Robustness Appendix for "Deconstructing Lifecycle Expenditure"
Mark Aguiar and Erik Hurst

This appendix presents results documenting the robustness of the primary results of "Deconstructing Lifecycle Expenditure," as well as several additional results not included in the paper. In Section 1, we present cohort-by-cohort results to augment the pooled results presented in the text. In Section 2, we test whether the results are stable across time. Specifically, we divide the sample in half by year of survey: 1980—1991 and 1992—2003. In Section 3, we discuss alternative controls for family size. In particular, we consider collapsing the dataset to demographic cells and tracking the within-cell lifecycle pattern of expenditure; using a crude "adult equivalence" scale popular in the literature; and using panel data from the PSID to estimate fixed effect regressions to test the validity of our benchmark controls in the case of food expenditures. In Section 4, we use the PSID to track households broken down by life cycle labor supply. This exercise documents the relationship between food expenditure and the opportunity cost of time, while controlling for selection into work status. In Section 5, we consider whether the results are sensitive to the paper's use of product-level deflators. In particular, in the robustness exercise we use a common deflator across all expenditure categories. In Section 6, we explore whether the results depend on whether we use log expenditure versus levels. The levels specification allows us to include households with zero expenditure on a category. Before proceeding with the details, we mention that the main results of the paper are robust across all of the alternative exercises.

To remind readers of our benchmark results, we depict mean log expenditure for several categories in Figure R1. Panel A plots categories that decline in the second half of the life cycle. These are the "Group 1" categories from Table 1 in the text, and include food, transportation, clothing and personal care, and other nondurables.[1] Panel B plots categories that show no declines over the life cycle (Group 2), namely, entertainment, housing services, utilities, domestic services, gambling receipts and charitable contributions. The figure depicts mean log expenditure relative to age 25, adjusted for family status and cohort effects, as discussed in the text.

Figure R2 depicts the benchmark cross-sectional dispersion of log expenditure over the life cycle relative to age 25. Panel A includes categories for which cross-sectional inequality increases over the life cycle: alcohol and tobacco, other nondurable, clothes and personal care, transportation, and gambling/charity. Panel B depicts those categories for which there is relative constant or declining inequality over the life cycle: food, domestic services, housing services, utilities, and entertainment. The dispersion measure is the "residual" standard deviation adjusted for cohort effects discussed in the text. Specifically, for each cohort-age pair, we compute the standard deviation of the residuals from a regression of log expenditure on age, family status, and cohort controls. We then regress the standard deviations on age and cohort dummies. The figure plots the coefficients on the age dummies.

[1] We omit alcohol and tobacco from the plots because its large decline over the lifecycle overshadows the movements of the other categories. However, in none of the robustness exercises does alcohol and tobacco exhibit a life cycle pattern that differs markedly from the benchmark.

1. Cohort Analysis

Recall from the paper, our primary specification for mean household expenditure is:

$$\ln(C_{it}^{k}) = \beta_0 + \beta_{age} Age_{it} + \beta_c Cohort_{it} + \beta_{fs} Family_{it} + \varepsilon_{it}^{k}, \quad (1)$$

where C_{it}^{k} is expenditure (in 2000 dollars) of household *i* during year *t* on consumption category *k*, Age_{it} is a vector of 50 one-year age dummies (for ages 26-75), $Cohort_{it}$ is a vector including eleven five-year age of birth cohort dummies, and $Family_{it}$ is a vector of family structure dummies that include a marital-status dummy and 10 household size dummies. The paper focused on the coefficients on the age dummies, β_{age}, which represent the impact of the lifecycle conditional on cohort and family size fixed effects.

The cohort dummies controlled for the shifts in the level of expenditure over the lifecycle. However, the specification imposes a common shape over the lifecycle across cohorts. That is, β_{age}, is constant across cohorts. In this section, we run an alternative to specification (1):

$$\ln(C_{it}^{k}) = \beta_0 + \beta_{age} Age_{it} * Cohort_{it} + \beta_{fs} Family_{it} + \varepsilon_{it}^{k}, \quad (2)$$

where *Age*Cohort* indicates age-cohort interactions. That is, each cohort has its own set of age dummies. The age coefficients for food and entertainment are depicted in Figure R3 Panels A and B, respectively. Each cohort is anchored at the mean log expenditure for the youngest age (or age 25) of each cohort contained in the data.

While the cohort-by-cohort plots are naturally very noisy given the much smaller sample sizes, they indicate that the "hump" seen in food and the lack of any decline in entertainment expenditure is consistently present across cohorts.

One perhaps surprising feature is that recent cohorts spend less on food than earlier cohorts at similar ages. This runs counter to the fact that recent cohorts have higher income than earlier cohorts. However, as documented by Laitner and Silverman (2005), the CEX has increasingly underestimated expenditure relative to the national income accounts. That is, under-reporting of expenditure in the CEX appears to be more severe in the recent waves of the CEX. This is true for food, but even more so for recreational expenditures, which includes nondurable entertainment (Table 1 of Laitner and Silverman). As the latter part of the life cycle tends to (relatively) reflect the recent waves of the CEX for a given cohort, this feature of the survey suggests we may be over-estimating the decline in food expenditure and under-estimating the increase in entertainment expenditure observed after middle age.

2. Sensitivity to Survey Time

A related issue is whether the results differ across different waves of the survey. That is, are the patterns documented in the paper confined to particular years? The fact that the results hold consistently across years for food and entertainment can be gleaned from Figure R3, given that a cohort-age pair reflects a particular year of the survey. For additional clarity, Figure R4 depicts life cycle expenditure for each category separately for 1980—1991 (Panels A and C) and 1992-2003 (Panels B and D). Group 1 (Panels A and B) goods are those which decline after middle age in the benchmark analysis, while

Group 2 (Panels C and D) are those that exhibit no decline. These classifications correspond to those in Table 1 of the main text. Figure R4 indicates that the patterns are stable over survey year; while the magnitudes of the movements over the life cycle vary somewhat across years, the general shape of life cycle expenditure does not vary for any categories.

Figure R5 depicts the dispersion of life cycle expenditure for the two time periods. Panels A and B depict those goods which exhibited an increase in dispersion over the life cycle in the benchmark analysis presented in the text (See Table 2 of the main text). Panels C and D are those goods which exhibited no increase (or even a decrease) in cross-sectional inequality over the life cycle. As with the cohort-by-cohort plots, there is some variation in magnitude across years. Most notably, alcohol and tobacco and housing services are not particularly stable across the sampling periods. Other than these categories, the evidence indicates the over all shape of life cycle dispersion is consistent across sampling periods.

3. Alternative Controls for Family Size

As seen in equation (1) above, our benchmark controls include a set of family status dummies to account for changing household composition over the lifecycle. In this section, we consider alternative approaches to family size.

We begin by simply normalizing expenditure by the square root of the number of household members. This approach has been frequently used in the literature.[2] We then regress log normalized expenditure on cohort and age dummies. The results for mean

[2] Unfortunately, the CEX does not allow us to adjust by the age of the family members. However, we discuss this issue in more detail below and argue that our results are not sensitive to this omission.

expenditure are depicted in Figure R6. Panel A in Figure R6 depicts our "Group 1" categories, which show a decline in the second half of the life cycle in the benchmark analysis, while Panel B depicts "Group 2" categories. We see that the alternative controls for family size do not reverse the general pattern of expenditure.

Similarly, Figure R7 depicts the cross-sectional standard deviation of expenditure. Panel A are those goods that showed an increase in dispersion over the life cycle in the benchmark results, while Panel B are those goods which showed no increase. Again, the patterns are not sensitive to the alternative family size controls.

A second robustness experiment assigns households to 110 demographic cells defined by age (3-year ranges), family size (0 through 5^{+}), and cohort. (Note that not all potential cells defined by these interactions are represented in the data set). For each cell, we compute mean expenditure on each category. We then document the within-cell variation over the life cycle. Specifically, we run a fixed-effects regression of log expenditure on age dummies. The coefficients on the age dummies are depicted in Figure R8. As seen in Panels A and B of this figure, the shape of life cycle expenditure is invariant to alternative family size controls.

One issue with our use of age dummies is that family size may be correlated with omitted variables, such as permanent income, as well as age. This may bias our age coefficients. Some of this concern is addressed by the demographic cell approach, which in effect allows family size controls to vary by age. However, we are still left with the problem that family size is ultimately a choice variable which may be correlated with omitted variables.

Ideally, we would like to control separately for a household fixed effect as well as changing family size. This requires a panel. To test the sensitivity of our methodology to these issues, we use data from the Panel Study of Income Dynamics (PSID), a panel data set which tracks food expenditure over long periods of time. Specifically, we regress log food expenditure from the PSID on age dummies, family size controls, and a household fixed effect using a sample of household from the 1980-1992 PSID. We place little restrictions on the PSID sample aside from restricting the sample to households with heads aged between 25 and 70 (inclusive). Our analysis excludes the SEO over-sample of poor households.[3]

Using the PSID sample, we estimate two sets of regressions. The results are shown in Figure R9. The first is identical to specification (1) from the main text. Specifically, we regress log food expenditures on a series of age, family size, and cohort dummies (without individual fixed effects). The age and family size controls are nearly identical to the controls used in the CEX estimation reported in the text. The second specification is identical to the first except we include individual fixed effects in the regression and exclude the cohort controls.

As seen in Figure R9, we find that within the PSID data the results are consistent across the two specifications, indicating that the results (at least concerning food expenditure) are not sensitive to the presence of a fixed effect control (relative to including cohort controls). Given that the PSID only has panel information on food expenditures (out of the broad class of non-durable consumption), we are forced to use the CEX to compute the lifecycle profile of expenditures across multiple non-durable

[3] The PSID did not collect food expenditure data in 1988 and 1989. Obviously, these two years are excluded from our sample. Our total sample size included 32,749 observations (with 4,960 distinct households).

consumption categories. Therefore, we find it very comforting that – at least for food – controlling for individual fixed effects and controlling for cohort fixed effects yields quantitatively similar results.

Additionally, the PSID also allows for more flexibility in controlling for the age of household members. Specifically, we include dummies indicating the age of children in the household. This accounts for the potential that younger family members have a different impact on desired household expenditure than older family members. Again, as seen in Figure R10, we find the life cycle pattern of food is not sensitive to these additional controls.

4. Additional Results on Employment Status and Food Expenditure

Our preferred explanation for the heterogeneity across categories in life cycle expenditure emphasizes the changing cost of time. A primary claim on time is market work. In this section, we use the PSID to examine life cycle expenditure on food across households with differing employment status. Such an analysis cannot be done in the CEX given that employment status at a given age is correlated with other factors that determine household consumption (like permanent income). Households that leave the labor force in their 50s (in many cases due to disability or other health shocks) are much more likely to have low permanent income. The panel dimension of the PSID, however, allows us to control for other features of a household which are correlated with employment (like permanent income) when assessing the degree to which household spending evolves differentially over the lifecycle for households who continuously work versus those that do not.

