

**Working Paper Series
Congressional Budget Office
Washington, D.C.**

**ECONOMETRIC FLEXIBILITY IN MICROSIMULATION:
AN AGE-CENTERED REGRESSION APPROACH**

John Sabelhaus (E-mail: john.sabelhaus@cbo.gov)
Lina Walker (E-mail: lina.walker@cbo.gov)
Congressional Budget Office
Washington, D.C.

January 2007
2007-02

Working papers in this series are preliminary and are circulated to stimulate discussion and critical comment. These papers are not subject to CBO's formal review and editing processes. The analysis and conclusions expressed in them are those of the authors and should not be interpreted as those of the Congressional Budget Office. References in publications should be cleared with the authors. Papers in this series can be obtained at <http://www.cbo.gov/publications/>.

Abstract

This paper describes a strategy for estimating predictive equations that has been shown to work well in microsimulation modeling. The technique, referred to here as “age-centered regression,” is particularly useful when the available data set for estimating a model equation is limited and the marginal effect of one or more explanatory variables might be expected to vary systematically by age. The examples used here to describe how age-centering works are taken from the labor supply equations in the Congressional Budget Office Long-Term (CBOLT) dynamic microsimulation model. By switching from a traditional single-equation approach to age-centered regression, we show that marginal effects of independent variables can vary significantly across age groups. The comparison also reveals that improvements in mean predictions by age can be achieved with little if any loss in statistical precision of coefficient estimates.

I. Introduction

A fundamental goal of microsimulation modeling is to replicate empirically observed heterogeneity in person-, household-, or firm-level outcomes. If individuals with certain characteristics are more likely to have certain outcomes, a good model will have transition equations that generate those correlations in simulations. Sample size limitations in the data sets used to estimate model equations can make achieving this goal a challenge, however, as they make it difficult to statistically sort out all of the covariance needed to replicate the desired heterogeneity. This paper discusses an estimation strategy, referred to as “age-centered regression,” that has proved useful for mitigating the types of problems in microsimulation that are typically associated with small sample sizes.

There are several conditions under which age-centered regression might be useful in microsimulation modeling. First, the behavioral processes being estimated vary systematically with age; these processes include marital status transitions, fertility, mortality, and labor market outcomes. Second, these processes are such that the effect of any given explanatory variable could also vary by age; for example, higher educational attainment might lower the probability of marriage for very young singles, but increase the probability of marriage for middle-aged singles. Third, the data set available for estimating the model equation is limited, so that estimating separate equations for each age is infeasible and therefore some sort of grouping is required.

A typical approach in microsimulation when these conditions arise is to estimate separate equations for two or more age groups. For instance, in the marital transition equation, one might estimate separate marriage probability equations for young, middle-aged, and perhaps older

individuals, so the model would have three marriage equations. Age centering takes this grouping approach a step further. The idea is to estimate a separate equation for each unique (or “reference”) age, but include every observation in the sample whose age is within a preset range (or “bandwidth”) around the reference age being estimated. Thus, if the sample being analyzed is 25-year-olds and the bandwidth is 4 years, the estimation phase would use all observations in the data set with ages 21 through 29. The equation for 26-year-olds would use every observation with ages 22 to 30, and so on. The end result is that the model will have separate equations for each reference age, but the sample used to estimate the equation for any given age overlaps with the data used to estimate nearby ages.

There are two benefits when using age-centered regressions, in terms of econometric flexibility. First, there is flexibility in the shape of the functional form with respect to age itself—one does not have to rely on polynomial terms or linear splines to specify how the process in question varies with age. Second, the effect of any given independent variable in the equation is no longer constrained to be equal for all age groups (or even a subset of age groups, as when one splits the sample). If the data suggest that the magnitude or even sign of an explanatory variable systematically varies with age, the differential relationship will show up in the estimates.

The examples used here to illustrate the age-centering approach are taken from the labor force modules in the U.S. Congressional Budget Office Long-Term dynamic microsimulation, also known as CBOLT.¹ The equations we consider sequentially predict labor force

¹ See Congressional Budget Office (2006). The age-centering technique is also applied to marital transitions in CBO’s model; see O’Harra and Sabelhaus (2002) and Harris and O’Harra (2001). For a general overview of the CBOLT model see O’Harra, Sabelhaus, and Simpson (2004). The model has been applied in several academic papers, such as Harris, Sabelhaus, and Simpson (2005), Harris and Sabelhaus (2005), Harris and Simpson (2005),

participation, full-time versus part-time employment, and hours worked for part-time employed persons. All of the equations are univariate or multinomial logit specifications with standard controls: the explanatory variables include age, marital status, educational attainment, in-school status, number of children under 6 years of age (for women), receipt of social insurance benefits, and cohort/time effects. The equations are all estimated using about thirty years of data from the March Current Population Survey (CPS), which is a large, annual, nationally representative cross-section.

