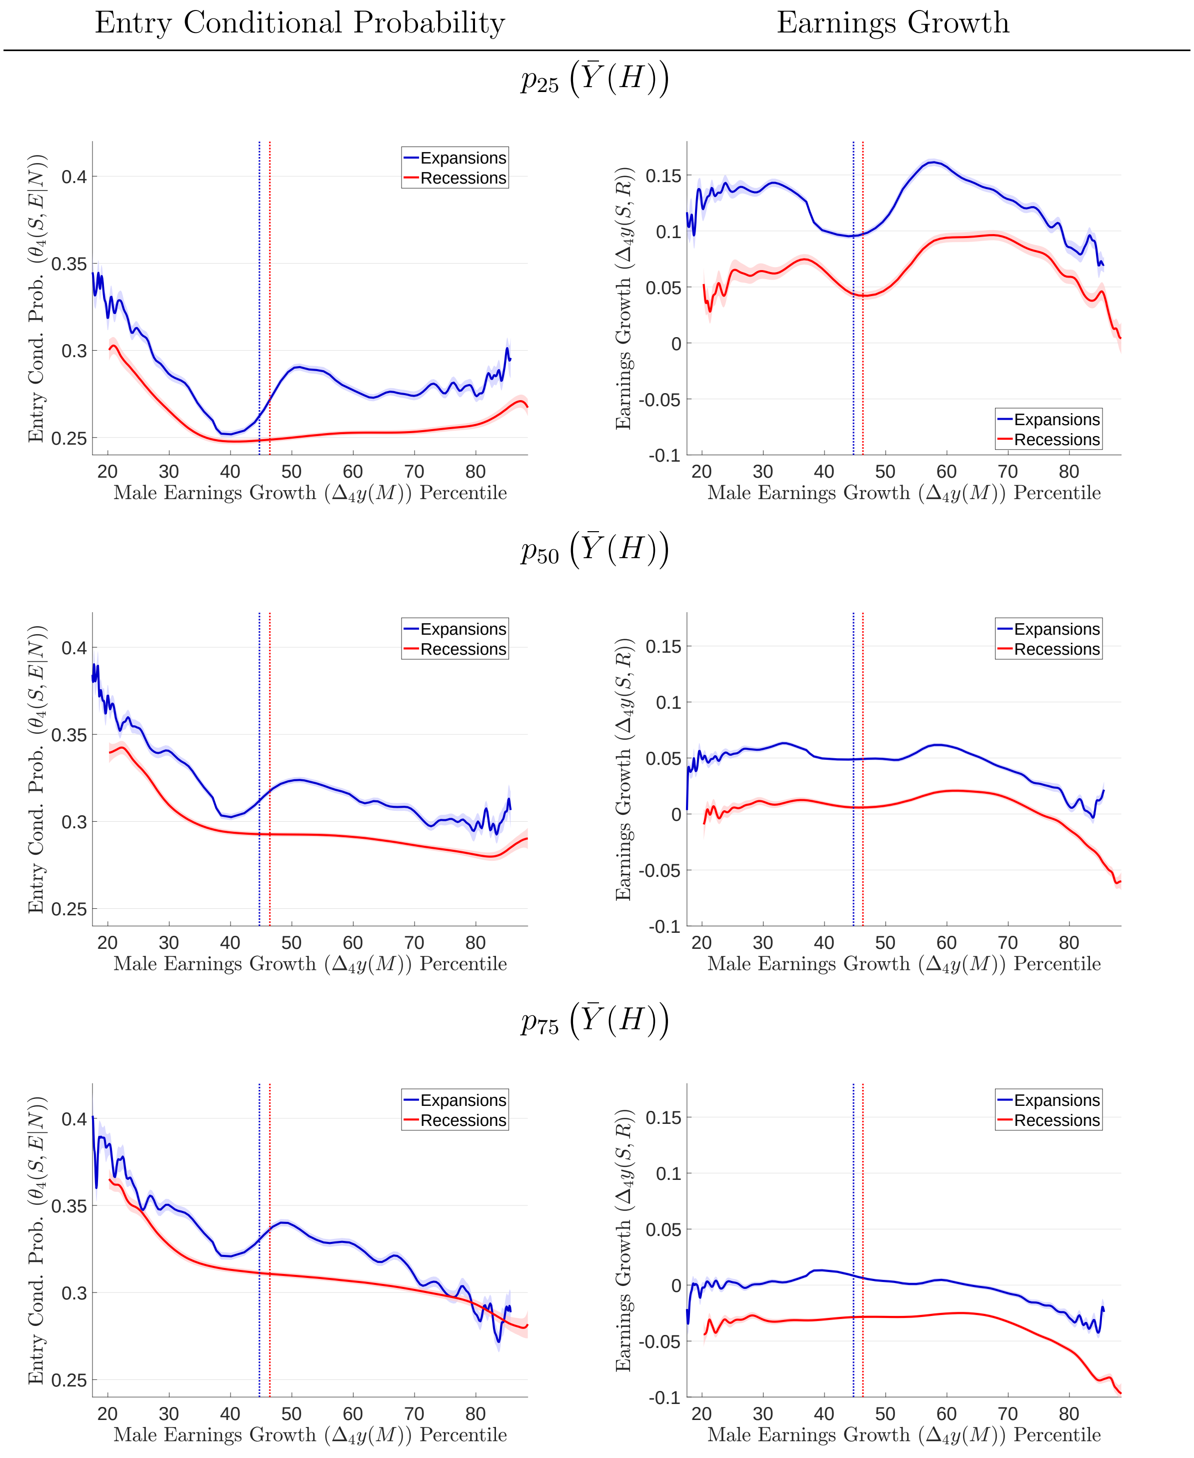

Figure OA10: Spousal Responses: Percentile Scale

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression, for generic spousal response variable $x(S)$. Confidence bands shaded. The left column shows results for $x(S) = \theta_4(S, E|N)$ the probability of non-employed spouses entering employment. The right column shows results for $x(S) = \Delta_4 y(S, R)$ the earnings growth of spouses who remain employed. The three rows report the nonparametric surface at the first, second, and third quartiles of the household earnings level distribution. The vertical dashed lines show where $\Delta_4 y(M) = 0$ for either recessions or expansions. Because of the countercyclical skewness we document, the solid lines do not cover the same horizontal values and the vertical lines are distinct, with the red recessionary vertical line to the right.

