American Economic Review: Papers & Proceedings 2009, 99:2, 380–386
http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.2.380

The Effects of Population Aging on the Relationship among Aggregate Consumption, Saving, and Income

By Karen E. Dynan, Wendy Edelberg, and Michael G. Palumbo*

As is well known, the US population has grown much older and is expected to continue to age. The share of adults age 65 or older—16 percent in 1980—has been rising steadily and is projected by the US Census Bureau to reach 22 percent by 2020. Given differences in saving rates and marginal propensities to consume (*mpc*'s) over the life cycle, demographic shifts could materially affect the relationship among macroeconomic aggregates such as income, consumption, and saving. A key question is how large these effects might be. Some past studies have used time series data to try to quantify these effects (for example, see Alan S. Blinder 1975), but demographics change more gradually than other determinants of aggregate spending, making it difficult to obtain precise estimates of the effects from such data. Our paper contributes to the literature by drawing lessons on how population aging might change the relationship between macroeconomic aggregates from household-level data. Household-level data enhance our ability to identify the effects because of the rich variation in consumption, saving, and income that we observe across households.

I. Using Cohort Methods in the Consumer Expenditure Survey

Our household-level data are drawn from the Consumer Expenditure Survey (CEX), which provides information about spending and income since the early 1980s. The public-use datasets include information from up to four interviews per household, spaced three months apart. Our sample is drawn from CEX data files corresponding to the period 1983:I (earlier data had problems with quality) through 2007:I (the latest quarter available at the time this paper was prepared).

Our main measure of consumption is derived from the detailed CEX expenditure files and is constructed to correspond as closely as possible to Personal Consumption Expenditures (PCE) in the US National Income and Product Accounts (NIPA). For each household, we measure annual consumption as the sum across all four interviews, dropping households with fewer than four interviews. Our income measure is the household's income after taxes for the previous 12 months, taken from the CEX family file corresponding to the household's final interview, so it covers the same period as the consumption measure. To adjust consumption and income for general price inflation, we construct household-specific deflators by taking a weighted average of the aggregate PCE chain-weighted price indexes for each major category of consumption, where the weights correspond to the share of each household's consumption represented by that category.[1]

The short time dimension of the CEX panel precludes tracking individual households over long periods. Instead, we follow an extensive literature, pioneered by Martin Browning, Angus Deaton, and Margaret Irish (1985), which groups households of similar birth years and then tracks how variables associated with these "synthetic cohorts" evolve over the life cycle. We group households by the birth year of the head, with the earliest of the ten-year cohorts born between 1910 and 1919. To limit the influence of top-coded values in the CEX on our results, for

* Dynan: Division of Research and Statistics, Federal Reserve Board, Washington, DC 20551 (e-mail: kdynan@frb.gov); Edelberg: Executive Office of the President, Council of Economic Advisors, Washington, DC 20502 (e-mail: Wendy_M._Edelberg@cea.eop.gov); Palumbo: Division of Research and Statistics, Federal Reserve Board, Washington, DC 20551 (e-mail: mpalumbo@frb.gov). We thank Annamaria Lusardi for helpful comments. The views expressed are those of the authors and not necessarily those of the Federal Reserve Board or the other members of its staff.

[1] We followed standard practices for dropping households with incomplete data and were left with a sample of 80,214. See the Data Appendix on the *AER* Web site (http://www.aeaweb.org/articles.php?doi=10.1257/aer.99.2.380)for more details.