For each of the three labor market equations, we compare the results using a standard (linear in age) specification with age-centered results. We show that estimated coefficients in the age-centered equations do vary systematically by age, and those differences in coefficients imply very different marginal effects by age. One example is the extent to which lagged labor force participation is correlated with current labor force participation; the effect varies systematically over the life cycle and the age-centered equation is better able to capture that pattern than the more standard equation. A second example is the effect of marital status on expected part-time hours (conditional on working part-time). The extent to which being married affects part-time hours worked also varies significantly across age groups.

A further benefit of using the age-centered approach in microsimulation modeling is that, by its nature, age-centered regressions are better at capturing any differences in mean predicted outcomes by age. In general, microsimulation modelers rely on the fact that the single equation approach will work if the underlying outcomes by age are smooth and if one uses an appropriate polynomial in age. Using the predictions from the labor force module, we show that deviating

and Sabelhaus and Topoleski (forthcoming). There are also several published CBO reports and studies using

from those conditions can lead to biased predictions by age in some cases.

II. The Mechanics of Age-Centered Regression

Age-centered regression is useful for estimating microsimulation model equations in cases where the underlying process varies systematically by age and one is trying to achieve maximum flexibility in the estimated econometric relationships. In microsimulation modeling, researchers typically split the sample by age when estimating equations, although it is well understood that the effect of certain independent variables differs across age groups. Age-centered regression takes this logic a step further by estimating different equations for each unique “reference” age. Estimating the equation for each age, however, limits the regression sample and reduces the statistical precision of the estimation. In order to overcome the small-sample-size problem, age-centered regression borrows a principle from kernel density analysis and includes all observations that are within a certain “bandwidth” of the reference age group in question. The technique allows statistical precision to be maintained at the same time that improved flexibility in estimated coefficients by age is achieved.

For any given bandwidth, the actual estimation strategy for an age-centered regression is straightforward. Assume the equation being estimated is for a reference group that is 25 years old and the bandwidth is set to five years. The equation estimation will include every observation in the data set with ages 20 through 30. However, the observations are not weighted equally. As in kernel density estimation, declining weights are applied to observations that are farther from the reference age group. A simple triangular weighting pattern is used, so if the

reference group is 25-year-olds with a 5-year bandwidth, then 25-year-olds have weights of one, persons who are one year plus or minus 25 (24- and 26-year-olds) have weights of 0.8, persons who are two years plus or minus 25 (23- and 27-year-olds) have weights of 0.6, persons three years plus or minus 25 (22- and 28-year-olds) have weights of 0.4, persons four years plus or minus 25 (21- and 29-year-olds) have weights of 0.2, and persons five years plus or minus 25 (20- and 30-year-olds) have weights of 0.0.

The decision about how wide a bandwidth to use when estimating an age-centered regression will depend on the tradeoff between desired flexibility and statistical precision. If the bandwidth is set low (with bandwidth one, the only observations included are in the reference group itself) the estimation process has the most possible flexibility. There is no smoothing of the estimated effect of independent variables across ages; therefore, if people of different ages are really very different in terms of mean outcomes or the effects of some independent variable, those differences will come through in the estimation. The tradeoff is the loss of statistical precision when the bandwidth is set too low; there may be too few observations exactly at or near some ages and one cannot estimate the coefficients of interest with any reliability.

The approach in kernel density estimation is to explicitly test for optimal bandwidth, but the decision on bandwidth ultimately depends on how much weight one puts ex ante on flexibility versus statistical precision. In the examples here (and other equations in the CBOLT dynamic microsimulation model) the bandwidth is set to five years, which balances the goals of flexibility and precision in these types of equations. The five-year bandwidth is large enough that the coefficients are precisely estimated, with significance levels on most coefficients that are comparable to those from the standard single equations using the entire data set. At the same

time, the five-year bandwidths are small enough that one can observe any systematic differences in estimates across the age distribution. Capturing these differences will improve the capacity of the microsimulation model to reflect the heterogeneity in the underlying data.

There are a few mechanical observations about using age-centered regression worth noting. First, in the estimates discussed in this paper (see Section III.) all of the age-centered equations include an age term as well as an intercept. Note that if the bandwidth is set to one, the intercept and age will be perfectly correlated. However, when the bandwidth is greater than one, the age term can be estimated and captures any systematic age differences (within the bandwidth range) not captured by other independent variables. It is probably easiest to think of this coefficient as the derivative by age of the process being estimated, evaluated at that particular reference age. In model simulations, when predicting outcomes for a given observation, the age term is effectively combined with the constant term for each reference age group.²

A second mechanical observation concerns how one actually estimates and uses age-centered regressions. At first glance it might seem that age centering is somewhat more cumbersome, because it replaces one equation for the entire population with separate equations for each age group. However, in practice the actual estimation and implementation are both straightforward. The equations here are all estimated using a simple looping feature in a standard software package, so the actual code for estimating, evaluating, and outputting multiple