Robustness Figure R11 shows the lifecycle profile of food spending using data from the PSID for three samples: all households, all married households, and all married households where both spouses worked continuously over the sample period. Otherwise, the PSID sample restrictions were similar to those discussed above (sample years 1980 – 1992, etc.) although we restrict our analysis to only those households with heads between the ages of 25 and 64 (inclusive). For married households, we restrict the sample to households whose members maintained the same marital status throughout their time in the PSID sample (between the ages of 25 and 64 and between the years of 1980 and 1992). For the continuously working married sample, we impose the additional restriction that both the husband and wife remained continuously employed during their time in the sample, and that in the year that food consumption is measured, both the husband and wife were working full time (each working more than 1900 hours per year).

The results are shown in Robustness Figure R11. For the unrestricted married sample and the total sample, food expenditures follow the standard lifecycle profile documented in the text and in other robustness figures. However, for households in which both spouses were working continuously, there is no decline in food spending with age. In other words, the decline in food spending for the total married sample reflects one or both spouses exiting the labor force. According to the model proposed in the text, exiting the labor force would free up more time to home produce and shop – both of which could decrease food expenditures while keeping food consumption (and total marginal utility) fairly constant.[4]

[4] The relevant sample sizes for the three regressions summarized in Figure R11 are: 30,307 observations (4,683 distinct households) for the total sample, 21,371 observations (3,234 distinct households) for the continuously married sample, and 1,151 observations (375 distinct households) for the continuously married and both spouses continuously employed sample. There were only 50 distinct households in the

5. Alternative Deflators

In the benchmark analysis, we deflate each category's expenditure using a product-level deflator. This accounts for the fact that relative prices may change over the sample period. To see whether this is an issue in practice, we have re-done the main CEX analysis deflating all expenditure categories by a common price index, namely, the CPI-U deflator. Figures R12 and R13 present the results for mean and standard deviation, respectively. Again, the main patterns remain.

6. Levels versus Logs

In the benchmark CEX analysis we analyze log expenditure. This necessitates dropping households that record zero expenditure in a category. In this section, we use the level of expenditure as our unit of analysis, and then convert to percentage changes using mean expenditure for households whose head is 25 years old. Specifically, we run regression from equation (1), but replace log expenditure with the level of household expenditure on that category. The sample includes households who report zero expenditure on the category. The coefficients on the age dummies now represent the difference in levels between households of the respective age and those aged 25. We calculate percentage deviation from age 25 by dividing the estimated coefficient on the age dummy with mean expenditure for 25 year olds. The results are depicted in Figure R14.

---

continuously married sample where the husband continuously worked and the wife continuously did not work over the sample period. For this reason, we did not report the results for this sample – the point estimates on the age dummies were extremely noisy. We could not rule out whether the lifecycle profile of food expenditure for this sample increased dramatically, decreased dramatically or stayed the same.

Figure R1: Benchmark Life Cycle Expenditure

Panel A: "Group 1" Mean Expenditure

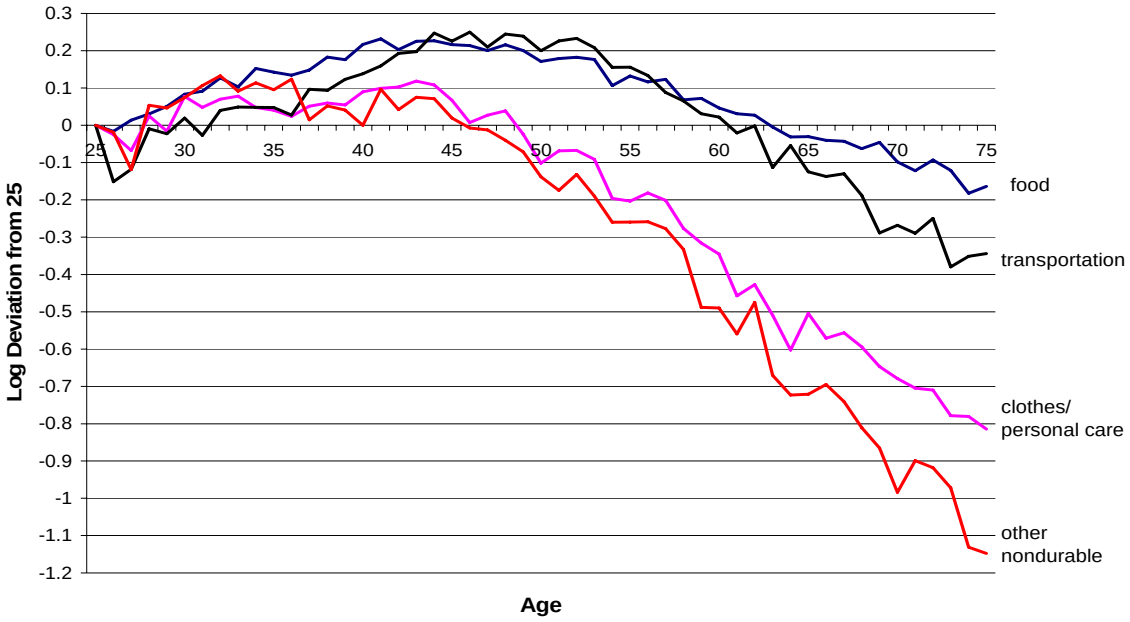

Panel B: "Group 2" Mean Expenditure

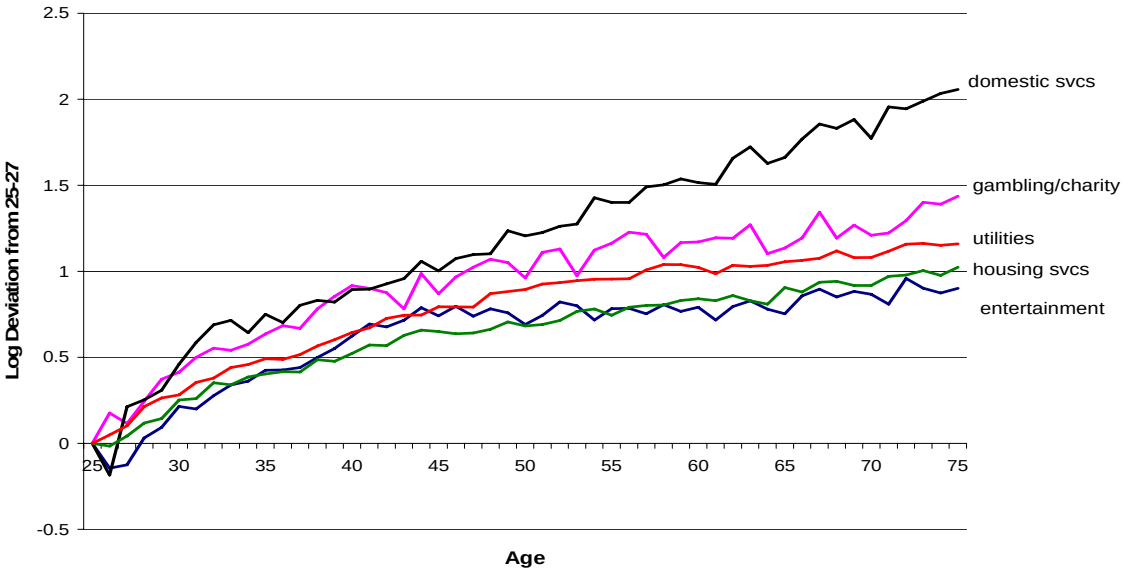

Notes: This figure depicts mean life cycle expenditure by category using the benchmark methodology from the text. The depicted profiles correspond to Figures 2 and 3 in the text.

Figure R2: Benchmark Life Cycle Expenditure -- Dispersion

Panel A: Increasing Categories

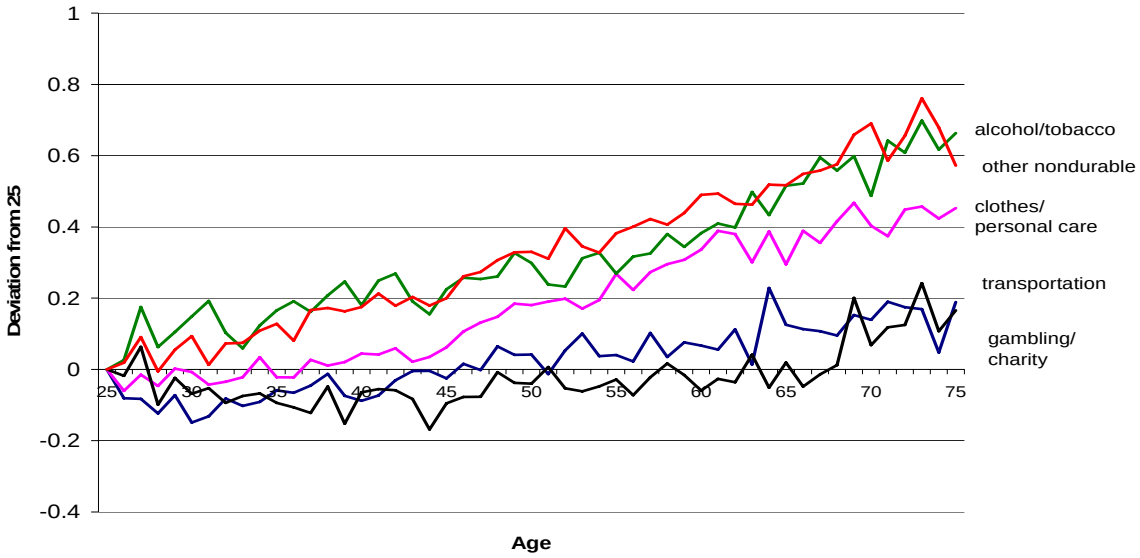

Panel B: Non-increasing Categories

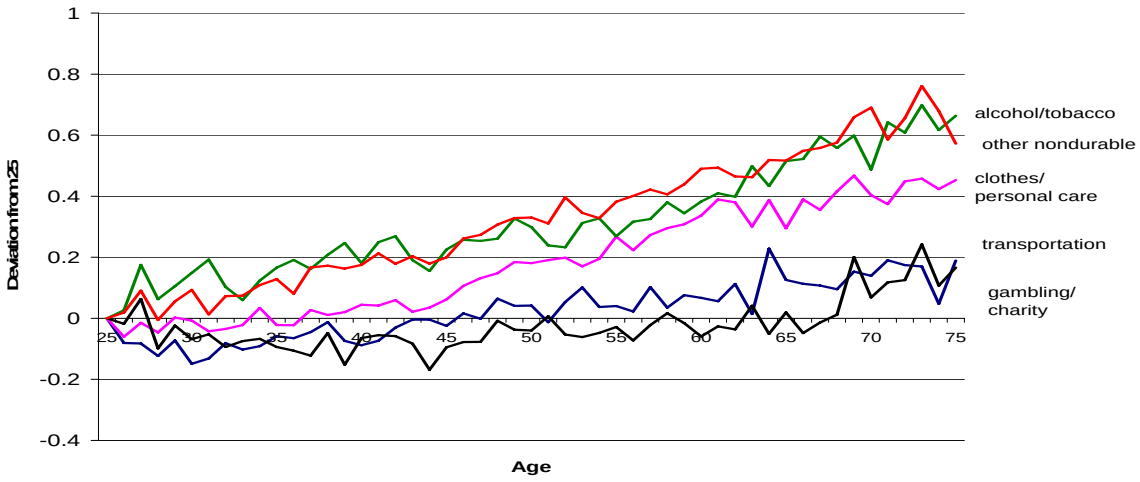

Notes: This figure depicts residual cross-sectional dispersion over the life cycle by category using the benchmark methodology from the text. The depicted profiles correspond to the methodology used to generate Figure 4 in the text.