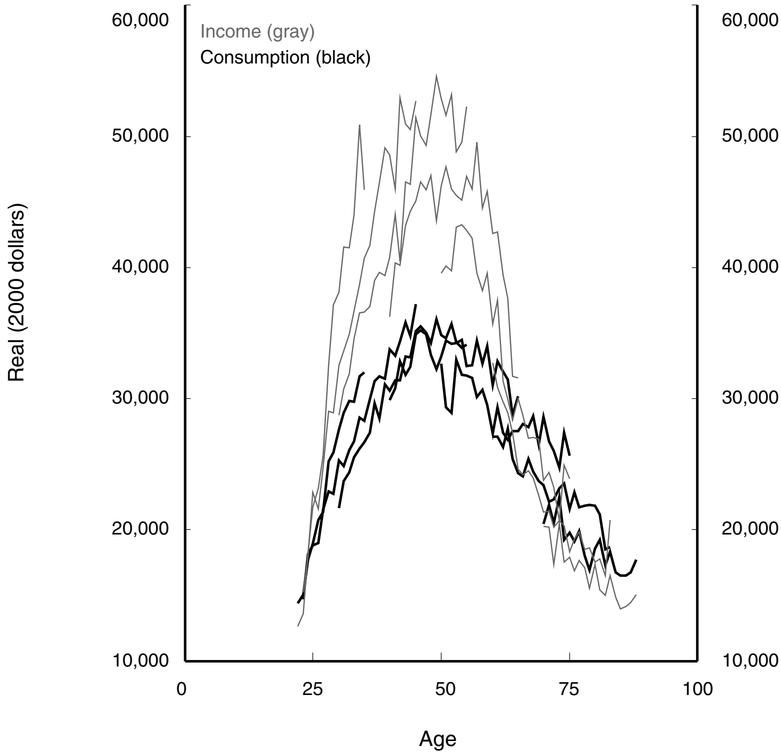

Figure 1. Real Consumption and Income by Age (Ten-year Cohorts)

each cohort, we calculate medians for the variables of interest at every age observed.[2]

Figure 1 shows the raw age-profiles for median income (thin gray) and consumption (thick black), with the separate line segments corresponding to different birth cohorts. As is well understood, the lines for each cohort embody an underlying relationship between age and income or consumption, as well as cohort-specific effects and business-cycle effects; we will use techniques that isolate the first of these. However, even without such adjustments, one can see that consumption is much flatter over the life cycle than income, suggesting some consumption smoothing, as predicted by the Life Cycle Hypothesis (LCH). That said, consumption has a hump over the life cycle, broadly consistent with models that include borrowing constraints or buffer-stock saving behavior (see Pierre-Olivier Gourinchas and Jonathan A. Parker 2002) or the influences of children and life-cycle labor supply on household preferences (see Orazio Attanasio et al. 1999).

Net saving over the life cycle looks to be highest between ages 35 and 60, although it is much larger than one would expect given the saving rates reported in the NIPA. We do not view this result as reflecting the true saving rate, but rather as reflecting a tendency for CEX respondents to underreport expenditures more than income. In any event, we make adjustments to reflect the discrepancy when we draw lessons for the NIPA saving rate from our CEX analysis.

II. Using the CEX to Account for Demographic Effects on the Aggregate Saving Rate

The aggregate US saving rate has shown a pronounced downtrend over the past two decades, falling from an average of 9 percent in the 1980s to 5 percent in the 1990s, and to 1½ percent so far this decade. We use our CEX cohort data to explore what effect the aging of

[2] We dropped the medians for the very youngest and oldest households observed in each cohort, as they are based on many fewer observations and tend to be less representative of the cohort.

the population might have had on aggregate saving and how it should be expected to affect saving going forward, should current projections of the age structure turn out to be accurate.

The key challenge in using synthetic cohort data to estimate the effect of population aging on the aggregate saving rate is distinguishing among the three determinants of the saving of a given cohort at a particular age noted earlier. The first is life-cycle influences on saving. One might also expect different cohorts to save more or less over the life cycle, depending, for example, on differing tastes, shocks to wealth, or changes in government programs providing income or health insurance in retirement. In addition, one would expect saving to vary with the business cycle as households seek to smooth their consumption in the face of job losses or other transitory shocks to income.

To distinguish between these factors, we follow the approach of Attanasio (1998), regressing median cohort saving rates on a fifth degree polynomial in age and on dummy variables for cohorts and for years. The year dummies are included to capture business-cycle effects, but, if left unrestricted, the coefficients will not be identified because year can be expressed as a linear function of age and cohort. Identification is achieved, however, by restricting the year-effects to average zero and to be orthogonal to a linear time trend. These identifying assumptions are strong: they mean that all linear trends in the cohort data are attributed to age or cohort effects. However, we found the estimated relationship between age and saving, the focus of our analysis, to be fairly robust across alternative specifications of the cohort- and year-effects.

