

Permanent-income inequality

BRANT ABBOTT

Economics Department, Queen's University and HCEO

GIOVANNI GALLIPOLI

Vancouver School of Economics, The University of British Columbia, CEPR, HCEO, and RCEA

Through certainty equivalent consumption (CE) measures, we show that dispersion of current earnings, expenditures, and net worth overstate welfare inequality. This is largely due to the unaccounted value of future earnings, which we call human wealth. The latter mitigates permanent-income inequality, though its influence is diminished by the growing importance of assets in lifetime wealth. Average expenditures and CE inequality roughly doubled between 1983 and 2016 and, to weigh these offsetting forces, we decompose aggregate welfare changes into contributions from the level and dispersion of consumption, as well as uncertainty and demographic composition. Rising inequality has offset about 1/4 of the welfare gains from higher consumption, with most of the losses accruing after 2000.

KEYWORDS. Wealth, human capital, permanent income, consumption, inequality.

JEL CLASSIFICATION. D31, E2, E21, I24.

1. INTRODUCTION

Average consumption expenditure in the United States has increased significantly over recent decades; for example, real consumption outlays per capita more than doubled between 1980 and 2020.¹ At the same time, a historically large rise in income and wealth inequality has occurred, and there is growing evidence of a pass through into greater consumption inequality. Establishing how these opposing secular trends have impacted societal well-being, and quantifying their relative contributions, is a nontrivial exercise because the cross-sectional distributions of income, wealth, or consumption do not directly reveal the prevailing patterns of individual welfare. Put differently, it is not obvious to what extent the costs of rising consumption inequality have offset the welfare gains from higher average consumption.

A vast literature examines the relationship between earnings and consumption inequality (among others, [Storesletten, Telmer, and Yaron \(2004\)](#), [Krueger and Perri](#)

Brant Abbott: abbottbrant@gmail.com

Giovanni Gallipoli: gallipol@mail.ubc.ca

We thank Kartick Athreya, Paul Beaudry, Mark Bills, David Green, Mark Huggett, Narayana Kocherlakota, Dirk Krueger, Costas Meghir, Joseph Mullins, Fabrizio Perri, Luigi Pistaferri, Victor Rios-Rull, and Gianluca Violante for feedback on earlier drafts. Davide Alonzo, Amedeus D'Souza, Denis Kojevnikov, and Xianya Zhou provided valuable research assistance. Financial support from the SSHRC is acknowledged. A replication file is posted ([Abbott and Gallipoli \(2022\)](#)).

¹See <https://fred.stlouisfed.org/series/A794RX0Q048SBEA>.

(2006), Blundell, Pistaferri, and Preston (2008), Heathcote, Storesletten, and Violante (2010), Attanasio, Hurst, and Pistaferri (2014), Aguiar and Bils (2015)).² There is also growing recognition that, while current consumption is related to the notion of welfare, it may reflect confounding influences due to credit market frictions, life-cycle variation, heterogeneous risk profiles, and temporary random fluctuations (Attanasio and Pistaferri (2016)). Thus, assessments of economic inequality should be based on welfare metrics that subsume lifelong processes, and account for differences in idiosyncratic uncertainty and demographic shifts. In fact, the point that one should account for lifetime values is a recurring theme in the empirical literature on earnings inequality (see Lillard (1977), Keane and Wolpin (1997), Geweke and Keane (2000), Bowlus and Robin (2004), Guvenen, Kaplan, Song, and Weidner (2017), Curtis, Garn, and Lester (2021)).³

This paper considers two alternative representations of the lifetime welfare of individuals and households, and develops an approach to map these theoretical constructs to data. The first mapping, called the “consumption representation,” delivers the expected present value of lifetime utility by directly employing expenditure data. The second, referred to as the “permanent-income representation,” estimates a theoretically equivalent quantity based on lifetime earnings and net worth. The latter measure is reminiscent of permanent income as defined by Friedman (1957), with the qualification that stochastic discount factors are applied in place of a risk-free discount factor.

Each of the two welfare representations can be recast in terms of certainty equivalent consumption, for which we characterize the cross-sectional distribution at different points in time. The certainty-equivalent measures of welfare provide a transparent way to account for life-cycle variation, and for the heterogeneous burden of uncertainty across households. Perhaps most importantly, the estimates can be used to assess how aggregate welfare has changed, given the underlying growth in both the level and dispersion of consumption.