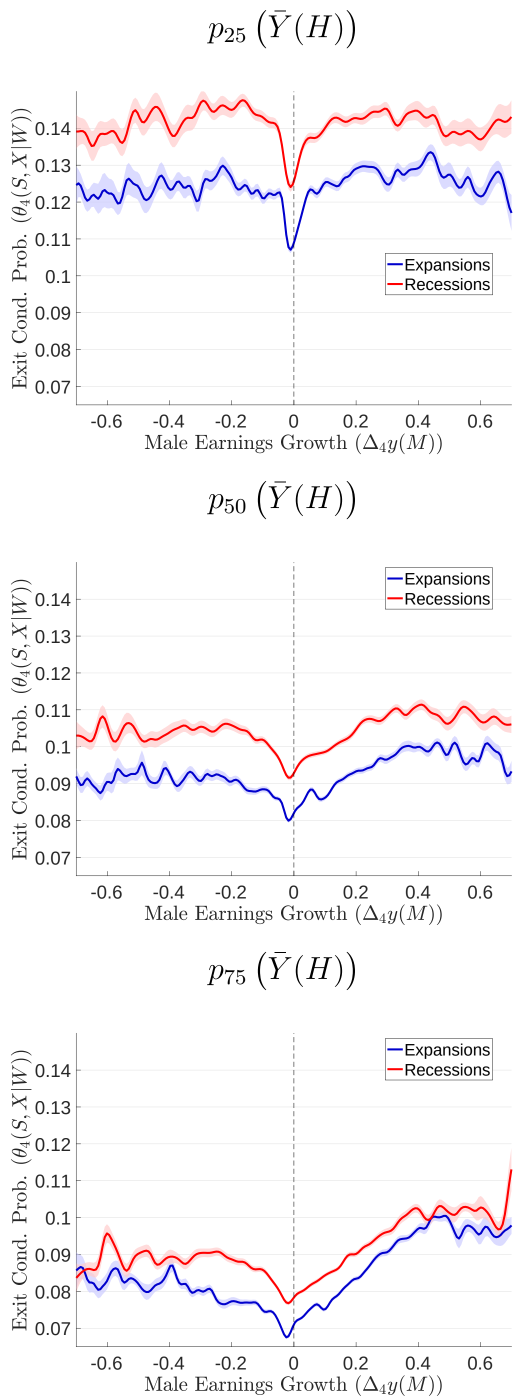

Figure OA11: Spousal Responses: Conditional Exit

Notes: Fitted values of $E(\theta_4(S, X|W)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression, where $\theta_4(S, X|W)$ is the probability of working spouses exiting employment. Confidence bands shaded. The three rows report the nonparametric surface at the first, second, and third quartiles of the household earnings level distribution.