Figure R3: Mean Log Expenditure by Cohort

Panel A: Food

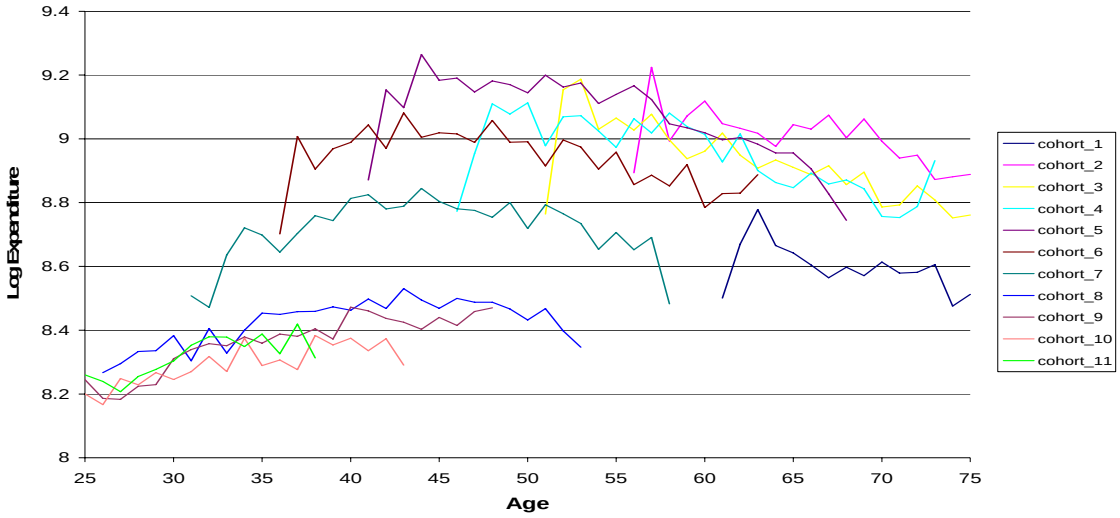

Panel B: Entertainment

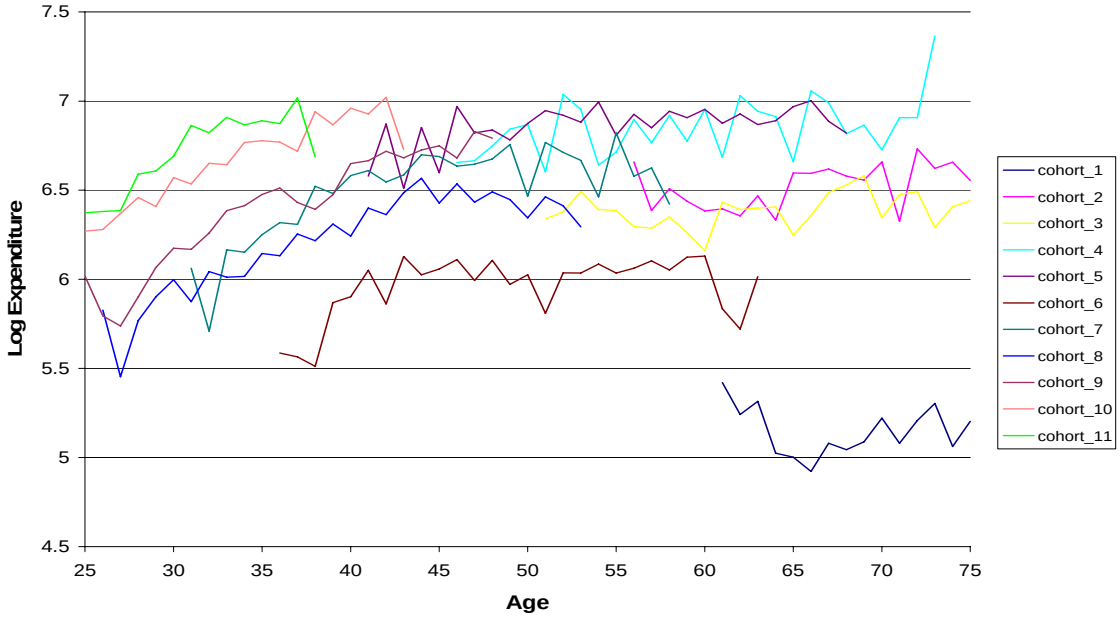

Notes: This figure depicts mean log expenditure on food (Panel A) and entertainment (Panel B), adjusted for family size. Each cohort line is anchored at unconditional mean log expenditure for the first age at which the cohort appears in the sample. The remainder of the line is computed using the coefficients on age dummies from a regression of log expenditure on age-cohort interactions and family status controls. Each cohort contains a five-year birth cohort, starting with those born between 1915 and 1919, inclusive, as cohort 1. The last cohort (cohort 11) is a four-year cohort consisting of those born between 1965 and 1968 inclusive.

