

**Robustness Appendix for “Deconstructing Lifecycle Expenditure”**  
**Mark Aguiar and Erik Hurst**

This appendix shows a variety of additional results that accompany our paper "Deconstructing Lifecycle Expenditure." In section 1, we discuss the standard errors associated with Figures 1a, 1b, 2a, and 2b of the main paper. In section 2, we discuss the covariances between core nondurables, food-at-home, and work-related expenditures over the lifecycle. In section 3, we explore the robustness of our results to alternate controls for family size. In section 4, we discuss whether good specific measurement error in the Consumer Expenditure Survey can be influencing our results. In section 5, we show the detailed results from our demand system estimates referenced in the text. In section 6, we show the robustness of our main results to unrestricted time effects and zero cohort effects when identifying the lifecycle patterns. Finally, in section 7, we explore the robustness of our disaggregated results to when we exclude individuals with zero spending on the category from our analysis. Broadly, the main conclusions of our paper are robust to these alternate specifications.

**1. Standard Errors for Figures 1 and 2**

Figures R1a, R1b, R2a, and R2b are the same as Figures 1a, 1b, 2a, and 2b from the text, respectively, but with the inclusion of two standard-error confidence bands around the lifecycle profiles. In particular, the dotted lines in Figures R1 and R2 are the lifecycle profile of the regression coefficient plus and minus twice the standard error. Given that the lifecycle profile of the variance is essentially identical with and without housing services, we depict only the case including housing to minimize on clutter. As seen from these robustness figures, the lifecycle profiles documented in the main text are estimated with a great deal of precision. On average,

the coefficients for the mean expenditure over lifecycle have standard errors of about 0.02 implying a two standard deviation range of the estimate of about 8 percentage points (Figures R1a and R2a). The lifecycle coefficients for the cross sectional variance have standard errors of about 0.002 implying a two standard deviation range of the estimate of about 1 percentage point (Figures R1b and R2b). The small standard errors reflect that we have a large sample size and are estimating a series of age-specific means and variances.

## **2. Covariances**

The main text highlights the lifecycle profile of age-specific expenditure inequality; that is, the variance of expenditure across households. The use of disaggregated expenditure also allows the computation of covariances. In Figure R3 we depict the lifecycle profile of the covariances for the three-good disaggregation from Figure 2 in the text. Panel (a) is the unconditional covariance, depicted as a deviation from the value at age 25. Panel (b) depicts the residual covariance, where we take the residuals from the regression of log expenditure on our age, cohort, year, and family size controls and compute the covariances, also as deviations from age 25. This is the same procedure used for Figure 2b in the text. For reference, we report the age-25 levels in the figure notes. In both panels, we see that the covariance across the goods is increasing over the lifecycle, particularly for the covariances involving work-related expenses.

A natural question is how much of the rise in total expenditure inequality is due to the increase in covariances versus the increase in good-specific inequality. Figure R4 provides some insight. In particular, consider the following decomposition of the unconditional total-expenditure variance. We can approximate the log of total expenditure as a weighted sum of its components:

$$\ln\left(\sum_k C_{i,t}^k\right) - \ln\left(\sum_k \bar{C}_t^k\right) \approx \sum_k s_t^k \left(\ln(C_{i,t}^k) - \ln(\bar{C}_t^k)\right).$$

Here, a bar indicates an average over households  $i$  at age  $t$  and for good  $k$ . We have used the approximation that log deviation from the average expenditure is approximately equal to the percentage deviation from the average. Using this expression, we can write the variance across households as

$$\text{Var}\left(\ln\left(\sum_k C_{i,t}^k\right)\right) \approx \sum_k (s_t^k)^2 \text{Var}\left(\ln(C_{i,t}^k)\right) + 2 \sum_{k,k'>k} s_t^k s_t^{k'} \text{Cov}\left(\ln(C_{i,t}^k), \ln(C_{i,t}^{k'})\right)$$

The first term on the right captures age-varying shares and variances, while the last terms are the covariances. In Figure R4 we plot the variance of the total (the right hand side above), which differs from Figure R1b in the text due to the absence of cohort, year, and demographic controls. We then decompose the first-term on the right hand side into (i) the “variance” component, which is constructed by holding the shares constant over the lifecycle at their average value and allowing the cross-sectional variance for each good to vary and (ii) the “shares” component, which is constructed by holding the variance of each good constant at the average value over the lifecycle and allowing the shares to vary. The figure also depicts the sum of the covariances (the last term in the above expression), as well as the approximation error, which is the difference between the right hand and left hand sides above. As before, all series are in differences from age 25.

Figure R4 indicates that changing shares and approximation error play a negligible role in explaining the change over the lifecycle in expenditure inequality. The main drivers are the

individual disaggregated variances and the covariances, with the covariances contributing more early in the lifecycle, and then the two series playing roughly equal roles post age 60.

### **3. Alternative Controls for Family Size**

In this section we explore an alternative scheme for translating household expenditure into per capita expenditure. The benchmark analysis presented in the text used a regression approach with a rich set of family size controls to convert household expenditure into “individual” consumption. The advantage of this approach is its flexibility. In particular, it allows for different scales for different goods, reflecting the fact that the returns to scale for some category (e.g. clothing) may be different from some other category (e.g., housing). However, this approach may lead to biased lifecycle profile if family size is correlated with permanent income, and family size varies over the life cycle. This may induce an artificial age profile to expenditure. A popular alternative - although one that still takes family size as exogenous- is to adjust household expenditure using an adult equivalent scale. Three such measures are popular in the literature. The first is the original or “old” OECD equivalence scales (see <http://www.oecd.org/dataoecd/61/52/35411111.pdf>). Specifically, this scale weights the first adult as 1, each additional adult in the household 0.7, and each child 0.5.<sup>1</sup> A second scale, the “modified” OECD scale, places a weight of 0.5 to additional adults and 0.3 to children. A third scale is the square root of total household size.

We have investigated all three alternatives. As we show below, the quantitative difference between the lifecycle profiles between core and work related expenditures hold regardless of how we control for family size. There is a difference in the levels of each profile individually. The fact that the levels of the profiles for both of the categories change similarly is

---

<sup>1</sup> Note, when implementing these procedures below, we define children as being under the age of 15.

what ensures that the difference between the profiles is held constant regardless of the family size controls. As a general rule, the higher the weight given to additional family members, the smaller the decline in mean expenditure post middle age and the smaller the increase in cross-sectional variance over the life cycle.

Consider Robustness Appendix Figure R5 which corresponds to Figure 1 in the main text. Specifically, for each panel, we construct adult equivalents for each household using the old OECD scale and the modified OECD scale. We then subtract log adult equivalents from log expenditure to form an adjusted expenditure measure. We then regress log adjusted expenditure on cohort and normalized year dummies. This equation is analogous to equation (1) in the main text aside from the fact that we are using the family size adjusted expenditure as our dependent variable and that we dropped all other family size controls, including marital status, from the regression. Figure R5a depicts our benchmark result for the mean profile over the lifecycle for nondurables with housing services alongside those using the two OECD equivalence scales. The red dashed line is the results using the old OECD scale and the black dotted line is the results using the modified OECD scale. Figure R5b depicts the results for nondurables excluding housing services. We see that the OECD scales generate a flatter profile early in the lifecycle and a larger increase (panel A) or a smaller drop (panel B) in the latter half of the life cycle. The fact that the hump can be mitigated with certain family size adjustment is emphasized in Browning and Ejrnaes (*Review of Economics and Statistics* 2009). Similar lifecycle profiles in consumption for which the family size adjustment is made using OECD equivalence scales can be found in the surveys of lifecycle consumption by Attanasio (1999) and Attanasio and Weber (2010). Fernandez-Villaverde and Krueger (2006) also document that the hump in lifecycle

expenditures is quantitatively sensitive to how researchers control for the lifecycle profile of family size.

In Figure R5 we also replicate Figure 1b from the text, which is the cross-sectional variance of log expenditure over the life cycle. These results are shown in panels (c) and (d) of Figure R5. Specifically, we take the residuals of the regressions from the first row of Figure R5 and compute a cross-sectional variance for each age-cohort pair. We then regress the variances on age and cohort dummies and plot the coefficients on the age dummies. Again, we present the results with and without housing services. From Figures R5c and R5d, we see that the OECD scales also lead to a smaller increase in variance over the life cycle. The old OECD scale indicates an increase that is one half of our benchmark results (or 0.08 points), while the modified scale lies halfway between the benchmark and the old OECD scale.

We now turn to whether the alternative family size controls affect our results on disaggregated consumption categories. In Figure R6 we replicate Figure 2 from the text with both the old OECD scale (panels (a) and (c)) and the modified OECD scale (panels (b) and (d)). Figures R6a and R6b indicates that mean expenditure on our core consumption categories all have an increasing profile over the life cycle, as was the case in our benchmark analysis. However, now the increase is even steeper than it was in our baseline case. Moreover, the figures also show that work related goods and food at home do not decline as much during the latter part of the lifecycle. In other words, even with the alternate family size controls, there is a substantial and systematic difference in the lifecycle profile of mean expenditures between core nondurables and work related expenditures and food. This shows that even with the different equivalent scales, the main results of our paper still hold. The behavior of work related goods and food behave very differently relative to core nondurables over the lifecycle. In panels (c)

and (d) of Figure R6, a similar story can be seen for the lifecycle profile of the cross sectional variances.

