Working Papers

FEDERAL RESERVE BANK OF BOSTON™

No. 23-5

Sustainable Consumption and the Comprehensive Economic Well-Being of American Households

Daniel H. Cooper, Barry Z. Cynamon, and Steven M. Fazzari

Abstract:
This paper develops a comprehensive measure of household economic well-being. The "sustainable consumption" concept accounts for income, assets, debt, transfer payments, and asset returns to estimate a consumption path that balances resources with expenditure over a household's lifetime. Calculating sustainable consumption using Panel Study of Income Dynamics data demonstrates that it acts as an anchor for actual household spending. Results show that following a period of rapid growth from the mid-1980s to the early 2000s, sustainable consumption stagnated on average. In the aftermath of the Great Recession, the decline in sustainable consumption exceeded the fall in actual consumption due in part to a decline in real asset returns. Decomposing sustainable consumption reveals the relative importance of different household resources in determining well-being and how these factors evolve over time—insights that would be missed when resources such as income or wealth are considered separately. Taxable income supports the majority of sustainable consumption; however, as a share of households' lifetime resources, taxable income has decreased on average while the Social Security share has grown.

JEL Classifications: E21, I31
Keywords: household spending, sustainable consumption, economic well-being

---

Daniel H. Cooper (daniel.cooper@bos.frb.org) is a vice president and economist in the Federal Reserve Bank of Boston Research Department. Barry Z. Cynamon (barry.cynamon@gmail.com) is an economist and former research fellow at the Federal Reserve Bank of St. Louis. Steven M. Fazzari (fazz@wustl.edu) is the Bert A. and Jeanette L. Lynch Distinguished Professor of Economics and a professor of sociology at Washington University in St. Louis.

Cynamon and Fazzari thank the Institute for New Economic Thinking, the Private Debt Project, and the Weidenbaum Center at Washington University for financial support. The authors are grateful to Nishan Jones, Sarah Morse, Hannah Rhodenhiser, and Zorah Zafari for research assistance. They thank J.W. Mason and Steven Waldman for comments on earlier drafts. They also thank participants in seminars at the Federal Reserve Bank of Atlanta, Colorado State University, and the Forum for Macroeconomics and Macroeconomic Policies Conference (Berlin) for helpful comments.

This version: July 2023 **https://doi.org/10.29412/res.wp.2023.05**

1. Introduction

Economists, policymakers, pundits, and the public focus much attention on the evolution of economic variables such as income, wealth, and debt at both the aggregate and household levels. Reasons vary, but an often implicit objective is to quantify and analyze the deeper concept of economic well-being. Undoubtedly, these variables correlate with well-being, but they provide only a partial picture.[1] This paper takes a comprehensive approach to analyzing household economic conditions—one that consolidates income, assets, debt, and other economic factors that affect well-being and its evolution over time. We call this concept *sustainable consumption*. In brief, sustainable consumption is based on a household's current balance sheet, a projected path for its income flows (including transfers), and a path for spending that exhausts the household's financial resources at the expected end of its lifetime, that is, the expected mortality date of the longest-surviving household member. A primary motivation for our work is summarized by Wolff et al. (2005, p. 1073): "The official measure of the level and distribution of economic well-being in the US is gross money income. . . . However, it is well known that this measure does not reflect households' command over, or access to, the products of a modern capitalist economy."

Over a household's life, its economic well-being is determined by its consumption path, subject to its intertemporal budget constraint.[2] We acknowledge that not every household will fully exhaust its resources and presume none will exceed its resources. In this sense, a consumption path is sustainable if it leads to a projected net worth of zero at the expected end of

[1] Many studies focus on a single dimension of household economic circumstances, often in the context of analyzing dynamics of inequality. The seminal work of Piketty and Saez (2003) explores income. Wolff (2014) considers wealth inequality alone. The Realtime Inequality website (https://realtimeinequality.org) presents data on income and wealth separately.
[2] Research using consumption measures to analyze household well-being includes Attanasio and Pistaferri (2014), Aguiar and Bils (2015), Fisher et al. (2015), Meyer and Sullivan (2017), and Meyer et al. (2022).

the household's lifetime. We estimate resources from current values of assets and liabilities, income, and taxes, employing assumptions about labor income growth along with projected rates of return on financial and housing assets. Projected transfers, especially Social Security, also add to lifetime resources.

For any time path of financial resources, and with the option for households to borrow and save, there is an infinite number of consumption paths that lead to projected zero net worth at the expected end of the household's lifetime. To measure sustainable consumption, we choose a path for household consumption that corresponds to the typical lifetime pattern of spending behavior observed in our data. This path is hump shaped and follows the similarly shaped path of income over a household's working years, which includes a significant decline during retirement. We do not claim such a path is necessarily optimal, but it is realistic, and thus sustainable consumption reflects how households, on average, spend their resources over time.

In addition to integrating all household resources and accounting for rates of return on assets, our sustainable consumption concept has several advantages over other measures of economic well-being. Importantly, it is comparable across households of different ages. All else being equal, the present value of future income is high when households are young, but their sustainable economic well-being may not be high because young households have to use their resources over more remaining years of life compared with older households. Also, one would expect older households to have more retirement savings than young households and therefore greater wealth, but older households are not necessarily systematically better off because they are closer to retirement, when their labor income will decline. Sustainable consumption adjusts for how such age differences affect available resources and therefore provides a more reliable measure of economic well-being over long periods of time as the age distribution shifts. In addition, because sustainable consumption is defined as an annual flow, it is directly comparable to actual consumption spending in a way that the present value of total household resources is not. Indeed, comparing actual and sustainable consumption at a point in time has important implications for future household well-being: If sustainable consumption is less than actual consumption, then household net worth will be less than projected along the current sustainable path, and future sustainable consumption will fall. Finally, our sustainable consumption measure allows one to analyze, on an annual flow basis, the impact of changes in macroeconomic variables, such as the interest rate, on economic well-being. This feature enables a straightforward assessment of the degree to which changes in macroeconomic conditions affect an average household.

We construct sustainable consumption and compare it with actual household expenditure using data from the Panel Study of Income Dynamics (PSID). The PSID has extensive data on income, wealth, demographics, and other characteristics at the individual and household levels. The panel structure of the data is critical for tracking household-level balance sheet changes across time and for measuring total household expenditure. The results show that, on average, sustainable consumption tracks the lifetime pattern of actual household spending. We also show that household-level shocks to sustainable consumption have large effects on actual spending, and when households spend more (or less) than their sustainable consumption in one period, they tend to adjust actual consumption toward the sustainable level in the subsequent period. This evidence supports our interpretation that sustainable consumption affects actual household spending over long horizons.

We also examine the source of differences across time in the fundamentals that determine household spending and economic well-being. Specifically, we can decompose sustainable consumption into the components determined by income, income growth, asset values, transfer payments, or asset returns. This decomposition allows us to analyze, for instance, how changes in the real rate of return affect the relative contribution of income flows arriving at different points during a household's lifetime to its measured well-being.

Our definition of sustainable consumption is related to the concepts developed by Auerbach et al. (2023), who take "a life-cycle perspective" and estimate the amount that a household can spend at a single point in time as determined by current net wealth and projected future labor earnings, taxes, and transfers.[3] By contrast, *evolution* of sustainable consumption over time is a major focus of this paper. Furthermore, whereas Auerbach et al. (2023) project the (constant) level of real consumption necessary to smooth individual living standards over future periods, we base our future consumption projection on the actual hump-shaped pattern of observed lifetime household spending. This approach facilitates comparisons of sustainable consumption across time periods that have different age distributions in the population. Furthermore, a realistic profile of consumption over households' lifetimes better reflects the effect of interest rates and asset returns on how they deploy financial resources over their life course. Given our data, we also can compare sustainable consumption with actual consumption. This approach relates to the study by Scholz et al. (2006) that estimates household wealth predictions based on a life-cycle model and compares them to actual wealth.

Our main results show that average household sustainable consumption grew rapidly from the mid-1980s to the early 2000s but then stagnated for several years before declining nearly 10 percent in the aftermath of the Great Recession. Although actual consumption expenditures follow a somewhat similar pattern, the shift in sustainable consumption was more

[3] See also McFall (2011), who uses the term "sustainable consumption" (with a definition similar to ours) to analyze the effect of wealth shocks during the Great Recession on retirement timing.

dramatic, due in large part to falling real asset returns. On average, sustainable consumption was less than actual consumption from 2004 through 2018, when our data end. This finding implies that American households typically did not accumulate enough wealth during the 2000s and 2010s to maintain the sustainable consumption path projected at the beginning of the century.

In addition, decomposing sustainable consumption reveals the extent to which income dominates relative to wealth in supporting household expenditures. The decomposition also highlights the importance of government transfer payments. Indeed, in 2018, projected transfer payments (excluding government health payments and dominated by Social Security) accounted for 20.3 percent of sustainable consumption on average.[4] By contrast, housing and financial net worth accounted for just 6.6 percent. The relative importance of income, wealth, and transfers also changes over our sample period. These changes arise in large part due to declining real rates of return that increase the relative value of Social Security and other income flows arriving later in a household's lifetime. Identifying and quantifying this effect on the annual level of sustainable consumption is, to our knowledge, an original contribution of our approach and is not possible using more common (independent) income or wealth measures of household well-being.[5]

The remainder of the paper proceeds as follows. Section 2 derives sustainable consumption from household balance sheet identities. Section 3 describes the PSID data we use to calculate sustainable consumption. Section 4 shows how our projected life-course consumption profile tracks the actual hump-shaped pattern of household spending and how average sustainable consumption evolves over time. Section 5 presents our decomposition of

[4] The PSID data include cash transfers to households but do not include non-cash transfers such as, most notably, Medicare and Medicaid benefits.
[5] Sabelhaus and Volz (2022) and Catherine et al. (2022) identify this effect for the present value of Social Security. Our approach generalizes the impact of declining real rates of return on household well-being.

sustainable consumption into its constituent components. Section 6 concludes and discusses the implications of our results for understanding the evolution of households' economic well-being in recent decades.

2. Sustainable Household Consumption

This section defines sustainable consumption—our measure of household economic well-being, hereafter denoted by $\hat{C}$. We assume a household consists of two individuals. Income variables with a *k* subscript are individual-specific and take on different values for each partner in the household (*k*=1, head of the household; *k*=2, spouse/partner). For a single-person household, the second income variable is set to zero.[6] We use these variables to define $\hat{C}$:

A_t: Assets, excluding owner-occupied homes, measured at the beginning of period *t*
H_t: Value of owner-occupied housing (if applicable) at the beginning of *t*
L_t: Liabilities (debt) at the beginning of *t*
r_t^A: Nominal rate of return on assets[7]
r_t^L: Nominal interest rate on liabilities
$Y_{k,t}^{NA}$: Disposable labor and transfer income for each partner *k*, *excluding* asset income
C_t: Consumption spending of the household (*excluding* interest payments)[8]
r_t^A: Nominal rate of return on financial assets
r_t^H: Nominal rate of return on housing assets
r_t^D: Nominal interest rate on debt

We derive $\hat{C}$ beginning with the law of motion for the value of household assets:

[6] In our empirical work, we account for income of other family members in addition to the two primary partners, such as children with income or an additional adult member of the household. Income of other family members is positive in 34.9 percent of the observations, but it accounts for just 4.9 percent of total household income.
[7] While we account for taxes on labor and transfer income, no taxes are imposed on asset returns, as if all asset accumulation for retirement occurs in tax-advantaged accounts, such as a Roth IRA. This approach biases the calculations in favor of higher $\hat{C}$ for households that accumulate savings outside of tax-sheltered retirement accounts.
[8] We use the term "consumption" to refer to cash expenditures by households that we compute from PSID survey data. This expenditure measure differs from the National Income and Product Accounts (NIPA) definition of personal consumption expenditure because it excludes imputations, such as the estimated rental value of owner-occupied housing. Also, it excludes transfer payments households never received as cash, most obviously government medical payments made directly to providers through Medicare and Medicaid. The consumption concept here corresponds closely to the adjusted consumption data presented in Cynamon and Fazzari (2017).

$$A_{t+1} + H_{t+1} = A_t(1 + r_t^A) + H_t(1 + r_t^H)$$
$$+(Y_{1,t}^{NA} + Y_{2,t}^{NA}) - C_t - r_t^L L_t + (L_{t+1} - L_t).$$

In addition, we define financial net worth as the difference between non-housing assets and total liabilities (including mortgage debt):

$$NW_t^{Fin} = A_t - L_t.$$

Combining these two identities yields an equation for the evolution of total household net worth (financial assets plus housing assets):

$$NW_{t+1}^{Fin} + H_{t+1} = A_t(1 + r_t^A) + H_t(1 + r_t^H) + (Y_{1,t}^{NA} + Y_{2,t}^{NA}) - C_t - (1 + r_t^L)L_t. \quad (1)$$

To make the solution for $\hat{C}$ tractable, assume the return on financial assets is the same as the interest rate on debt ($r_t^A = r_t^L = r_t$).[9] Then equation 1 reduces to a law of motion for household net worth that does not depend on the value of assets and liabilities separately:

$$NW_{t+1}^{Fin} + H_{t+1} = NW_t^{Fin}(1 + r_t) + H_t(1 + r_t^H) + (Y_{1,t}^{NA} + Y_{2,t}^{NA}) - C_t. \quad (2)$$

Note that debt service is implicit in equation 2. If financial liabilities exceed financial assets, NW_t^{Fin} is negative.

We further assume home values appreciate at a rate r_t^H and that a household owning a home at a given point in time continues to own the home for the remainder of its life. However, homeowners implicitly have access to a reverse mortgage at rate r_t that allows them to use their home equity to help finance consumption in retirement.