Figure R4: Stability over Sample Periods--Means

Panel A: "Group 1" Mean Expenditure 1980--1991

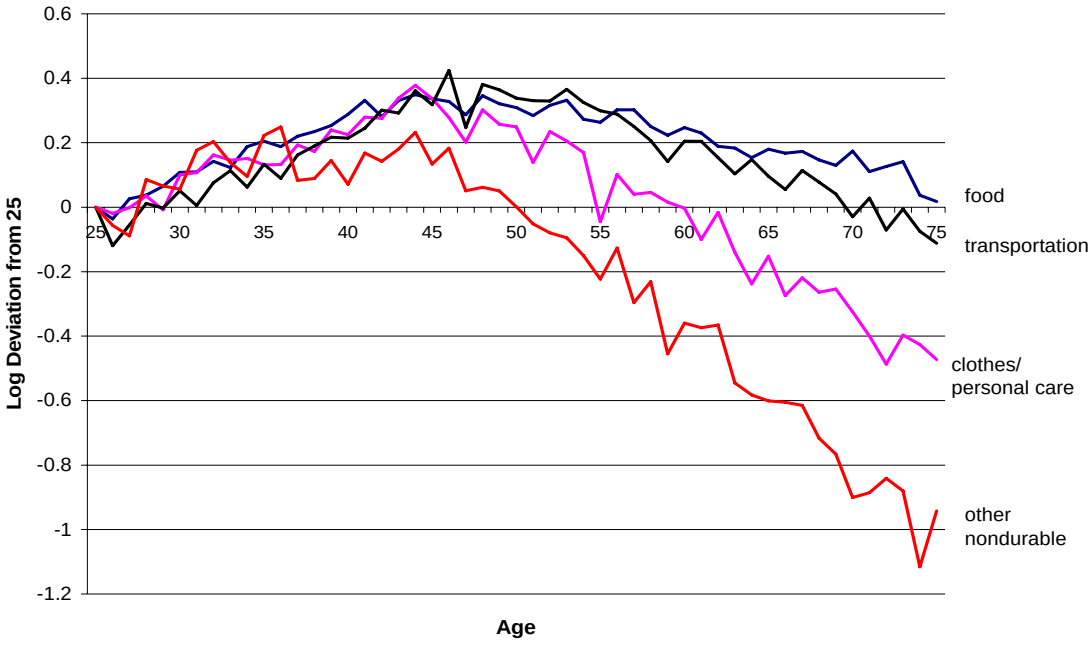

Panel B: "Group 1" Mean Expenditure 1992--2003

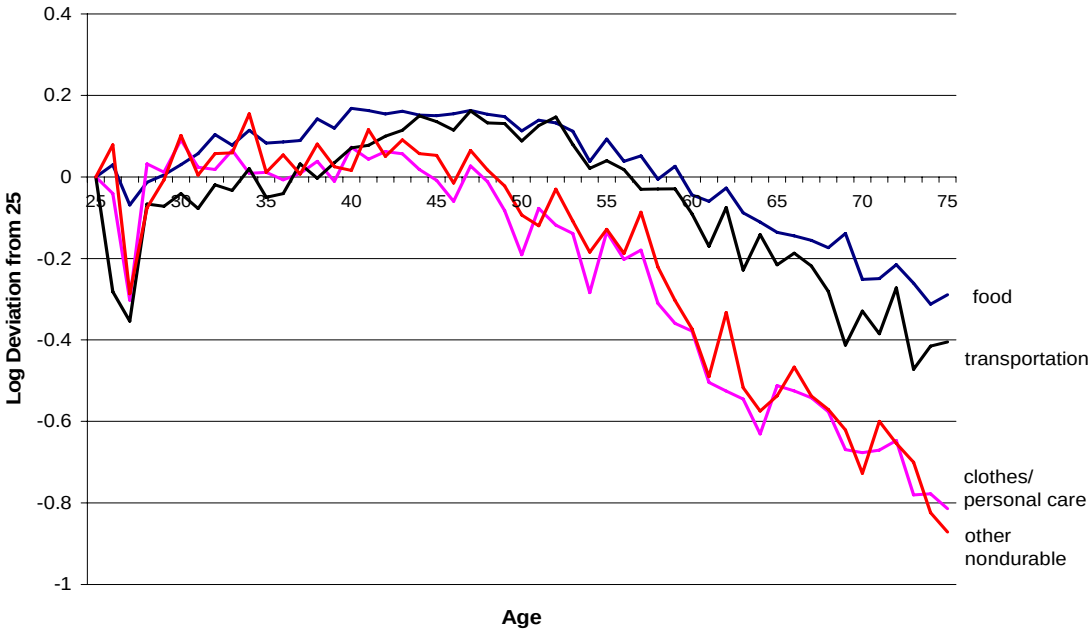

Panel C: "Group 2" Mean Expenditure 1980--1991

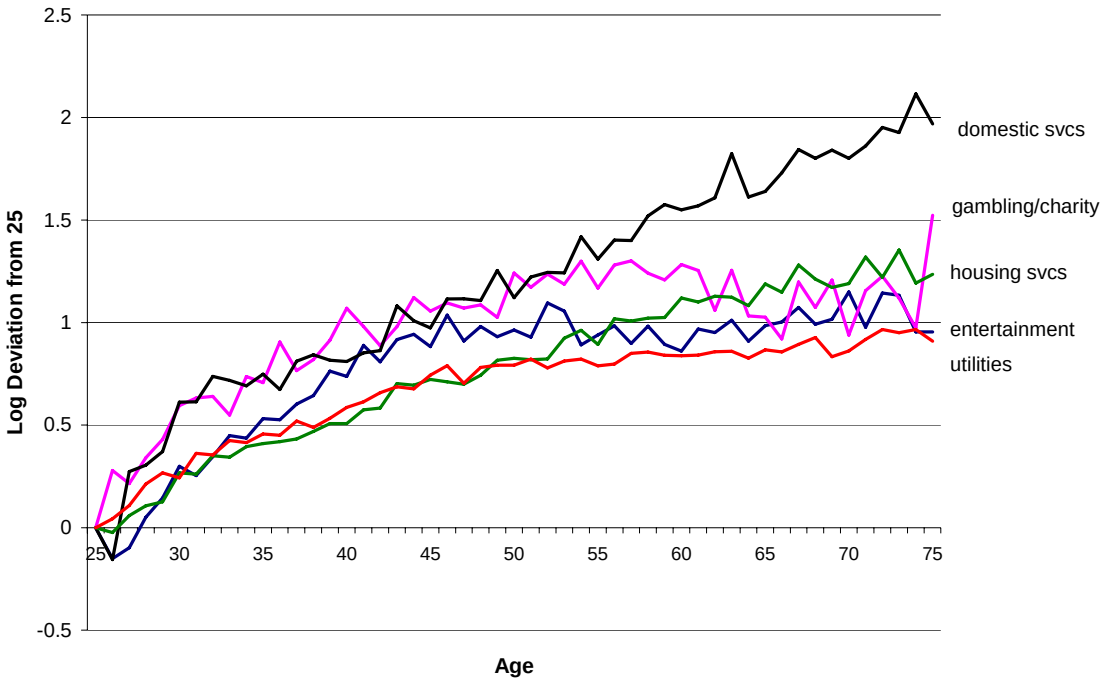

Panel D: "Group 2" Mean Expenditure 1992--2003

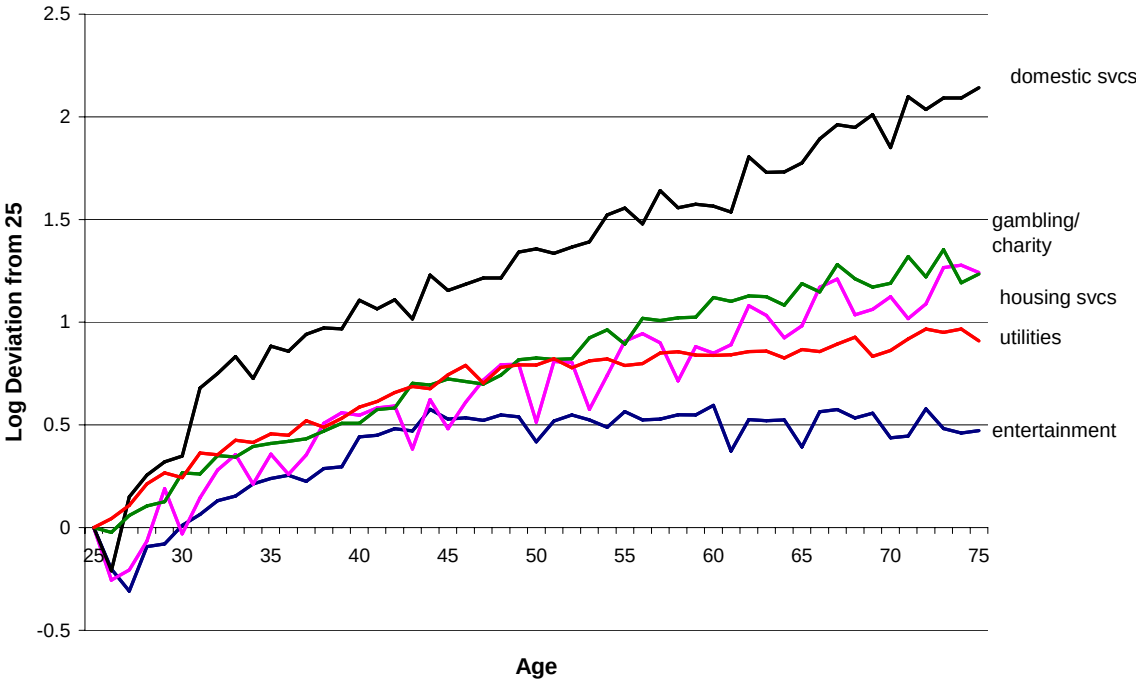

Notes: This figure uses the methodology of Figure A1, but splits the sample into two periods: 1980-1991 (Panels A and C) and 1992-2003 (Panels B and D).