² It is worth noting that this also can lead to a lot of apparent imprecision in the estimated constant and age terms, because if the slope of the process being estimated is relatively flat over a particular age range, there is too little variation by age to separately identify a constant and age coefficient. Indeed, in some processes, one observes a pattern of estimated age and constant term coefficients that appear volatile, but the linear combination of the two actually input to the microsimulation model is much more stable.

equations is only slightly more cumbersome than estimating a single equation.³ In the actual CBOLT model code, switching from a single equation to age-centered equations simply involves adding an extra dimension (age) in the coefficient arrays used in the simulation. In sum, the benefits of adopting age centering exceed any computational burdens the technique introduces.

III. Age Centering in the CBOLT Labor Force Equations

The overarching goal of CBOLT is the same as many other dynamic microsimulation models: to simulate demographic, labor market, and government tax/transfer outcomes for a representative sample of the population forward through time. Although age centering is used in several CBOLT equations, for reasons of brevity the analysis in this paper is focused on the labor force participation and hours worked equations in the CBOLT labor force module. The approach used here to demonstrate age centering is to compare and contrast age-centered results with a more standard single-equation approach for each of the three equations in the CBOLT labor market module.

CBOLT Labor Force Equations

In CBOLT, annual hours worked for each individual are estimated using a sequence of three equations. For each individual, the model solves for (1) labor force participation, (2) full-time versus part-time work for those in the labor force, and (3) part-time hours worked for those who work part-time. The first two equations are univariate logits, and the resulting probability is compared to a random number draw to determine the actual model outcome. The third equation

³ The Stata code used to estimate all of the equations in the next section is available from the authors upon request.

is a multinomial logit with seven possible outcomes (or “bins”) for annual part-time hours outcomes.⁴ As with the first two equations, a random number draw is used to place individuals into each of the seven annual part-time bins.

Each of the three equations in the labor force module is estimated separately for men and women. The explanatory variables used in the CBOLT labor force equations include age, marital status, educational attainment, in-school status, number of children under 6 years of age (for women), receipt of social insurance benefits, and cohort/time effects. The labor force participation equation also includes a lagged independent variable in order to capture the observed persistence in the data.⁵

The three CBOLT labor force equations are estimated using pooled Current Population Survey (CPS) data. The March CPS collects information on about 60,000 households each year. The data sets are cross-sectional and contain a wide variety of economic and demographic information on the individual, family, and household. The data used to estimate the equations in the labor force module are for calendar years 1975 through 2003. In the equations discussed here, the sample is restricted to individuals ages 25 through 61, which includes nearly 1 million observations over the 29 year period.⁶

The CBOLT approach to estimating annual hours worked may seem overly complicated—one could use a single equation that directly predicts annual hours worked—but

⁴ The bins for the multinomial logit are for annual part-time hours of 125, 375, 625, 875, 1,125, 1,375, and 1,625. Everyone who works full-time (determined by the second equation in the module) has annual hours worked of 2,080.

⁵ The other two equations—full-time versus part-time work, and part-time hours conditional on working part-time—exhibit the same longitudinal persistence as labor force participation, but the CPS data used to estimate the equations do not have the requisite lagged information. CBOLT introduces that persistence into the simulation *ex post* using a technique described in Congressional Budget Office (2006).

⁶ CBOLT uses different equations and/or bandwidth limits for people under age 25 and over 61. For the young, the effects of schooling dominate labor force decisions. For people over 61 the effects of social insurance dominate

there are good reasons to separate the process into the three steps. Because each process has different dynamics and marginal impacts from changing independent variables, separately identifying each equation improves the simulations of actual population heterogeneity. Also, separating the module into these three logical steps makes it easier to build in other features of labor market outcomes. For example, unemployment incidence and spell lengths differ for part- and full-time workers. The CBOLT approach also makes it feasible to introduce behavioral responses into the model. In particular, retirement in CBOLT is modeled as a decision to start collecting Social Security, which does not necessarily end labor market activity—it just changes the intensity.⁷

Estimated Coefficients in Single-Equation and Age-Centered Regressions

In order to draw out how age-centered regression differs from traditional single-equation estimates we compare versions of each labor force equation, estimated using the same data set for the same time period. Table 1 shows six sets of single-equation estimates, three each for men and women. For each sex group, the first column shows estimated coefficients for the labor force participation equation, the second column for the full-time work equation, and the third column for part-time hours worked.⁸

The signs of the coefficients in Table 1 make sense, and most of the parameters of interest are significant at the 1 percent level. In general, more education is associated with

labor force decisions, because the eligibility age for Social Security is 62 in the United States.