Following Attanasio, we compute the saving rate as the difference between after-tax income and consumption relative to consumption. Table 1 presents the estimated coefficients and standard errors on the age terms and the cohort dummies. Lower-numbered cohorts correspond to those born more recently (represented by the leftmost line segments in Figure 1). The estimated coefficients imply that earlier cohorts in the CEX had *lower* saving rates than their later counterparts.

The result stands in contrast with that of Attanasio. Using a much shorter sample from the CEX (1980–1991), he concluded that the earliest cohorts showed some tendency toward having higher saving rates. Our results imply that this pattern has reversed since the early 1990s. While this could reflect a trend toward greater saving among the later cohorts, we suspect that measurement issues play a significant role. Indeed, the aggregate saving rate implied by the CEX data has actually risen a bit since the mid-1980s, even as the NIPA personal saving rate has trended strongly down.

Assuming the cohort dummies in the regression account for such measurement issues, we can use the estimated coefficients on age—together with information about the changing age distribution in the economy—to infer how population aging has, and will, affect the aggregate saving rate. We emphasize that our approach captures the effects through the life-cycle–saving channel alone. A broader analysis of demographic effects might also capture the interaction between any cohort-specific components of saving and population aging, but the measurement problems just discussed preclude such a focus.

Arthur Kennickell (1990) showed that the effects of demographics on the saving rate between two periods can be expressed approximately as the weighted average of the predicted initial saving rate for each age, multiplied by the change in the share in the population represented by that age. We use US Census Bureau estimates (for 1980, 1990, and 2000) and projections (for 2010 and 2020) to calculate changes in the shares of the population between ages 20 and 80 for five-year age groups, relative to their values in 1980. We calculate initial saving by multiplying CEX consumption for each age as of the beginning of our sample by the value of the consumption-based saving rate for that age implied by the coefficient estimates on age. We then scale CEX saving for each group by the ratio of aggregate CEX saving to NIPA saving (to adjust for the differences in the levels) and divide by initial NIPA disposable personal income.

The results are shown in Table 2. The movement of the baby boomers into their prime saving years is estimated to have increased the aggregate saving rate by 0.8 percentage point in the 1980s and then by an additional 1.1 percentage points in the 1990s, for a two-decade total of close to 2 percentage points. Thus, life-cycle influences on saving behavior do not help explain the decline in the saving rate between 1980 and 2000, but instead leave more to be explained by other factors. With baby boomers increasingly reaching retirement age after 2000, the pattern reverses: the contribution of

Table 1—Age and Cohort Effects on Saving Rates, Ten-year Birth Cohorts

	Coefficient	s.e.
Age	45.99	10.32
$\text{Age}^2/100$	−183.64	44.91
$\text{Age}^3/1{,}000$	37.74	9.29
$\text{Age}^4/10{,}000$	−3.90	0.92
$\text{Age}^5/100{,}000$	0.16	0.04
Cohort 2	−2.74	2.40
Cohort 3	−6.85	3.13
Cohort 4	−14.00	4.05
Cohort 5	−21.18	4.94
Cohort 6	−21.42	5.89
Cohort 7	−17.19	6.76

Note: Regression also includes a constant and year dummies that are constrained to sum to zero and to be orthogonal to a linear trend.

Table 2—Change in the NIPA Saving Rate Related to Population Aging

	Change relative to 1980 value (percentage points)
1990	0.8
2000	1.9
2010	1.7
2020	1.0

Notes: Estimates reflect changes in saving related to life-cycle factors. They do not reflect cohort effects.

population aging to the saving rate is predicted to decline 0.2 percentage point over the current decade and to fall by an additional 0.7 percentage point over the next decade. However, we do not estimate a full reversal of the demographic effect. In isolation, the shifting age distribution points to an aggregate NIPA personal saving rate 1 percentage point higher in 2020 than in 1980. Although the older baby boomers are likely to be dissaving (spending more than their incomes) by 2020, the youngest baby boomers will still be at their peak saving ages, and the children of the older baby boomers will be entering their higher-saving ages.