There are advantages and limitations associated with each measure. The consumption representation is more direct; however, the permanent-income representation allows one to decompose lifetime wealth into its human and financial components, which aids in understanding the long-term trends. To the extent that future transfers, for example, bequests, enter utility but are not captured by current expenditures; it is valuable to complement the welfare analysis with information about income and wealth. On the other hand, the consumption representation will perform better in situations where income and taxes are not measured precisely. However, either representation offers an advantage over welfare measures based on current expenditures and earnings if the latter contain measurement error that can be averaged out through present value calculations.

²For related literature on earnings, see references in Guvenen, Karahan, Ozkan, and Song (2021) and De Nardi, Fella, and Paz-Pardo (2016), De Nardi, Fella, and Paz-Pardo (2020), who consider rich nonsymmetric income processes. On wealth inequality and its evolution, see Kopczuk and Saez (2004), Piketty and Saez (2006), Saez and Zucman (2016), Bricker, Henriques, Krimmel, and Sabelhaus (2016), Kaymak and Poschke (2016), and Kuhn and Ríos-Rull (2016). See also Athreya, Ionescu, and Neelakantan (2021), who document covariation of human capital and financial portfolios.

³Sanders and Taber (2012) and Abbott and Gallipoli (2020) review this extensive literature.

The empirical framework is based on a standard life-cycle incomplete-markets model, from which we derive each distinct measure of certainty equivalent consumption. Both approaches indicate that (i) welfare inequality (i.e., certainty equivalent consumption inequality) is considerably lower than income or wealth inequality, but also that (ii) welfare inequality increased substantially since the early 1980s, albeit less than wealth inequality. A breakdown based on the permanent-income representation shows that (iii) human wealth mitigates inequality and accounts for the lower welfare dispersion, but also that (iv) this mitigating influence has waned over time as net worth has become a larger contributor to lifetime wealth and permanent income.

The aggregate implications of the changing level and dispersion of consumption are quantified through the lens of a utilitarian welfare function. Changes in aggregate expenditure, uninsurable uncertainty, inequality, and demographic composition can each independently affect aggregate welfare; crucially, the estimates of certainty equivalent consumption allow us to map each of these moving parts into contributions measured in consumption equivalent units.⁴ We find that, between 1983 and 2016, aggregate consumption increased by 88%; however, using isoelastic utility and a baseline CRRA coefficient of 2, the consumption equivalent value of the average welfare change is only 3/4 as much. This discrepancy is entirely due to the costs of rising inequality of certainty equivalent consumption. The losses are especially conspicuous as they occur in the face of stable or marginally improved cross-sectional insurance. Inequality patterns over time are similar when we consider alternative curvatures of the utility function and account for the value of labor supply.

The initial part of the paper outlines the model and derives the estimation approaches, illustrating how they can be leveraged in small survey samples that, unlike administrative data, have the advantage of reporting consumption expenditures and a rich set of individual characteristics. An important consideration, irrespective of data sources, is that only one realization of the future state of the world is observed for each person and time period. Because we cannot run an individual's life multiple times, we do not observe the entire distribution of their possible future outcomes, on which certainty equivalent consumption depends. This fundamental data limitation is addressed by estimating the distribution of possible outcomes from those observed for individuals who are *ex ante* similar in terms of current consumption, income, and various other characteristics. This approach works under the assumption that individuals who are *ex ante* equivalent, in terms of a broad enough set of variables, face the same distribution of *ex post* outcomes.

Both the consumption and permanent-income representations feature state-dependent stochastic discounting to account for the ease with which resources can be shifted across time periods, and for uncertainty about future earnings and consump-

⁴The welfare decomposition builds on original methods developed in Benabou (1996) and Floden (2001). See also Abbott, Gallipoli, Meghir, and Violante (2019).

tion.⁵ Being constrained by a credit limit, or facing a great deal of risk, reduces a household's valuation of their expected lifetime resources.

Sections 5 and 6 summarize the trends in certainty equivalent consumption inequality and explore the implications of these trends for aggregate and cross-sectional welfare. Finally, in Section 7 we overview numerous robustness checks and test the sensitivity of our results.

2. LIFETIME WEALTH AND WELFARE

Mappings from data observations to estimates of certainty equivalent consumption are obtained using a standard life-cycle model. First, we derive a risk-adjusted version of the lifetime budget constraint with state-dependent discounting (stochastic discount factors replace the risk-free rate). Then we recover the elements of this equation from data and employ them to form estimates of indirect utilities (i.e., value functions). Such estimates are made comparable across individuals by expressing them in terms of certainty equivalent consumption.