A natural question arises as to which control for family size is better? The main drawbacks of the OECD scales are that (1) they are determined rather arbitrarily, (2) they do not differ across consumption goods, and (3) they still do not deal with the endogeneity of family size over the lifecycle. The main drawback of the inclusion of family size controls directly into the regression (as we do in the main body of the text) is just they do not account for the endogeneity of family size over the lifecycle. If family size over the lifecycle is correlated with permanent income, our lifecycle profiles purged of family size effects could be biased.

We can test whether such an issue is a concern for our analysis – at least for food expenditures – using the Panel Study of Income Dynamics (PSID). The PSID is a longitudinal survey that tracks multiple observations of income, demographics, and food expenditure of the same household across time.<sup>2</sup> Given the panel nature of the survey, we can test whether the family size controls are biasing our lifecycle profiles by replicating our main empirical specification from the text and then comparing it to the same analysis but including individual fixed effects. In the latter case, all the estimation comes from within family variation. Permanent income variation across individuals will be proxied by the individual fixed effects.

This analysis is shown in Figure R7. As seen from Figure R7, the estimates of the lifecycle profile of food expenditure are nearly identical regardless if we use our method from the text or if we include individual fixed effects in the estimation. This suggests that our procedure does not result in any significant bias from the potential correlation of family size controls and permanent income. We realize that this analysis is only for food expenditure and

---

<sup>2</sup> For this analysis, we restricted our attention to the 1977-1999 waves of the PSID excluding the 1988 and 1989 waves. We exclude the 1988 and 1989 waves because those waves do not include information on food expenditures.

does not imply that there will not be bias with other categories. However, we find it reassuring that for a category where we can include individual fixed effects, our method provides very similar results.

Given the fact that we can estimate different returns to scale for different category and the fact that there does not seem to be much bias in the lifecycle profiles resulting from the potential correlation of family size and permanent income, we strictly prefer our method to the method of deflating expenditures by constant and arbitrary equivalence scales.

#### **4. Measurement Error in the CEX**

It has been documented that the Consumer Expenditure Survey has measurement error that has increased over time. In particular, the ratio of CEX expenditure to NIPA Personal Consumption Expenditure (PCE) has fallen sharply over the period of our analysis. Also, it has been well documented that the deterioration has differed by categories. For example, the BLS reports that ratio of CEX to PCE expenditure for entertainment falls from 0.66 in 1988 to 0.52 in 1999, while the corresponding ratio for rent, utilities, and public services declines from 0.99 to 0.95 and the ratio for food is roughly constant at 0.74. These ratios are reported in various issues of “Consumer Expenditure Survey Compared with Personal Consumption Expenditures,” available on the BLS website: [www.bls.gov/cex.cecomparison.htm](http://www.bls.gov/cex.cecomparison.htm).

To assess whether such discrepancies bias our results, we perform the following exercise. For each of our three-goods, food at home, work-related expenses, and core expenditure, we calculate the ratio of CEX to PCE expenditures. To do this, we use the disaggregated ratios reported by the BLS and sum up using the CEX shares reported by the BLS. The concordance provides a rough approximation of the under-reporting for each of our disaggregated goods,



although the coverage across goods and years is not complete. For each household in the respective years we divide reported expenditure by the BLS ratio. That is, we assume that the under-reporting is proportional across households in a given year. Using these adjusted series, we recomputed Figure 2 from the text (although given the missing years in the BLS ratio data, we drop the year controls from the regressions, which has a negligible affect on the series in the presence of cohort and age dummies).

The results are depicted in Figure R8, with panel (a) depicting means and panel (b) depicting variances. Each figure includes a solid, adjusted series as well as a corresponding dashed line, which is the benchmark data without adjustment. The adjusted series tracks the benchmark data closely. The main difference is that when adjusted, mean core expenditure increases more over the entire lifecycle, and work-related expenses increase slightly more early in the lifecycle. Overall, the adjustment does not change the conclusions drawn from the benchmark exercises reported in the text.

## **5. Demand System Estimates**

Given the potential importance of work related expenses in driving changes in expenditure over the life cycle, a natural approach would be to directly control for work status when estimating the lifecycle profile of mean expenditures or dispersion.<sup>3</sup> A difficulty with simply adding controls for employment status to the estimation of equation (1) from the main text is the fact that labor supply is closely associated with permanent income. For example, lower wage workers in the time frame of our sample tend to work fewer hours than high wage

---

<sup>3</sup> This whole section used to be part of the main text of the paper. Given the comments of the editor and referees, we added much more material to the main text of the paper during the revision process. To conserve space, we moved this section to the robustness appendix. We provide a summary of these results in the main text. Some of the summary overlaps closely with the exposition in this section.

workers (see, for example, Aguiar and Hurst (2007)). Absent a panel, controls for labor supply will also proxy for permanent income.<sup>4</sup> However, using the standard tools of demand system analysis, we can explore the effect of labor supply on how expenditure is allocated across different goods, conditional on a given level of total expenditure. That is, by including total expenditure, we can isolate the effect of labor supply from variation in permanent income across households.

Specifically, we estimate the following:

$$s_{it}^k = \omega_0 + \omega_{age} Age_{it} + \omega_c Cohort_{it} + \omega_t D_t + \omega_{fs} Family_{it} + \sum_k \omega_p^k \ln P_t^k + \omega_p \ln P_t + \omega_x \ln X_{it} + \omega_L L_{it} + \varepsilon_{it}^k, \quad (R1)$$

where  $X_{it}$  is our measure of total nondurable spending and is defined as the sum of spending across all the categories shown in Figures 3a and 3b of main text (i.e., total nondurable spending including housing services and excluding alcohol and tobacco) for household  $i$  in period  $t$ .  $s_{it}^k$  is the share of spending on consumption category  $k$  out of  $X_{it}$  for household  $i$  in period  $t$ . By definition, for each household the shares across the different consumption categories sum to 1. The age, cohort, year, and family status controls are the same as in equation (1) of the main text. We include as additional controls the log price index of each of our sub-aggregates ( $\ln P^k$ ) as well as the overall price index ( $\ln P$ ). These variables, together with the normalized year dummies, control for changes in relative prices across the consumption categories. We compute the category specific price indices by using the weighted share of the price indices for the goods that comprise the consumption category. Finally, we include a vector of controls describing household labor supply ( $L$ ). We discuss these controls below.

---

<sup>4</sup> We discuss estimates using panel data from the PSID below.

Given the fact that expenditures on the individual consumption categories are determined simultaneously and the fact that any measurement error in one category will lead to measurement in  $X$ , we follow the standard practice of instrumenting  $X$  with log total household family income and education dummies.<sup>5</sup> Our measure of total household family income includes labor and transfer income of both husbands and wives.

Note that equation (R1) is a close parallel to the almost ideal demand system (AIDS) of Deaton and Muelbauer (1980), conditioned on work status, family size, cohort, and age. We impose the restriction that the overall price index is given by the CPI-U, but do not impose restrictions related to consumer optimization such as symmetry and homogeneity. The inclusion of work status controls to form a conditional demand system follows the important work of Browning and Meghir (1991) and Blundell et al (1994).

Using equation (R1), we answer two different questions. First, among younger households (those under the age of 50), how is working associated with spending on different consumption goods? If there are work related consumption needs, we would predict that, all else equal, an increase in household labor supply would be positively associated with spending on those categories. By estimating (R1), we assess whether transportation spending, clothing spending, and food away from home are positively associated with household labor supply. Second, we use (R1) to assess how much of the decline in spending post middle age on work related consumption categories can be attributed to changes in household labor supply. In particular, we estimate (R1) both with and without controls for labor supply and see how the age coefficients change.

---

<sup>5</sup> Specifically, the instruments consist of log income (where the 217 households with zero income are bottom coded at 1 dollar), a dummy for zero income, income squared, income cubed, and four educational dummies indicating the head's educational attainment (<12, 12, 13-15, 16+).

The results of the first question are shown in Table R1. To avoid issues of changes in household formation and its effect on labor supply, we restrict our analysis sample to include only married households. Similarly, to avoid the issue of retirement, we restrict our analysis to only those households where the head is 50 years old or younger. This leaves us with a sample of 21,034 households. Specifically, Table R1 shows the results from estimating (R1) when our measure of household labor supply ( $L$ ) consists of two dummy variables; one indicating whether the husband is currently employed and another indicating whether the wife is currently employed.

Table R1 shows that there are only three consumption categories for which the share of spending is positively associated with household labor supply. These three categories are nondurable transportation, food away from home, and clothing. Specifically, the unconditional mean for the share of spending ( $s$ ) allocated to nondurable transportation is 13 percent. Households where both spouses work spend an additional 1.4 percentage points (or an additional 11% above the average share) on nondurable transportation compared with an otherwise similar household where only the husband works. Having a working wife increase the share spent on food away from home by 8 percent. Both of these differences are statistically significant at the 1 percent level.

Table R1 also shows that there is no positive relationship between the share of spending allocated to any of the other main consumption categories and household employment status. Rather, given the adding up constraint, the share of spending on most of these other consumption categories is negatively related to employment status. Our simple demand system estimates confirm what we discussed above: spending on clothing, nondurable transportation, and food away from home are positively associated with household labor supply.

Figures R9a-R9c show the results from our second exercise. In these results, we plot the age coefficients from our estimation of (R1) where the sample includes all married households between the ages of 25 and 75 (34,195 households). We focus our attention on the share expenditures allocated to the specific work related categories: transportation (Figure R9a), food away from home (Figure R9b), and clothing (Figure R9c). We then ask how much of the declining share of spending on these goods post middle age can be explained by changing work status. To do this, we estimate (R1) with (dotted line) and without (solid line) work controls.<sup>6</sup> As always, each point represents the deviation from age 25, with the units for Figure R9 being share of total nondurable expenditures.