Figure R5: Stability over Sample Periods—Residual Inequality

Panel A: Increasing Categories 1980--1991

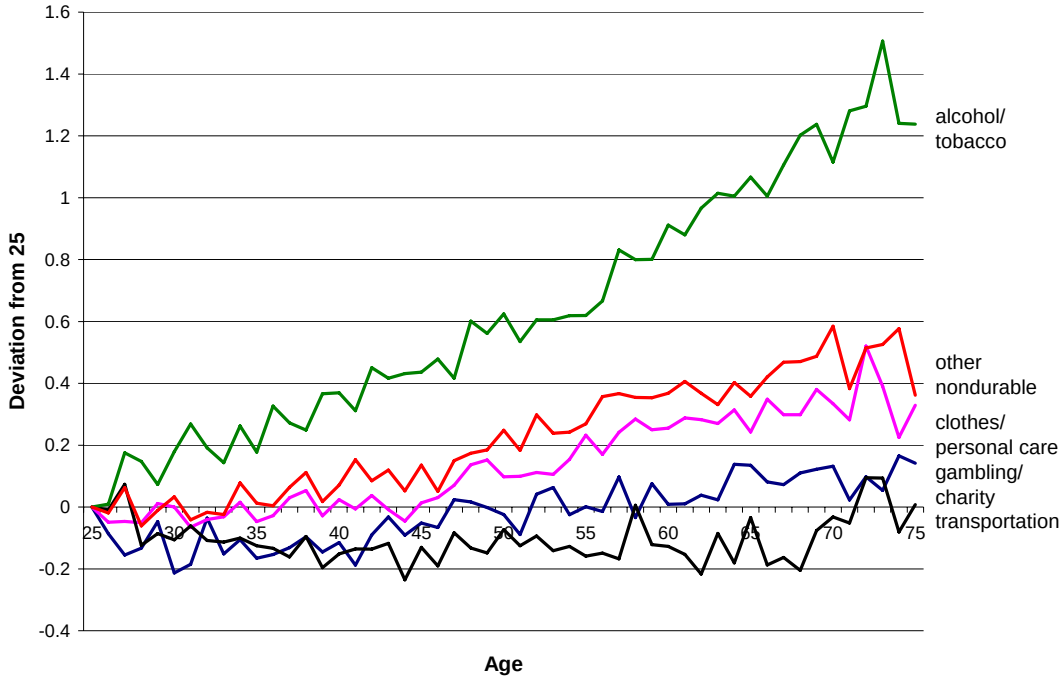

Panel B: Increasing Categories 1992—2003

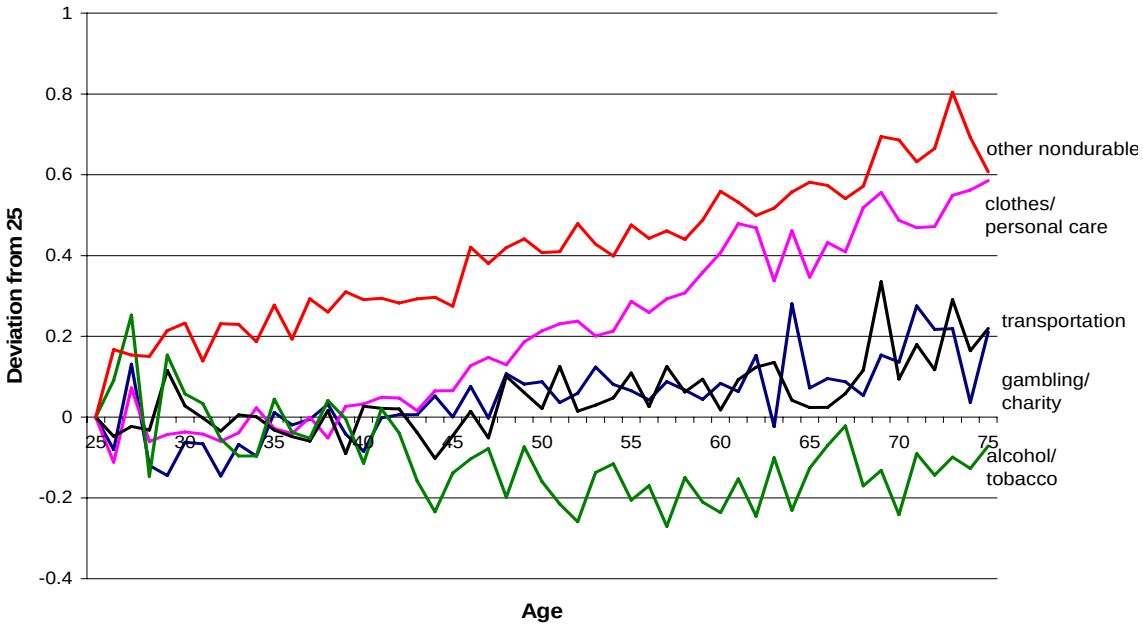

Panel C: Non-increasing Categories 1980--1991

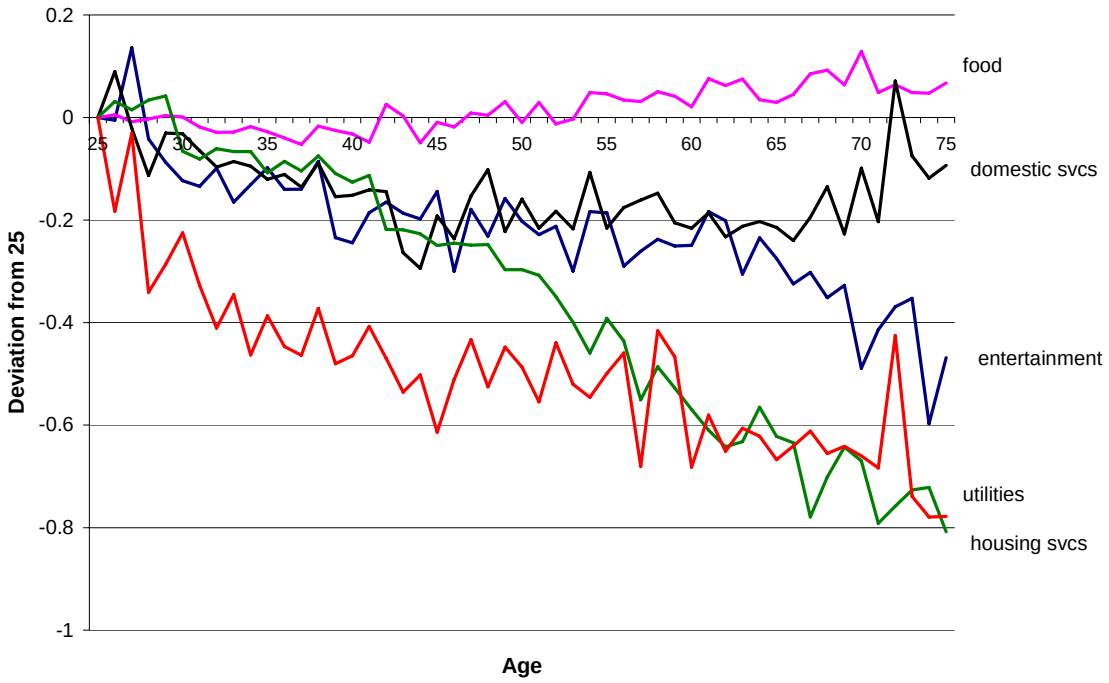

Panel D: Non-increasing Categories 1992—2003

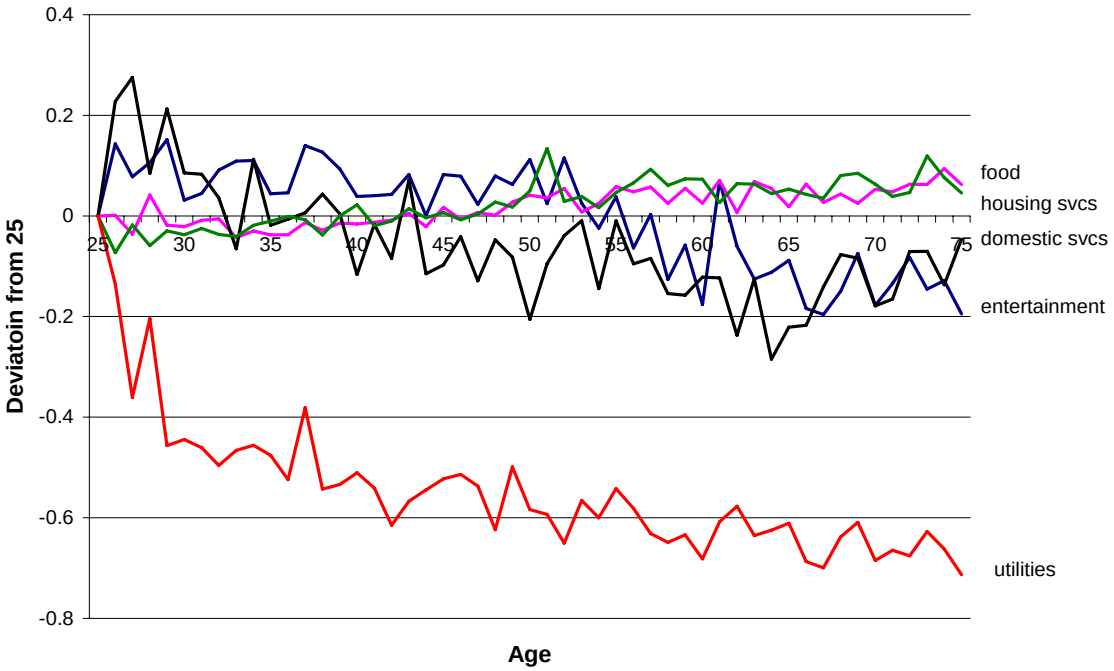

Notes: This figure uses the methodology of Figure A2, but splits the sample into two periods: 1980-1991 (Panels A and C) and 1992-2003 (Panels B and D).

Figure R6: Alternative Family Size Controls – Square Root

Panel A: "Group 1" Mean Expenditure

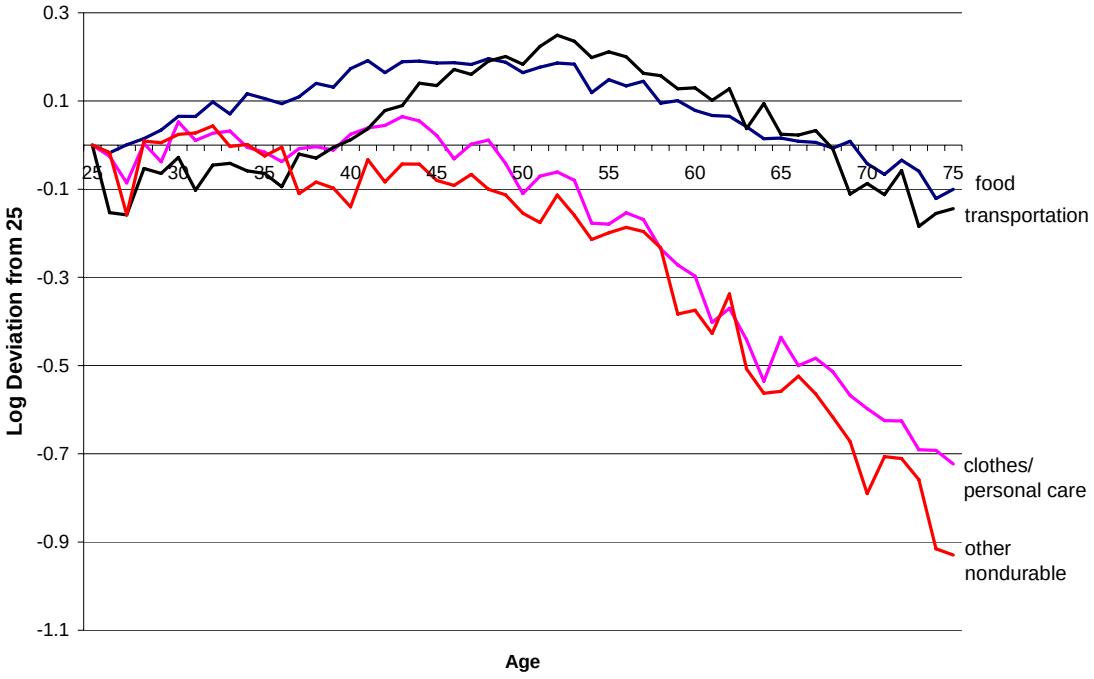

Panel B: "Group 2" Mean Expenditure

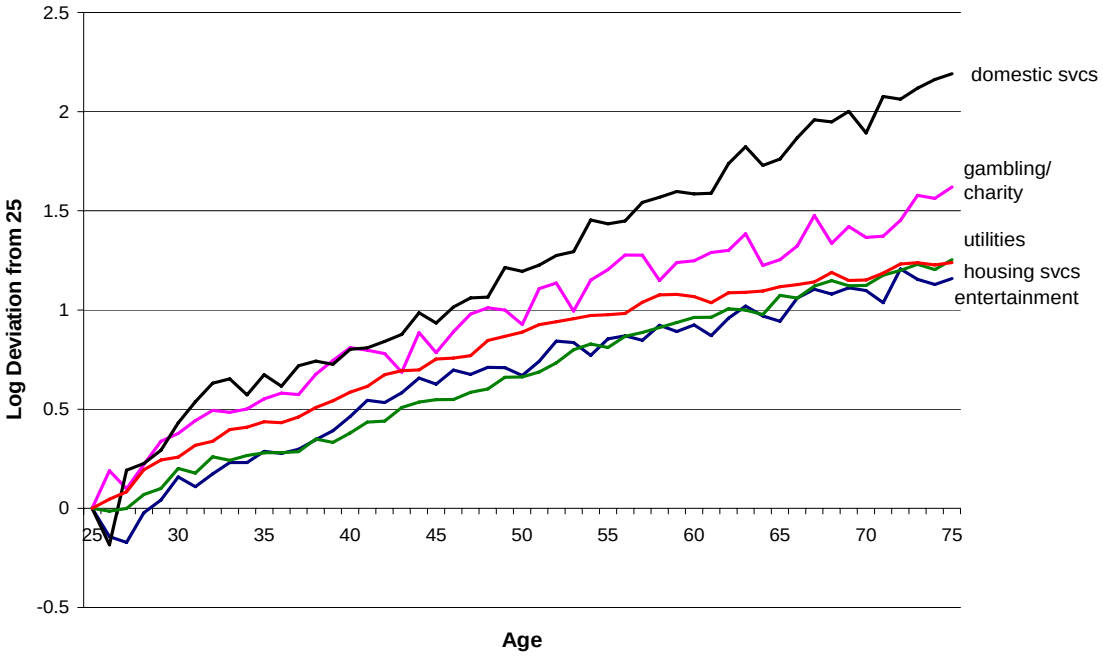

Notes: This figure depicts mean log expenditure normalized by the square root of family size and adjusted for cohort effects.