⁷This is not the only way to achieve desired heterogeneity in labor market outcomes. For other examples, see the description of the labor force modules for the Urban Institute’s DYNASIM model in Favreault and Smith (2004) or for the Social Security Administration’s Modeling Income in the Near Term (MINT) model in Toder et al. (2002).

⁸ The estimated thresholds or “cut” parameters are not shown for the part-time multinomial logits. Those are available from the authors upon request.

higher labor supply for both men and women, while marriage is associated with higher labor supply for men but lower labor supply for women. Receipt of Social Security income has a strong negative effect on labor supply for both men and women. The most dominant effect in the labor force participation equation is from lagged labor force participation, which we focus on below when computing marginal effects. This coefficient shows that persistence in labor supply within the population is a first-order effect that should be accounted for, clearly dominating the magnitude of the other zero-one dummy variables.

The CBOLT age-centered versions are shown for two reference ages (age 30 and 55) in Table 2a (for men) and Table 2b (for women).⁹ The first observation about Tables 2a and 2b is that the estimated coefficients often vary significantly between the two age samples. For example, the coefficient on lagged labor force participation for 55-year-olds is about 30 percent higher than the coefficient for 30-year-olds, and that holds for both men and women. What that suggests is that persistence in labor supply is much stronger at age 55 than it is at age 30, everything else constant.

The estimated coefficients for the 30- and 55-year-old reference age samples diverge from the single-equation estimates to varying degrees across independent variables and equations, but the presumption that using age centering can reveal differences in estimated effects by age is clearly borne out in Tables 2a and 2b. It is worth noting that single-equation parameter estimates (in Table 1) generally fall in the range spanned by the age-centered estimates in Tables 2a and 2b, which is expected given the way the data are used to estimate the equations. Finally, there also appears to be no significant loss of statistical precision when

⁹ The entire list of coefficients for reference ages 16 through 70 can be found in Congressional Budget Office

shifting to age centering in these equations. Almost every variable that is significant in Table 1 remains so in Tables 2a and 2b.

Differences in Marginal Effects by Age

Differences in estimated coefficients across age groups are one way to show how age centering can affect predicted outcomes, but a better way to see the implications for model simulations is to compute marginal effects. The marginal effect of a given independent variable is computed by applying the estimated equation coefficients within-sample. The value of the independent variable being investigated is varied and the difference in the corresponding predicted outcomes is the marginal effect.

The first set of marginal effects computed for the CBOLT labor force equations applies to the lagged dependent variable in the labor force participation equations. For each person in the full sample the probability of working is computed using the estimated coefficients, with all independent variables set to actual values except lagged labor force participation. In the first set of calculations, the value of lagged labor force participation is set to zero (no work in the previous period) for every observation and the mean predicted labor force participation probability by age is computed. In the second set of calculations, the value of lagged labor force participation is set to one for every observation, and again the mean participation probability by age is computed. Because lagged labor force participation has a positive effect on current-year participation, the second set of mean probabilities is higher at every age. The marginal effect of lagged labor force participation is the gap between these two sets of mean probabilities.

Figure 1 shows four sets of marginal effects for lagged labor force participation. There is one set each for men and women and for the two equations (single-equation and age-centered). For both men and women, the marginal effects in the single-equation estimates show much less variation by age than in the age-centered regressions. This makes sense because the single-equation estimate provides the weighted-average marginal effect of lagged labor force participation across the age distribution. Indeed, the only variation by age in marginal effects is associated with underlying variation in the other independent variables in the data set; there is only one coefficient on lagged labor force participation that applies to every age group, so the range of marginal effects is limited.

The variation in marginal effects by age is much more pronounced in the age-centered regressions, which is consistent with observations about the estimated coefficients by age made in the previous section. The coefficient on lagged labor force participation rises systematically by age in the age-centered regressions for both men and women, and this rise is reflected in significant increases in marginal effects as age increases. The fact that both age-centered marginal effects lines cross their respective single-equation lines is also consistent with the observations made about the coefficients in Tables 1 and 2a, 2b.

These observed differences in estimated marginal effects could be of first-order importance in the microsimulation. The lagged labor force participation coefficient is a key to establishing longitudinal persistence in labor force participation within the micro sample, and in a single-equation model that persistence will be biased up for the young and biased down for workers nearing retirement. One often-used technique in microsimulation—calibration factors by age—could be used to match predicted outcomes by age, but unless one addresses the

underlying cause of the bias calibration will not fix the problems with the equation.

A second comparison of estimated marginal effects using single-equation estimates and age-centered regression leads to the same basic conclusion. Figure 2 shows the marginal effect of being married on part-time hours worked (conditional on working part-time). In this case, the marginal effect of being married for both men and women is fairly similar between the age-centered and single-equation estimates, with the notable exception of the youngest and oldest age ranges, where the bias is noticeable.