As seen from Figures R9a – R9c, without work status controls, the share of spending allocated to nondurable transportation, food away from home, and clothing falls between the age of 50 and 75 by roughly 3.3 percentage points, 0.7 percentage points, and 2.1 percentage points, respectively. The dashed lines indicate that the inclusion of labor supply controls reduce the post-middle age decline in expenditure shares dramatically. In particular, controlling for the labor supply of both the husband and wife explains essentially all the decline in the share of spending allocated to clothing and personal care and food away from home, and roughly 40 percent of the decline in the share of spending allocated to non-durable transportation.

In Table R2, we exploit the panel dimension of the PSID to document similar results.<sup>7</sup> In columns 1 and 2 of Table R2, we summarize the lifecycle profile of food expenditure estimating the regression with the inclusion of individual fixed effects so as to exploit within

---

<sup>6</sup> We found that labor supply has a nonlinear effect on consumption for these different consumption categories, so when we include work status controls, we include a much more extensive set of controls than simply whether or not the husband and wife were employed. Specifically, our work status controls in Figure R9 also include: a vector of 7 dummies indicating the number of weeks worked by husbands during the previous year, a vector of 7 dummies indicating the number of weeks worked by the wife during the previous year, a vector of 9 dummies indicating the number of hours usually worked during the week by the head, and a vector of 9 dummies indicating the number of hours usually worked during the week by the wife.

<sup>7</sup> The PSID sample we used for this analysis is the same sample we used in section 3 of this robustness appendix.

person variation over the lifecycle. The regression also includes the same demographic controls as in Figure R7 (i.e., controls for family size, marital size, and the sex and age of household children). We estimate this for all individuals in the PSID (column 1) and for a sample of married households (column 2). In column 3, we restrict our analysis to those married households where both the husband and the wife worked full time continuously over the time they were in the PSID. As seen from the sample sizes at the bottom of Table R2, this includes only about 10 percent of the married sample. However, for this sample, food expenditure does not decline post middle age.<sup>8</sup>

Collectively, the demand system results and the panel estimates from the PSID suggest that the lifecycle profiles for food, clothing, and nondurable transportation post middle age are to a large extent determined by changing work status over the lifecycle.

## **6. Cohort Versus Time Effects**

All the analysis in text controlled for cohort effects and restricted time effects to be cyclical (that is, have no trend). Controlling for cohort heterogeneity is important when using cross-sectional data given that younger cohorts tend to have higher lifetime resources than older households. If lifetime resources for successive cohorts are increasing, excluding cohort effects will bias down estimated age effects of older households relative to younger households. However, there may be trends in consumption patterns that argue for time effects. Some of this may result from changing relative prices of different goods over time. However, it may also reflect changing aggregate economic conditions and changes in preferences or technology.

---

<sup>8</sup> Given the smaller sample size, these coefficients are estimated with less precision. But, these estimates are consistent with the demand system estimates shown in Figures R8a-R8c,

Given the co-linearity of age, year of birth, and time, we cannot identify independent trends in all three effects. To assess how sensitive our main results are to the choice between cohort versus time effects, in this section we re-estimate (1) and (2) from the main text with unrestricted year effects and no cohort effects for total expenditures as well as the three sub-aggregates. The results are shown in Figures R10, R11, and R12. Figures R10 and R11 are akin to Figures 1 and 2 where the estimation is done with unrestricted year effects and no cohort controls. The main take away from our analysis is that for nearly all categories, it makes little difference whether we estimate the lifecycle profiles with cohort effects or with year effects. The only difference is with housing services. The housing services results are shown in Figure R12. With the cohort effects, housing services is upward sloping throughout the lifecycle. However, with year effects, housing services decline slightly after middle age. Any differences with year effects in Figures R9 and R10 are attributable to housing. The essential ambiguity is whether the systematic increase in expenditure on housing is a cohort or time effect.

We want to emphasize a few additional points. First, the lifecycle profiles of the cross sectional variance change very little regardless of whether or not we use cohort or year effects. This is seen clearly from Figure R11b. The reason is that the big increase in lifecycle variance can be attributable to work related expenses and this changes little over the lifecycle. Second, with respect to the mean patterns, the profile of core nondurables has changed while the lifecycle profile of work related expenses remained relatively constant (Figure R11a). Instead of core nondurables rising sharply from age 25 through age 50 and then rising further from age 50 through age 75, core nondurables with year affects rise less from age 25 through age 50 and then actually declines slightly between age 50 and 75.

With regard to housing services, the identification issue is how to interpret the fact that a 35-year-old in 1980 purchases more housing services as a 55-year-old in 2000 than a 55-year-old in 1980 (adjusting for housing inflation and family size). Our benchmark follows a cohort over time to trace out the lifecycle, so it interprets this as an increase in expenditure over the lifecycle. An alternative is that the cross-section at a point in time is a better measure of the lifecycle, and so the difference between young and old in 1980 (averaged with the differences in each year of the sample) indicates a true hump in housing expenditure. We favor the former interpretation (as does the vast majority of the lifecycle consumption literature). However, if we allow for a decline in housing services that starts in middle age, the differences between core consumption, work-related, and food-at-home are mitigated. In our quantitative exercise, this brings the implications of a multigood model slightly closer to the one-good model. In particular, the conditional and unconditional differences in risk at a 20-year horizon are now 27 and 19 log points, rather than 41 and 28, respectively.

## **7. Excluding the “Zeros”**

There are a number of households who record zero expenditures for our smaller categories (food away from home, domestic services, and the residual category “other nondurables”). When looking at these categories in isolation, looking at log deviations implies dropping observations with zero expenditure. The benchmark results in the text bottom coded all expenditure on these three categories at one dollar. The rationale for this was that the fraction of households with strictly positive expenditure on these categories changes over the life cycle. Dropping the zeros meant dropping the low spending households at various points in the life cycle, potentially giving a misleading picture of the life cycle pattern. While these small



categories do not figure prominently in the analysis, for completeness we report how our handling of this issue changes the profiles for these categories. Note that the log zero issue does not arise for (1) the other disaggregated categories, (2) our three composite goods of food at home, work related expenditures, or core nondurables or (3) for our aggregate nondurables.

Figure R13 plots the fraction of households with non-zero expenditure at each age for the three goods in question: food away from home, domestic services, and the residual other category. The figure indicates that young households are less likely to consume positive amounts of domestic services and other nondurables, and are more likely to consume food away from home.

In Figure R14, we show how dropping the zeros alters the results depicted in Figure 3 of the text. The main difference between Figure R14a and Figure 3a in the text is in regard to other nondurables, which now has a flat profile and had an increasing profile in the benchmark figure. This reflects the fact that many young households report zero expenditures on other nondurables. Dropping these zeros raises young expenditure relative to older expenditure. In regard to cross-sectional variance, dropping the zeros has an ambiguous effect on the life cycle profile of consumption inequality. We see that the increase in the cross-sectional variance in consumption of food away from home is much smaller if we drop the zeros. Also, dropping the zeros also removes the increase in cross-sectional variance for other nondurables.

## References

Aguiar, Mark and Erik Hurst (2007). "Measuring Trends in Leisure: The Allocation of Time Over Five Decades", *Quarterly Journal of Economics*.

Attanasio, Orazio (1999). "Consumption", *Handbook of Macroeconomics*, in J Taylor and M. Woodford (ed.), *Handbook of Macroeconomics*, edition 1, volume 1, chapter 11, pages 741-812, Elsevier.

Attanasio, Orazio and Guglielmo Weber (2010). "Consumption and Saving: Models of Intertemporal Allocation and Their Implications for Public Policy", *Journal of Economic Literature*, 48(3), 693-751.

Blundell, Richard, Martin Browning, and Costas Meghir (1994), "Consumer Demand and the Life-Cycle Allocation of Household Expenditures," *Review of Economic Studies*, 61, 57-80.

Browning, Martin and Mette Ejnaes (2009), "Consumption and Children", *Review of Economics and Statistics*, 91(1), 93-111.

Browning, Martina and Costas Meghir (1991), "The Effects of Male and Female Labor Supply on Commodity Demands," *Econometrica* 59(4): 925-951.

Deaton, Angus and John Muelbauer (1980), "An Almost Ideal Demand System", *American Economic Review*, 70, 312-26.

Fernandez-Villaverde, Jesus and Dirk Krueger (2006), "Consumption Over the Lifecycle: Some Facts from CEX Data", *Review of Economics and Statistics*, 2006, 89, 552-65.