---

[9] If the interest rate on liabilities is different from the return on assets, the derivation of sustainable consumption will depend on the timing of debt principal payments. This timing is a function of other variables in the model such that the solution for sustainable consumption is highly nonlinear. Reported interest rate data in the PSID (available from 2000 onward) suggest that the average interest rate on debt (r_t^L) is 1.3 percentage points higher than the rate of return on assets (r_t^A) that we use in the $\hat{C}$ calculations. Appendix A presents an alternative approach to modeling (r_t^L) that provides an approximate upper bound on the effect on $\hat{C}$ if we assume that the interest rate on debt is greater than the nominal return on financial assets. The results in Appendix A show that this alternative calculation reduces $\hat{C}$ by an average of just 2.4 percent over time and has very little effect on qualitative trends in $\hat{C}$.

Our objective is to find a solution for $\hat{C}$ that is sustainable in the sense that the present value of a household's projected consumption path over its remaining lifetime equals its current net worth plus the present value of future income flows. That is, a household that consumes at $\hat{C}$ in the current period and stays on the sustainable consumption path going forward, while realizing its projected income and asset returns, would have terminal wealth of zero.[10] The final period for each household is the latest expected death year across the two partners denoted by $N = \max(N_1, N_2)$, where N_k is the expected death year of partner *k*. Terminal wealth is the sum of final-period financial assets and house values ($NW_{N+1}^{Fin} + H_{N+1}$) and is dated *N+1* because stock variables are defined at the beginning of each period. We assume the nominal rates of return on financial assets and owner-occupied housing are constant across years without loss of generality because only the geometric averages of time-varying returns affect the results. We solve for terminal wealth through backward iteration:

$$NW_{N+1}^{Fin} + H_{N+1} = NW_t^{Fin}(1+r)^{N-t} + H_t(1+r^H)^{N-t} + \sum_{j=t}^{N_1} Y_{1,j}^{NA}(1+r)^{N-j} + \sum_{j=t}^{N_2} Y_{2,j}^{NA}(1+r)^{N-j} - \sum_{j=t}^{N} C_j(1+r)^{N-j}. \quad (3)$$

To determine sustainable consumption at time *t* ($\hat{C}_t$), we replace the future values of $Y_{k,j}^{NA}$ and C_j with projections based on period *t* data from a given PSID survey wave. We can then solve equation 3 for a level of $\hat{C}_t$ that results in zero terminal wealth for each household.

We project non-asset income for each partner k ($Y_{k,t}^{NA}$) from data in year *t* by assuming it grows at a nominal rate $g_{k,i}^{Y}$ between each future period i and $i+1$ until partner k's normal retirement age. Projected future income in period *j* can therefore be written as:

[10] This approach assumes zero bequests, which, we propose, is consistent with the concept of sustainable consumption. If a household plans to leave a bequest, it might consume less than its sustainable level.

$$Y_{k,j}^{NA} = Y_{k,t}^{NA} \prod_{i=t}^{j} \left(1 + g_{k,i}^{Y}\right) \qquad t + 1 \leq j \leq M_k \,, \tag{4}$$

where M_k is partner k's normal retirement age. We allow $g_{k,i}^{Y}$ to vary depending on each partner's education level and age as described in the next section. Because $g_{k,i}^{Y}$ captures nominal income growth, it will also vary due to changes in projected inflation. Similar to that in Auerbach et al. (2023), this specification assumes that current income, $Y_{k,t}^{NA}$, represents "permanent income" in the sense that an additional dollar of income in period t translates into additional real income over partner k's remaining working years. In addition, we assume individuals receive income in retirement from Social Security, which we incorporate in present-value terms (*PVSS*) and describe further in the next section and in Appendix B.

Solving equation 3 for $\hat{C}_t$ also requires projecting a path of consumption over a household's lifetime (C_j where $j=t$ ... N). The most straightforward household consumption profile would assume complete expenditure smoothing as in Auerbach et al. (2023), where real consumption for the household remains constant over time. However, this assumption does not match actual consumption patterns over households' lifetimes observed in the data (see Section 4). Instead, the mean and median of real household consumption tend to mirror the hump-shaped path of real income during households' working years before declining substantially during households' retirement years. We do not take a stand on whether this time profile of consumption is optimal, but it is consistent with previous research.[11]

Given the spending profile observed in our data, we assume consumption grows at the same rate as household income until the household head reaches age 61.[12] After age 62, we

[11] Aguiar and Hurst (2013) present evidence of hump-shaped consumption profiles over households' lifetimes and provide references to related findings. Carroll and Summers (1991) find that consumption follows income over the life cycle. Also, see Attanasio and Weber (2010) for additional evidence and a survey of related research.

[12] We choose age 61 to avoid complications created by early retirement and calculate household income growth as the growth in the sum of both partners' labor and business income plus transfer income and the income of other

assume, consumption declines at a constant real rate calibrated to match households' spending reductions in their later years (see Section 4 and Appendix C for more details).[13] More specifically, we define the projected growth rate of consumption in future year *i* as g_i^C. Then, consumption (C_j) in equation 3 can be replaced by consumption in year *t* (C_t) times a product of future growth rates:

$$C_j = C_t \prod_{i=t}^{j} \left(1 + g_i^C\right) \qquad t \leq j \leq N. \tag{5}$$

To further clarify the exposition, we define the present value of projected future consumption implied by an extra dollar of current consumption at time *t* as:

$$V_t^C = \sum_{j=t}^{N} \left[\prod_{i=t}^{j} (1 + g_i^C) \right] \left[\frac{1}{(1+r)^{j-t}} \right], \tag{6}$$

and, similarly, we define the present value of future income implied by an extra dollar of current income for partner *k* at time *t* as:

$$V_{k,t}^Y = \sum_{j=t}^{M_k} \left[\prod_{i=t}^{j} (1 + g_{k,i}^Y) \right] \left[\frac{1}{(1+r)^{j-t}} \right]. \tag{7}$$

Recall that sustainable consumption based on data at time *t* $(\hat{C}_t)$ is the level of current consumption, given projected future resources, that results in zero terminal household net worth. Therefore, we can derive $\hat{C}_t$ by dividing equation 3 by $(1+r)^{N-t}$ and then substituting in the present-value terms from equations 6 and 7:

family members. Labor and business income account for 89 percent of household income on average. We also generated results assuming total household consumption drops by one-third in the expected death year of the shortest-lived partner. This alternative assumption has no material effect on our results.

[13] Scott et al. (2020) argue that substantial reductions in retirement consumption could be optimal for many households under realistic assumptions about time preference rates, asset returns, and health and mortality risks. In addition, Aguiar and Hurst (2013) identify reduced needs for some consumption categories in retirement, while Chen and Munnell (2021) find large annual consumption declines in retirement, especially for low-wealth households.

$$0 = NW_t^{Fin} + H_t\left(\frac{1+r^H}{1+r}\right)^{N-t} + PVSS_t + Y_{1,t}^{NA}V_{1,t}^Y + Y_{2,t}^{NA}V_{2,t}^Y - \hat{C}_t V_t^C$$

$$\hat{C}_t = (V_t^C)^{-1}\left[NW_t^{Fin} + H_t\left(\frac{1+r^H}{1+r}\right)^{N-t} + PVSS_t + Y_{1,t}^{NA}V_{1,t}^Y + Y_{2,t}^{NA}V_{2,t}^Y\right]. \quad (8)$$

Given equation 8, $\hat{C}_t$ is the level of current consumption that equates the present value of a household's lifetime projected resources with its lifetime projected consumption.

The *V* terms in equation 8 also have relevant interpretations. In particular, if r, r^H, and the assumed growth rates of consumption (g_i^C) and income ($g_{k,i}^Y$) in equations 6 and 7 are all zero, the expressions simply reflect years of life remaining (V_t^C) or years of work remaining ($V_{k,t}^Y$). For younger people, $V_{k,t}^Y$ will be larger because an extra dollar of income today, which we assume implies additional (permanent) income over the household's remaining working life, adds more to the household's lifetime resources because its members have more working years remaining.

Therefore, sustainable consumption ($\hat{C}_t$) is a forward-looking estimate of the fundamental economic factors and household-specific circumstances that affect households' decisions regarding their actual spending in year *t*. Specifically, $\hat{C}_t$ combines households' current balance sheet and income information with realistic projections of future consumption and income flows (including Social Security retirement benefits) to yield a time path of sustainable consumption. It is also important to recognize that $\hat{C}_t$ projects sustainable consumption for a given household at a point in time. If household fundamentals change between *t* and *t*+1 (for example, income, asset returns, or house prices), or if the household consumes more or less than $\hat{C}_t$ in the current period, then actual sustainable consumption in *t*+1 will differ from the level projected at time *t*.

3. Data and Implementation

To operationalize our approach for measuring sustainable consumption, we need household-level data on income, assets, and debt along with projections of asset returns and interest rates on household liabilities. To estimate life-cycle consumption dynamics and to validate our approach and compare sustainable and actual consumption, we also need data on household expenditures.

Expenditures

Microdata on household income and balance sheet variables are available from several sources, but only the PSID provides the data we need to estimate total household expenditures. Indeed, while the Consumer Expenditure Survey (CEX) contains comprehensive household expenditure data, it lacks adequate balance sheet information and has been well known to understate total household consumption, most obviously for higher income households (Cooper 2010; Sabelhaus 2010; Carroll et al. 2015).[14] While the reported expenditure data in the PSID have problems similar to the CEX data, the wealth data and the panel structure allow us to estimate total household outlays (defined as consumption plus personal transfers and debt service).

Following the approach of Ziliak (1998) and Cooper (2010), we use the PSID data to estimate household outlays as follows:[15]

$$Outlays = Income - Taxes - Active\ \Delta NW^{Fin}. \tag{9}$$

The active change in financial net worth (final term) excludes capital gains resulting from the change in asset prices. If positive, it equals the net accumulation of financial and housing assets;

[14] Passero et al. (2015) show that the understatement of consumption in the CEX increased significantly from 1992 to 2010.
[15] See Browning et al. (2013) and Crawley and Kuchler (2023) for a similar approach using Danish data.

if negative, it is the value of outlays financed by a net increase in debt or a drawing down of financial assets. The panel structure of the PSID allows us to compute the change in net worth within households—a variable that cannot be calculated from cross-sectional household surveys such as the Survey of Consumer Finances (SCF).[16] The PSID does not record information on household taxes, but following the existing literature, we estimate tax payments using the National Bureau of Economic Research (NBER) TAXSIM software, which calculates federal and state tax liabilities from the detailed household characteristics and income data available in the PSID.

The PSID did not collect household balance sheet data until 1984, and from that year until the 1999 wave, the data were recorded only every five years. Since 1999, the PSID has been conducted biennially, and each wave contains household balance sheet information. In addition, every wave of the PSID contains detailed income data measuring receipts in the preceding calendar year. The latest survey available as of this writing was the 2019 wave, for which respondents reported their 2018 income. Differencing the balance sheet data between waves and taking into account the timing of reported household income yields household outlay estimates for three five-year periods (1984 through 1988, 1989 through 1993, and 1994 through 1998) and 10 two-year periods (1999 and 2000 through 2017 and 2018). For simplicity of exposition, we denote the five-year expenditure periods using their midpoint date (1986, 1991, and 1996), and we identify each two-year period using its second year (2000, 2002, . . . 2018) when presenting our results. We estimate real, annualized outlays (equation 9) using the average of the personal consumption price index over each period.

[16] Indeed, while the SCF contains excellent household balance sheet data and added expenditure data in recent waves, it is not a panel, and thus one cannot calculate active savings (and total household outlays).

The primary difference between outlays and consumption (excluding imputed service flows) is that outlays include interest payments on debt. To generate a measure of household consumption, we remove debt service costs from outlays (see Appendix D for further details).[17]

Income

To calculate the present value of future household income, we follow a two-step process. First, we determine non-asset (labor plus business) income for each household partner ($Y_{k,t}^{NA}$) in every period for which we have data.[18] Second, we project each partner's income growth forward until the age they become eligible for full Social Security benefits (65 to 67, based on their birth year, rounded to the nearest year). Projected future growth rates for each partner's labor and business income consist of three components: inflation, aggregate real-wage growth, and an estimated profile of real income changes across individuals' working years that accounts for differences based on age and education. We treat labor and business income symmetrically and discount each partner's income growth profile back to the current period to calculate $V_{k,t}^{Y}$ as defined in equation 7. (Appendix F provides further details on our projection approach.)

In addition, the difference between permanent and transitory income is an important issue when studying household financial sustainability because transitory income fluctuations will introduce unnecessary volatility into our projected income profiles and in turn into $\hat{C}_t$. We address this issue by using an algorithm to estimate (after-tax) permanent income for each household in every survey year in our sample, as described in Appendix E.

[17] Personal transfers, such as charitable contributions, are treated as consumption for our purposes. These items are likely small, as they have averaged about 1 percent of aggregate household disposable income in recent years, according to National Income and Product Accounts (NIPA) data.
[18] Non-asset income data for each partner come from individuals' reported labor and business income, which the PSID tracks separately. However, from 1983 to 1991, the PSID combined household partners' business income. For these years, we split business income between partners based on each partner's share of household labor income.

As noted earlier, our measure of household income also includes two additional sources of potential resources for financing consumption. The first is reported transfer income obtained by the household prior to its normal retirement age. This source of income as a share of total household resources is zero or small for most observations in our sample of working-age households. Indeed, transfer income comprises just 5.5 percent of total income included in our $\hat{C}$ calculations. However, transfer income is substantial for some households (the 90th percentile of the distribution of the share of transfer income in total household income is 34 percent, and the 95th percentile is 59 percent). For simplicity, we assume transfer income is nontaxable and grows at the projected inflation rate until the household head's normal retirement age.[19] We also include taxable income earned by other members of the household if the household reports non-zero "other family member" income. We assume this income will grow at the projected inflation rate until the household head's normal retirement age[20] and that taxes on it are levied at the household's average tax rate.