Differences in Mean Predicted Outcomes by Age

In addition to eliminating bias by age in estimated marginal effects, age centering can also improve the mean predicted outcomes by age in the microsimulation. The examples used here to draw out this point are somewhat contrived, because we estimate the single-equation versions using only a linear age term. In many applications, microsimulation modelers will examine patterns by age for the types of processes we are considering and estimate higher-order age polynomials.

Figure 3 shows the mean probability of working by age for men and women, evaluated within-sample using the single-equation and age-centered regression approaches. Both equations clearly capture the concave shape in labor force participation between ages 25 and 61. The age-centered predictions will, by virtue of maximum likelihood principles, more closely track the actual patterns of labor force participation by age, so we can characterize the deviation of the single-equation from the age-centered line as bias. Although the deviation is not too extreme, it is systematically biased upwards for women. For men, the single-equation estimate is

biased up for young workers and biased down for older workers, in generally the same direction of bias as the marginal effects of lagged labor force participation noted above. Adding higher-order age terms to the equation may help the single-equation track better, but it is unlikely to eliminate the bias.

Figure 4 shows the mean of part-time hours conditional on working part-time, and for these outcomes the bias is even stronger. The effect of imposing a linear coefficient on age comes through clearly for men, as the functional form induces a predicted relationship with age that is linear. As with labor force participation, the higher-order age-terms could improve predicted outcomes but the predictions will only asymptotically approach the age-centered values as higher-order terms are added.

IV. Conclusions

The technique described here as “age-centered regression” analysis is a useful way to extract information from a limited data set when estimating equations for dynamic microsimulation models. The situation in which age centering can help is very common: when the behavioral process in question varies systematically by age and the effect of one or more independent variables might also differ across age groups. The basic idea is to estimate separate equations for each age sample, but to include observations that are close to the reference age being estimated. Using the extra information from nearby observations makes it possible to statistically identify how marginal effects and predicted mean outcomes differ across the age distribution under consideration.

The examples used here to illustrate age centering are fairly simple labor supply

equations—the sequence of participation, full- versus part-time, and part-time hours given part-time employment—used in the Congressional Budget Office Long-Term (CBOLT) dynamic microsimulation. We show that moving from a single-equation to an age-centered approach involves little loss of statistical precision but has a significant impact on estimated marginal effects and mean predicted outcomes by age.

V. References

Congressional Budget Office, *Projecting Labor Force Participation and Earnings in CBO's Long-Term Microsimulation Model*. Background Paper, Congressional Budget Office, Washington, DC (October 2006).

Favreault, Melissa M., and Karen E. Smith, "A Primer on the Dynamic Simulation of Income Model (DYNASIM3)." Discussion paper, Retirement Project, The Urban Institute, Washington, DC (December 2004).

Harris, Amy Rehder, and Josh O'Harra, "The Impact of Marriage and Labor Force Participation Trends on the Social Security Benefits of Women." *Proceedings of the National Tax Association* (November 2001).

Harris, Amy Rehder, John Sabelhaus, and Michael Simpson, "Uncertainty About OAI Worker Benefits Under Individual Accounts." *Contemporary Economic Policy*, 23:1, 1-16 (January 2005).

Harris, Amy Rehder, and John Sabelhaus, "How Does Differential Mortality Affect Social Security Progressivity and Finances?" Technical Paper 2005-5, Congressional Budget Office, Washington, DC (May 2005).

Harris, Amy Rehder and Michael Simpson, "Winners and Losers Under Various Approaches to Slowing Social Security Benefit Growth," *National Tax Journal*, Vol. LVIII, No. 3, pp 523-544 (September 2005).

O'Harra, Josh, John Sabelhaus, and Michael Simpson, "Overview of the Congressional Budget Office Long-Term (CBOLT) Policy Simulation Model." Technical Paper 2004-1, Congressional Budget Office, Washington DC (January 2004).

O'Harra, Josh, and John Sabelhaus, "Projecting Longitudinal Marriage Patterns for Long-Term Policy Analysis" Technical Paper 2002-2, Congressional Budget Office, Washington, DC (October 2002).

Sabelhaus, John, and Julie Topoleski, "Uncertain Policy for an Uncertain World: The Case of Social Security" *Journal of Policy Analysis and Management* (forthcoming, 2007).

Smith, Karen E., David B. Cashin, and Melissa M. Favreault, "Modeling Income in the Near Term 4." The Urban Institute, Washington DC (April 2005).

Toder, Eric, and others, "Modeling Income in the Near Term: Revised Projections of Retirement Income Through 2020 for the 1931-1960 Birth Cohorts." The Urban Institute, Washington DC (June 2002).