**Table R1: The Relationship Between Work Status and Spending, by Consumption Category**

<i>Disaggregated Consumption Group</i>	<i>Mean Share Out of Total Nondurable Expenditures</i>	<i>Coefficient on Dummy Variable: Husband Working</i>	<i>Coefficient on Dummy Variable: Wife Working</i>
Transportation	0.13	0.014 (0.002)	0.014 (0.001)
Food Away From Home	0.06	0.008 (0.001)	0.005 (0.001)
Clothing and Personal Care	0.08	0.003 (0.001)	0.001 (0.001)
Housing Services	0.34	-0.009 (0.003)	-0.012 (0.002)
Utilities	0.12	-0.005 (0.001)	-0.003 (0.001)
Entertainment	0.04	0.000 (0.001)	0.000 (0.001)
Food At Home	0.18	-0.016 (0.002)	-0.013 (0.001)

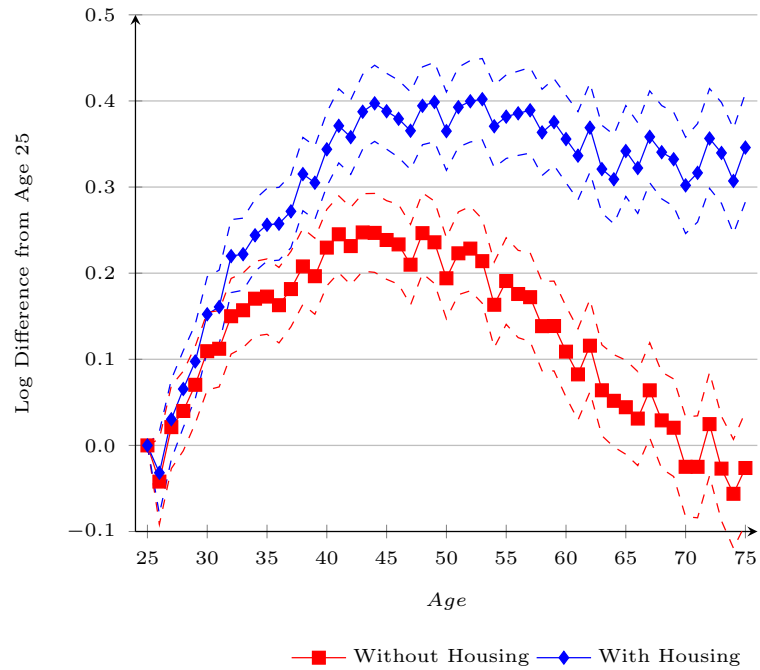
Notes: This table reports the coefficients on dummy variables indicating whether the husband is working (column 3) and/or the wife is working (column 4) in a regression in which the dependent variable is the share of expenditure out of total nondurable spending (including housing services but excluding alcohol and tobacco) allocated to the category indicated in column 1. Each regression includes the standard vector of controls used in tables 1 and 2, the full set of (log) price indices for each category indicated in column 1 and the overall CPI deflator, and log total nondurable expenditure (see equation 3 in the text). We instrument total nondurable expenditure with a flexible function of total household income and dummies indicating the educational attainment of the head (<12, 12, 13—15, >=16). For the regression analysis, we restrict our sample to include only married households between the ages of 25 and 50 (inclusive), which is a sample of 21,034 households. Standard errors are in parentheses. Column 2 includes the unconditional mean of the share of spending out of total nondurable spending excluding alcohol/tobacco for each of the consumption categories. For this column, we use the full sample of 53,412 households.

**Table R2: Lifecycle Profile of Food Expenditure,  
Panel Estimates From the PSID With and Without Conditioning on Work Status**

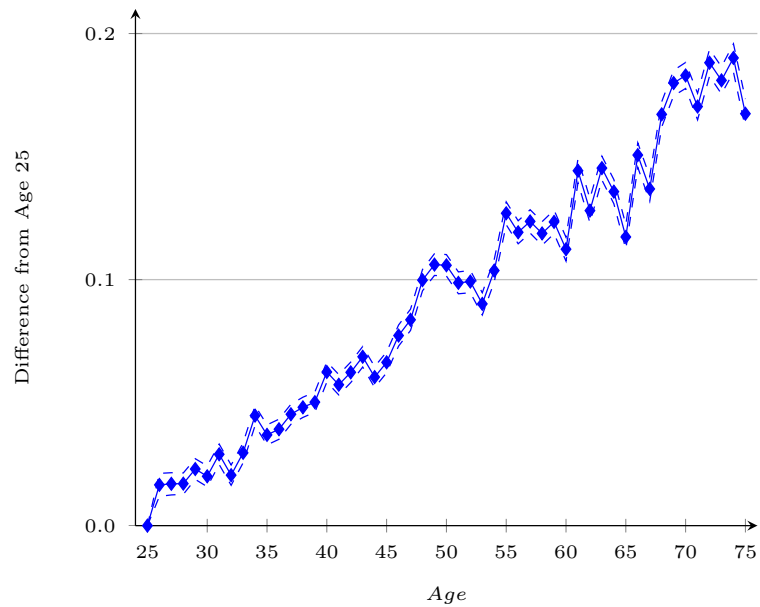
<i>Log Expenditure, Relative to Age 25-29</i>	<i>All, Without Conditioning on Work Status</i>	<i>Married, Without Conditioning on Work Status</i>	<i>Married, Conditioning on Constant Full Time Employment</i>
Dummy: Age 30-34	0.04 (0.01)	0.03 (0.01)	0.01 (0.05)
Dummy: Age 35-39	0.06 (0.01)	0.06 (0.01)	0.13 (0.06)
Dummy: Age 40-44	0.07 (0.01)	0.08 (0.01)	0.17 (0.06)
Dummy: Age 45-49	0.03 (0.01)	0.07 (0.01)	0.16 (0.07)
Dummy: Age 50-54	-0.01 (0.02)	0.040 (0.02)	0.21 (0.08)
Dummy: Age 55-59	-0.08 (0.02)	-0.02 (0.02)	0.18 (0.09)
Dummy: Age 60-64	-0.17 (0.02)	-0.10 (0.02)	0.15 (0.12)
Difference: Age 60-64 - Age 45-49	-0.20	-0.17	-0.01
p-value of Difference:	<0.01	<0.01	0.93
Sample Size	30,307	21,371	1,151

Notes: Estimates of the age effects from a regression of log food expenditures on individual fixed effects, detailed family size controls, and five year age dummies. In columns 1 and 2, we include all individuals and married individuals, respectively. In column 3, we restrict the sample to include all married households where (1) the head and wife are currently working, (2) the head and wife worked at least 1900 hours in the prior year, and (3) the head and wife worked every year they were in the PSID. For the regressions in column 3, we also included controls for hours worked in the prior year.

Figure R1: Lifecycle Profiles of Nondurable Expenditures with Standard Error Bands



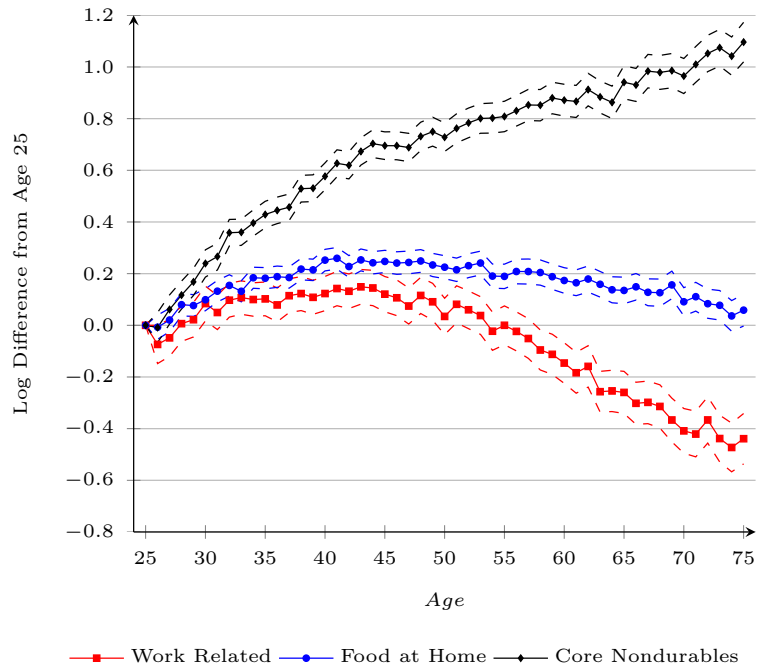
(a) Mean Nondurable Expenditures over the Lifecycle



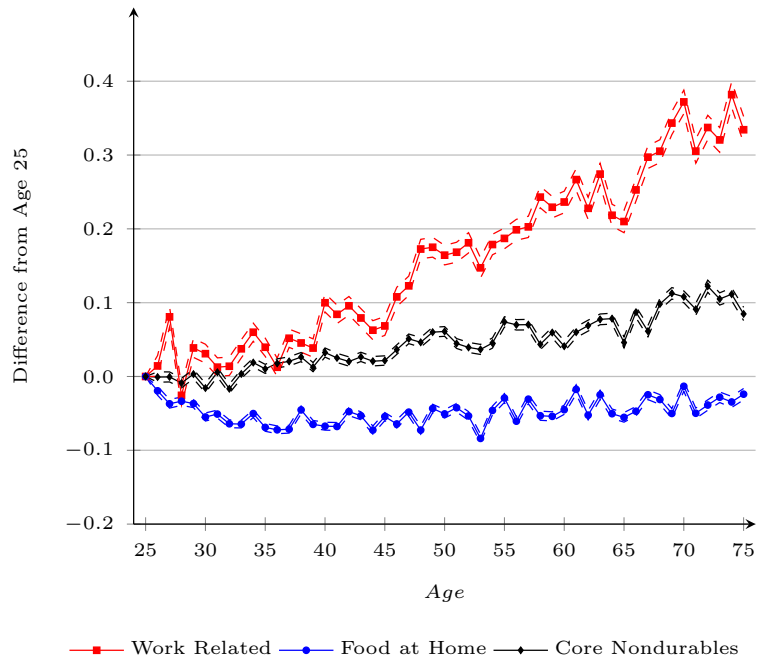
(b) Cross Sectional Variance of Nondurable Expenditures over the Lifecycle

Notes: This figure replicates figure 1 from the text with the addition of standard error bands. Specifically, each point represents the coefficient on an age dummy of the regression of log expenditure on our controls. The dashed lines around each point represent the regression coefficient plus and minus two times its standard error. In panel (b) we omit nondurables without housing given how close the point estimates are in figure 1 (b).