Household Net Worth

Another key variable in the definition of $\hat{C}$ is household net worth. The PSID includes values for households' cash, stocks, bonds, retirement accounts, and real estate holdings. Liabilities consist of the remaining principal on mortgages, vehicle debt, and other debt (including credit card balances and student loans). We treat all financial assets as being held in a

[19] Because tax rates are usually low for households with substantial transfer income, treating transfer income as taxable makes very little difference for our results. Imposing average household tax rates on transfer income reduces total household resources by less than 1 percent for 90 percent of the observations in our sample and less than 5 percent for 99 percent of the observations. The logic behind stopping transfer income at normal retirement age is that we incorporate a separate estimate of Social Security income at that time.

[20] This assumption will overstate resources for households with income-earning children. However, "other family member" income for households with one or more children is just 3.3 percent of total family income.

tax-advantaged account, such as a Roth IRA, and thus financial returns in our calculations are nontaxable.

Social Security Benefits

Finally, we use benefit formulas from the Social Security Administration to estimate $PVSS_t$ given information about household structure and each partner's income (if applicable). First, we obtain labor income data from each partner's earliest recorded income observation in the PSID through the current time period *t*. Second, we project each partner's income up to their full Social Security retirement age. Lastly, we employ an algorithm that chooses an approach to claiming Social Security benefits that maximizes their present value for the household. (See Appendix B for a full description of our approach.)

Although the PSID includes data on assets in retirement accounts, it does not include defined-benefit (DB) pension information. Therefore, our estimates of $\hat{C}$ are potentially biased downward. owever, this bias is declining over time, as DB pension plans have become much less prevalent for workers. At the same time, the lack of DB pension information induces a positive bias in the *trend* of $\hat{C}$ because omitted DB pension resources are shrinking over time.[21]

Nominal Rates of Return and the Inflation Rate

Future rates of return play a significant role in our sustainability calculations. The horizon for these returns is long term, as they are relevant for the remainder of households' lifetimes. We project financial rates of return using the 10-year-ahead forecasts in the Survey of Professional Forecasters (SPF) for equity, bond, and cash returns. The overall financial asset return for each household, which functions as the discount rate in equation 8, is a weighted average of the SPF

[21] Butrica et al. (2009) discuss and evaluate BLS data and show that the share of wage and salary workers with access to a DB pension fell from 38 percent in 1980 to 20 percent in 2008. Access to a DB pension fell further to 16 percent in 2019, with only 12 percent of workers actually participating, according to our own calculations of updated BLS data.

forecasts for the different asset returns. Portfolio weights are determined by household-specific wealth holdings, as described in detail in Appendix G.

Long-term SPF projections for housing returns are limited. Instead, we use a fixed long-term average of real housing returns (1.45 percent) estimated from the Case-Shiller house-price index. Projected nominal house price growth is this real return plus forecasted inflation. (See Appendix G for further details).

Finally, expected inflation is based on the SPF 10-year-ahead inflation forecast, which declined from 4.8 percent in the mid-1980s to near 2 percent by 2008. Importantly, projected nominal asset returns declined more than forecasted inflation over time, and thus real returns are lower at the end of our sample period than in the beginning. This reality has important implications for movements in average $\hat{C}$ over time, as we discuss further in sections 4 and 5.

Estimation Sample

Our analysis sample contains 60,764 observations for 14,370 unique households. Due to data availability, the number of observations per period is lower early in our sample period, although even the smallest sample count—3,140 households for the 1994–1998 period—is substantial.

In addition, as discussed earlier, the consumption data cover periods between waves of the survey when the necessary household balance sheet data are available for calculating household outlays and beginning-of-period stocks of assets and debts. We impose several rules, explained in detail in Appendix H, to ensure consistent data over time. The most important restriction limits the sample to households headed by someone aged 25 to 61 years old. This restriction is designed to capture individuals when they have completed most of their education and before they retire.[22] As a result, our sample is not fully representative of the US population; however, our data are broadly consistent with aggregate trends in household financial and spending data.[23]

4. The Evolution of Household Well-Being

This section begins with evidence supporting the assumption in our calculations that households' actual consumption over their lifetimes follows the lifetime profile of income through most of a household's working years, and then consumption declines sharply when the household head passes age 61. We then show how sustainable consumption evolves over time.

Figure 1 depicts mean real consumption and income in the PSID by age of the household head and shows a clear hump-shaped pattern of consumption over the life cycle. (To ensure adequate sample size, the figure reports mean real consumption and income for three-year age ranges.) Indeed, real consumption for a household in its late 40s is two times larger than it is for a household in its late 20s. Average consumption then stays roughly constant for households with heads in their 50s before it declines significantly during typical retirement ages.

This consumption pattern closely follows real, after-tax income (dashed line) during households' normal working years.[24] From age 30 to 60, the ratio of consumption to income remains in a narrow range between 0.875 and 0.915. Patterns for subgroups are similar. In

---

[22] Altig et al. (2022) show that approximately half of Americans are working fewer than 20 hours per week by age 62.

[23] The PSID does not contain enough information on the highest-income households to adequately represent the top 1 percent of the income distribution. For example, in 2018, the average income of the top 1 percent of households in our data was $697,000 compared with an average in the population of $2 million, according to data from the Congressional Budget Office ("The Distribution of Income, 2018," Congressional Budget Office [2021], exhibit 1, page 6. See https://www.cbo.gov/system/files/2021-08/57061-Distribution-Household-Income.pdf).

[24] The income profile in Figure 1 is qualitatively similar to the income-age profile derived from the SCF (see Feiveson and Sabelhaus 2019, figure 4). Our observed consumption profile is similar to theirs over individuals' standard working years, but the spending decline we observe over individuals' retirement years is more substantial. One possible explanation for this difference is that the SCF oversamples high-income and high-wealth households, and such households are better able to maintain their consumption during retirement.

particular, the mean income and consumption profiles are hump-shaped and proportional across education-level groups (not shown). This evidence supports our approach of projecting consumption as proportional to after-tax income during households' normal working years and then declining at an annual rate after age 61.

Figure 1: Mean Real Consumption and After-Tax Income across Age Groups

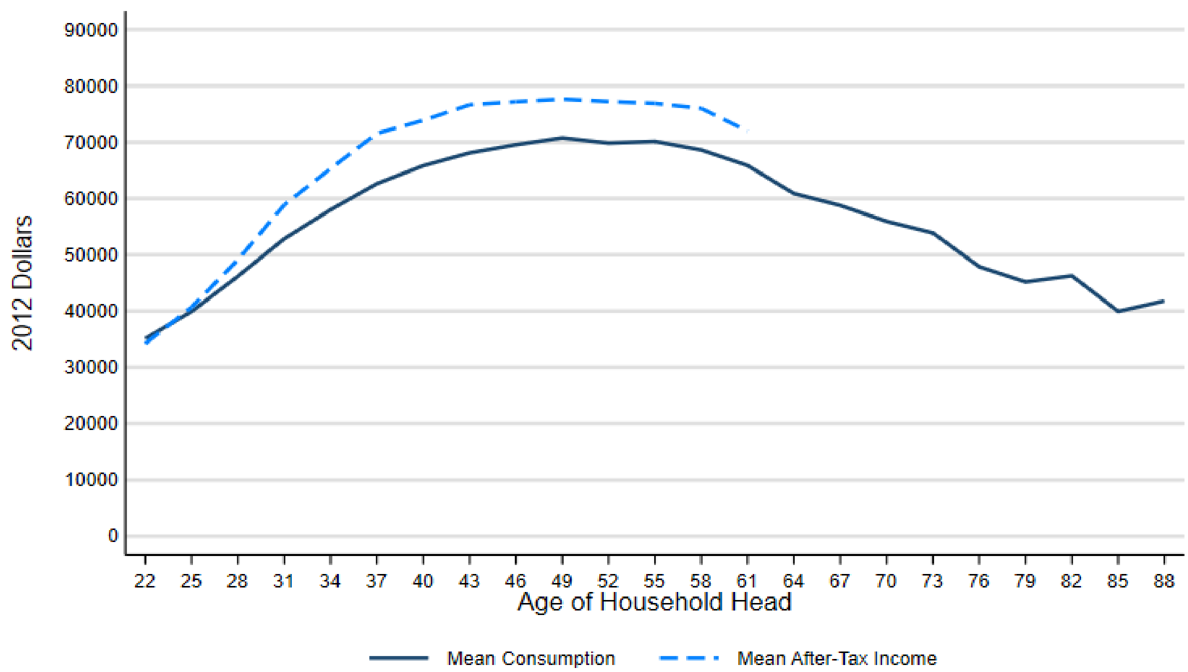

Note: The displayed age covers a three-year interval of household head ages to ensure adequate sample size (for example, age 25 includes households with heads aged 23 to 25 and so on).
Source: Authors' calculations using PSID data.

To compute sustainable consumption at any point in time ($\hat{C}_t$), we need to calibrate the projected annual change in real consumption when household heads reach age 62. We set this growth rate to –1.4 percent, the median within-household annual growth rate of real consumption for households with heads aged 62 to 90. (See Appendix C for further information about this calibration.)

Figure 2 presents the mean and median sustainable consumption ($\hat{C}$) and actual consumption (C) by age group. Median $\hat{C}$ tracks median C closely (the average difference between medians is 0.8 percent across all ages), and mean $\hat{C}$ is almost identical to mean C for middle-aged households, although the gap widens for the oldest age groups (the average difference between means across all ages is 3.1 percent).

Figure 2: Sustainable Consumption ($\widehat{C}$) and Actual Consumption (C) by Age

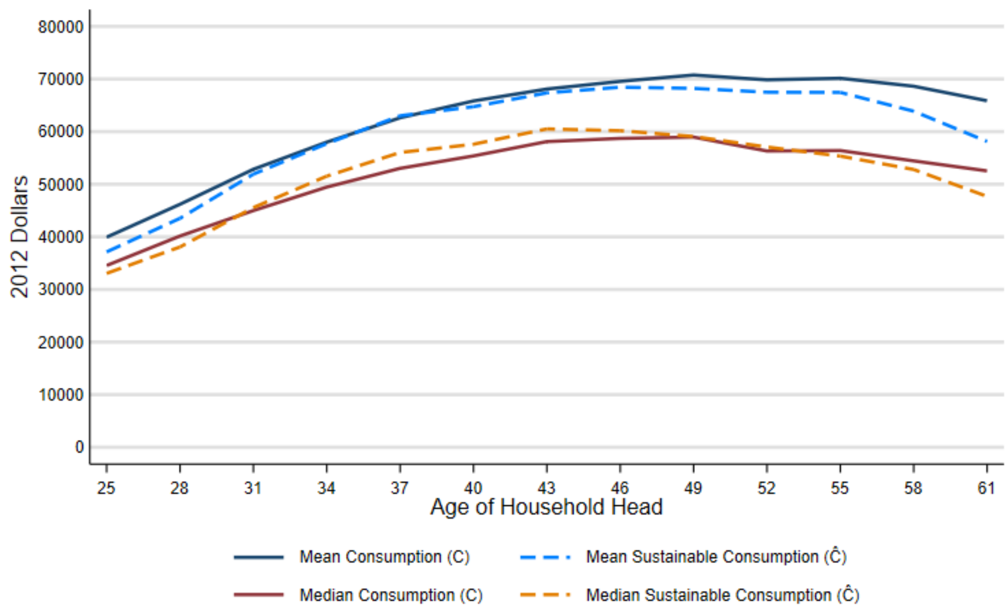

Note: Sustainable consumption calculated based on equation 8. The displayed age covers a three-year interval of household head ages to ensure adequate sample size (for example, age 25 includes households with heads aged 23 to 25 and so on).
Source: Authors' calculations using PSID data.

The results in Figures 1 and 2 suggest that households do not completely smooth real consumption over their lifetimes. Understanding the source of the hump-shaped spending pattern would be interesting, but it is not particularly important for our purposes. What is relevant for this analysis is that the pattern of consumption we project describes how households actually allocate expenditures over time. Indeed, the figures show that our measure of sustainable consumption across the age distribution effectively captures the central tendency of actual consumption over households' lifetimes using very parsimonious assumptions about the age distribution of actual expenditures.

In addition, at a disaggregated level, there is evidence that $\hat{C}$ anchors C. Consider the simple regression explaining consumption for household *j* at time *t*:

$$C_{jt} = \beta_{0j} + \beta_1 \hat{C}_{jt} + \beta_2 \left(C_{jt-1} - \hat{C}_{jt-1} \right) + \epsilon_{jt}, \tag{10}$$

where β_2 captures a correction factor for past shortfalls or overspending relative to the households' sustainable expenditure level. The idea is that if C_{jt-1} exceeds $\hat{C}_{jt-1}$ for household *j* at *t*-1, actual consumption for the household at time *t* will be lower, all else being equal, and vice versa. Table 1 shows estimates of equation 10 with household fixed effects and in first differences.

Table 1: Household-Level Consumption Regression (Levels)

Estimation Method	$\boldsymbol{\beta_1}$	$\boldsymbol{\beta_2}$
Fixed Effects	1.204 (0.152)	–0.172 (0.042)
First Differences	1.214 (0.141)	–0.580 (0.020)

Note: Standard errors clustered at the household level are in parentheses.
Source: Authors' calculations using PSID data.