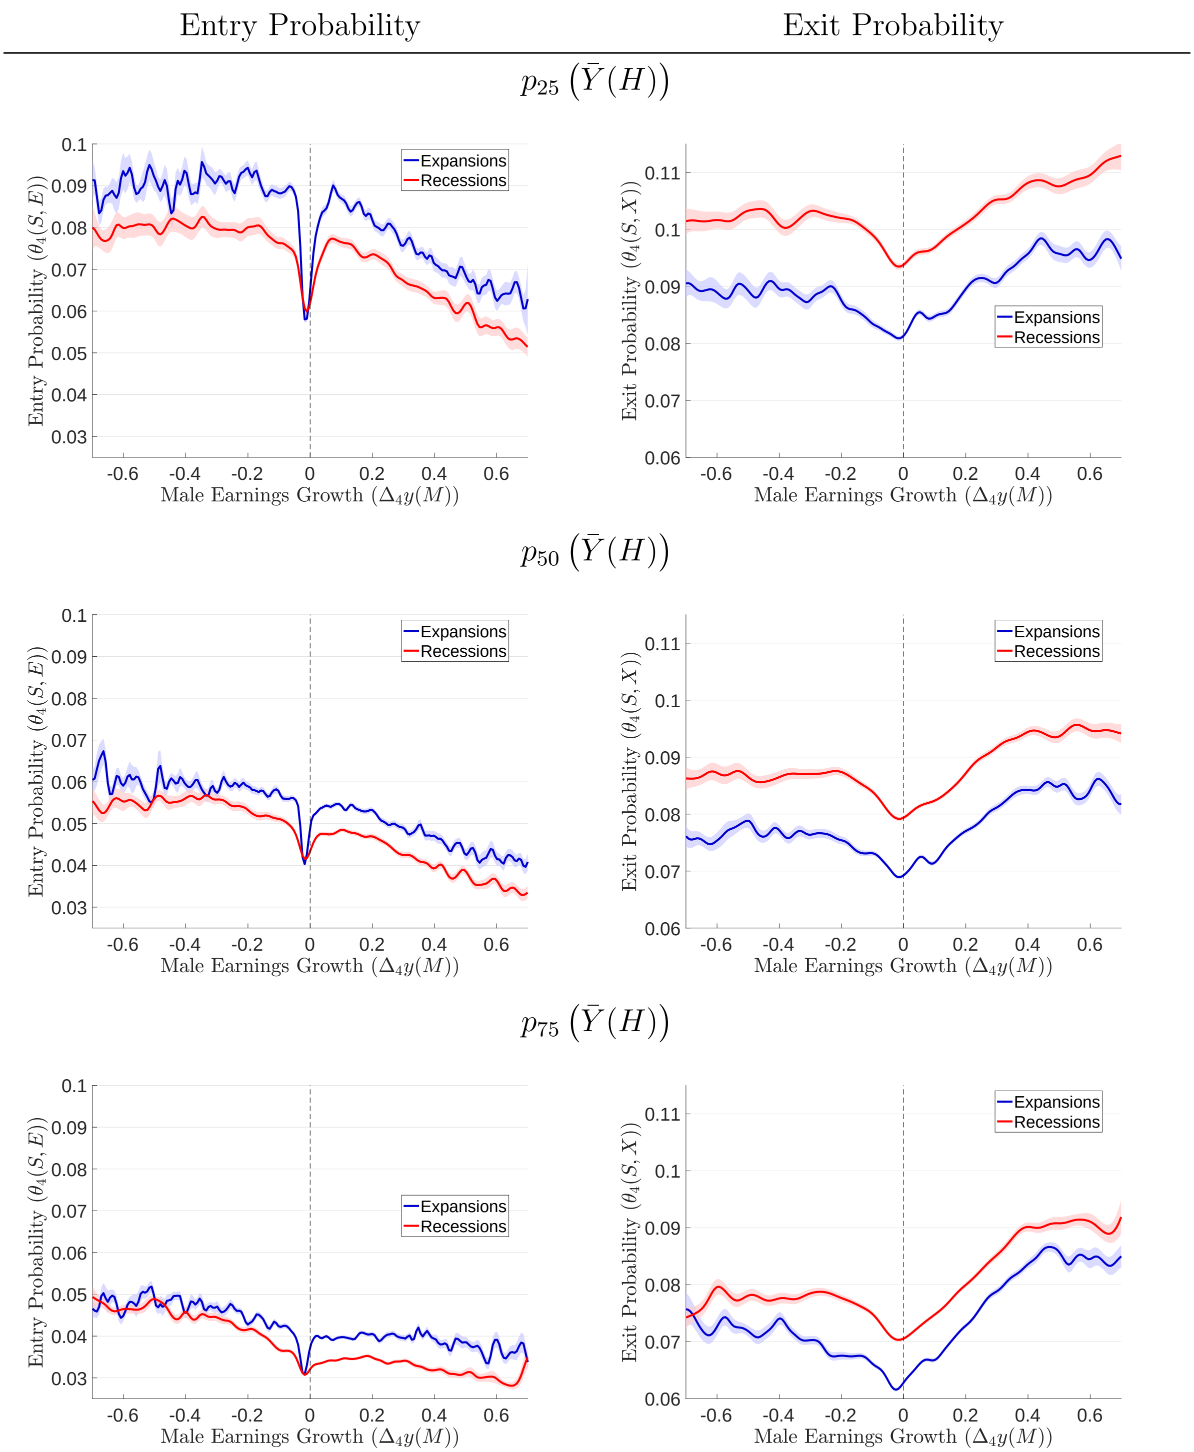

Figure OA12: Spousal Responses: Entry and Exit

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. The left column shows results for $x(S) = \theta_4(S, E)$ the unconditional entry probability. The right column shows results for $x(S) = \theta_4(S, X)$ the unconditional exit probability. The three rows report the nonparametric surface at the first, second, and third quartiles of the household earnings level distribution.