The model features a general endowment process that depends on the state of the world s_j at age j , denoted as $y(s_j)$. Assuming a maximum life-cycle length of J and a single risk-free asset a with return r , the household's recursive optimization problem is

$$V_j(a_j, s^j) = \max_{\{c_j, a_{j+1}\}} \{u(c_j) + \beta E_{s_{j+1}|s_j}[V_{j+1}(a_{j+1}, s^{j+1})]\} \quad (1)$$

$$\text{s.t. } c_j + a_{j+1}/(1+r) = a_j + y(s_j) \quad \text{and} \quad a_{j+1} \geq \underline{a}_j(s_j),$$

where value functions depend on histories $s^j = \{s_0, s_1, \dots, s_j\}$. Households can borrow up to an amount $\underline{a}_j(s^j)$, which they can repay with certainty given their age and history. For ease of exposition, we temporarily employ a constant discount parameter β ; however, as will be detailed, the main results allow for mortality risk and age-varying β_j . To obtain closed-form solutions for lifetime utility, we posit isoelastic utility $u(c_j) = c_j^{1-\gamma}/(1-\gamma)$. The household's lifetime budget constraint is

$$\sum_{k=j}^J \left(\frac{1}{1+r}\right)^k E_j[c_k] = a_j + \sum_{k=j}^J \left(\frac{1}{1+r}\right)^k E_k[y(s_k)]. \quad (2)$$

This accounting identity connects expected lifetime wealth and consumption. It does not, however, describe how uncertainty affects the valuation of resources. A mean-preserving spread in the distribution of consumption outcomes would not change this equation; yet, such a change would unambiguously change household welfare. This observation highlights the need for an approach that recasts the budget identity in terms of valuations based on state-dependent stochastic discount factors (SDFs). By applying stochastic discount factors to the sequence of expected consumption and income

⁵Discounting future earnings at the risk-free rate overstates the value of human capital (Huggett and Kaplan (2016)). Mechanically discounting income flows ignores state-dependent valuations of earnings and other forms of heterogeneous discounting (Gabaix and Laibson (2017)).

realizations, we establish an equivalence between risk-adjusted present values of optimal life-cycle consumption and lifetime wealth. The latter result is derived in two steps (see Appendix A for a formal derivation where we also show how this approach can accommodate cross-sectional heterogeneity in rates of return across households). First, we multiply each potential realization of the age $j + 1$ budget constraint in problem (1) by the corresponding stochastic discount factor $\pi(s_{j+1}|s^j)\beta^{\frac{u'(c_{j+1})}{u'(c_j)}}$, where $\pi(s_{j+1}|s^j)$ denotes the conditional transition probability. Next, we sum across these probability-weighted realizations of s_{j+1} to define an expected constraint at each age; these age-specific constraints are sequentially added up and, after imposing standard intertemporal optimality, the following lifetime relationship is obtained:⁶

$$\sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} c_k \right] = \sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} y(s^k) \right] + a_j. \tag{3}$$

The expression on the left is analogous to a Lucas (1978) asset pricing relationship and describes the value of an asset yielding an uncertain flow of consumption to the household. In our model, this value equates to that of an asset that pays the household an uncertain stream of endowment income plus their current net worth. In this sense, the asset value of one’s consumption equals the asset value of their human capital (their risk-adjusted human wealth) plus their financial wealth.⁷ Using $\theta_j^h = \sum_{k=j}^J E_j [\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} y(s^k)]$ to denote human wealth and $\theta_j^c = \sum_{k=j}^J E_j [\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} c_k]$ to denote the asset value of future consumption, one can write (3) as $\theta_j^c = \theta_j^h + a_j$.

2.1 From lifetime wealth to welfare

Intuitive measures of welfare can be derived using the intertemporal budget identity. With CRRA preferences $u(c_j) = \frac{1}{1-\gamma} u'(c_j) \times c_j$, so the value function in (1) can be written as

$$V_j(a_j, s^j) = \frac{u'(c_j)}{1-\gamma} \theta_j^c \tag{4}$$

$$= \frac{u'(c_j)}{1-\gamma} (\theta_j^h + a_j). \tag{5}$$

The first equality follows from the isoelasticity property noted above and by applying the definition of θ_j^c . The second equality follows from the intertemporal constraint (3). Equations (4) and (5) are alternative representations of lifetime values; they are what we refer to as, respectively, the consumption and permanent-income representation. Below, we suggest procedures to separately estimate them.

⁶Both our assumption of natural borrowing limits and its opposite extreme of no borrowing ($a = 0$) deliver (3) exactly. In other cases, (3) holds approximately. However, any approximation error would only affect the permanent-income representation, not the consumption representation.

⁷See Huggett and Kaplan (2016) for a related derivation under general conditions.