The estimates of β_1 are larger but not significantly different from one, suggesting that actual consumption closely tracks sustainable consumption on average. In addition, the correction-term estimates are strongly negative, implying that deviations of actual consumption from the sustainable consumption path predict future consumption. Overall, the findings in Table 1 support our interpretation of $\hat{C}$ as a measure of the fundamental factors that determine how much households can sustainably consume, and therefore it serves as a comprehensive measure of economic well-being.[25]

[25] The results in Table 1 are robust to including controls for household age, the presence of children, and any time variation in education and marital status. Including fixed effects or differencing consumption within households captures any relevant (unobserved) time-independent differences across households.

Figure 3 shows mean real $\hat{C}$ and C across households over time. Over our entire sample, the two measures have the same shape, but there are some important differences across subperiods. From 1986 to 2000, mean $\hat{C}$ grows somewhat faster than mean C (1.7 percent versus 1.3 percent). This trend turns a small negative gap between $\hat{C}$ and C (3.5 percent) at the start of the period into a small positive gap (1.9 percent) by 2000. However, in 2000 and 2002, mean C growth slows as the economy experiences a mild recession, and then beginning in 2004, growth in $\hat{C}$ stagnates due largely to slower income growth.[26] The growth of C also stagnates starting in 2004, but strong growth from 2002 to 2004 restores actual consumption close to its longer-term trend. As a result, actual consumption runs above our estimate of its sustainable level for the remainder of our sample period.

[26] The fact that $\hat{C}$ does not decline until 2004 despite the mild recession that occurs in the early 2000s is a result of our approach that incorporates the permanent (smoother) component of household income.

Figure 3: Mean Sustainable Consumption ($\widehat{C}$) and Mean Actual Consumption (C) across Survey Years

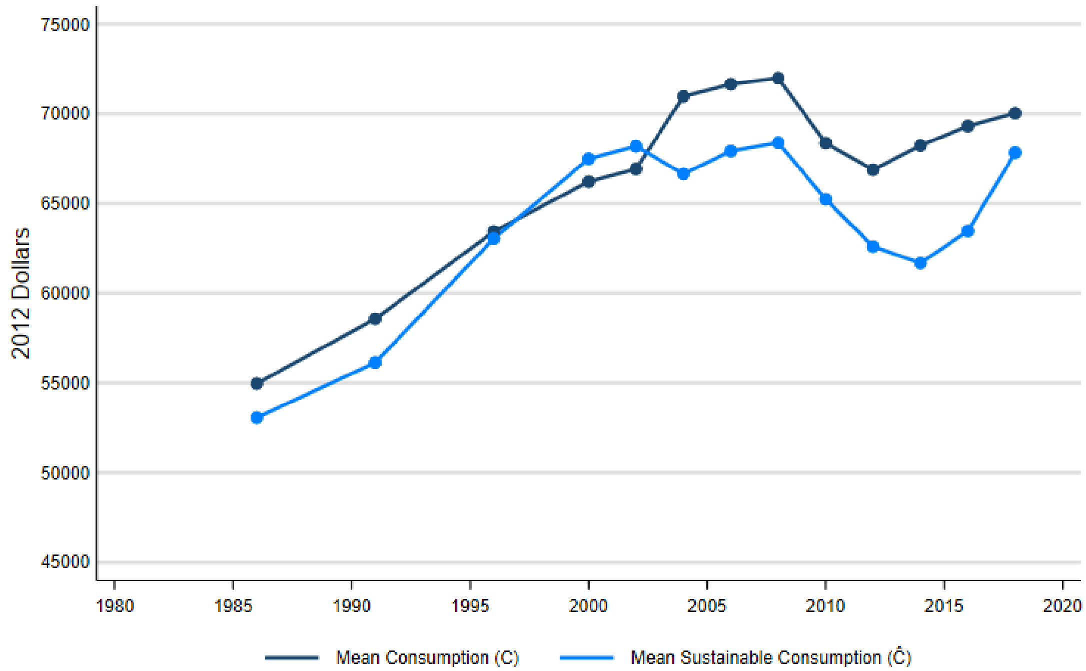

Note: The years shown represent the midpoint of the multi-year expenditure periods as discussed in Section 3.
Source: Authors' calculations using PSID data.

Figure 3 also shows actual consumption falling noticeably during and after the Great Recession. Indeed, mean real C declines by 7.1 percent from 2008 to 2012. The size of this drop may seem exceptionally large considering that real personal consumption expenditures (PCE) per capita from the national accounts fell just 3.6 percent peak to trough over those years. However, PCE includes large expenditures that are not captured in the PSID, such as government-financed health care and the imputed service flow from owner-occupied housing.[27] Figure 3 further highlights that mean $\widehat{C}$ decreased more severely than C in the Great Recession, dropping 9.8 percent from 2008 to 2014. As a result, the gap between actual and sustainable

[27] Our PSID consumption concept is more closely aligned with the adjusted consumption concept from Cynamon and Fazzari (2017), which measures the actual cash flows of spending from the household sector, than it is with PCE from NIPA data. This cash-flow measure (real, per capita) declined 6.8 percent in the Great Recession, peak to trough, similar to our observed decline in mean expenditures in the PSID.

(mean) consumption reached its largest point (9.6 percent) over our sample period in 2014. However, this gap subsequently narrowed as actual C rose slowly from 2012 through the end of the sample, but $\hat{C}$ rose noticeably in 2018. The resulting 3.1 percent negative gap between sustainable and actual consumption in 2018 is similar in size to the gap that preceded the Great Recession.

Overall, following two decades (1984 to 2004) in which mean $\hat{C}$ was nearly equal to or even somewhat higher than C, the data from 2004 to 2018 show actual consumption running persistently above sustainable consumption ($C > \hat{C}$). This pattern implies households, on average, were not accumulating net worth at a rate consistent with typical lifetime consumption patterns. If C exceeds $\hat{C}$ year after year, the cumulative impact reduces the level of sustainable household expenditures. More specifically, households spending in excess of the sustainable level from 2004 through 2018 results in roughly $31,000 less net worth accumulation, on average, relative to what would be expected if average C equaled average $\hat{C}$ over this period. Going forward, this loss in average available household resources implies lower $\hat{C}$ at the beginning of 2020 by about $1,100, or 1.7 percent of 2018 $\hat{C}$.[28] This implied reduction in sustainable consumption is economically meaningful for a single year, but it is especially striking because it is *permanent* if households maintain actual spending at the current (lower) level of $\hat{C}$ going forward. All else being equal, the only way for households to make up for lost resources and raise $\hat{C}$ would be to spend less than $\hat{C}$ for an extended period.

[28] To estimate the effect of lower wealth on $\hat{C}$, we use an average "wealth effect" of 3.6 percent (the average of $1/V_t^C$ from equation 8 over the 2006–2018 period).

5. Decomposition of Resources Supporting Sustainable Consumption

The previous section describes how the evolution of average $\hat{C}$ over time compares with C. To better understand the drivers of $\hat{C}$ over our sample period and what they imply for the economic condition of American households, we use equation 8 to decompose $\hat{C}$ into the five fundamental components of household resources that support it. These include projected taxable income, projected transfer income, the value of owner-occupied homes, financial net worth, and projected Social Security benefits. At the household level, the contribution of each resource is the amount of the household's annual $\hat{C}$ supported by the resource.[29]

Table 2 shows the results of this decomposition for our full sample when we average contributions across households and over time. It is not surprising that taxable income accounts for the largest share of household resources supporting $\hat{C}$. Among the remaining components, projected Social Security benefits dominate, with a share nearly three times larger than the contribution of housing and financial net worth combined.[30] Transfer income for households with a head aged 25 to 61 is a relatively small contributor, and financial net worth plays only a small supporting role on average.[31]

[29] For example, from equation 8, the financial net worth component for each household is NW_t^{Fin}/V_t^C; the other components are defined similarly.

[30] Altig et al. (2022) and Sabelhaus and Volz (2022) also emphasize the importance of Social Security as a share of household resources.

[31] Note that our approach in calculating $\hat{C}$ includes house values in the housing component, while mortgage debt reduces households' financial net worth. We use this approach because debt service on mortgages is more closely related to current financial conditions than to the real growth of house prices. If mortgages were shifted to the housing component of $\hat{C}$, the housing share would fall and the financial net worth share would rise, but the sum of the two components would be unchanged.

Table 2: Decomposition of Resource Support for $\widehat{C}$ (Full Sample)

Resource Component	Mean Contribution to $\widehat{C}$
Taxable Income	73.1%
Social Security	16.0%
Owner-Occupied Housing	5.3%
Transfer Payments	4.3%
Financial Net Worth	1.3%

Note: The numbers in the table are the mean of each component across households in a given year divided by the mean of $\hat{C}$ for the same year then averaged over all years.
Source: Authors' calculations using PSID data.

Figure 4 shows the decomposition of $\hat{C}$ over time. The contribution of taxable income drops noticeably from a peak of about 79 percent in the early years to just under 70 percent by the end of our sample, while the Social Security contribution grows from about 12 percent in the 1990s to nearly 19 percent by 2018. In addition, the housing share approximately doubles during the housing boom from the late 1990s up to the Great Recession, although it never exceeds 7 percent of total resources supporting $\hat{C}$, on average. During the Great Recession, the housing contribution falls to about 6 percent—less than one might expect considering the collapse in home prices, as we discuss further below.

The observed changes over time in the composition of resources supporting sustainable consumption are explained to a large extent by the decline in the real return on assets. More specifically, let ρ be the real rate of return, defined as the difference between the mean weighted average nominal return on assets across households and the expected inflation rate (see Appendix G for further details). From 1986 through 2000, ρ is nearly constant, averaging 3.3 percent, but it falls below 3 percent when interest rates decline during the 2001 recession before decreasing noticeably further during the Great Recession. From 2014 to 2018, ρ averages just 1.6 percent.

Figure 4: Contribution of Household Resources to Sustainable Consumption

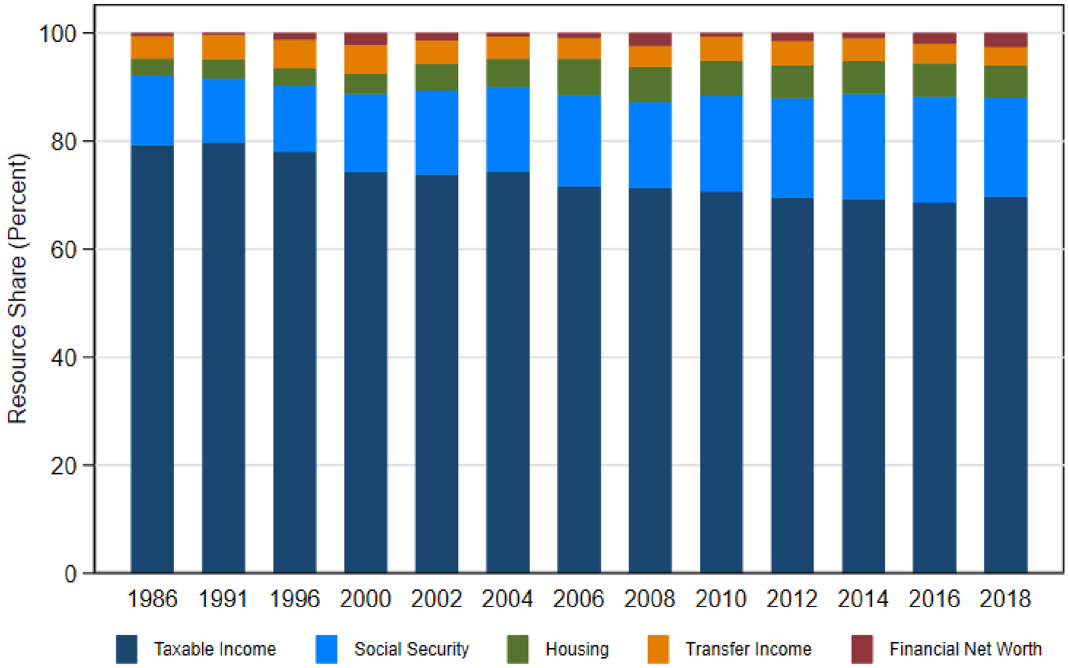

Note: The figures are the mean of each component across households divided by the mean of $\hat{C}$ in a given year.
Source: Authors' calculations using PSID data.

To better understand how a decline in ρ affects $\hat{C}$, first consider an unrealistic special case in which consumption is financed entirely by labor income and projected income and consumption flows are perfectly aligned over a household's lifetime. In this case, households would simply consume what they earn period by period, and the level of ρ would have no effect on $\hat{C}$. But planned retirement imposes a timing mismatch between income and consumption that increases the effect of ρ. One way to think about this effect is that a household's consumption stream has a longer financial duration than its labor income stream. That is, over a household's lifetime, the flow of labor income arrives, on average, earlier than consumption occurs. This difference in duration makes the present value of a household's projected consumption stream more sensitive to ρ than to the present value of its stream of income.

To evaluate the size of this effect, consider the ratio $V_{k,t}^{Y}/V_{t}^{C}$ (based on equations 6 and 7), which can be thought of as capturing the income capitalization rate. In particular, using equation 8, the ratio shows how much current sustainable consumption an additional dollar of current income earned by one partner "buys" given our projections of income and consumption over a household's lifetime. This ratio is less than one because of retirement. (If the income and consumption real growth rates and ρ were all zero, the income capitalization ratio would just be the number of the partner's remaining working years divided by the number of projected remaining years of life.) As discussed in the previous paragraph, a decline in ρ raises the present value of the longer duration household consumption stream (V_{t}^{C}) more than the present value of its income stream ($V_{k,t}^{Y}$), so the ratio declines when ρ falls.