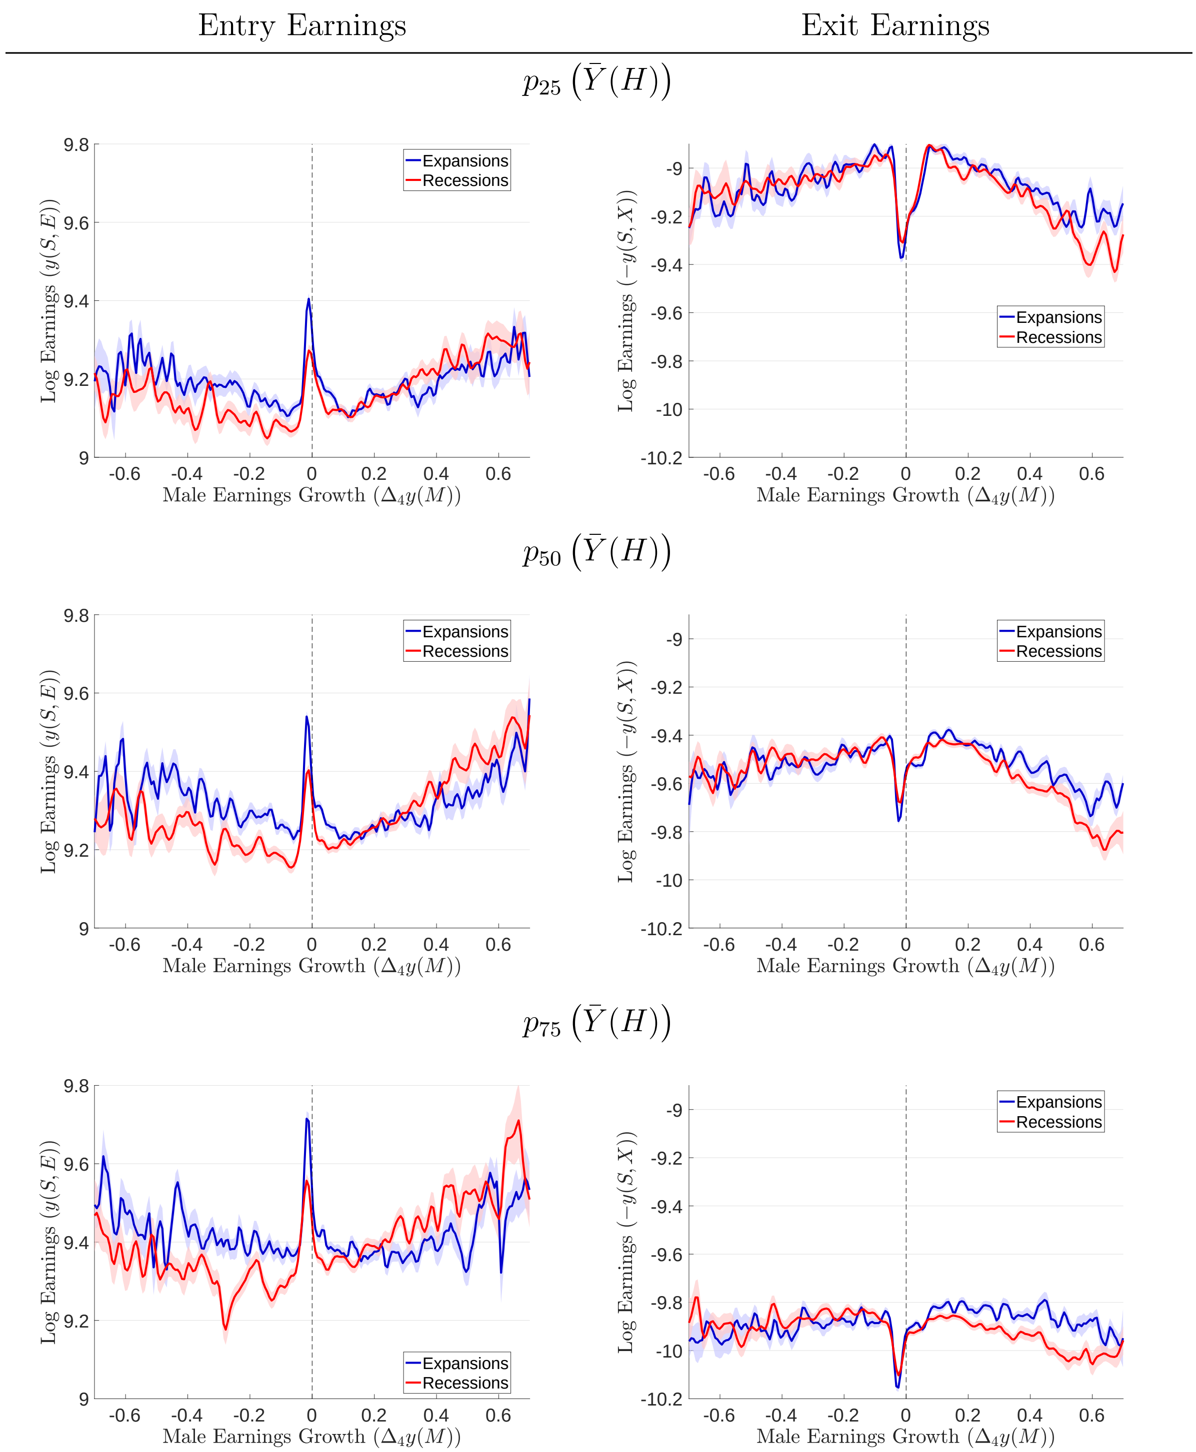

Figure OA13: Spousal Responses: Entry and Exit Earnings

Notes: Fitted values of $E(x(S)|\Delta_4 y(M), \bar{y}(H))$ from nonparametric regression. Confidence bands shaded. The left column shows results for $x(S) = y(S, E)$ the entry earnings. The right column shows results for $x(S) = -y(S, X)$ the exit earnings. The three rows report the nonparametric surface at the first, second, and third quartiles of the household earnings level distribution.