Two considerations must be made when using the $V_j(a_j, s_j)$ representations in (4) and (5) for welfare analysis: first, utility comparisons are not cardinal; second, it is hard to make welfare comparisons between individuals of different ages. To remedy these limitations, we compute the certainty equivalent consumption values, $\bar{c}(a_j, s^j)$, which solve

$$\sum_{k=j}^J \beta^{k-j} \frac{(\bar{c}(a_j, s^j))^{1-\gamma}}{1-\gamma} = V_j(a_j, s^j). \quad (6)$$

The certainty equivalent (CE) estimates are expressed in expenditure terms, which facilitates comparisons across individuals, households, and age groups. We use \bar{c}_{ij}^c and \bar{c}_{ij}^p to denote CE consumption based on (4) and (5), respectively.⁸

3. EMPIRICAL ANALYSIS

3.1 Identification and estimation

Equations (4) and (5) make clear that the main measurement problem involves identification of lifetime consumption and earning values at each age j (θ_j^c and θ_j^h). The latter are stochastically discounted present values so that the methods to recover them are analogous. For brevity, we present only the steps for θ_j^h estimation and note that θ_j^c estimation follows a similar procedure with y_k replaced by c_k . The most pertinent details of identification and estimation are provided here, while a detailed discussion is in Appendix B.

We begin by observing that the value of θ_{ij}^h can be expressed in recursive form as

$$\theta_{ij}^h = y_{ij} + E_{ij} \left[\beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \theta_{ij+1}^h \right], \quad (7)$$

where subscript i is now introduced to index an individual observation. The age-specific discount parameters β_j capture age-varying mortality risk and are taken as given.

The expectation in (7) must be discussed before proceeding any further. Only one observation of the age $j + 1$ outcome is observed for any individual at age j . Therefore, one has to assume that the distribution of possible age $j + 1$ outcomes for an individual can be estimated using the outcomes observed for similar individuals at age j . The notion of similarity is based on a vector \mathbf{z} of characteristics that includes both individual and aggregate information. The idiosyncratic expectation E_{ij} can then be mapped into an age-specific conditional operator such that the following relationship holds:

$$E_{ij} \left[\beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \theta_{ij+1}^h \right] = E_j \left[\beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \theta_{ij+1}^h | \mathbf{z} = \mathbf{z}_{ij} \right]. \quad (8)$$

In practice, the condition holds if the vector \mathbf{z}_{ij} is sufficient to span the current information set of an individual i at age j . Given (8), human wealth can be written recursively,

$$\theta_j^h(\mathbf{z}) = y_{ij} + \mathbb{E}_j \left[\beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \theta_{j+1}^h(\mathbf{z}') | \mathbf{z} \right]. \quad (9)$$

⁸The superscript p stands for “permanent income.”

Here, \mathbf{z}' is the age $j + 1$ realization of \mathbf{z} and, because current earnings y_{ij} are in the conditioning vector $\mathbf{z} = \mathbf{z}_{ij}$, they do not imply any heterogeneity beyond \mathbf{z} itself.

To flexibly estimate (9), we employ the Nadaraya–Watson method. Denote as $\xi_{ij}(\mathbf{z})$ the weighting function (kernel), which depends on how similar the evaluation point \mathbf{z} is to the value of \mathbf{z}_{ij} observed for individual i at age j . Then the empirical counterpart to (9) is a recursive weighted summation:

$$\hat{\theta}_j^h(\mathbf{z}) = y_{ij} + \sum_{i=1}^{N_j} \beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \hat{\theta}_{j+1}^h(\mathbf{z}_{ij+1}) \xi_{ij}(\mathbf{z}), \tag{10}$$

where the $\hat{\theta}_j^h(\mathbf{z})$ on the left-hand side is the unknown object we need to solve for. Some remarks can be made about (10). First, $\hat{\theta}_j^h(\mathbf{z})$ delivers an estimate of the whole function, rather than a set of point estimates; this facilitates evaluation of human wealth at data points that are not included in the estimation. Second, while the function $\hat{\theta}_j^h(\mathbf{z})$ depends on estimates of $\hat{\theta}_{j+1}^h(\mathbf{z})$, and thus is not yet determined, estimation of the $\hat{\theta}_j^h(\mathbf{z})$ on the left-hand side of (10) only requires that $\hat{\theta}_{j+1}^h(\mathbf{z})$ be known at the points $\{\mathbf{z}_{ij+1}\}$ where data are observed (as opposed to being fully identified).