Figure 5: Real Rate of Return (Left Scale) and Income Capitalization Ratio (Right Scale)

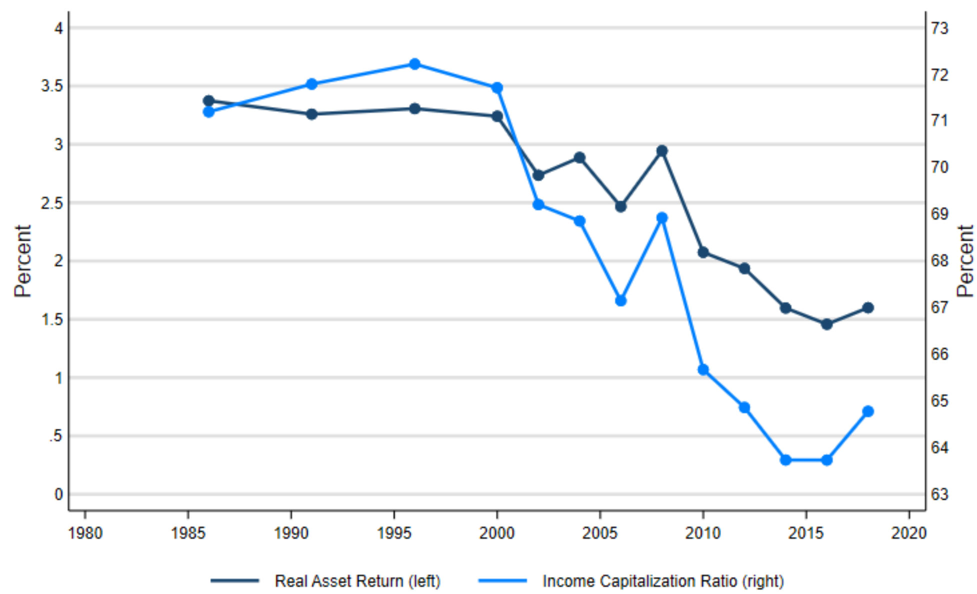

Note: The real rate of return (ρ) is the average return on household asset portfolios less the average inflation rate. The income capitalization ratio, $V^{Y}_{ki,t}/V^{C}_{t}$, is averaged across households.
Source: Authors' calculations using PSID data.

This effect is evident in Figure 5, which plots the average value of ρ and the $V_{k,t}^{Y}/V_{t}^{C}$ ratio over our sample period. The income capitalization ratio declines steeply from 2000 to 2014.

Over the same period, the change in the income component of $\hat{C}$ reduces average $\hat{C}$ by \$7,537 (15 percent). The fall in the capitalization ratio accounts for approximately *70 percent* of this decline in sustainable consumption, with the remaining 30 percent coming from a decrease in current mean real after-tax household income. While lower real returns clearly affect the income of households that hold financial assets, they also matter for households with minimal financial wealth as of a given survey date. For low-wealth households, lower real returns reduce $\hat{C}$ because, along any given sustainable path, almost all households will need to save during their remaining working years to finance projected retirement expenditures. These results suggest that the low-rate-of-return environment since the early 2000s helps explain the stagnation in the average economic well-being of American households through 2018 based on our $\hat{C}$ measure.

The decline in ρ also affects the extent to which homeownership supports $\hat{C}$. In particular, equation 8 identifies the value of a dollar of home value relative to a dollar of financial assets in the $\hat{C}$ calculation as $[(1+r^H)/(1+r)]^{N-t}$, where r^H and r are the projected nominal rates of return on housing and financial assets, respectively. If r falls relative to r^H, a given home value at time t will support more $\hat{C}$. The intuition for this effect is that our methodology allows a household to effectively borrow against the value of its home and pay off the debt by selling the home in the household's final year of life (N). The real cost of that implicit loan, effectively a reverse mortgage, is ρ (inflation affects both r^H and r symmetrically; see Appendix G). Figure 6 shows average reported beginning-of-period home values and the average housing contribution to $\hat{C}$ for homeowners. Until 2008, the $\hat{C}$ housing component closely follows home prices. But although home values in our data collapsed during the Great Recession, the real (dollar) contribution of owner-occupied housing to $\hat{C}$ was roughly constant because of the decline in ρ.

Figure 6: Mean Real Home Prices (Left Scale) and Mean Real Housing Contribution to $\hat{C}$ (Right Scale) for Home-Owning Households Only

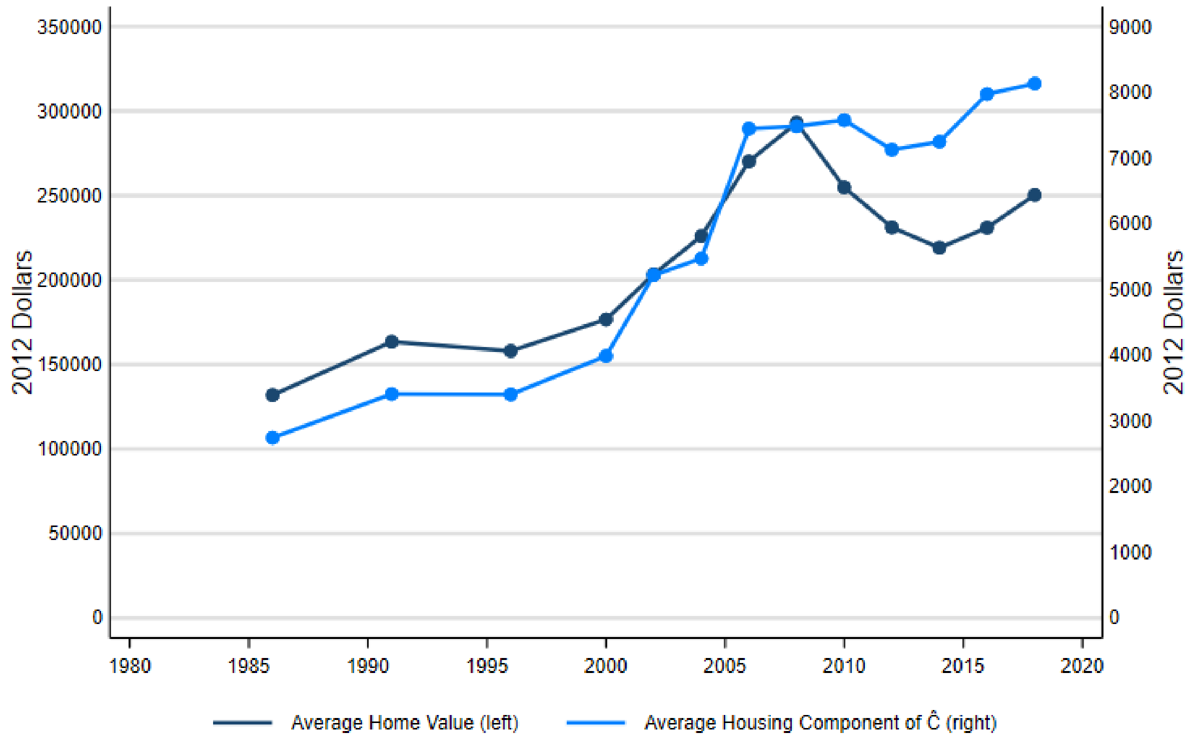

Note: Home values are self-reported by homeowners in the PSID. Averages across home-owners by sample period.
Source: Authors' calculations using PSID data.

Finally, the downward trend in real asset returns also has an important impact on the contribution of Social Security to $\hat{C}$ over time (see also Sabelhaus and Volz 2022 and Catherine et al. 2022). Indeed, the real value of the Social Security component of average $\hat{C}$ nearly doubles over our sample period (from $6,821 in 1986 to $12,378 in 2018), and the share of Social Security in resources supporting $\hat{C}$ rises from about 12 percent to 19 percent. Through 2000, these gains are driven by rising real incomes that increase the wage base for Social Security. After 2000, however, incomes are stagnant, but the Social Security component of $\hat{C}$ continues to rise as ρ declines. In contrast to labor income receipts, projected Social Security payments for working-age households occur later in life compared with the time profile of consumption.

Therefore, a fall in the real return raises the present value of Social Security relative to the present value of consumption.

Quantifying how real asset returns affect $\hat{C}$ and how changes in the relative importance of different sources of household resources impact household well-being over time is, to our knowledge, an original contribution of our approach measuring sustainable consumption. The results in this section demonstrate how the low-interest-rate environment in the aftermath of the Great Recession had economically meaningful effects on the well-being of American households. The analysis also reveals how changing asset returns matter for the relative value of different components of households' lifetime resources: They reduce the value of income relative to financial and housing assets and substantially increase the importance of Social Security benefits.

6. Discussion and Conclusions

This paper proposes a comprehensive measure of household economic well-being that we call sustainable consumption, denoted by $\hat{C}$. In our analysis, the resources that support $\hat{C}$ are a combination of households' current wealth and estimates of both after-tax income and cash transfers over their lifetimes. We measure $\hat{C}$ empirically using PSID data for working-age American households from the mid-1980s through 2018.

Our analysis provides insights into the economic well-being of American households and how it has changed in recent decades. First, after about two decades of consistent growth in both $\hat{C}$ and actual consumption (C), both measures stagnated in the early 2000s. At the same time, American households' financial fragility increased in the sense that C exceeded $\hat{C}$. Second, when the Great Recession hit, both C and $\hat{C}$ declined, but the drop in $\hat{C}$ was much larger than the fall in C. Third, because C tends to return to $\hat{C}$ over time, the large negative gap between $\hat{C}$ and C in the aftermath of the Great Recession is broadly consistent with the historically slow recovery of household demand during this period. Fourth, the downward trend in real interest rates over time, and the effect of that decline on projected real asset returns, has several important effects on $\hat{C}$. Declining real rates reduce the value of savings out of current and future labor income, but they also cushion homeowners' sustainable consumption from a collapse in home values and significantly increase the importance of Social Security in supporting $\hat{C}$. Identifying these effects adds an important dimension to debates about both monetary policy and fiscal policy in periods of economic slowdown.

Calculating and evaluating sustainable consumption over households' lifetimes also provides insights into the well-being of American households in retirement. Our data show that average consumption declines significantly in households' later working years and during retirement. Our methodology, which integrates lifetime income, wealth, and Social Security benefits can help identify the reasons behind this decline and inform debates about effective policy for supporting households' well-being in retirement.

Finally, we recognize that the sustainable consumption concept developed in this paper requires assumptions about future income growth, the time allocation of future consumption, and future returns to financial and housing assets, and other factors. Certainly, these projections involve a degree of uncertainty, and further research can likely help refine the analysis and explore the impact of alternative assumptions. However, the choices we made enable us to draw informed conclusions about how household economic well-being has evolved over time. Importantly, our approach is comprehensive and combines information about multiple household resources in studying households' financial sustainability. In doing so, it provides a more holistic picture of household well-being than an analysis that focuses on the impact of only one or two dimensions of households' lifetime resources.

References

Aguiar, M., and M. Bils (2015). "Has Consumption Inequality Mirrored Income Inequality?" *American Economic Review* 105: 2725–2756.

Aguiar, M., and E. Hurst (2013). "Deconstructing Life-cycle Expenditure." *Journal of Political Economy* 121(3): 437–492.

Altig, D., L.J. Kotlikoff, and V.Y. Ye (2022). "How Much Lifetime Social Security Benefits Are Americans Leaving on the Table?" National Bureau of Economic Research Working Paper 30765.

Attanasio, O., and L. Pistaferri (2014). "Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure." *American Economic Review* 104(5): 122–126.

Auerbach, A.J., L.J. Kotlikoff, and D. Koehler (2023). "U.S. Inequality and Fiscal Progressivity: An Intergenerational Accounting." *Journal of Political Economy* 131(5): 1249–1291.

Blanchard, O. (2022). *Fiscal Policy under Low Interest Rates*. Cambridge, MA: MIT Press.

Browning, M., M. Gørtz, and S. Leth-Petersen (2013). "Housing Wealth and Consumption: A Micro Panel Study." *Economic Journal* 123(May): 401–428.

Butrica, B.A., H.M. Iams, K.E. Smith, and E.J. Toder (2009). "The Disappearing Defined Benefit Pension and Its Potential Impact on the Retirement Incomes of Baby Boomers." *Social Security Bulletin* 69(3): 1–27.

Carroll, C.D., T.F. Crossley, and J. Sabelhaus. (2015). "Introduction to *Improving the Measurement of Consumer Expenditures*." In *Improving the Measurement of Consumer Expenditures* (Vol. 74), ed. C.D. Carroll, T.F. Crossley, and J. Sabelhaus, 1–20. Chicago: University of Chicago Press.

Carroll, C.D., and L.H. Summers (1991). "Consumption Growth Parallels Income Growth: Some New Evidence." In *National Saving and Economic Performance,* ed. B.D. Bernheim and J.D. Shoven, 305–348. Chicago: University of Chicago Press.

Catherine, S., M. Miller, and N. Sarin (2022). "Social Security and Trends in Wealth Inequality." Working paper (SSRN-id3546668).

Chen, A., and A.H. Munnell (2021). "Do Retirees Want Constant, Increasing, or Decreasing Consumption?" Center for Retirement Research Working Paper 2021-21.

Chetty, R., D. Grusky, M. Heil, N. Hendren, R. Manduca, and J. Narang (2017). "The Fading American Dream: Trends in Absolute Income Mobility Since 1940." *Science* 356(April 28): 398–406.

Congressional Budget Office (2021). "The Distribution of Household Income, 2018."

Cooper, D.H. (2010). "Imputing Household Spending in the Panel Study of Income Dynamics: A Comparison of Approaches." Federal Reserve Bank of Boston Research Department Working Paper No. 10-12.

Cooper, D.H., B.F. Lutz, and M.G. Palumbo (2015). "The Role of Taxes in Mitigating Income Inequality across the US States." *National Tax Journal* 68(4): 943–974.

Crawley, E., and A. Kuchler (2023). "Consumption Heterogeneity: Micro Drivers and Macro Implications." *American Economic Journal: Macroeconomics* 15(1): 314–341.