III. Are These Results Consistent with Time Series Evidence on Aggregate Saving?

We now explore how our results on the effect of aging on the saving rate based on the CEX cohorts line up with a time series analysis of aggregate data. We use a typical forecasting equation for aggregate personal consumption expenditures, which assumes a long-run cointegrating relationship in (log) levels of real personal consumption expenditures (c), labor income (y), transfer income (τ), and household wealth (w), where (stationary) deviations (η) between actual c each period and the target level given by the long-run coefficients are assumed to be proportional to the magnitude of spending. An error-correction specification relates quarterly percentage changes in c to the lagged deviation η, the level of real interest rates (r), lagged growth of c, income growth, an indicator of the economy's potential growth rate (p), an index of consumer sentiment (S), and the change in the unemployment rate (u). The model consists of two regression equations, which are estimated sequentially:

$$(1)\quad \ln c_t = \alpha_0 + \ln y_t + \alpha_1(\tau_t/y_t) + \alpha_2(w_t/y_t) + \varepsilon_t$$

$$\Delta \ln c_t = \beta_0 + \Delta \ln p_t + \beta_1\widehat{\varepsilon_{t-1}} + \beta_2(\Delta \ln y_t - \Delta \ln p_t) + \beta_3 r_{t-1} + \beta_4(\Delta \ln c_{t-1} - \Delta \ln p_t) + \beta_5 \Delta u_t + \beta_6 S_t + \mu_7,$$

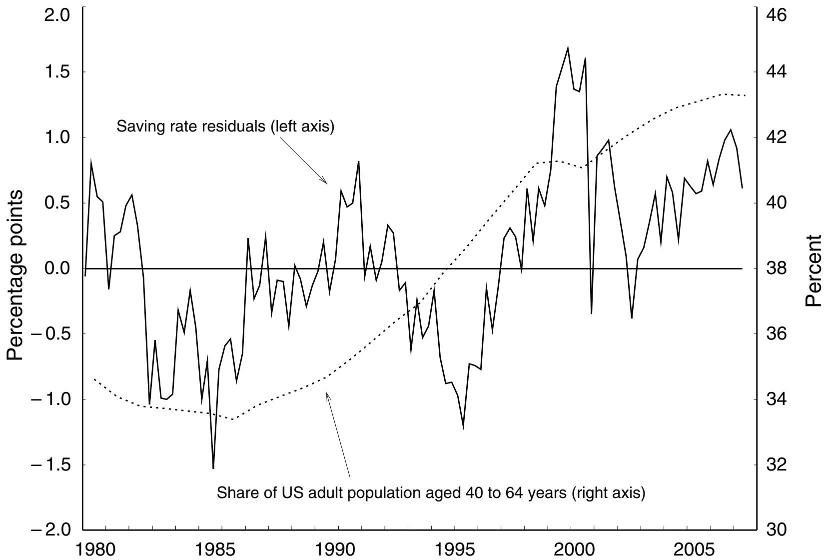

Figure 2. Demographic Shifts and Unexplained Changes in the NIPA Personal Saving Rate

using quarterly aggregate data from 1963 to 2008.[3] $\widehat{\varepsilon_{t-1}}$ is the difference between actual log consumption and the predicted value from the first regression. We then dynamically simulate the level of the personal saving rate beginning in 1980—roughly the start of its steep downtrend—and compute residuals as the difference between the actual NIPA personal saving rate and the simulated value from the model.[4] These residuals are plotted (in percentage points against the left axis) as the solid line in Figure 2.