Figure R7: Alternative Family Size Controls – Square Root – Residual Inequality

Panel A: Increasing Categories

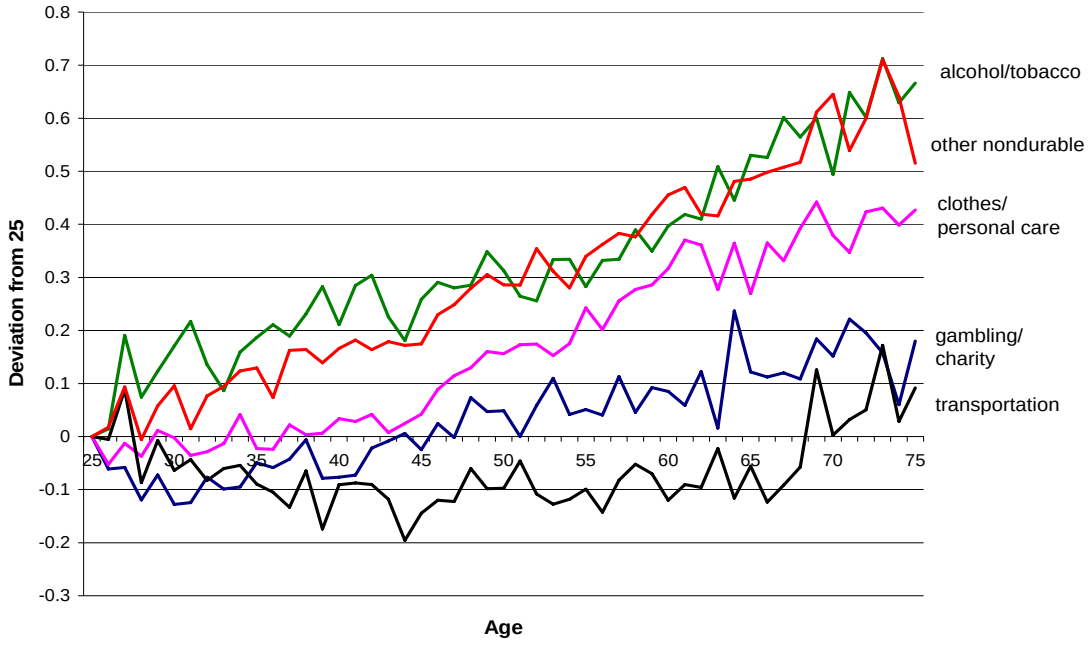

Panel B: Non-Increasing Categories

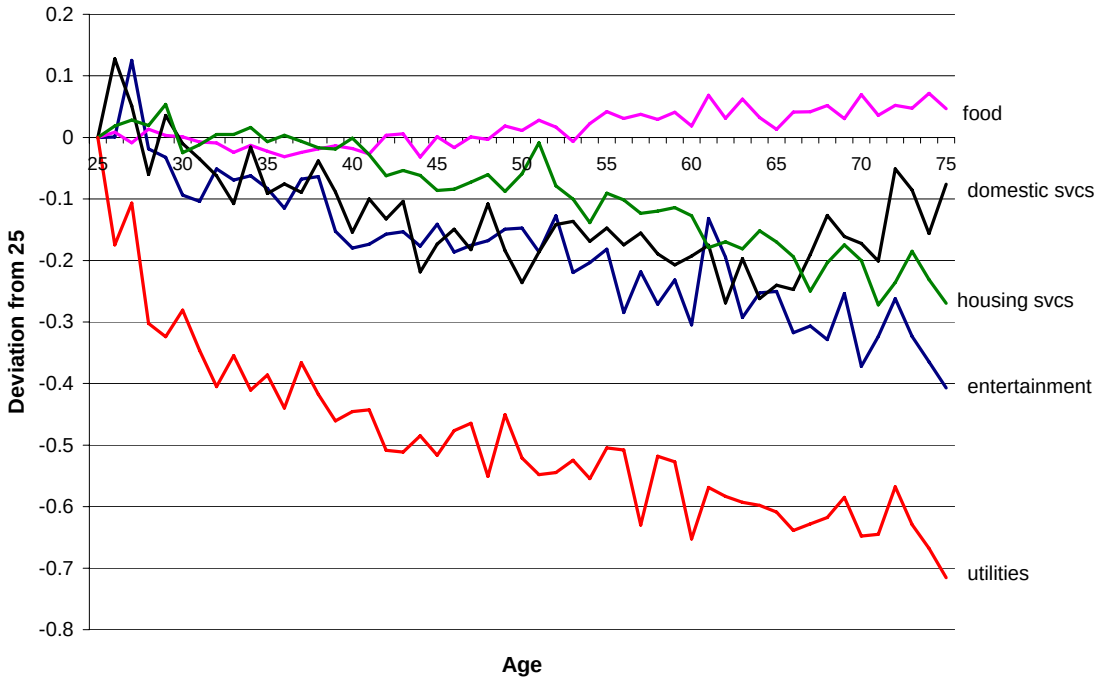

Notes: This figure depicts the life cycle pattern of residual dispersion (as calculated in Figure A2), with log expenditure normalized by the square root of family size rather than adjusted using family size dummies.

Figure R8: Alternative Family Size Controls – Demographic Cells

Panel A: "Group 1" Mean Expenditure

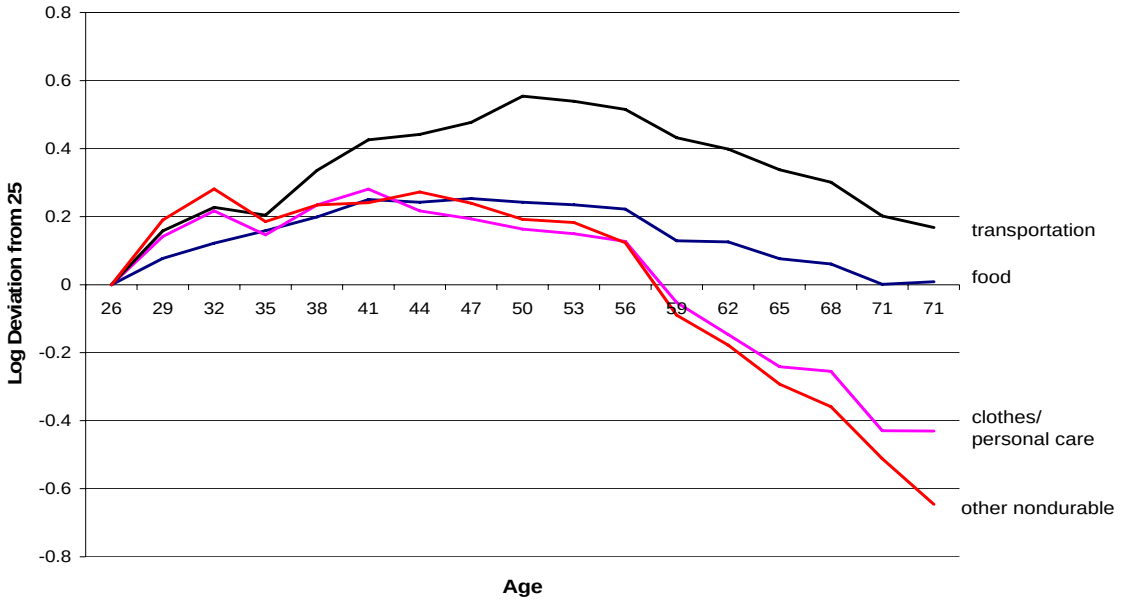

Panel B: "Group 2" Mean Expenditure

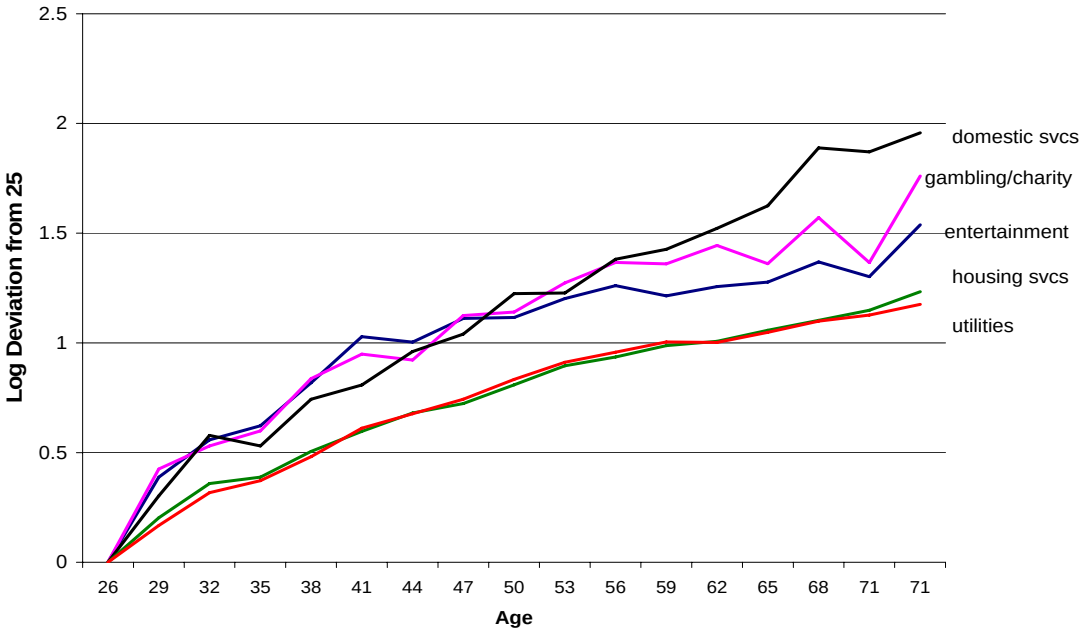

Notes: This figure depicts the evolution of mean log consumption within demographic cells. Demographic cells are defined by age, cohort, and family status. See text for details.