Cynamon, B.Z., and S.M. Fazzari (2017). "Household Income, Demand, and Saving: Deriving Macro Data with Micro Data Concepts." *Review of Income and Wealth* 63(1): 53–69.

Dynan, K., J. Skinner, and S.P. Zeldes (2004): "Do the Rich Save More?" *Journal of Political Economy* 112(2): 397–444.

Fisher, J.D., D.S. Johnson, and T.M. Smeeding (2015). "Inequality of Income and Consumption in the U.S.: Measuring the Trends in Inequality from 1984 to 2011 for the Same Individuals." *Review of Income and Wealth* 61: 630–650.

Fisher, J.D., D.S. Johnson, T.M. Smeeding, and J.P. Thompson (2021). "Inequality in 3-D: Income, Consumption, and Wealth." *Review of Income and Wealth* 68(1): 16–42.

Feiveson, L., and J. Sabelhaus (2019). "Lifecycle Patterns of Saving and Wealth Accumulation." Finance and Economics Discussion Series 2019-010. Board of Governors of the Federal Reserve System. https://doi.org/10.17016/FEDS.2019.010.

Halikias, D., and R.V. Reeves (2016). "How Many People Are Better Off Than Their Parents? It Depends How You Cut the Data." Social Mobility Memos, Brookings Institution, August 10. https://www.brookings.edu/blog/social-mobility-memos/2016/08/10/how-many-people-are-better-off-than-their-parents-depends-on-how-you-cut-the-data/

McFall, B.H. (2011). "Crash and Wait? The Impact of the Great Recession on the Retirement Plans of Older Americans." *American Economic Review: Papers and Proceedings* 101(3): 40–44.

Meyer, B.D., and J.X. Sullivan (2017). "Consumption and Income Inequality in the U.S. since the 1960s." National Bureau of Economic Research Working Paper 23655. http://www.nber.org/papers/w23655.pdf

Meyer, B.D., C. Murphy, and J.X. Sullivan (2022). "Changes in the Distribution of Economics Well-being during the COVID-19 Pandemic: Evidence from Nationally Representative Consumption Data." University of Chicago Becker Friedman Institute Working Paper 2022-46. https://bfi.uchicago.edu/wp-content/uploads/2022/03/BFI_WP_2022-46.pdf

Passero, W., T.I. Garner, and C. McCully (2015). "Understanding the Relationship: CE Survey and PCE." In *Improving the Measurement of Consumer Expenditures* (Vol. 74), ed. C.D. Carroll, T.F. Crossley, and J. Sabelhaus, 181–203. Chicago: University of Chicago Press.

Petev, I., L. Pistaferri, and I.S. Eksten (2011). "An Analysis of Trends, Perceptions, and Distributional Effects in Consumption." In *The Great Recession*, ed. D.B. Grusky, B. Western, and C. Wimer, 161–195. New York: Russell Sage Foundation.

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

Sabelhaus, J. (2010). "The Joint Distribution of Consumption and Income." University of Maryland Working Paper.

Sabelhaus, John, and Alice Henriques Volz (2022). "Social Security Wealth, Inequality, and Lifecycle Saving." In *Measuring and Understanding the Distribution and Intra/Inter-Generational Mobility of Income and Wealth*, ed. Raj Chetty, John N. Friedman, Janet C. Gornick, Barry Johnson, and Arthur Kennickell, 249–286. Chicago: University of Chicago Press.

Scholz, J.K., A. Seshadri, and S. Khitatrakun (2006). "Are Americans Saving 'Optimally' for Retirement. *Journal of Political Economy* 114(4): 607–643.

Skinner, J. (2007). "Are You Sure You're Saving Enough for Retirement? National Bureau of Economic Research Working Paper 12981.

Summers, L.H. (2014). "U.S. Economic Prospects: Secular Stagnation, Hysteresis, and the Zero Lower Bound." *Business Economics* 49(2): 65–73.

Wolff, E.N. (2014). "Household Wealth Trends in the United States, 1983–2010." *Oxford Review of Economic Policy* 30(1): 21–43.

Wolff, E.N., A. Zacharias, and A. Caner (2005). "Household Wealth, Public Consumption, and Economic Well-being in the United States." *Cambridge Journal of Economics* 29(6): 1073–1090,

Ziliak, J.P. (1998). "Does the Choice of Consumption Measure Matter? An Application to the Permanent-Income Hypothesis." *Journal of Monetary Economics* 41(1): 201–216.

Appendices

A. Alternative Debt Interest Rate Assumption

The results presented in the main text are derived under the assumption that the interest rate on household liabilities (r^L) is the same as the weighted average return on financial assets (r). Here, we consider the robustness of our results to this assumption by instead allowing the interest rate on debt to be higher than the interest rate on assets given information from the Panel Study of Income Dynamics (PSID).

Starting with 2000, we can estimate interest expenses from the methods described in Appendix D and compute a household-specific average interest rate on a household's outstanding debt. We then calculate the spread (premium) between this rate and the household's average return on financial assets. This spread averages 1.3 percentage points across households in our sample.[32] The spread varies across wealth groups (as defined in Appendix H) from an average of 2.5 percentage points for the lowest wealth group to –0.6 percentage point for the top wealth group. The average spread also declines over time, falling from 2.0 percentage points in 2000 to 0.7 percentage point in 2018. For years prior to 2000, we set the debt interest rate premium at its 2000 median by wealth group.

While incorporating endogenous timing of debt payoffs would greatly complicate our sustainable-consumption calculations, we can check the robustness of our results under the assumption that household debt is repaid at the end of a household's life by the household's housing or financial assets. In this context, debt appears similar to owner-occupied housing in the $\widehat{C}_t$ calculation but with a negative sign and the average interest rate on debt (r^L) in place of the return on housing assets. The alternative definition of sustainable consumption becomes:

$$\hat{C}_t^{Alt} = (V_t^C)^{-1}\left[Assets_t^{Fin} - D_t\left(\frac{1+r^L}{1+r}\right)^{NM-t} + H_t\left(\frac{1+r^H}{1+r}\right)^{NM-t} + \ PVSS_t + Y_{1,t}^{NA}V_{1,t}^Y + Y_{2,t}^{NA}V_{2,t}^Y\right].$$

Note that the alternative definition reduces to equation 8 in the text if $r^L = r$.

For the PSID observations that have non-zero liabilities and a positive spread between the interest rate on debt and the return on financial assets, assuming that debt is paid off only at the end of a household's lifetime likely overstates the negative effect of liabilities on $\hat{C}_t^{Alt}$ because most households will likely repay their (more costly) liabilities earlier. Therefore, it is reasonable to assume that sample averages of sustainable consumption with a fully endogenous treatment of the timing of debt repayment will lie between our main results and the estimates under the assumption that debt is not paid off until the end of a household's lifetime ($\hat{C}_t^{Alt}$).

Our results change only marginally when we allow the interest rate on debt to differ from the return on financial assets, as shown in Figure A1. Over our full sample, average real $\hat{C}_t^{Alt}$ is 2.4 percent lower than average $\widehat{C}_t$. The largest difference is 4.3 percent in 2002. Looking only at observations with positive debt, the difference averaged over all years is 2.9 percent, with the largest difference being 5.0 percent in 2000. Differences between the two sustainable consumption measures shrink to almost zero toward the end of our sample due to the very low debt interest rates at that time. Overall, the differences between the two measures are small, and all interpretations presented in the main text would be very similar if we used the alternative assumption about the interest rate on household debt.

[32] We capped this premium at 5.3 percent, the 99th percentile of the difference between household average interest rate on debt and the weighted average return on assets.

Figure A1: Comparison of Mean $\hat{C}$ with Alternative Debt Interest Assumptions

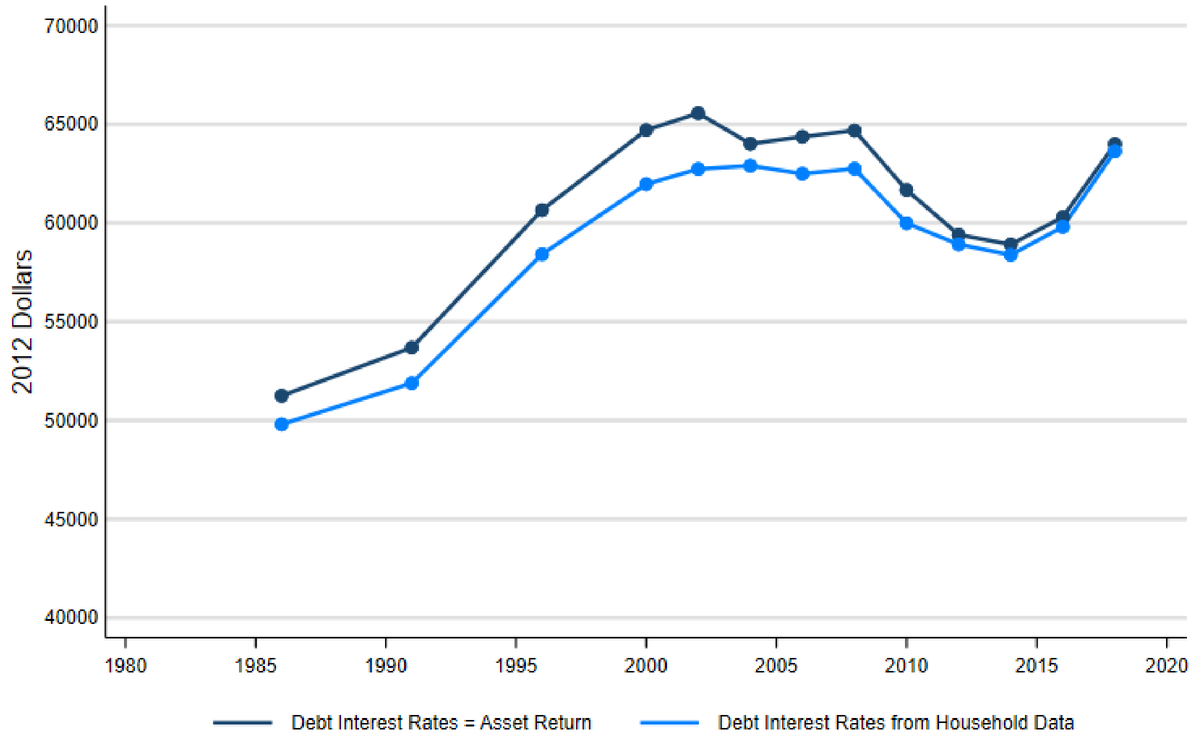

Note: The years represent the midpoint of the expenditure period used in the analysis, as discussed in Section 3.
Source: Authors' calculations using PSID data.

B. Estimating Social Security Benefits

As noted in the main text, the present value of lifetime Social Security benefits (PVSS) is an important component in determining households' sustainable consumption. This section outlines the data we use to calculate PVSS. Importantly, we account for the rules that determine spousal benefits as a function of the income records of both partners.

<u>Basic Social Security Benefit Information Used in Calculations</u>

We obtained a series of basic data needed to calculate Social Security benefits for households over time. The information comes from the Social Security Administration (SSA) unless otherwise noted:

- *Normal retirement age, rounded to the nearest integer.* This is 65 for individuals born before 1941, 66 for birth years 1941 through 1957, and 67 for birth years after 1957.
- *Life expectancy by sex and retirement age, 2019 values.* This is 83 years for men with normal retirement ages of 65 or 66 and 84 years for men with a normal retirement age of 67. For all women in our sample, life expectancy is 86 years. (See https://www.ssa.gov/oact/STATS/table4c6.html for further details.)

- *Yearly cost-of-living adjustments*. Historical adjustment data are available from the SSA through 2021 (see https://www.ssa.gov/oact/cola/colaseries.html). Projections for future years are assumed to equal projected inflation (2 percent per year in baseline calculations).
- *Wage indexing*. We determine Social Security benefits by first indexing individuals' historical earnings to the wage index two years prior to the first year they are eligible for Social Security to account for cost-of-living changes over their lifetimes. We obtain historical average wage index data through 2020 from https://www.ssa.gov/OACT/COLA/AWI.html. We calculate wage indexes for future years based on our projected aggregate real-wage growth assumptions and projected inflation. (These are the same assumptions we use to project individuals' future labor income; see Appendix F).
- *Benefit "bend points."* Social Security payments are computed from a formula that generates decreasing marginal payments as a function of an individual's historical income. The marginal effect of income on their benefits is piecewise linear with two breaks called "bend points." We obtain the historical bend points from (https://www.ssa.gov/oact/cola/bendpoints.html), and as with the wage index data, we project them forward for future retirement years based on our assumptions for real-wage growth and inflation.

Establishing Individuals' Social Security Income Record

The monthly Social Security benefit for an individual, called the "primary insurance amount" (PIA), is calculated from a formula based on the individual's annual income from their 35 years of highest indexed earnings. To determine the PIA for individuals in the PSID, we use their reported historical income data (excluding asset and transfer income) where available. However, we do not observe 35 years of historical earnings data for most individuals in our sample for several reasons:

- Starting with 1997, the PSID survey occurs every second year.
- Some individuals have holes in their earnings records because some households do not participate every survey year.
- Many individuals enter the survey at older ages, and therefore we do not observe income data from the beginning of their likely working years. Similarly, the PSID begins with 1968, and older individuals in the earlier years of our sample will have earnings in their 35-year historical record prior to this date.
- We estimate individuals' PIA at their normal retirement age, which means we need to include future earning years as part of the income record for individuals who have not reached this age by their final observation in our sample (which may be earlier than the final available survey wave if a household drops out of the PSID earlier).