Because demographics are not included in the equations, their effect might show up as an upward trend in the saving rate residuals for the periods when more households were entering peak saving ages. Notwithstanding the extremely low level of the NIPA saving rate this decade, the model views the actual saving rate as having been surprisingly high (given the estimated coefficients and values for the right-hand-side variables) since about 1999. That is, the model would have expected an even lower saving rate this decade, primarily because of the relatively high level of household wealth-income ratio, high readings on consumer sentiment, and low interest rates.

The dotted line (plotted against the right axis) corresponds to the share of the adult US population age 40 to 64 years. We plot this line as a rough benchmark, since saving rates in the CEX peak in this age range. Figure 2 suggests some relationship between the residuals and the fraction of households in the high saving years, with the residuals averaging below zero when the fraction was lower in the 1980s and early 1990s, and above zero more recently, when the fraction has been high.

IV. Aging and the Relationship between Changes in Consumption and Income

We now explore what can be learned from our household data about the aggregate marginal propensity to consume. To this end, we compare the relative magnitudes of typical changes in consumption, c, from age to age with typical changes in income, y. The differential of c to y, dc/dy, can be estimated as the ratio of the age-differentials for c, dc/da, and y, dy/da. This approach has limitations in that identification of the *mpc* comes only from the typical response of consumption to income changes over the life cycle without distinguishing anticipated changes in income from "shocks" or attempting

[3] See the online Data Appendix for information about the sources for the data.

[4] This simulation fixes all coefficients at the values estimated for the full sample, and then conditions on actual realizations of the determinants of consumption.

Table 3—Annual Changes in Consumption and Income over the Life Cycle

	Dependent variable		
Age range	Consumption	Income	Ratio
20–24	1,595.2	2,776.4	0.57
25–29	1,306.0	3,083.3	0.42
30–34	1,276.9	2,174.8	0.59
35–39	804.4	1,278.1	0.63
40–44	1,164.0	1,930.4	0.60
45–49	200.9	644.4	0.31
50–54	74.4	592.8	0.13
55–59	−407.1	−849.4	0.48
60–64	−323.9	−1,842.9	0.18
65–69	−10.6	−407.2	0.03
70–74	−289.1	−474.4	0.61
75–79	−219.9	−497.6	0.44

Note: Regressions included a constant, cohort dummies, and year dummies that are constrained to sum to zero and to be orthogonal to a linear trend.

to differentiate the spending response to permanent and transitory movements in income. Rather, this approach simply examines how closely consumption tracks income over the life cycle, in the spirit, for example, of Christopher D. Carroll and Lawrence H. Summers (1991).

We estimate dc/da and dy/da with separate regressions. To preserve the nonlinear age-consumption and age-income profiles apparent in the raw CEX cohort data, we use linear spline regression models with knots that correspond to the 12 five-year age groups between ages 20 and 80. Using household-level data, we regress real consumption (and then real income) on the spline terms, as well as cohort and time dummies, with the latter restricted as in our saving rate regressions. Our procedure estimates the typical change in consumption (and income) as households age; within five-year age groups, we impose a linear relationship with age, but slopes are allowed to vary across age categories.

Table 3 reports typical annual consumption changes for the various age groups in the left column, income changes in the middle column, and the ratios of changes in consumption to changes in income in the right column. The regression-based patterns are similar to those seen in Figure 1: consumption and income are steeply upward-sloping for households from their early 20s through their mid-40s; both variables reach a peak in the mid-50s.

The regression coefficients confirm that typical changes in consumption track those in income fairly closely across the life cycle—the age bins associated with the largest increases in income also show the largest increases in consumption, and, as increases in income tend to slow in the late 30s and early 40s, so do increases in consumption. The ratios of the coefficients suggest the closest tracking prior to the mid-40s, consistent with borrowing constraints or buffer-stock saving motives being important determinants of consumption at those stages.

The estimated age-income profile flattens out in the late 40s and early 50s, as does the age-consumption profile, and both age profiles become distinctly downward-sloping from the late 50s through the late 70s. On average, the ratios of the coefficients are smaller for these groups than for the younger groups. Possible interpretations would be that household incomes have risen sufficiently such that borrowing constraints no longer bind by these ages or that households have already built sufficient precautionary buffers of assets.