Figure R2: Lifecycle Profiles of Three Sub-aggregates with Standard Error Bands



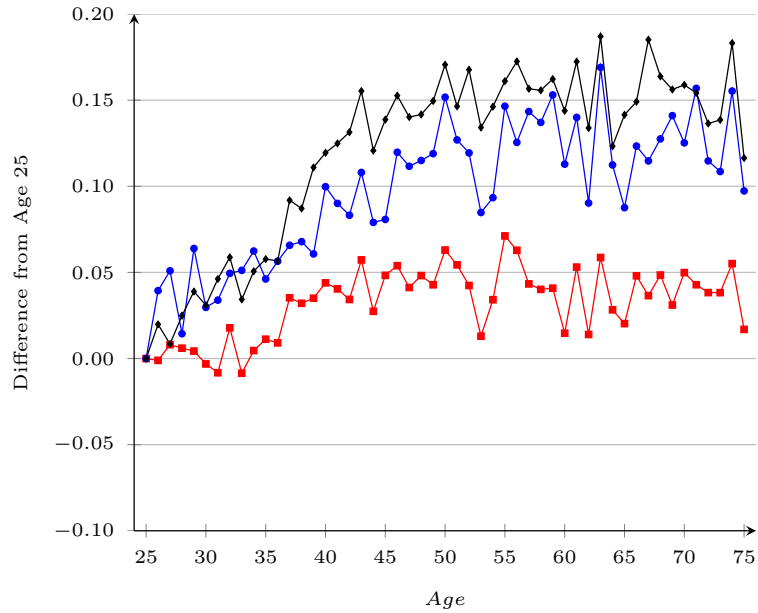
(a) Mean Expenditure



(b) Cross-Sectional Variance

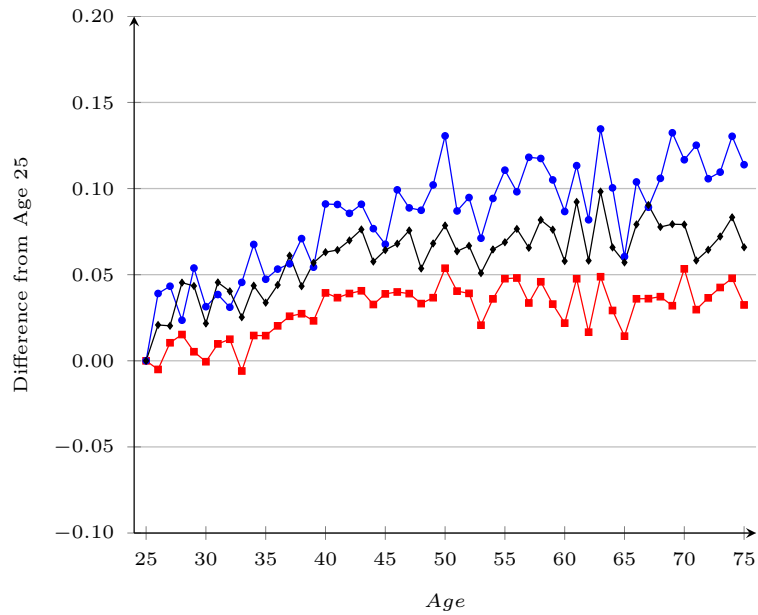
Notes: This figure replicates figure 2 from the text with the addition of standard error bands. Specifically, each point represents the coefficient on an age dummy of the regression of log expenditure on our controls. The dashed lines around each point represent the regression coefficient plus and minus two times its standard error.

Figure R3: Lifecycle Profiles of Three Sub-aggregates: Covariances



—■— Covar(Core,Food) —●— Covar(Core,Work Related) —◆— Covar(Food,Work Related)

(a) Unconditional Covariance

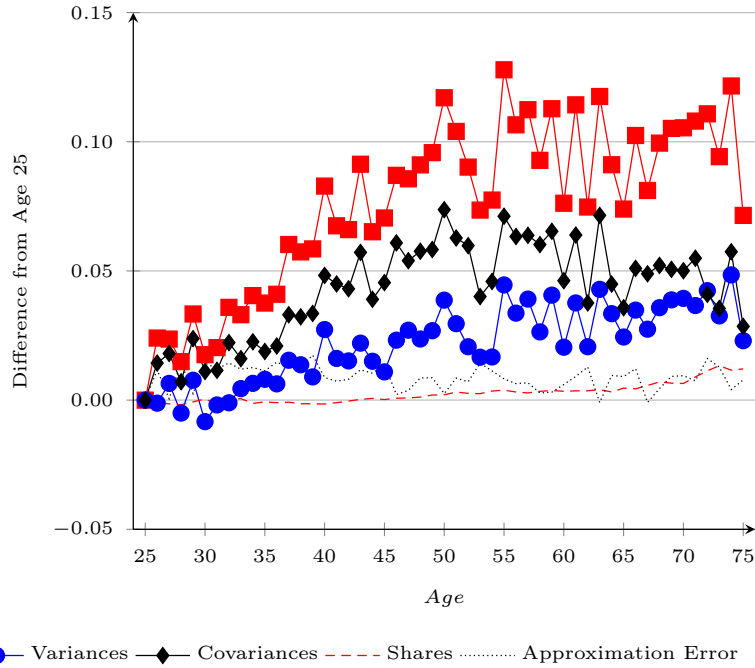


—■— Covar(Core,Food) —●— Covar(Core,Work Related) —◆— Covar(Food,Work Related)

(b) Residual Covariance

Notes: This figure plots the covariances between log core, log food at home, and log work related expenditures. Each point represents the cross-sectional covariance at that age minus the respective covariance at age 25. Panel (a) depicts the raw covariances and panel (b) depicts the residual covariance after regressing on our cohort, age, year, and family size controls. The age-25 unconditional covariance between core and food is 0.08, between core and work-related is 0.15, and between work-related and food is 0.04. The respective residual age-25 covariances are 0.04, 0.12, and 0.02.

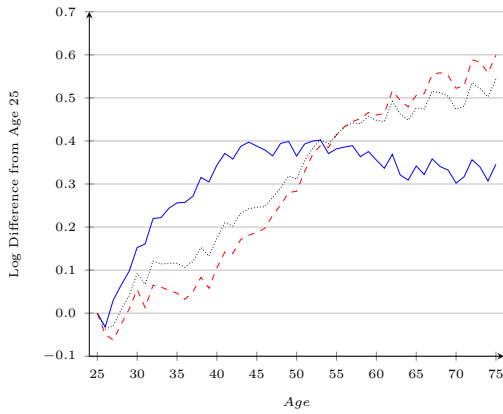
Figure R4: Components of Variance



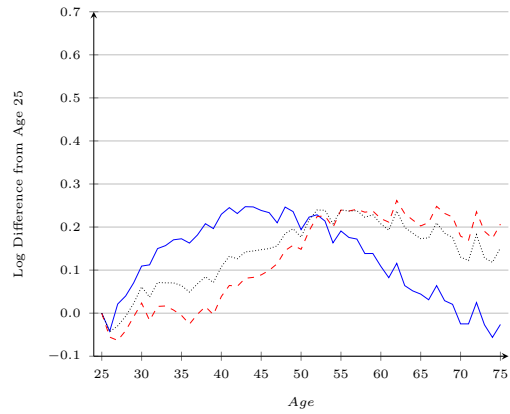
Notes: This figure explores the components of the lifecycle profile of consumption inequality. Total variance refers to the cross-sectional variance at each age of the log of the sum of core, work-related, and food-at-home expenditures. These are unconditional variances, without demographic, cohort, or year controls. The difference between this series and that depicted in figure 1b in the text is the absence of such controls. The remaining series are based on the approximate decomposition discussed in the appendix. The series labelled “Variances” depicts  $\sum_k (\bar{s}^k)^2 Var(\ln(C_{i,t}^k))$ , where  $\bar{s}^k$  is the average share of good  $k$  (=core, food, work-related) averaged over all ages. That is, this series is the contribution of age-dependent variances of each subagregates, weighted by a constant set of shares. The series labelled “Covariances” depicts  $\sum_k \sum_{k'} s_t^k s_t^{k'} Cov(\ln C_{i,t}^k, \ln C_{i,t}^{k'})$ . That is, this series represents the contribution of age-dependent covariance terms weighted by age-dependent shares. The series labelled “Shares” represents  $\sum_k (s_t^k)^2 \bar{V}(\ln C_{i,t}^k)$ , where the bar over the variance indicates we average the variance over all ages. That is, this series represents the contribution due to the fact that expenditure shares are shifting away from food and work-related and toward core over the lifecycle. The series “Approximation Error” is the approximation error; that is, the difference between the total variance and the sum of the depicted sub-components. All series are in deviations from age 25.