To fill in missing data and generate (at least) 35 years of earnings information for all individuals in our data, we project forward and backward individuals' observed labor income data as needed using income growth factors based on an individual's age and education. The growth factor for individual k from survey years t-1 to t is:

$$(1 + \pi_t)(1 + g_t)(1 + z_{k,a}),$$

where $z_{j,a}$ denotes the individual's education-age-specific component of income growth, g_t is aggregate real-wage growth, and π_t is aggregate inflation. Appendix F discusses how we determine the different components of income growth in detail. In short, we calculate g_t and $z_{j,a}$ for 1968 through 2018 using data from the CPS. For 1952 through 1967 and 2018 through 2020, we use the growth in the aggregate wage index from the SSA less the growth of the personal consumption expenditures (PCE) deflator. For projected income after 2020, we use our baseline real-wage growth assumption of 1.1 percent. For historical years with no other information (1930 through 1951), we assume real-wage growth of 1.0 percent. Note that very little historical income information before 1952 will have any effect on our results because the vast majority of the relevant working years for individuals in our sample are after 1952. For inflation, we use the annual rate of change in the PCE deflator and assume a 2 percent inflation rate (the Federal Reserve target rate) for future years.[33]

Income projections are made step by step from the earnings data we observe in the PSID so that each individual has earnings information for all ages from 18 to the individual's normal retirement age. When there is a hole (year gap) in the individual's reported earnings data (most obviously because the survey is biennial after 1997), we average the forward- and backward-looking income projections to obtain our income estimate for the missing year.

The key measure that determines individuals' Social Security benefits is their "average indexed monthly earnings" (AIME). We generate individuals' AIME as follows. First, we cap all earnings data (observed and estimated) for an individual's working years at the Social Security maximum for the relevant year (see https://www.ssa.gov/OACT/COLA/cbb.html). The capped data are then indexed to the year in which the individual turns 60, using the wage index information discussed earlier. (Working years beyond age 60 are not indexed for the AIME calculation.) Following Social Security eligibility rules, we average an individual's annual income from their 35 years of (capped and indexed) highest earnings. This average is the AIME estimate.[34]

To compute an individual's PIA, we enter their AIME estimate into the Social Security benefits formula that has declining marginal benefits as AIME rises. There are three benefit tiers that generate marginal increases to PIA of 0.90, 0.32, and 0.15, respectively, for each additional dollar of AIME. These break points between marginal benefit rates are the "bend points" discussed earlier in this appendix.

[33] See Altig et al. (2022) for a similar approach to using earnings records to determine Social Security benefits.

[34] An individual needs 40 quarters of positive Social Security earnings to qualify for any benefit. We proxy this requirement by setting an individual's AIME to zero if they have fewer than 10 years of positive income estimates. Note that this constraint applies to income *estimates*, not survey observations. For example, a person in our data aged 25 to 30 with positive income would have a sequence of positive projected income estimates in their record up through age 60, and so they would have a positive estimated AIME.

We determine the PVSS for a single individual *k* with data from time *t* in a straightforward way. We first inflation-adjust their PIA from age 62 to their normal retirement age (NRA). We denote this amount by $PIA_{k,t}^{NRA}$. For years through 2021, we use actual cost-of-living adjustment (COLA) values from the SSA. For 2022 and later years, we apply our assumption for projected aggregate inflation (baseline value of 2 percent). Using these adjusted data, we compute PVSS as follows:

$$PVSS_{k,t} = \left(12 * PIA_{k,t}^{NRA}\right) * \sum_{i=NRA+1}^{N} \frac{(1+\pi_t)^{i-NRA}}{(1+r_t)^{i-Age}},$$

where π_t and r_t are the forward projected inflation rate and the nominal weighted average return on financial assets, respectively, (see Appendix G), and *N* is the individual's demographically determined expected age of death. (Note that we assume individuals work through the calendar year in which they reach normal retirement age, so the Social Security benefit begins in year *NRA* + 1.)

To compute benefits for a couple, we could calculate *PVSS* for each partner in isolation and combine the two results. However, this approach ignores the possibility of spousal benefits having a large impact on the household's overall *PVSS* if the individual earnings of each partner are significantly different. A qualifying partner is entitled to 50 percent of their spouse's benefit if that amount is larger than the benefit they would obtain based on their own earnings record. In addition, a surviving spouse may take their deceased partner's full benefit if that amount is greater than their own benefit.

To account for spousal benefits, we compute the *PVSS* for a couple under all possible spousal-benefit choice and survivorship scenarios and assign the amount that maximizes *PVSS* for the household in our financial sustainability calculations. Table B.1 documents the possibilities considered for both the household head (as defined by the PSID, usually the male for a male/female couple) and their spouse/partner.

In particular, there are three survivorship scenarios based on which partner lives longer. Within each scenario, we choose the maximum *PVSS* from among three cases: (1) both partners take their own benefit; (2) the head takes their own benefit, and the spouse takes the spousal benefit based on the head's record; and (3) the spouse takes their own benefit, and the head takes the spousal benefit based on the spouse's record. If an individual takes their own benefit, payments start the year following the year in which they reach NRA and end at the individual's demographically determined expected death age (DA). In cases where an individual takes spousal benefits, payments begin at either the individual's NRA or the individual's age when the spouse reaches NRA, whichever is later.

To further clarify our approach, Table B.1 reports the benefit taken by each partner in the different survivorship scenarios and benefit choice cases.

Survivorship Scenarios:

A. Single-person household or both partners reach expected death age in the same year
B. Household head reaches expected death age later than spouse
C. Spouse reaches expected death age earlier than household head

Benefit Cases:

1. Single-person household, or both partners take their own benefits
2. Head takes own benefit, and the spouse takes the spousal benefit
3. Spouse takes own benefit, and the head takes the spousal benefit

Each cell in the table highlights a different possible scenario for the benefits claimed by the household head (H) and partner/spouse (S). The percentage in parentheses shows the share of the appropriate benefit received (either 100 percent or 50 percent according to Social Security rules). The phrase "H age when S reaches NRA" means the age of the head in the year in which the spouse/partner reaches their normal retirement age.

Table B.1: Possible Structure of Social Security Benefits

Scenario A: Both partners end benefits in the same year

	Head (H)	**Spouse (S)**
Case 1	H (100%) from NRA-H to DA-H	S (100%) from NRA-S to DA-S
Case 2	H (100%) from NRA-H to DA-H	H (50%) from Max [NRA-S, S age when H reaches NRA] to DA-S
Case 3	S (50%) from Max [NRA-H, H age when S reaches NRA] to DA-H	S (100%) from NRA-S to DA-S

Scenario B: Household head has later expected death year

	Head (H)	**Spouse (S)**
Case 1	H (100%) from NRA-H to DA-S; then Max [H (100%), S (100%)] from DA-S to DA-H	S (100%) from NRA-S to DA-S
Case 2	H (100%) from NRA-H to DA-H	H (50%) from Max [NRA-S, S age when H reaches NRA] to DA-S
Case 3	S (50%) from Max [NRA-H, H age when S reaches NRA] to DA-S; then S (100%) from Max [NRA-H, DA-S] to DA-H	S (100%) from NRA-S to DA-S

Scenario C: Spouse has later expected death year

	Head (H)	**Spouse (S)**
Case 1	H (100%) from NRA-H to DA-H	S (100%) from NRA-S to DA-H; then Max [H (100%), S (100%)] from DA-H to DA-S
Case 2	H (100%) from NRA-H to DA-H	H (50%) from Max [NRA-S, S Age when H reaches NRA] to DA-H; then S (100%) from Max [NRA-S, DA-H] to DA-S
Case 3	S (50%) from Max [NRA-H, H age when S reaches NRA] to DA-H	S (100%) from NRA-S to DA-S

Notes: NRA (normal retirement age); DA (expected death age); H age when S reaches NRA (age of the head in the year in which the spouse/partner reaches their normal retirement age); S age when H reaches NRA (age of the spouse in the year in which the head reaches their normal retirement age).

C. Projections of Household Consumption Growth

In calculating sustainable consumption ($\hat{C}$), we assume future consumption grows at the same rate as household income through the year the household head turns 61. After that year, we assume, real consumption decreases at an annual rate of 1.4 percent until the expected age of death of the longest-surviving partner.

To calculate the projected income growth rate for each household, we compute the age- and education-specific growth rate for each partner's income (as described in Appendix F) along with the projected path for income from transfers and other family members (assumed constant in real terms). Next, all projected income for the household is added together, and we calculate annual household income growth using these household-level income profiles. These results determine households' projected real consumption growth through the year the head turns 61. This estimated growth rate of total household income determines projected consumption growth until the household head is older than 61.

We estimate the annual rate of real consumption decline after age 61 for the household head as follows. Exploiting the panel structure of the data, we compute the annualized growth rate of real consumption within each household between all contiguous survey waves in which they appear. To prevent outliers from having an excessive effect on our results, we drop approximately 1 percent of the observations in which real consumption is greater than $300,000 (2012 dollars). We then calculate the median annual value of the within-household real consumption growth rate for households with a head aged 62 to 90. This rate is –1.4 percent. Regressing log consumption on age for household head aged 62 to 90 produces nearly identical results.

To check the robustness of assuming a common decline in consumption after age 61 across different types of households, we regress our observed data on the log-difference in real household consumption on the age of the household head, indicator variables for the head's education group (less than high school [HS], HS, some college, college or more) and time fixed effects. None of the regressors has predictive power for consumption growth—the lowest p-value for any estimate is 0.17 on a negative coefficient for the college-educated group. In fact, the coefficients on all the education groups above "less than high school" are negative (but not significant), perhaps suggesting that lower coverage of pre-retirement income with Social Security leads to greater consumption decline in later years. Overall, the results suggest that trends in consumption during households' retirement years do not depend on the age or education status of the household head.

We note that other research identifies large declines in consumption during households' retirement years. In particular, Chen and Munnell (2021) find smaller but still highly significant reductions of 0.7 to 0.8 percent per year in nondurables consumption.

D. Estimation of Annualized Outlays and Consumption

Cooper (2010) estimates household-level outlays in the PSID using household income and balance sheet (wealth) data. The approach harnesses the panel structure of the data and yields the cumulative amount of household cash spending between PSID survey waves that include full balance sheet data. For example, the first PSID wave with balance sheet data was conducted in

1984, and the second one was in 1989. The outlay estimate from these two waves therefore covers the five years from 1984 through 1988. As of 1999, when the PSID shifted to biennial surveys, balance sheet data are available every two years. For example, 2013 and 2015 survey waves allow for an estimate for cumulative household outlays in 2013 and 2014.

We follow the approach in Cooper (2010) and calculate household outlays for the available time periods in the PSID. We define "consumption" as outlays less interest expenses. This definition of consumption corresponds to actual household cash flows (along the lines described in Cynamon and Fazzari 2017). This definition does not cover owners' equivalent rent and other imputed service flows that are included in personal consumption expenditures (PCE) in the national accounts. It also does not include expenditures in PCE not paid for by households (for example, medical services paid for by employer-sponsored health insurance or Medicare/Medicaid). We annualize consumption, dividing by two or five as appropriate, and convert it to real terms using the average of annual PCE inflation over the relevant time periods.

Note that we calculate household interest expenses by multiplying beginning-of-period household debt by an appropriate interest rate. If households take on more debt during the period over which we calculate outlays, we will underestimate interest payments. However, this effect would be offset by households that pay off some debt during the period and thus may lower their interest expenses. To calculate mortgage interest expenses, we use the household-specific mortgage interest rate data available in the PSID starting in 1996. For earlier years, we obtain mortgage interest rate information from the Freddie Mac primary mortgage market survey data (from the FRED database, variable MORTGAGE30US). In addition, the PSID reports interest rates on vehicle loans starting in 1998, and we use the 48-month auto loan rate in other years (from the FRED database, variable TERMCBAUTO48NS). There are no data in the PSID for interest rates on other household debt. We use the 48-month auto loan rate for these liabilities.[35] Annual interest rates are averaged over the periods between survey waves.

E. Estimating Permanent Income

Our method for calculating household financial sustainability uses individuals' non-asset income for each wave of the PSID as the basis for estimating their future income path until they reach their normal retirement age. Transitory fluctuations of individual income will introduce volatility into our measures of sustainable household income. We therefore use individuals' (historical) earnings histories in the PSID to generate a proxy for each person's permanent income.

To calculate permanent income, we first generate complete income records for each individual (see appendices B and F). This yields a series of estimated and actual income observations, $Y_{k,t}$, for each individual k in year t back to the individual's first year in the PSID. (Note that even though a lack of balance sheet data prevents us from calculating outlays and hence sustainable consumption before 1984, individuals' income records go back as far as 1969.)

[35] Other debt includes credit card debt, student loans, loans from relatives, and any other household debt that does not have an asset pledged against it.

An individual's estimated permanent income is a weighted average of their past income with adjustments to reflect both real growth and inflation. More specifically, for individual *k,* we denote past income at time $t - i$ (determined as discussed in Appendix F) by $Y_{k,t-i}$. The adjusted value of this income observation that enters the permanent income estimate is then:

$$Y_{k,t}^{i} = \begin{cases} Y_{k,t} \text{ for } i = 0 \\ Y_{k,t-i} \prod_{j=1}^{i} \left(1 + g_{k,t-j}^{Y}\right) \text{ for } i > 0 \end{cases},$$

where $g_{k,t-j}^{Y}$ is the nominal income growth rate determined by an individual's age and education.

We define permanent income for individual *k* at time *t* as the weighted average of past observed and interpolated income up to *t*:

$$Y_{k,t}^{P} = \sum_{i=0}^{S_{k,t}-1} \omega_i Y_{k,t-i}^{i},$$

where $S_{k,t}$ is the number of income realizations in the data for individual *k* through time *t.* Weights are chosen to decay geometrically using a parameter $\delta < 1$ and are normalized to sum to one for any given horizon *N*:

$$\omega_i = \delta^i \left(\frac{1-\delta}{1-\delta^x}\right) \quad i = 0,1, \dots, N-1,$$

where $x = S_{k,t}$. Note that if the household has only a single available income observation, then $S_{k,t} = x = 1$, and $i = 0$, resulting in a weight of one. For relatively large $S_{k,t}$, the weight on the most recent income observation (period *t, i*=0) converges to $1 - \delta$. We set $\delta = 0.7$ for our analysis.