Given the limitations of our dc/dy measure noted above, we do not try to quantify the effects of population aging on an aggregate *mpc* as we did for the saving rate. However, the CEX data are at least suggestive that the aggregate *mpc* might have been expected to decline as the baby boomers moved out of their early working years and into the years in which their spending patterns were not so dominated by borrowing constraints or precautionary motives.

V. Time Series Data and the Aggregate *mpc* out of Disposable Personal Income

We can also use the reduced-form model in equation (1) to see whether the estimated aggregate *mpc* out of real current income (β_2) has fallen over time, as might be suggested by the analysis above. Estimating the model with rolling 20-year windows from 1963 to 2008, we find a decline in the point estimates of β_2: they average 0.20 in the early 1980s and 0.04 thereafter. This pattern is consistent with the results from household data in the last section, given the rise in the share of the population between ages 40 and 64 shown in Figure 2 (and the implied decline in the share of households in their early working years). However, trends in the population share and the estimated aggregate *mpc* are not perfectly aligned, and other changes over time (such as broader access to household credit) may well have also contributed to a lower *mpc*. Our household-level and time series analysis seems generally consistent with a relationship between demographics and the aggregate *mpc* out of current income over this period, but our analysis does not establish a strong or causal link.

VI. Conclusion

Our cohort analysis of household-level data in the CEX suggests that the movement of the baby boom into its higher saving years should have provided an increasing boost to the aggregate saving rate in the 1980s and 1990s. We expect some reversal of this pattern as this cohort moves into retirement in the current and next decade. We show with time series data that these effects appear to be at least roughly consistent with the realized pattern of the aggregate saving rate, after controlling for other determinants. We are hesitant to draw firm conclusions about the *mpc*, but our analysis suggests that demographics may have contributed in part to the trend toward a smaller response of aggregate consumption to income changes in the past two decades.

REFERENCES

Attanasio, Orazio P. 1998. "Cohort Analysis of Saving Behavior by U.S. Households." *Journal of Human Resources*, 33(3): 575–609.

Attanasio, Orazio P., James W. Banks, Costas Meghir, and Gugliemo Weber. 1999. "Humps and Bumps in Lifetime Consumption." *Journal of Business and Economic Statistics*, 17(1): 22–35.

Blinder, Alan S. 1975. "Distribution Effects and the Aggregate Consumption Function." *Journal of Political Economy*, 83(3): 447–75.

Browning, Martin, Angus Deaton, and Margaret Irish. 1985. "A Profitable Approach to Labor Supply and Commodity Demands over the Life-Cycle." *Econometrica*, 53(3): 503–43.

Carroll, Christopher D., and Lawrence H. Summers. 1991. "Consumption Growth Parallels Income Growth: Some New Evidence." In *National Saving and Economic Performance*, ed. B. D. Bernheim and J. B. Shoven, 305–43. Chicago: University of Chicago Press.

Gourinchas, Pierre-Olivier, and Jonathan A. Parker. 2002. "Consumption over the Life Cycle." *Econometrica*, 70(1): 47–89.

Kennickell, Arthur B. 1990. "Demographics and Household Savings." Board of Governors of the Federal Reserve System Finance and Economics Discussion Series 123.

This article has been cited by:

1. Marianne Bertrand, Adair Morse. 2016. Trickle-Down Consumption. *Review of Economics and Statistics* **98**:5, 863-879. [CrossRef]
2. Mark Vink. 2016. Intergenerational developments in household saving behaviour. *New Zealand Economic Papers* **50**:1, 3-28. [CrossRef]
3. Daniel L. Tortorice. 2014. CREDIT CONSTRAINTS, LEARNING, AND AGGREGATE CONSUMPTION VOLATILITY. *Macroeconomic Dynamics* **18**:02, 338-368. [CrossRef]
4. Z. He. 2012. Dynamic Compensation Contracts with Private Savings. *Review of Financial Studies* **25**:5, 1494-1549. [CrossRef]