Figure R9: Cohort vs. Fixed Effects in the Lifecycle Profile of Food Expenditures Using PSID Data

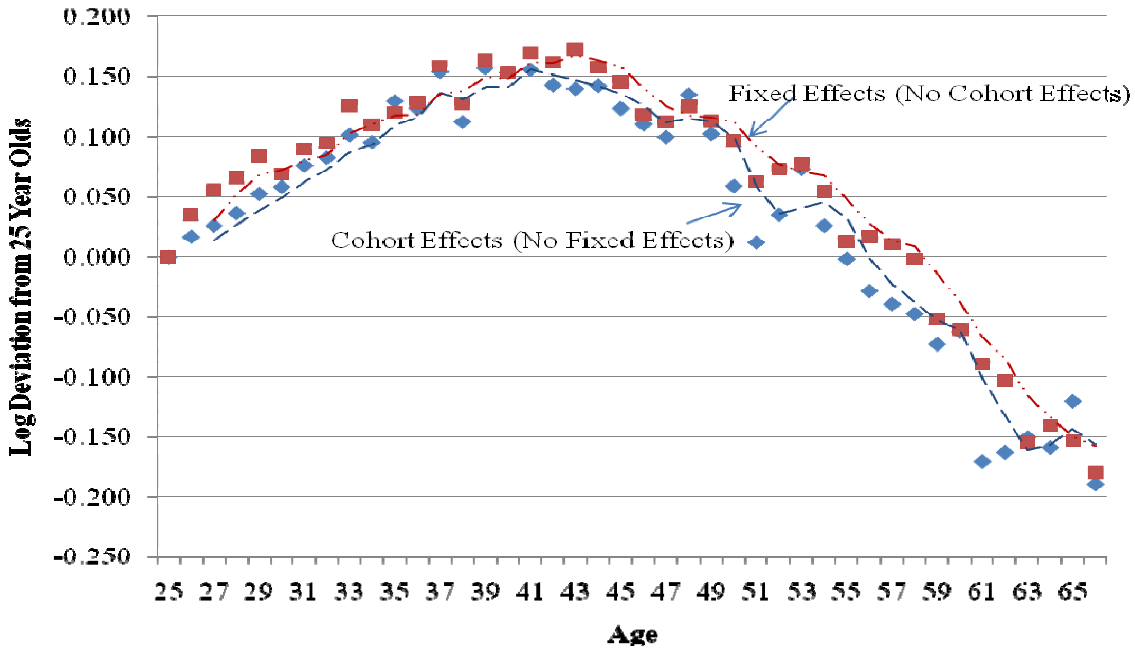

Notes: Figure shows the lifecycle profile of food expenditures in the PSID either with individual fixed effects (and no cohort effects) or with cohort effects (and no individual fixed effects). Both specifications control for family size dummies

Figure R10: Controlling for Child Age Effects on Lifecycle Profile of Food Expenditures Using PSID Data

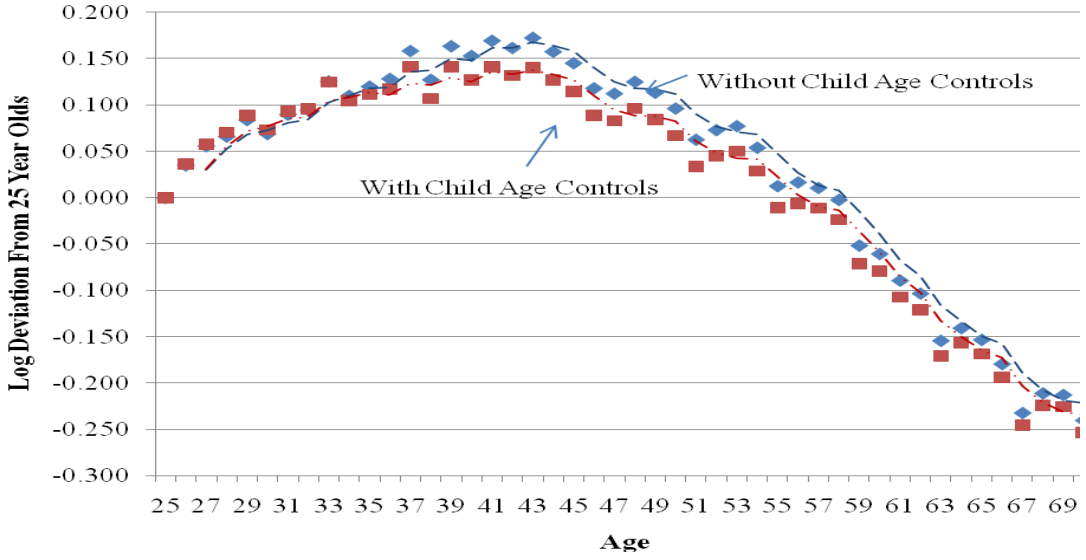

Notes: Figure shows the lifecycle profile of food expenditures in the PSID with and without controls for child age. The age coefficients come from a regression of log food sepnding on age dummies, individual fixed effects, and family size dummies (with and without child age controls).

Figure R11: Working vs. Nonworking Food Expenditure in the PSID

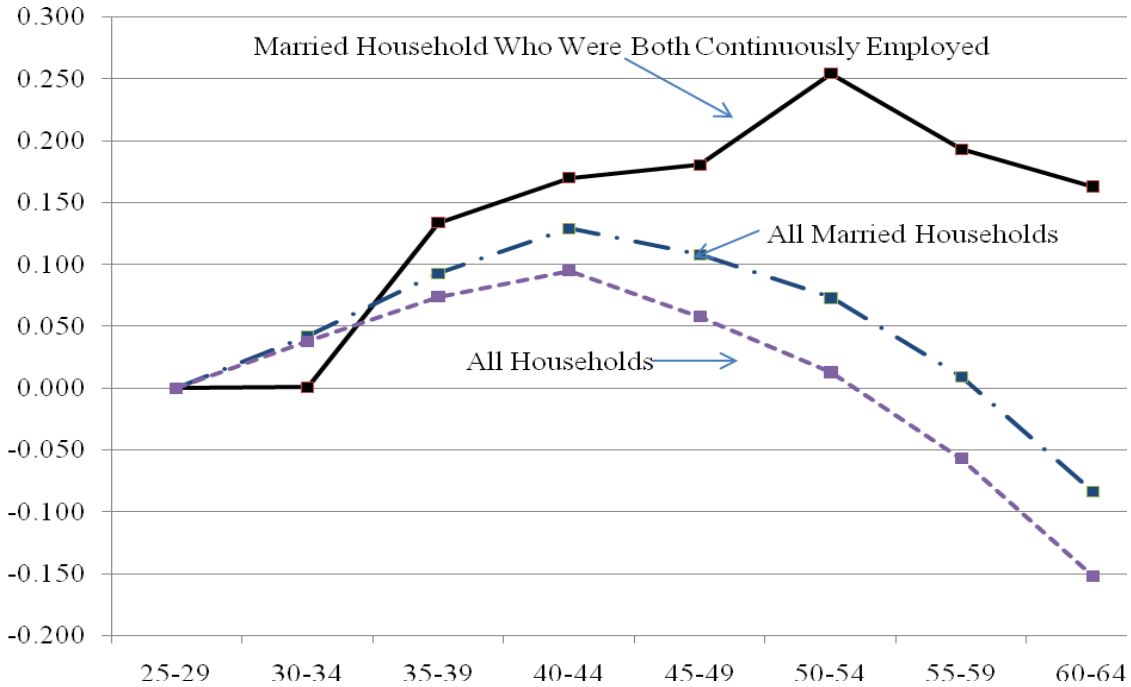

Notes: Figure shows the lifecycle profile of food expenditures in the PSID for three different samples: all households, all married households, and all marrried households were both spouses were continuously employed full time throughout the sample. The latter two groups also condition on being continuously married throughout the sample period. The age coefficients come from a regression of log food sepnding on age dummies, individual fixed effects, and family size dummies. For the continuously employed married sample, we also restricted the sample to only households where both the head and wife continuously worked throughout the sample and that there are currently working full time (more than 1900 hours per year).

Figure R12: Alternative Deflator – A Single Deflator for All Categories

Panel A: "Group 1" Mean Expenditure

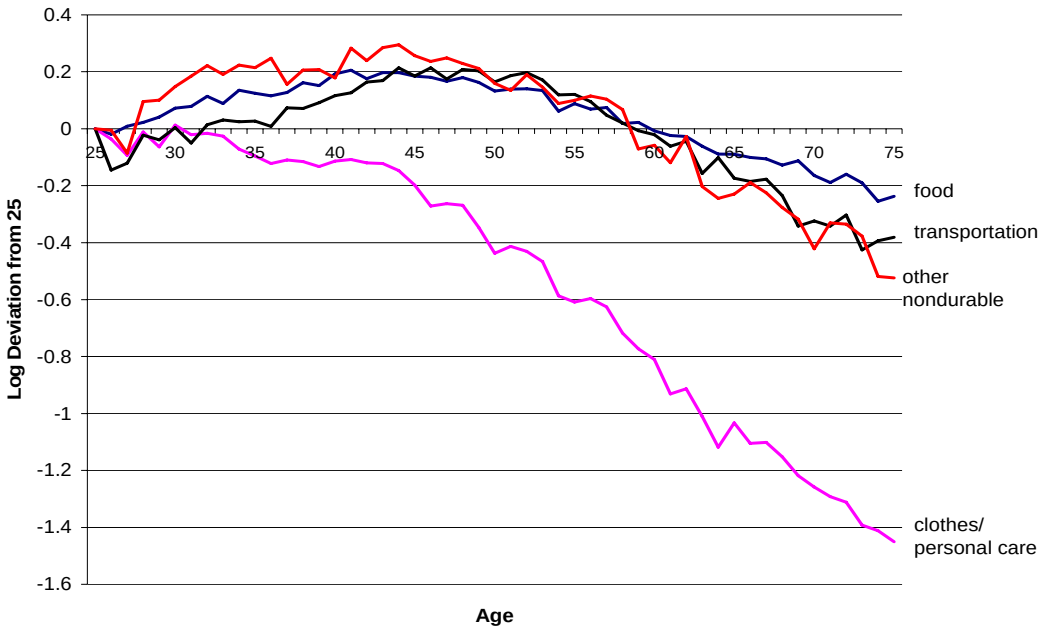

Panel B: "Group 2" Mean Expenditure

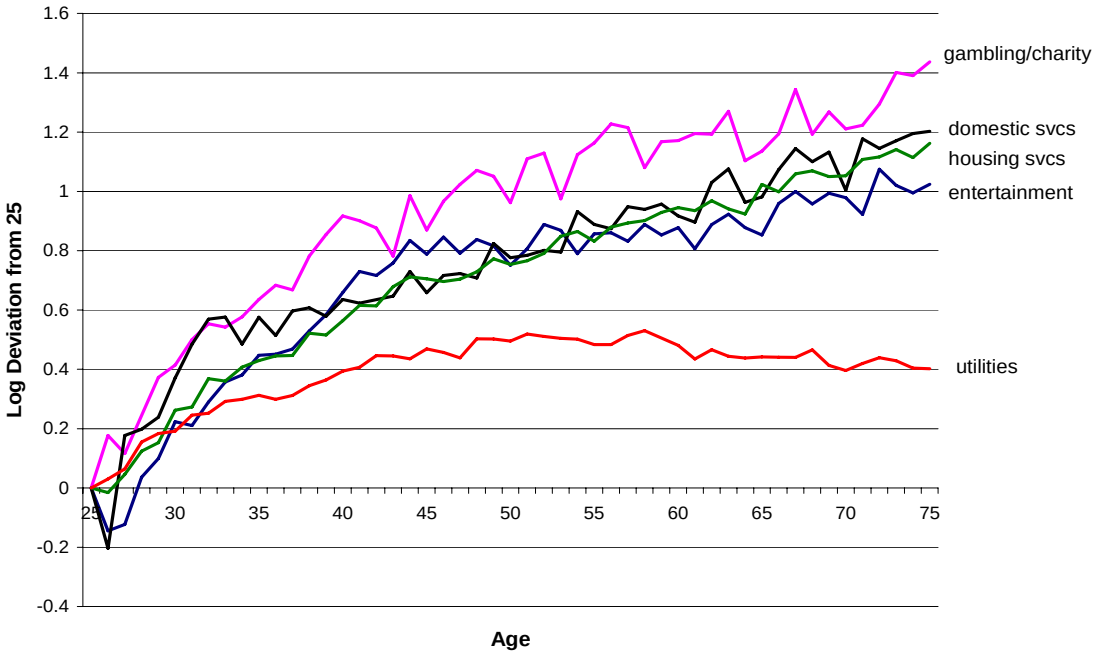

Notes: This figure replicates Figure A1 but using an alternative deflator for expenditures. Specifically, all categories are deflated using a common deflator, namely, the CPI-U. The benchmark Figure A1 used separate category-level deflators.