Our estimates of households' permanent (non-asset) income are fairly closely aligned with their actual income. In particular, over our sample, real non-asset income (head plus partner) has a mean of \$79,516 and median of \$59,998 (2012 dollars). The permanent income data are very similar overall, with a mean that is 2.7 percent lower and a median that is 2.1 percent higher. However, the standard deviation of real permanent non-asset income is 20 percent lower than the standard deviation of reported non-asset income, which is consistent with the latter measure containing a lot of transitory fluctuations in individuals' income.

F. Projecting Individuals' Income Growth Profiles

The sustainability calculation requires projecting individuals' labor income through their remaining working years (years following the observed data point in the PSID). We do so using estimated income growth rates from the Annual Social and Economic Supplement (ASEC, released in March) to the Current Population Survey (CPS) stratified by age and educational group. Altig et al. (2022) use a similar approach.

Specifically, consider individual k in a given household who is a years old and is a member of education group j. Every individual is assumed to retire at age M, and we decompose the growth in their annual real-wage and salary income (w) from sample year t to a future year $t+i$ into two parts. The first component depends on projected aggregate real income growth (g), while the second component is age- and education-specific ($z_{j,a}$): [36]

$$w_{k,j,a,t+i} = w_{k,j,a-1,t+i-1}(1+g)\left(1+z_{j,a}\right) \text{for } i = 1, \dots, M-a.$$

We calculate $z_{j,a}$ from the ASEC by first restricting the sample to CPS respondents who are in the labor force and have positive income. We then divide the remaining sample into four education groups: (1) less than high school, (2) high school diploma, (3) associate degree or some college, and (4) bachelor's degree or above. Because the education question in the CPS was changed in 1991, we calculate education-specific wage growth using data from 1991 through 2019. Nominal wages in the data are converted to real wages using the PCE deflator.

Next, we average (using CPS sample weights) log real-wage income (defined as the inflation-adjusted "incwage" variable from the ASEC) across individuals within each education (j) and age group (a) for each year (t) to obtain average log real income $\bar{w}_{j,a,t}$. We define $x_{j,a,t}$ as:

$$x_{j,a,t} = \bar{w}_{j,a,t} - \bar{w}_{j,a-1,t-1}.$$

That is, $x_{j,a,t}$ is average real income growth from ages a-1 to a for education group j from times t-1 to t. To isolate age-specific income growth for a given education group j, we remove time-specific aggregate CPS wage growth, g_t,[37] from $x_{j,a,t}$ and then average the resulting growth rate over time. That is,

$$z_{j,a} = (1/n) \textstyle\sum_t \left(x_{t,j,a} - g_t\right),$$

where n is the number of CPS years in the sample. To further smooth education-age-specific wage growth, we average the year-by-year values of $z_{j,a}$ across three-year age groups.

Under the assumption that age-education-specific wage growth remains the same in the future, real labor income growth for individual k from the current sample year t forward i years is:

$$w_{k,j,a,t+i} = w_{k,j,a,t} \textstyle\prod_{y=1}^{i} (1+\hat{g})(1+z_{j,a+y}) \text{ for } y = 1, \dots, M-a.$$

The constant $\hat{g}$ represents the future real aggregate wage growth. To estimate this parameter, we average g_t from 1980 through 2019. Based on this estimation, we set $\hat{g}$ equal to 1.1 percent.

While the primary use of these wage growth estimates is to project individuals' future labor income growth, the implied wage profiles are interesting in themselves. Figure F.1 shows the estimated profiles of mean real incomes by education group and age. In particular, the profiles

[36] We considered whether the z estimates varied across time by comparing results from several subsamples of the CPS. The age-education components seemed reasonably stable over the sample period we use.

[37] We define g_t as average real income growth in the CPS:

$$g_t = \bar{w}_t - \bar{w}_{t-1},$$

where $\overline{w_t}$ is average log real income for all individuals in the CPS in year t.

begin at age 27 with the initial values set at the mean real non-asset income of household heads aged 25 through 29 in each education group in our PSID sample.

The figure shows a noticeable difference in income levels by educational attainment, as one would expect. In addition, there is a pronounced decline in income during individuals' later working years for all education groups, but the drop is larger, both in levels and percentages, for more highly educated individuals. Since we include only currently working individuals in our CPS calculations, the drop in wages when individuals are in their late 50s and 60s is not due to higher rates of unemployment or lower labor force participation. However, the results could be driven in part by job changes, voluntary or involuntary, that lead to wage reductions or fewer working hours.

Figure F.1: Lifetime Projected Real Median Wage Profiles (by Education Group)

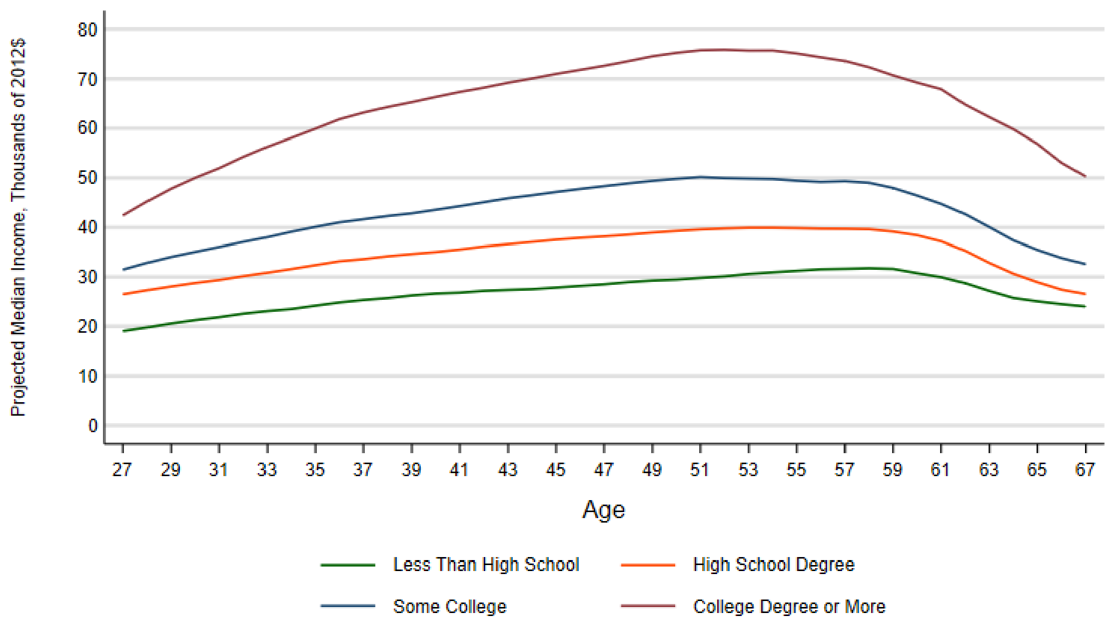

Note: The income profiles begin at age 27 with the initial values set at the mean real income of household heads aged 25 through 29 in each education group in our PSID sample. These initial income levels are based on reported PSID non-asset income from 2014 to 2018, while the remaining data points derive from projecting forward those values across the age distribution using our calculated age-education growth rates from the CPS.

As a validation check of our income projection approach with the CPS data, we compare forecasted future income h years ahead for individuals in the PSID at a given point in time with their actual reported income (for individuals who are still in the sample at that time). Overall, the results are reasonable, although the forecast errors can be large for any given individual.[38] For example, the 10-year forecast from 1988 income for 1998 income was below actual 1998 income by 1.5 percent at the sample median. Conversely, the 10-year forecast from 2002 for 2012 was

[38] The forecast errors have a clear cyclical pattern, with median actual income modestly less than projected income in recession years and the opposite in years with strong macroeconomic growth.

8.2 percent above actual 2012 income at the median. The average of the forecast error for median income over all 10-year periods we can analyze was 3.4 percent (higher forecasted income than actual income).

G. Forecasts of Inflation, Financial Returns, and Housing Returns

Our sustainability calculations require estimates of the projected rate of return on assets over the remainder of a household's lifetime.

Financial Wealth Returns and Inflation

We base projected returns for cash, bonds, and equities on the median 10-year ahead expectations of nominal returns from the Survey of Professional Forecasters compiled by the Federal Reserve Bank of Philadelphia (https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters/data-files). We assume these forecasts apply beyond their stated 10-year horizon since our calculations require projecting asset returns over the remaining years of households' expected lifetimes.

To compute portfolio weights, we divide the households in our sample into four groups based on their total financial wealth: bottom 50 percent, 50th to 75th percentile, 75th to 90th percentile, and top 10 percent. These group definitions are determined using all PSID observations before any sample restrictions are imposed. We calculate the distribution cut points for each year with available household balance sheet data. Households can change positions in the wealth distribution across years.

For each wealth group, we then calculate the average share of assets held in cash (or equivalents), bonds, and equities. Since retirement assets are not allocated to specific asset classes in the PSID, we assume 40 percent of households' retirement assets are held in bonds and 60 percent in equities. Table G.1 shows these estimated portfolio shares for the four wealth groups. Not surprisingly, households with greater wealth have financial wealth portfolios that are more heavily invested in equities.

Table G.1: Portfolio Shares by Wealth Group

Wealth Group	Cash Share	Bond Share	Equity Share
1	90.5%	4.8%	4.6%
2	52.8%	21.3%	25.9%
3	30.2%	27.2%	42.6%
4	14.7%	23.7%	61.6%

Table G.2 shows the forecasted rates of return for financial assets in each consumption period along with the forecasted inflation rate. Projected inflation is based on the Survey of Professional Forecasters median 10-year ahead forecast for the PCE price index. (Note that annual data are averaged across the years included in each period.) The data reflect forecasts as of a given date and therefore reflect information available at that time. Unanticipated changes—for example the acceleration of inflation during and after the COVID-19 pandemic—will change the projections and therefore affect the estimation of sustainable consumption.

Table G.2: 10-year Projections of Financial Returns and Inflation (Averages across households for each time period)

Period	Equity Return	Bond Return	Cash Return	Inflation
1984–88	11.38%	8.88%	6.38%	4.80%
1989–93	10.33%	7.93%	5.35%	3.93%
1994–89	9.05%	6.32%	4.79%	2.94%
1999–00	8.00%	5.75%	4.75%	2.52%
2001–02	7.25%	5.50%	4.13%	2.41%
2003–04	7.88%	5.34%	3.75%	2.12%
2005–06	7.00%	5.00%	3.98%	2.44%
2007–08	7.00%	5.00%	4.25%	2.10%
2009–10	6.75%	4.90%	3.00%	2.15%
2011–12	7.03%	4.44%	2.75%	2.13%
2013–14	6.06%	4.09%	2.45%	2.00%
2015–16	5.41%	3.68%	2.58%	1.99%
2017–18	6.00%	3.78%	2.63%	2.05%

Housing Returns

The Survey of Professional Forecasters does not provide long-term projections for nominal home-price growth. Instead, it provides one-year ahead forecasts, and they begin in 2010, long after the earliest years in our data. We therefore do not use SPF data to project the return on home ownership. In addition, our sample period encompasses times of substantial volatility in home prices, and we assume that households expect more typical (and stable) house-price returns in the post-sample years. We therefore set the real return on owner-occupied homes equal to the long-run growth rate in inflation-adjusted home-sale prices of 1.45 percent. This average return is calculated based on annualized growth of the 10-city Case-Shiller house-price index deflated by the personal consumption expenditure price index from 1988 (when the Case-Shiller data begins) through 2018. This growth rate is almost the same as the annualized real growth rate of the US home-price index from the Federal Housing Administration for our full sample period (1.50 percent, 1984 through 2018). Nominal rates of return on housing equal 1.45 percent plus the projected inflation by period shown in Table G.2.

H. Data Restrictions and Adjustments to Determine Analysis Sample

We employ several restrictions with the available PSID data to ensure a consistent analysis sample and to reduce the influence of outliers or inconsistent survey data on our results.

- Since we are interested in studying households during their main working years after they have completed any full-time education, we delete observations from when the household head is younger than 25. We also drop observations if the household head is older than 61 to avoid including data that reflect early retirement. (Individuals become eligible for Social Security at age 62.)

- We drop observations (households) in the PSID that are tied to the (short-lived) over-sampling in the mid-1990s designed to increase the survey coverage of under-represented groups.
- We drop observations with missing, negative, or bottom-coded household (family) income.
- We delete observations with negative after-tax non-asset income.
- We drop observations for which we cannot estimate household outlays and consumption.
- We also drop observations with negative estimated consumption values that most likely arise from inconsistencies in households' reported balance sheet data across survey years. (For example, if an asset value was substantially under-reported in year 1 and then corrected to a much higher value in year 2, the asset accumulation inferred from the reporting error could lead to a large negative calculated outlay.) Such errors are also likely to affect observations (in the opposite directions) for the household in other sample periods. We therefore delete observations if calculated expenditures in the sample period before or after the negative consumption year is large enough to generate a ratio of consumption to after-tax income that exceeds 1.5.
- Some households also have estimated consumption that appears inconsistent with their income, most likely due to reporting errors. All observations for a household are dropped if the sum of their total consumption for all years they are in the survey are more than three times the sum of the household's after-tax income.
- We cap reported mortgage interest rates at 20 percent per year.
- We winsorize estimated household tax rates (both federal and state) if they fall in the top or bottom 1 percent of the tax rate distribution for each year in the data.