Table 1: Regression Results Using Single Equation with Linear Age Term, Men and Women

	Men			Women		
	Pr(LFP=1)	Pr(FT=1)	PT Hours	Pr(LFP=1)	Pr(FT=1)	PT Hours
Lagged Labor Force Participation	4.082 (0.012)**			4.361 (0.008)**		
Married	0.166 (0.012)**	0.822 (0.009)**	0.236 (0.015)**	-0.525 (0.009)**	-0.708 (0.006)**	-0.371 (0.010)**
Age of Person	-0.021 (0.001)**	-0.002 (0.001)**	-0.011 (0.001)**	-0.004 (0.001)**	0.012 (0.000)**	0.013 (0.001)**
Receiving Social Security Income	-2.081 (0.026)**	-2.243 (0.035)**	-1.218 (0.042)**	-1.085 (0.020)**	-1.110 (0.022)**	-0.606 (0.025)**
High School Education	0.111 (0.015)**	0.552 (0.012)**	0.111 (0.019)**	0.398 (0.011)**	0.299 (0.009)**	0.101 (0.012)**
Some College Education	-0.049 (0.017)**	0.491 (0.013)**	-0.074 (0.022)**	0.482 (0.013)**	0.325 (0.010)**	0.093 (0.014)**
College Education	0.241 (0.016)**	0.646 (0.013)**	-0.116 (0.021)**	0.637 (0.013)**	0.490 (0.010)**	0.071 (0.014)**
Number of Children Under 6 Years				-0.481 (0.008)**	-0.516 (0.006)**	-0.299 (0.007)**
Birth Year: 1920-1929	0.208 (0.042)**	0.129 (0.041)**	0.074 (0.066)	0.175 (0.032)**	0.078 (0.033)*	0.055 (0.045)
Birth Year: 1930-1939	0.376 (0.041)**	0.202 (0.041)**	0.073 (0.064)	0.422 (0.032)**	0.157 (0.033)**	0.185 (0.044)**
Birth Year: 1940-1949	0.392 (0.042)**	0.331 (0.042)**	0.178 (0.065)**	0.64 (0.032)**	0.347 (0.033)**	0.337 (0.044)**
Birth Year: 1950-1959	0.341 (0.043)**	0.293 (0.042)**	0.151 (0.068)*	0.831 (0.034)**	0.543 (0.033)**	0.547 (0.045)**
Birth Year: 1960-1969	0.147 (0.046)**	0.266 (0.044)**	0.068 (0.070)	0.863 (0.035)**	0.700 (0.034)**	0.660 (0.046)**
Birth Year: 1970-1979	-0.001 (0.052)	0.115 (0.047)*	0.002 (0.074)	0.744 (0.040)**	0.732 (0.036)**	0.711 (0.049)**
Constant	-0.443 (0.063)**	1.125 (0.056)**		-2.169 (0.048)**	0.152 (0.041)**	
Observations	1,034,541	944,231	86,049	1,113,085	782,972	254,685

Robust standard errors in parentheses

* significant at 5%; ** significant at 1%

Table 2a: Regression Results Using Age Centering, Men Ages 30 and 55

	Men, Age 30			Men, Age 55		
	Pr(LFP=1)	Pr(FT=1)	PT Hours	Pr(LFP=1)	Pr(FT=1)	PT Hours
Lagged Labor Force Participation	3.451 (0.022)**			4.594 (0.025)**		
Married	-0.043 (0.018)*	0.793 (0.014)**	0.226 (0.024)**	0.335 (0.028)**	0.562 (0.022)**	0.091 (0.036)*
Age of Person	0.046 (0.004)**	0.064 (0.003)**	0.028 (0.005)**	-0.123 (0.005)**	-0.089 (0.004)**	-0.062 (0.006)**
Receiving Social Security Income	-1.687 (0.052)**	-1.905 (0.075)**	-1.211 (0.090)**	-2.293 (0.050)**	-2.555 (0.067)**	-1.248 (0.068)**
High School Education	-0.088 (0.026)**	0.617 (0.021)**	0.230 (0.034)**	0.152 (0.030)**	0.360 (0.024)**	0.012 (0.039)
Some College Education	-0.364 (0.028)**	0.475 (0.022)**	-0.060 (0.036)**	0.123 (0.037)**	0.371 (0.029)**	-0.078 (0.047)
College Education	-0.043 (0.028)	0.536 (0.022)**	-0.248 (0.035)	0.464 (0.035)	0.580 (0.027)**	0.047 (0.046)
Birth Year: 1920-1929				-0.065 (0.081)	-0.180 (0.067)**	-0.169 (0.103)
Birth Year: 1930-1939				-0.133 (0.082)	-0.382 (0.067)**	-0.291 (0.104)**
Birth Year: 1940-1949				-0.277 (0.082)**	-0.388 (0.068)**	-0.295 (0.105)**
Birth Year: 1950-1959	0.011 (0.030)	-0.124 (0.025)**	-0.008 (0.041)	-0.365 (0.098)**	-0.289 (0.080)**	-0.232 (0.126)
Birth Year: 1960-1969	-0.025 (0.030)	-0.058 (0.025)*	-0.044 (0.042)			
Birth Year: 1970-1979	-0.001 (0.034)	0.009 (0.028)	-0.035 (0.047)			
Constant	-1.578 (0.117)**	-0.549 (0.094)**		5.136 (0.302)**	6.613 (0.239)**	
Observations	308,015	290,385	29,024	188,934	160,656	16,547