The last point suggests an intuitive two-step approach: in a first step, construct estimates of $\hat{\theta}_j^h(\mathbf{z})$ at the observed data points for each j ; then use the latter to recover estimates of the function $\hat{\theta}_j^h(\mathbf{z})$ in its entirety. It should be noted that joint estimation across all ages is desirable in the first step because it improves efficiency relative to a recursive iterative procedure. The joint estimator is constructed by stacking all $\hat{\theta}_j^h(\mathbf{z}_{ij})$ point evaluations into a vector $\tilde{\Theta}_j$, and then stacking these age-specific vectors into a larger vector $\tilde{\Theta}$. Point observations of y_{ij} are similarly stacked into a vector \tilde{Y} so that the system of equations corresponding to all point evaluations of (10) is represented in compact form as $\tilde{\Theta} = \tilde{Y} + \Gamma \tilde{\Theta}$, where the elements of the matrix Γ are the $\beta_j \frac{u'(c_{ij+1})}{u'(c_{ij})} \xi_{ij}(\mathbf{z}_{ij})$ terms from the summation in (10).⁹ The nice thing about these arrangements is that point estimates of human wealth become straightforward to solve for in a single step using $\tilde{\Theta} = (I - \Gamma)^{-1} \tilde{Y}$. Once the point estimates $\tilde{\Theta}$ are attained, an age specific function $\hat{\theta}_j^h(\mathbf{z})$ can be estimated as in (10). A detailed description of this procedure, including formal nonparametric identification arguments and how to include biennial data periods, is in Appendix B.

It is worth emphasizing that, while workers cannot forecast the exact path of their future earnings, the information embodied in the evolution of \mathbf{z} helps shape their expectations. For example, \mathbf{z} may contain information about education effects: because the present value for a young person depends on yet-to-be realized returns to education, the estimator imposes that expectations about the distribution of such returns are consistent with what is later observed. This approach also allows for extrapolation, in the sense that present values for younger cohorts in late sample periods are based on the expectation that life-cycle profiles will exhibit growth similar to that of previous cohorts.

⁹The Γ matrix consists of age specific Γ_j submatrices such that $\tilde{\Theta}_j = \tilde{Y}_j + \Gamma_j \tilde{\Theta}_{j+1}$.

3.2 Data

Estimation requires panel data on earnings, consumption, wealth, and a sufficient set of conditioning characteristics. Unlike other data sources, the PSID satisfies these requirements allowing estimation to be carried out at the level of individual expenditures and income. To optimize panel length, we follow [Attanasio and Pistaferri \(2014\)](#) (AP) and estimate total expenditures as a function of observed food outlays, relative prices, and other variables using the richer consumption data in the post-1997 waves. Then we invert this demand model to impute total household expenditure from food expenditure, which is observed in most sample periods. For consistency, we use the imputed measure of consumption in our analysis even after 1997. For married couples, we attribute half of the predicted expenditure to each spouse to generate an individual level measure of expenditure. As AP show, this approach delivers log expenditures whose variance closely matches empirical observations.

A concern is that expenditure *levels* in survey data are lower than the average expenditures measured in aggregate data.¹⁰ To account for this discrepancy total consumption is scaled so that sample averages match expenditure-per-capita in NIPA data. The rescaling is consequential because the aggregate expenditure level is an important determinant of changes in welfare over time. However, scaling all observations by the same factor may generate inaccurate measures of expenditure because, as documented in [Aguar and Bils \(2015\)](#) and [Abbott and Brace \(2020\)](#), affluent households underreport their expenditures much more than poorer households. A pragmatic correction to account for this phenomenon is to scale consumption exponentially so that household expenditures are $c_{ij} = \tilde{c}_{ij}^{\alpha_t}$, where \tilde{c}_{ij} is the AP predicted expenditure and the year- t parameter α_t solves $E[\tilde{c}_{ij}^{\alpha_t} | t] = c_t^{\text{nipa}}$.¹¹ Real consumption expenditure per capita, c_t^{nipa} , is taken from the FRED database.

Data on household wealth is observed every five years from 1984 through to the 1999 PSID wave, and biennially thereafter until 2016. Income and consumption data are retrospective and we treat wealth the same way.¹² Our measure of earnings includes usual labor income as well as social security payments for retired households. This ensures that social security entitlements are directly accounted for in welfare calculations. Earnings are converted to a net-of-tax measure by subtracting the tax liabilities using NBER TAXSIM. The sample used for estimation of $\theta_j^h(\mathbf{z})$ and $\theta_j^c(\mathbf{z})$ includes 179,936 individual

¹⁰See [Parker, Vissing-Jorgensen, and Ziebarth \(2009\)](#), for example.