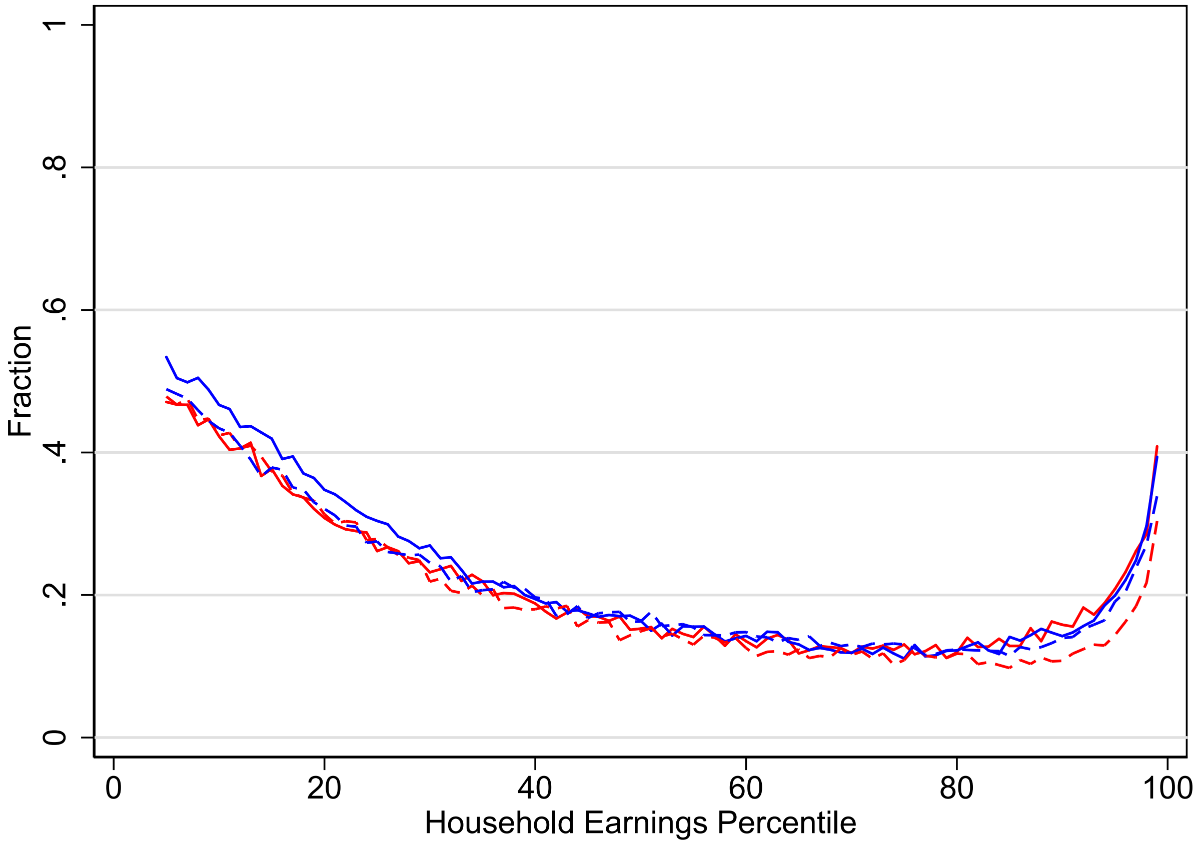

Figure OA14: Sample Selection

OA.6.4 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 exits employment, we restrict attention to households with a working spouse. Figure OA14 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-earnings households are more likely to have a non-working spouse relative to households in the middle of the earnings distribution. There is only modest variation across the male earnings shocks, conditional on earnings, as suggested by the figure at the 25th and 75th percentiles of the male earnings shocks distributions.