Robust standard errors in parentheses

* significant at 5%; ** significant at 1%

Table 2b: Regression Results Using Age Centering, Women Ages 30 and 55

	Women, Age 30			Women, Age 55		
	Pr(LFP=1)	Pr(FT=1)	PT Hours	Pr(LFP=1)	Pr(FT=1)	PT Hours
Lagged Labor Force Participation	3.717 (0.013)**			5.028 (0.020)**		
Married	-0.534 (0.015)**	-0.547 (0.011)**	-0.306 (0.016)**	-0.535 (0.022)**	-0.780 (0.016)**	-0.413 (0.024)**
Age of Person	0.029 (0.003)**	0.001 (0.002)	0.007 (0.003)*	-0.061 (0.004)**	-0.020 (0.003)**	-0.017 (0.004)**
Receiving Social Security Income	-0.882 (0.044)**	-0.794 (0.050)**	-0.515 (0.054)**	-1.110 (0.047)**	-1.368 (0.050)**	-0.718 (0.050)**
High School Education	0.436 (0.019)**	0.398 (0.018)**	0.107 (0.023)**	0.373 (0.024)**	0.272 (0.019)**	0.076 (0.026)**
Some College Education	0.526 (0.021)**	0.454 (0.019)**	0.132 (0.024)**	0.425 (0.031)**	0.290 (0.022)**	0.047 (0.032)
College Education	0.700 (0.021)**	0.731 (0.019)**	0.166 (0.024)**	0.598 (0.033)**	0.403 (0.023)**	0.066 (0.034)
Number of Children Under 6 Years	-0.435 (.008)**	-0.584 (0.007)**	-0.309 (0.009)**	-0.209 (0.084)	-0.201 (0.060)**	-0.201 (0.085)*
Birth Year: 1920-1929				-0.054 (0.062)	-0.094 (0.054)	-0.034 (0.074)
Birth Year: 1930-1939				0.096 (0.063)	-0.095 (0.054)	0.084 (0.075)
Birth Year: 1940-1949				0.220 (0.063)**	0.135 (0.055)*	0.220 (0.075)**
Birth Year: 1950-1959	0.299 (0.020)**	0.255 (0.016)**	0.293 (0.022)**	0.342 (0.076)**	0.244 (0.062)**	0.351 (0.086)**
Birth Year: 1960-1969	0.409 (0.021)**	0.449 (0.017)**	0.455 (0.023)**			
Birth Year: 1970-1979	0.393 (0.024)**	0.499 (0.019)**	0.510 (0.027)**			
Constant	-2.408 (0.086)**	0.530 (0.067)**		0.855 (0.247)**	2.184 (0.178)**	
Observations	330,794	239,891	81,777	205,062	127,326	40,052

Robust standard errors in parentheses

* significant at 5%; ** significant at 1%

Figure 1: Marginal Effect of Lagged LFP on Labor Force Participation
Comparing Single-Equation and Age-Centered Regressions
Men and Women, Ages 25-61