¹¹Relative to linear rescaling, the exponential correction delivers inequality trends in the consumption representation that are marginally closer to current consumption inequality. Thus, linear rescaling would make the use of present value calculations even more compelling. Second-order effects occur also in the permanent-income representation through stochastic discount factors. Overall, the choice of rescaling has only limited impacts on either representation, although we consider exponential scaling preferable as it is consistent with underreporting by the affluent. Appendix D provides the counterpart of Table 1 with linear scaling.

¹²One exception is the 1989 wave, when food expenditure was not surveyed and total expenditure cannot be imputed. Given the sparsity of wealth observations, we take steps to preserve the 1989 data point and combine the 1989 wealth records of the PSID with income and consumption from the 1990 wave. Since retrospective income and consumption data describe experiences in 1989, we set the year of these constructed observations to be 1989 as well.

observations spread over 32 years.¹³ Our main results employ those estimated functions, along with wealth data, for which we have 57,533 household level observations spread over 13 sample years. In Appendix D, we provide a more comprehensive table of summary statistics.

The conditioning vector \mathbf{z} includes the following variables: cohort (birth year), gender, education (less than high school, high school, some college, college), current earnings, current consumption, and aggregate GDP per capita. Gender and education are treated as discrete, while the remaining variables are modeled as continuous. Crucially, the nonparametric estimator implicitly allows for arbitrary interactions and higher-order terms for any of these variables.

There is a different valuation function for each age j . The estimation procedure described in Section 3.1 is carried out on individual-level data, although our main analysis is at the household level. For couples, this means that estimated values for the head and spouse are summed before a household level certainty equivalent consumption is computed.¹⁴ Data on net worth are observed at the household level. Unlike consumption, there is no need to divide these resources between spouses as wealth is not used in the estimation of θ_{ij}^c nor θ_{ij}^h . Rather, net worth is added to other household-level variables when constructing the right-hand side of the permanent-income representation of lifetime utility in equation (5).

By design, the two welfare representations differ in their treatment of housing as they leverage alternative data sources and formats. While the consumption representation converts the flow of housing services (rent equivalents, other expenditures) from a yearly frequency to a present value, the permanent-income representation directly accounts for the stock value of housing assets. If housing wealth measures diverge from the present value of housing service flows, this will be reflected in the gap between the two welfare measures.

3.3 Mortality and utility parameters

The estimation described above takes the utility parameters γ and $\{\beta_j\}$ as given. Risk aversion is set to $\gamma = 2$, which is a common choice in the literature. The $\{\beta_j\}$ parameters reflect both time discounting and mortality risk. We form these age-specific discount factors by multiplying estimated mortality rates, based on Pijoan-Mas and Ríos-Rull (2014), by 0.95. Robustness to these parameter assumptions, particularly the utility curvature γ , is examined below. We set maximum age $J = 86$ for the practical reason that sample sizes fall off substantially at that point.

4. ESTIMATES OF WEALTH OVER THE LIFE CYCLE

Estimates of θ_{ij}^h and θ_{ij}^c , along with observed net worth, are the primary contributors to the welfare metrics we consider. To illustrate their behavior, we provide several snapshots illustrating how they vary over the life cycle, across cohorts, and with certain dimensions of heterogeneity. Figure 1 begins by presenting age profiles of average human

¹³Food consumption is not reported for 1972, 1987, or 1988, and so observations based on these years are not used.

¹⁴The age of the head is used to compute the sum on the left-hand side of (6).