OA.7 Welfare analysis

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%).[12]

We therefore have seven Δy values, $p_0, p_{10}, p_{25}, p_{50}, p_{75}, p_{90}, p_{100}$, that define six distinct regions, denoted $p_{10} - p_0$, $p_{25} - p_{10}$, 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 $r > q$.[13] 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.

OA.8 Central tendency of earnings growth

To robustly 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 plot the median (solid) and mean (dashed) of household earnings shocks $\Delta y(H)$ over the business cycle in Figure OA15. 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 earnings and nearly independent of the business cycle. Consider a family at the median level of earnings ($54,500). 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 earnings distribution, the average growth rate implies a loss of 23 log points during

[12]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.

[13]Ie. if $u_1 \in [0, 0.1]$ then earnings growth is drawn from region $p_{10} - p_0$, if $u_1 \in (0.1, 0.25]$ then earnings growth is drawn from region $p_{25} - p_{10}$, etc.

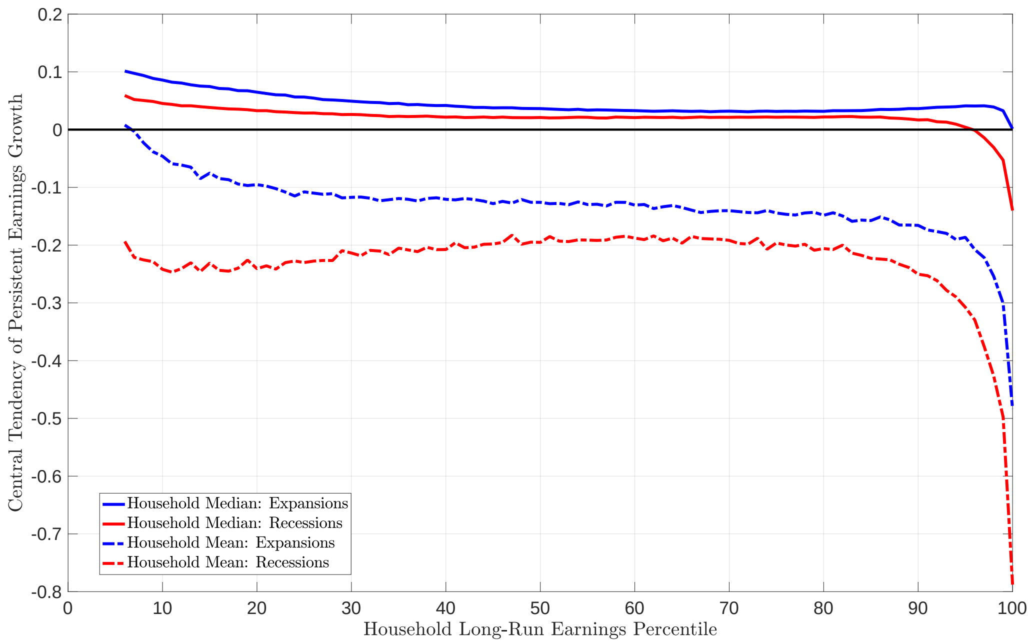

Figure OA15: Median and Mean Household Earnings Growth

Notes: The median and mean of household persistent earnings growth. Blue lines are for expansions and red lines are for recessions. Household earnings median shown as solid lines, and mean shown as dashed lines.

recessions and 13 log points during expansions. The average effect is substantially worse for higher earnings 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 growth for higher earnings 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.

References

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.

Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction.* MIT Press.

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

Johannesen, Niels, Patrick Langetieg, Daniel Reck, Max Risch, and Joel Slemrod. 2018. "Taxing Hidden Wealth: The Consequences of U.S. Enforcement Initiatives on Evasive Foreign Accounts." National Bureau of Economic Research Working Paper 24366.

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

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

Pruitt, Seth, and Nicholas Turner. 2019. "A Half Century of Female and Male Earnings Risk: Evidence from Millions of Individuals." Arizona State University and Federal Reserve Board of Governors.

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.