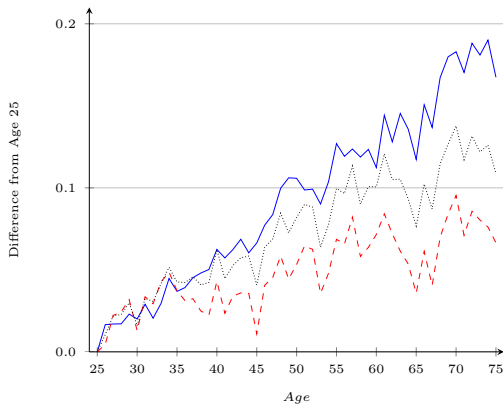
Figure R5: Lifecycle Profiles of Nondurable Expenditures: Alternative Family Size Controls



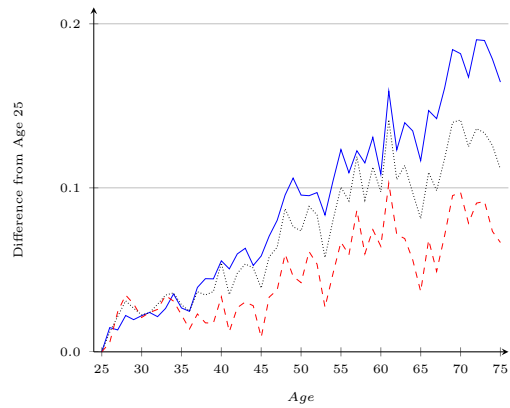
(a) Nondurable Expenditures with Housing Services



(b) Nondurable Expenditures without Housing Services



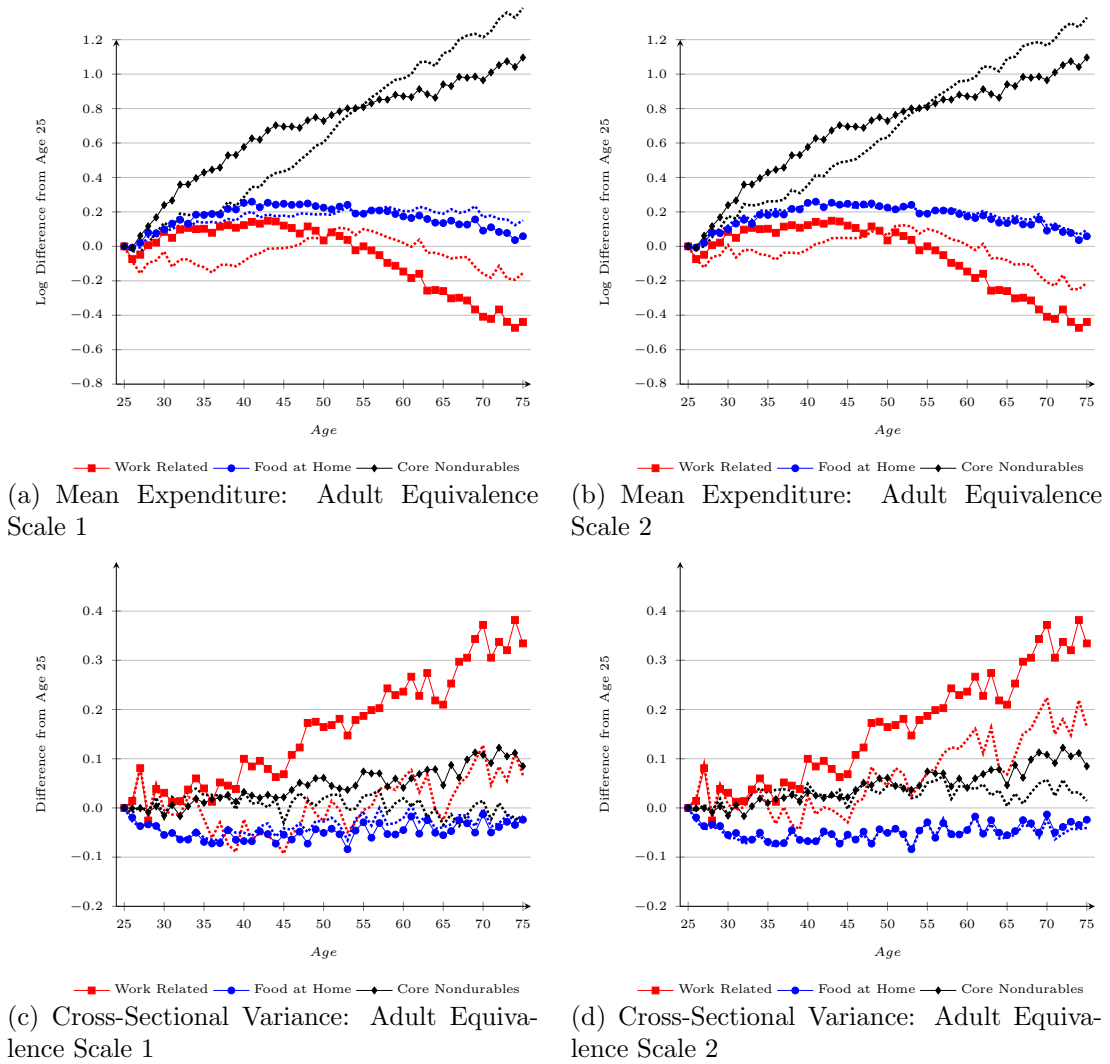
(c) Cross Sectional Variance of Nondurable Expenditures with Housing Services



(d) Cross Sectional Variance of Nondurable Expenditures without Housing Services

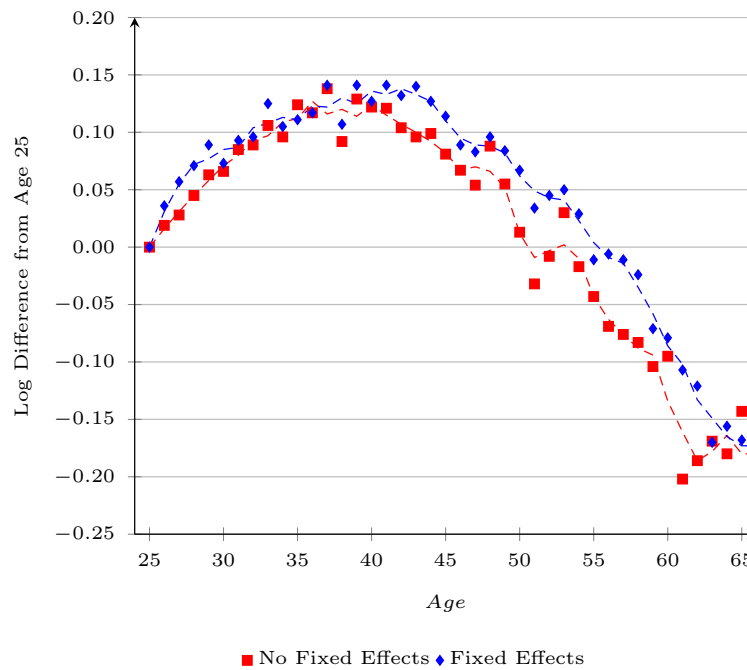
Notes: This figure explores alternative family size control. The solid line in each figure represents the benchmark controls depicted in figure 1 of the text; namely, the inclusion of a rich set of demographic dummies. The red dashed line uses an alternative “adult equivalence” scale; namely, household expenditure is divided by one plus 0.7 times the number of members age 15 or older plus 0.5 times the number of members less than 15. The black dotted line replaces (0.7,0.5) with (0.5,0.3), respectively. Panel (a) is mean nondurables including housing services, (b) is mean nondurables excluding housing services, (c) is the cross sectional variance of nondurables including housing services and (d) is the cross sectional variance of nondurables excluding housing services.

Figure R6: Lifecycle Profiles of Three Sub-aggregates: Alternative Family Size Controls



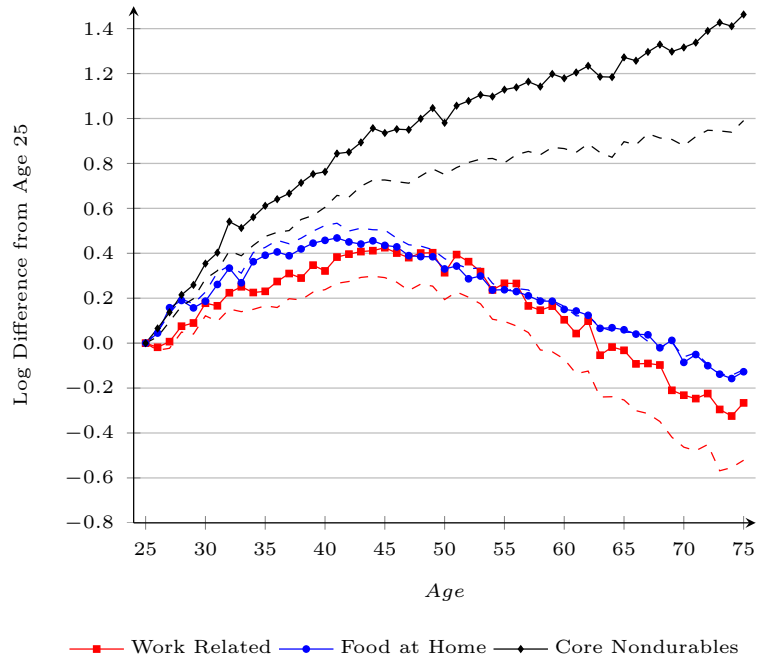
Notes: This figure explores alternative family size control using our three-subaggregates. The solid line in each figure represents the benchmark controls depicted in figure 2 in the text; namely, the inclusion of a rich set of demographic dummies. The dotted lines represent the corresponding series deflated by alternative adult equivalence scales. Scale 1 (left column, panels (a) and (c) ) deflates by one plus 0.7 times the number of members age 15 or older plus 0.5 times the number of members less than 15. Scale 2 (right column, panels (b) and (d)) replaces (0.7,0.5) with (0.5,0.3), respectively.