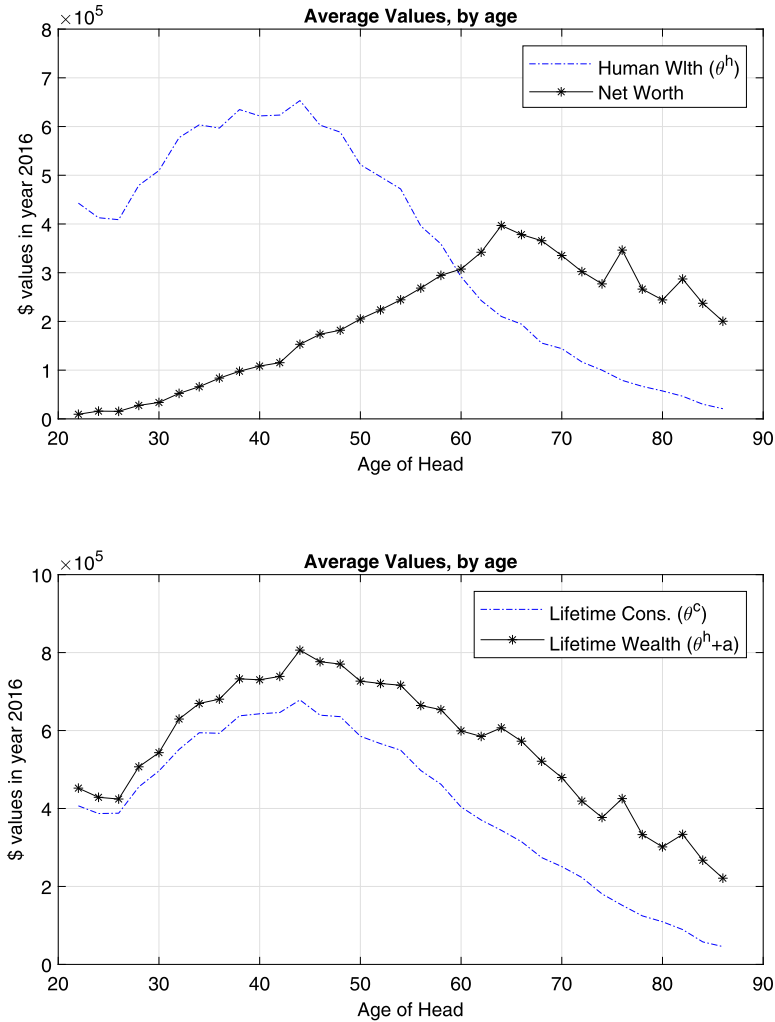


FIGURE 1. This figure illustrates the life-cycle evolution of human wealth (θ_{ij}^h) and net worth in the top panel, and the life-cycle evolution of lifetime consumption (θ_{ij}^c) and lifetime wealth ($\theta_{ij}^h + a$) in the bottom panel. The top panel illustrates how, early in life, the average household is rich in human wealth; the latter is converted into financial wealth as the life cycle unfolds. The bottom panel shows that the sum of human wealth and net worth (lifetime wealth) follows a hump-shape pattern over the life cycle; moreover, lifetime consumption exhibits a very similar shape, although the latter tends to zero toward the end of life. All values are reported in 2016 equivalent dollars.

wealth (θ_{ij}^h) and net worth (a_{ij}) in the upper panel. Observations are grouped into two-year age brackets, for example, 22–23, 24–25, etc., to reduce noise. It is apparent from this figure that the average young household is human wealth rich, and that this wealth is converted to financial wealth as the life-cycle progresses. Two factors explain the rising value of human wealth early in the life cycle: (i) very young households discount their

peak earnings more than households closer to middle age because of the longer time lag, and (ii) younger households experience rapid expenditure growth so that their discount factors tend to be smaller. The bottom panel of Figure 1 combines human wealth and net worth into a single measure of lifetime wealth ($\theta_{ij}^h + a_{ij}$), and plots this alongside the life-cycle profile of average lifetime consumption (θ_{ij}^c). Here, we see that the two measures of lifetime resources exhibit similar hump shapes over the life cycle, but that the θ_{ij}^c profile is generally lower and converging to zero at the end of life. This difference arises because households tend not to exhaust their net worth, and the lifetime consumption measure does not take into account any value associated with, for example, bequests.

Figure 2 illustrates how the life-cycle profiles of different variables have evolved across generations. In our estimates, growth across cohorts is captured by cohort effects, as well as general shifts in the distributions of variables in \mathbf{z} (e.g., earnings and consumption themselves). We plot segments of the life-cycle profiles of four birth-year groups. Increases across cohorts are evident for θ_{ij}^h and θ_{ij}^c (top and middle panels of Figure 2) with higher averages in younger cohorts. For net worth (bottom panel of Figure 2), substantial increases of wealth across cohorts are apparent among older households, while at younger ages there is little or no distinction between cohorts.

Figure 3 plots age-profiles for the two lifetime wealth measures (θ_{ij}^c and θ_{ij}^h +net worth) at different levels of education, current consumption, and current earnings. The first row of Figure 3 shows that, as expected, more education is associated with higher lifetime wealth. Comparing across the two representations of lifetime welfare, the consumption representation indicates somewhat lower values than those estimated using permanent income (based on human wealth and net worth). This is especially apparent close to retirement and late in the life cycle when average net worth becomes larger than the gaps between the age-specific averages of θ_{ij}^c and θ_{ij}^h . The latter observation recasts the well-known retirement savings puzzle (Banks, Blundell, and Tanner (1998)) in terms of present values and suggests that some of the net worth of richer households is never consumed by them and is possibly passed on to the next generation.

Notably, quartiles of current consumption (middle panels of Figure 3) offer more accurate predictions of lifetime wealth gaps than either education or earnings quartiles (top and bottom panels). This suggests that early consumption choices encompass private information that cannot be surmised from current earnings or education alone. This illustrates the importance of using data sources that combine expenditures with more traditional predictors of income. In practice, conditioning on expenditure data broadens the empirical value of the model as it allows it to capture additional layers of heterogeneity.