Figure R13: Alternative Deflator – Inequality

Panel A: Increasing Categories

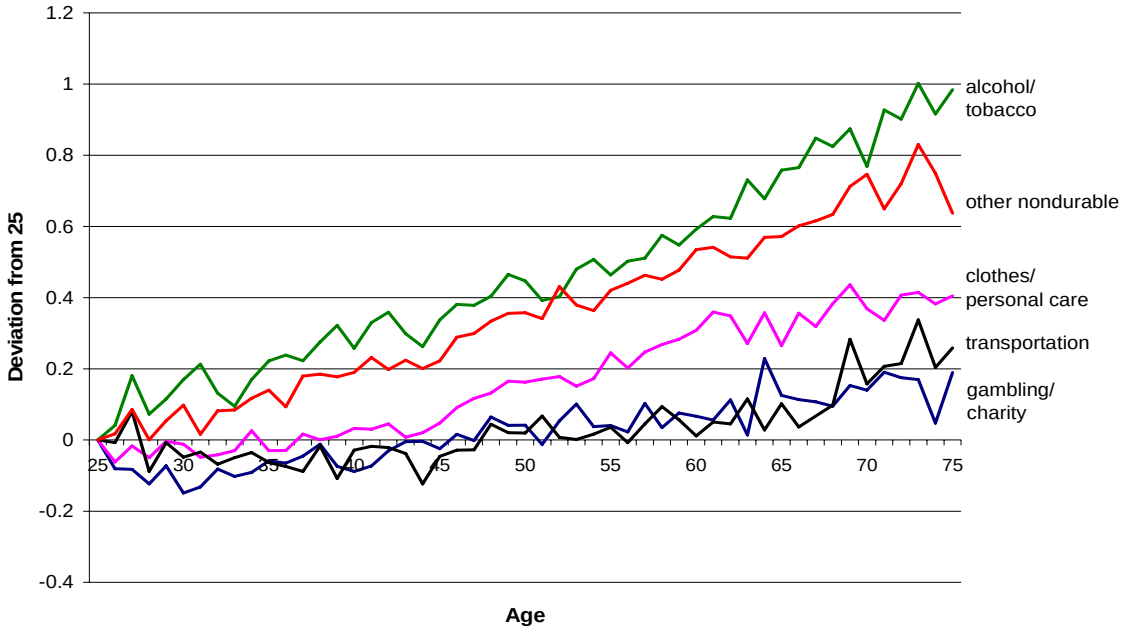

Panel B: Non-increasing Categories

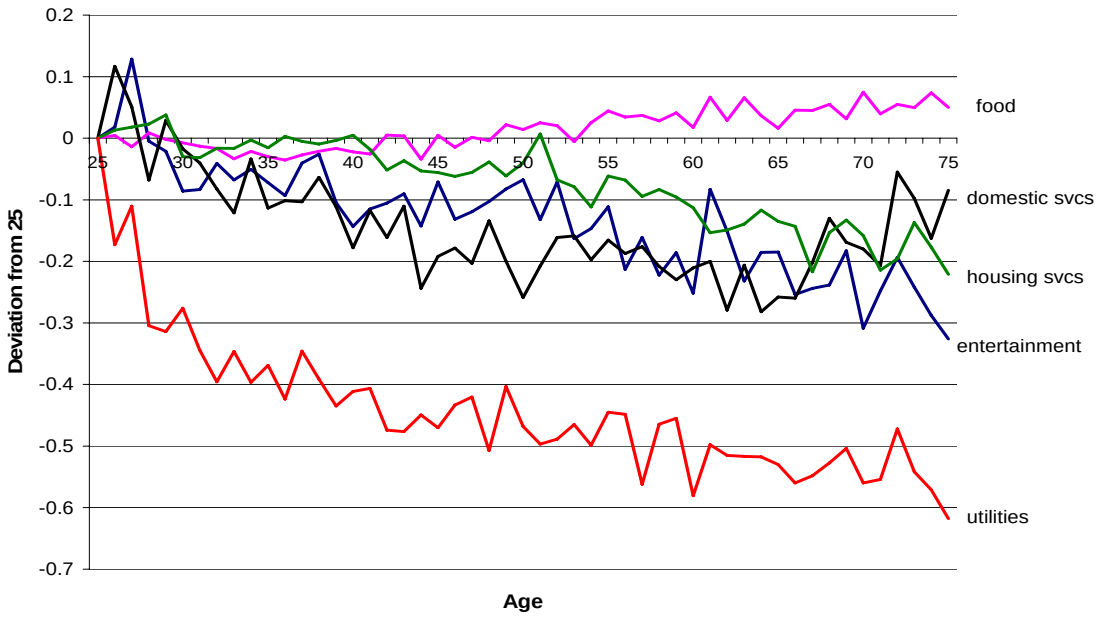

Notes: This figure replicates Figure A2 but using an alternative deflator for expenditures. Specifically, all categories are deflated using a common deflator, namely, the CPI-U. The benchmark Figure A1 used separate category-level deflators.

Figure R14: "Levels" Specification

Panel A: "Group 1" Mean Expenditure

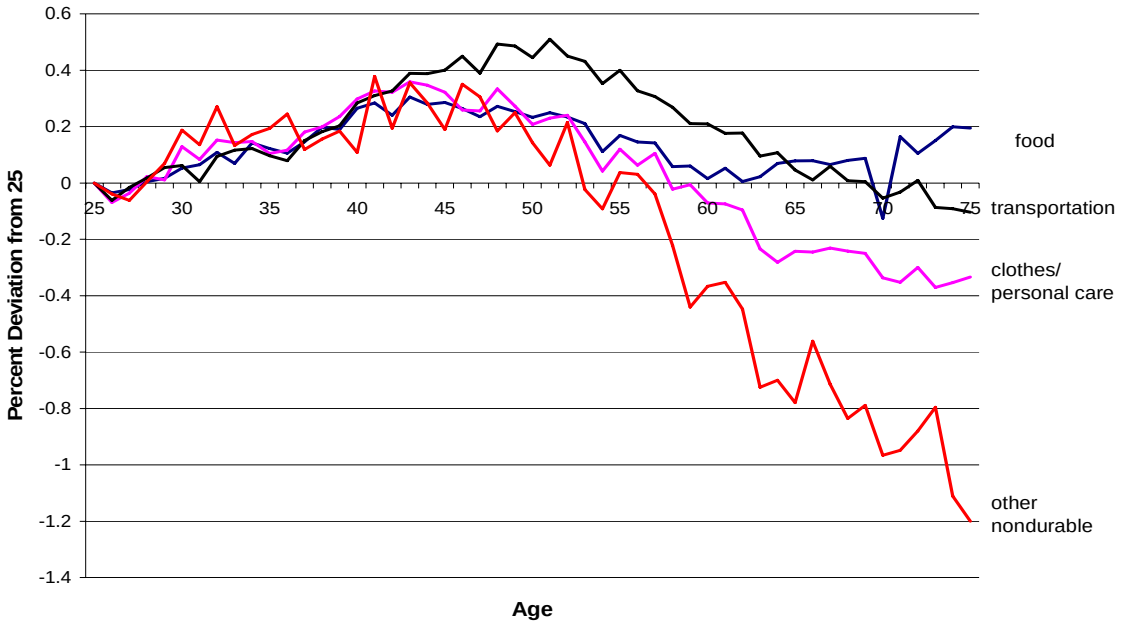

Panel B: "Group 2" Mean Expenditure

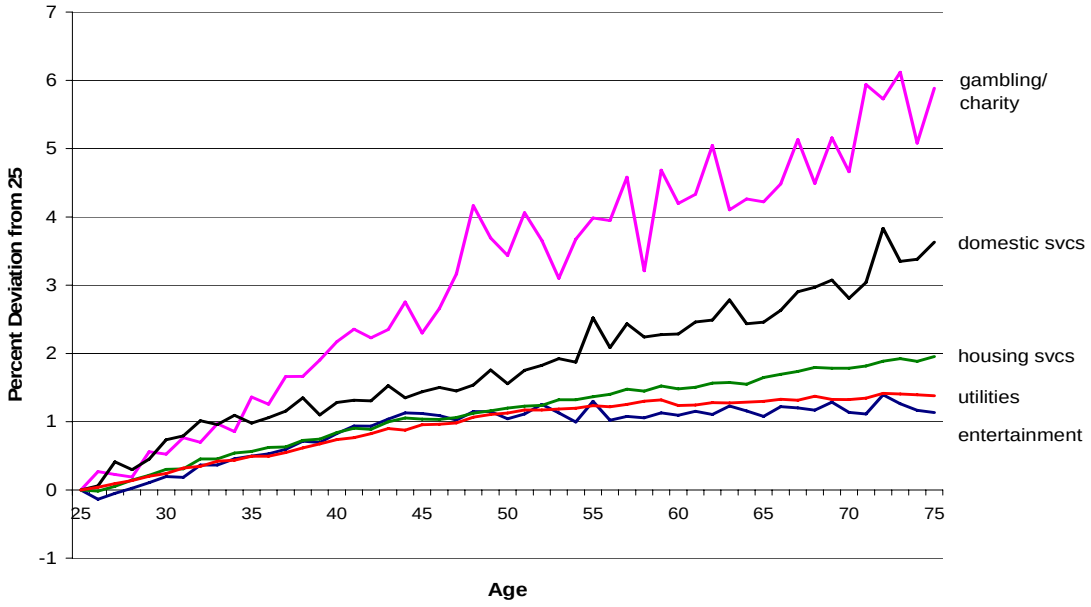

Notes: This figure replicates the benchmark analysis of Figure A1, replacing log expenditure with percentage deviations from 25. Specifically, we regress the level of expenditure on age, cohort, and family status dummies. We normalize the coefficient on each age dummy by the mean expenditure at age 25 (the omitted group). These normalize coefficients consitute the depicted series.