Figure R7: Endogeneity of Family Size Controls: A Comparison using PSID Food Expenditure

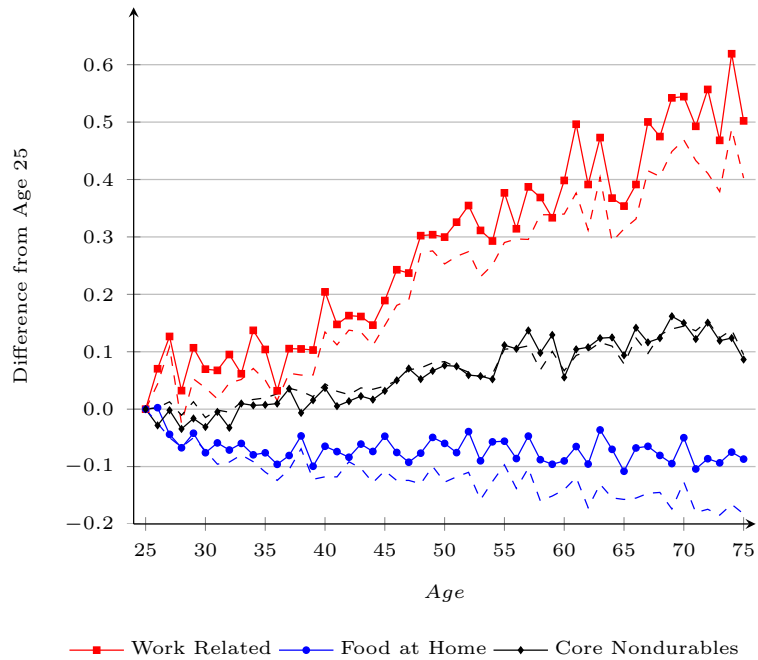


Notes: This figure uses a sample from the PSID to explore the endogeneity of family size controls. In particular, the series labelled "no fixed effects" plots the lifecycle profile of mean log expenditure on food using our family size and cohort controls, treating the data as repeated cross-sections as in the benchmark. The series labelled "fixed effects" uses the panel dimension of the PSID to control for household-specific fixed effects. The respective dashed lines are the corresponding three-year averages. The proximity of the two series suggests that, in the case of food expenditure over the lifecycle, a synthetic panel with cohort and family size controls has similar properties to a true panel with fixed effects.

Figure R8: Lifecycle Profiles of Three Sub-aggregates: NIPA Adjustments



(a) Mean Expenditure



(b) Cross-Sectional Variance

Notes: This figure replicates figure 2 from the text including a proportional adjustment to bring the CEX in line with the national income accounts personal consumption expenditure (PCE). Specifically, we rescale each series using the ratio of CEX to PCE expenditure reported at the subaggregate level by the BLS. The dashed lines are the unadjusted series. The BLS reports a disaggregated ratio for only a subset of years, and the benchmark is recalculated using only the available years as well. These series have cohort and demographic controls, but not year dummies given the restricted number of years available.

Figure R9: Effect of Work Status on Expenditure Shares



(a) Transportation



(b) Food Away From Home

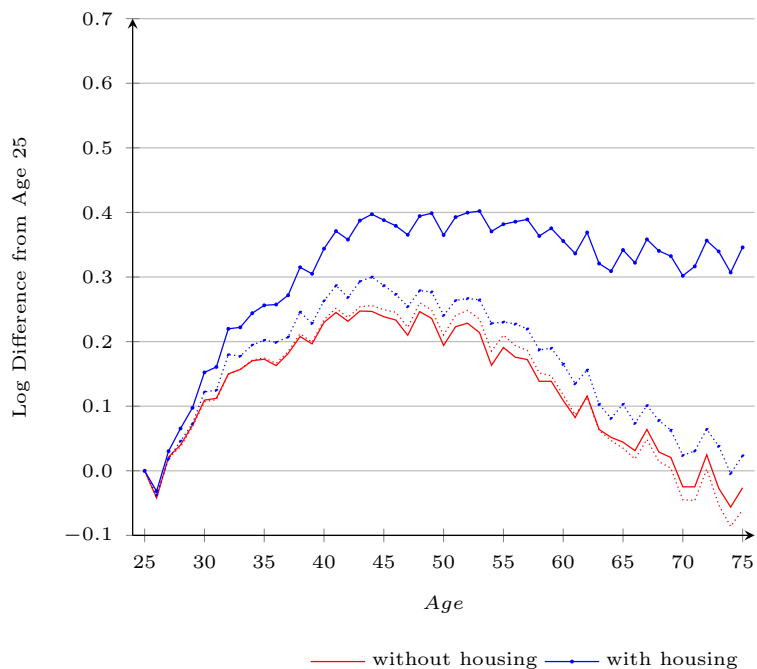
Figure R9: Effect of Work Status on Expenditure Shares (cont'd)



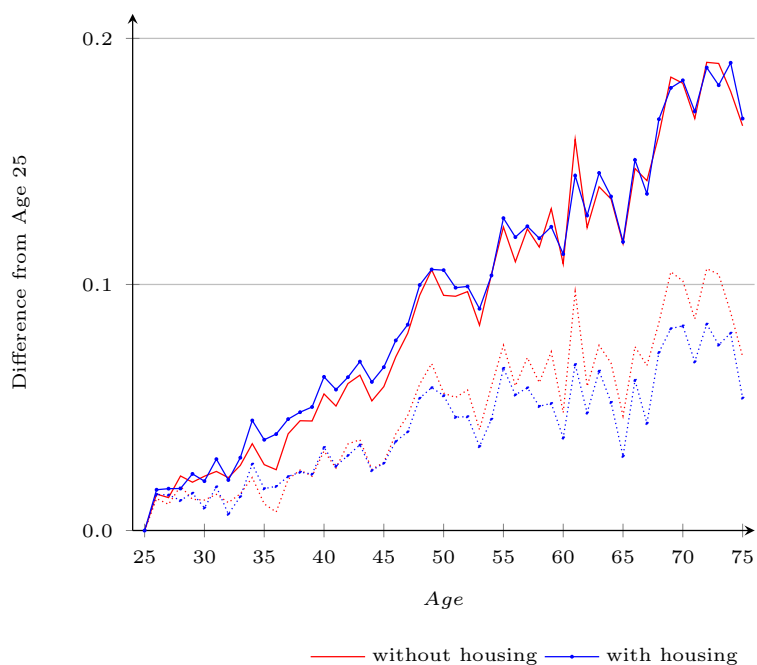
(c) Clothing

Notes: This figure depicts the age coefficients from a regression of expenditure shares on age, cohort, family size, price, and total expenditures, with (solid) and without (dashed) work status controls. Specifically, we regress the indicated category's share of expenditure out of total nondurables (including housing services and excluding alcohol and tobacco) on the standard vector of controls used in Figure 1 of the main text, plus the full set of (log) price indices for each category and the overall CPI deflator, as well as log total nondurable expenditure (see equation R1 in the appendix). We instrument total nondurable expenditure with a flexible function of total household income and dummies indicating the educational attainment of the head ( $\mathbb{1}_{12}$ ,  $\mathbb{1}_{12}$ ,  $\mathbb{1}_{1315}$ ,  $\mathbb{1}_{16}$ ). The figures depict the coefficients on the age dummies. The dashed line differs from the solid line in that the dashed line represents age coefficients in a regression that does not include a vector of work status controls consisting of dummies for the number of hours worked per week and dummies for the number of weeks work during the past year, for both spouses. For the regression analysis, we restrict our sample to include only married households, which is a sample of 34,195 households. Panels (a), (b) and (c), depict the share of expenditure spent on nondurable transportation, food away from home, and clothing, respectively.

Figure R10: Lifecycle Profiles of Nondurable Expenditures: Complete Year Controls and No Cohort Controls



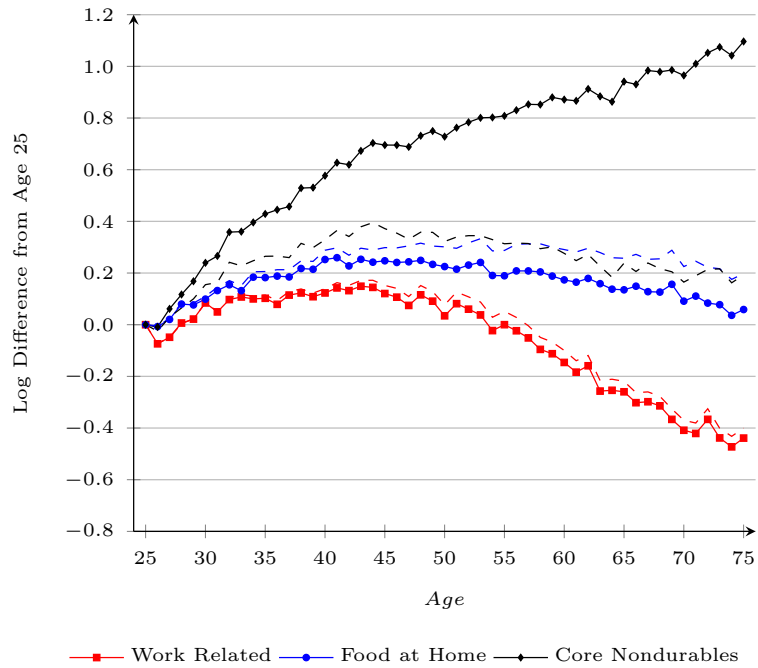
(a) Mean Nondurable Expenditures



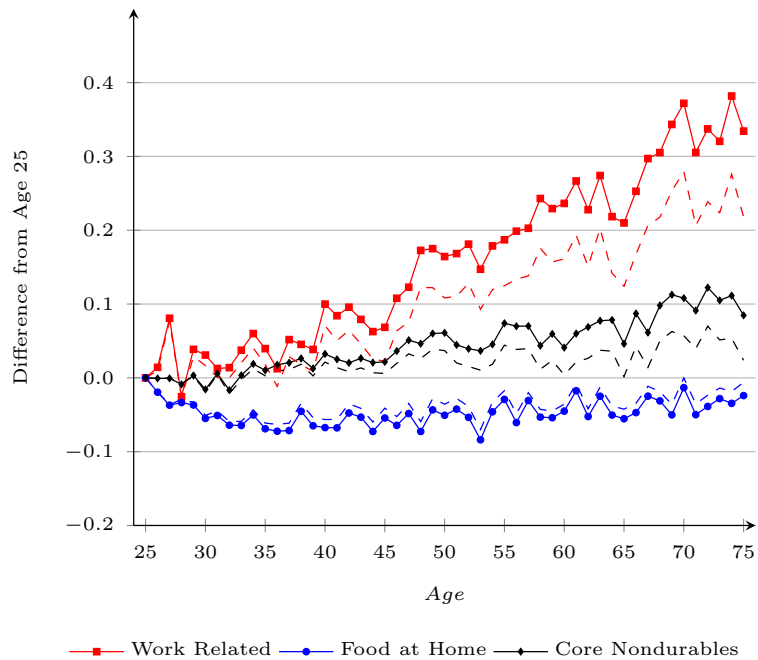
(b) Cross Sectional Variance of Nondurable Expenditures

Notes: This figure replicates figure 1 from the text with the full set of year controls and no cohort controls (dotted lines) as well as the benchmark series (solid lines).