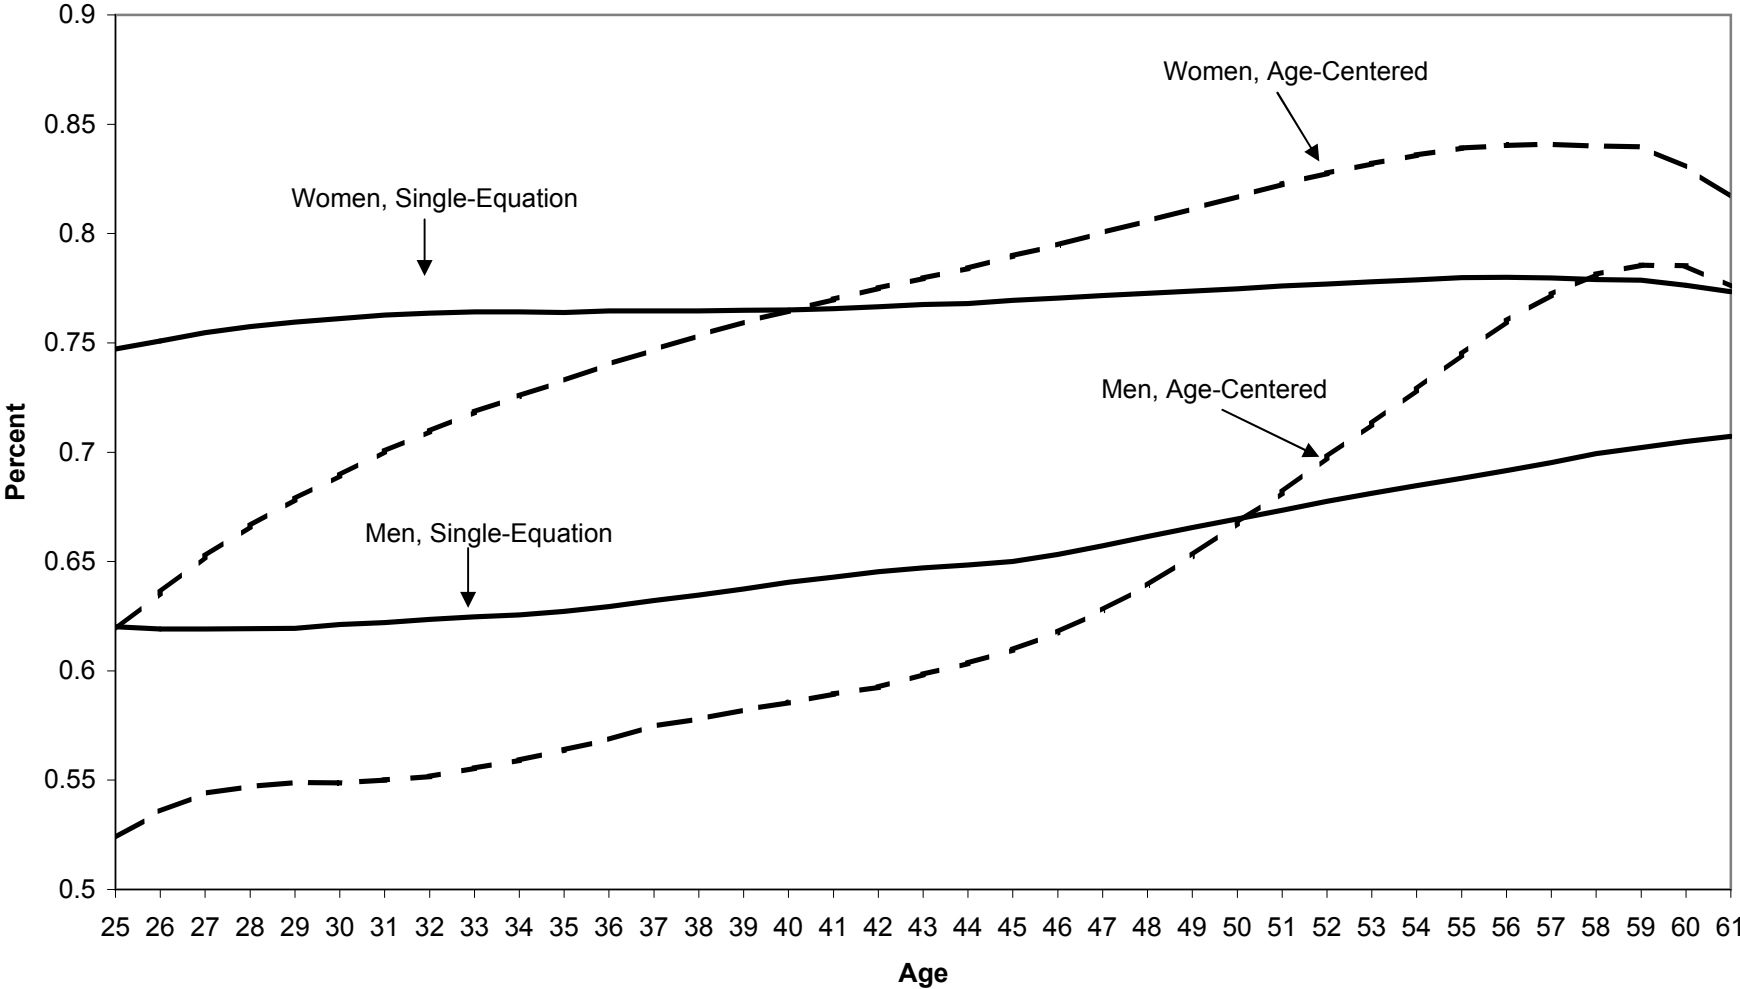


Figure 2: Marginal Effect of Being Married on Part-Time Hours Worked
Comparing Single-Equation and Age-Centered Regressions
Men and Women, Ages 25-61




Figure 3: Mean Predicted Labor Force Participation
Comparing Single-Equation and Age-Centered Regressions
Men and Women, Ages 25-61




Figure 4: Mean Expected Part-Time Hours
Comparing Single-Equation and Age-Centered Regressions
Men and Women, Ages 25-61