Figure R11: Lifecycle Profiles of Three Sub-aggregates: Complete Year Controls and No Cohort Controls



(a) Mean Expenditure

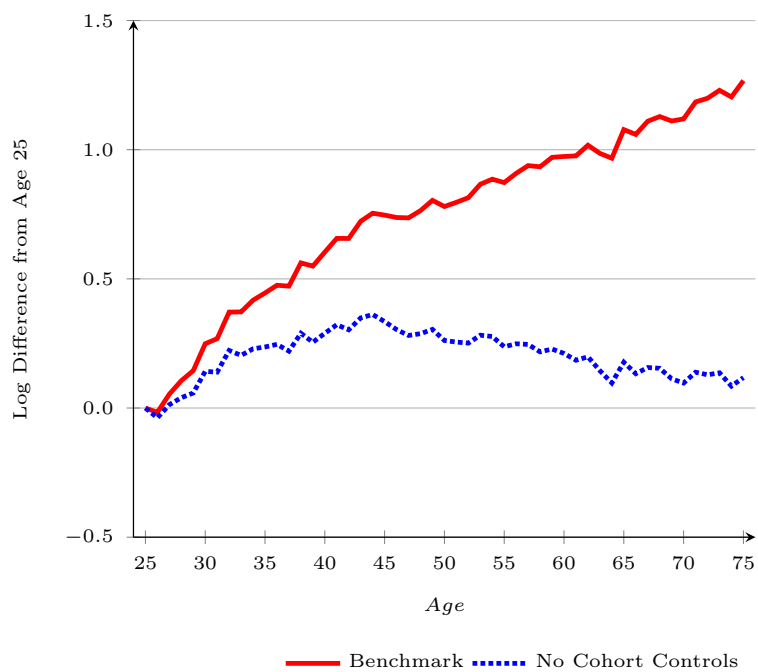


(b) Cross-Sectional Variance

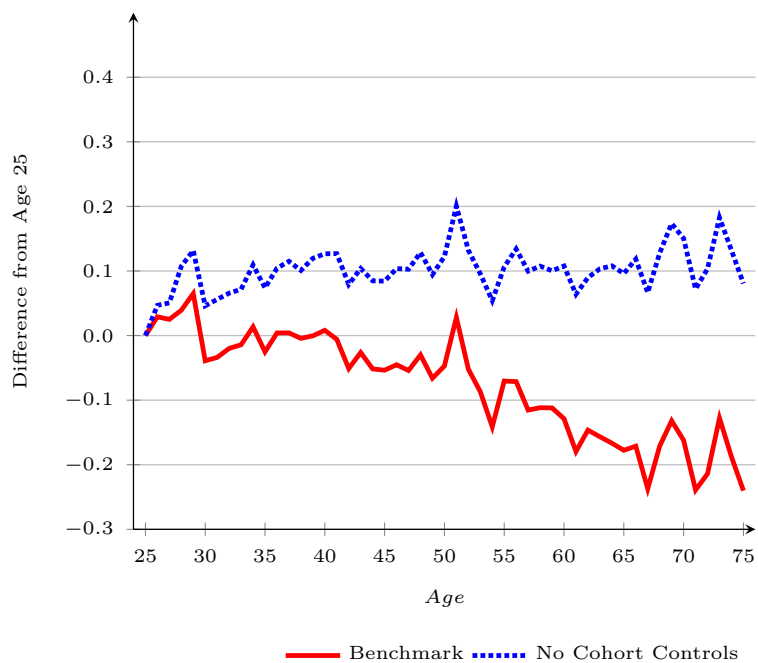
Notes: This figure replicates figure 1 from the text with the full set of year controls and no cohort controls (dotted lines) as well as the benchmark series (solid lines).



Figure R12: Lifecycle Profiles for Housing Services: The Importance of Cohort Controls



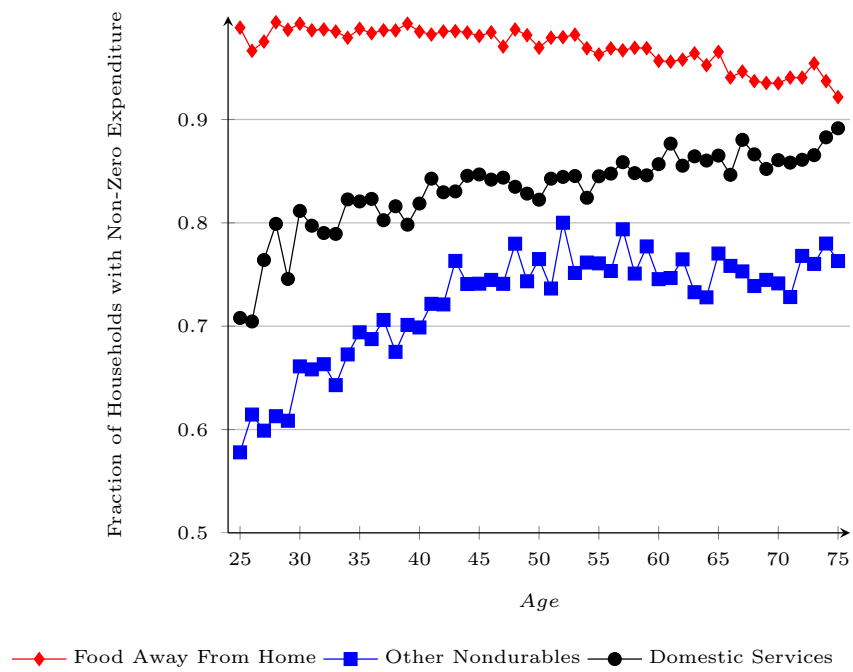
(a) Mean Housing Expenditures



(b) Cross-sectional Variance of Housing Expenditures

Notes: This figure depicts the lifecycle profile for housing services with the full set of year controls and no cohort controls (dotted lines) as well as the benchmark series (solid lines).

Figure R13: Fraction of Households with Non-Zero Expenditure on Three Small Categories

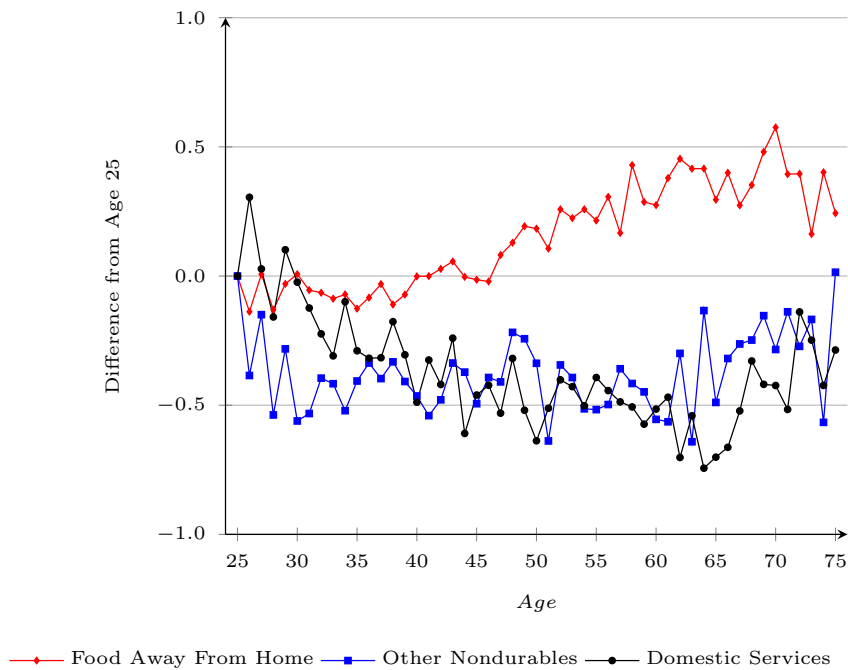


Notes: This depicts the fraction of households with non-zero expenditure at each age for food away from home (red, diamonds), other nondurables (blue, squares), and domestic services (black, stars).

Figure R14: Mean Life Cycle Expenditure for Small Categories: Excluding Zeros



(a) Mean Expenditure



(b) Cross-Sectional Variance

Notes: This figure replicates figure 2 in the text for the small categories which may have respondents recording zero expenditure in a year. Respondents with zero expenditure were dropped in this robustness exercise. They were bottom coded at one dollar in the benchmark.