5. THE EVOLUTION OF CROSS-SECTIONAL INEQUALITY

Table 1 displays the evolution of proportional variation for several variables between 1983 and 2016. The first two columns display year-specific log variances of the CE measures defined in (6). Both measures indicate a persistent increase in welfare inequality over the 34-year sample period, differing only in the magnitude of the increase. Dispersion in \bar{c}_{ij}^p rises by 12.6 log points while for \bar{c}_{ij}^c the growth is 9.9 log points. Over the same

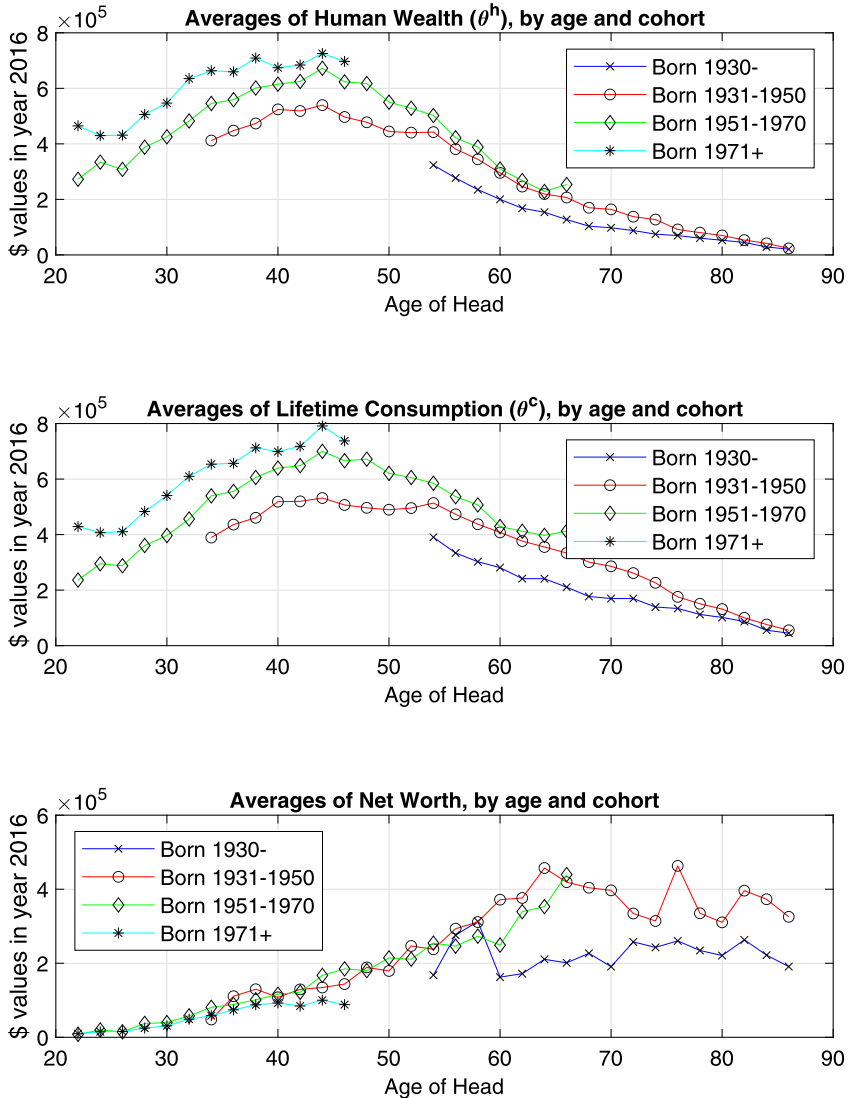


FIGURE 2. This figure illustrates the evolution of the average human wealth (θ_{ij}^h), lifetime consumption (θ_{ij}^c), and net worth over the life cycle of successive cohorts in our sample. The life-cycle profiles for younger cohorts lie above those of older ones, confirming that average expenditure and earnings have risen over time. The cross-cohort patterns of net worth are not as sharply ordered; however, for the two cohorts that we observe into their 80s, the more recent one clearly displays more wealth at the end of life. All values are reported in 2016 equivalent dollars.

period, the proportional variation of expenditures and earnings increased, respectively, by 10.4 and 11.7 log points, albeit from much larger base values. By comparison, the proportional dispersion of net worth is not only orders of magnitude larger but it also grew a lot faster. These patterns document a clear but uneven link between rising proportional variation in observable variables, for example, consumption, earnings and net

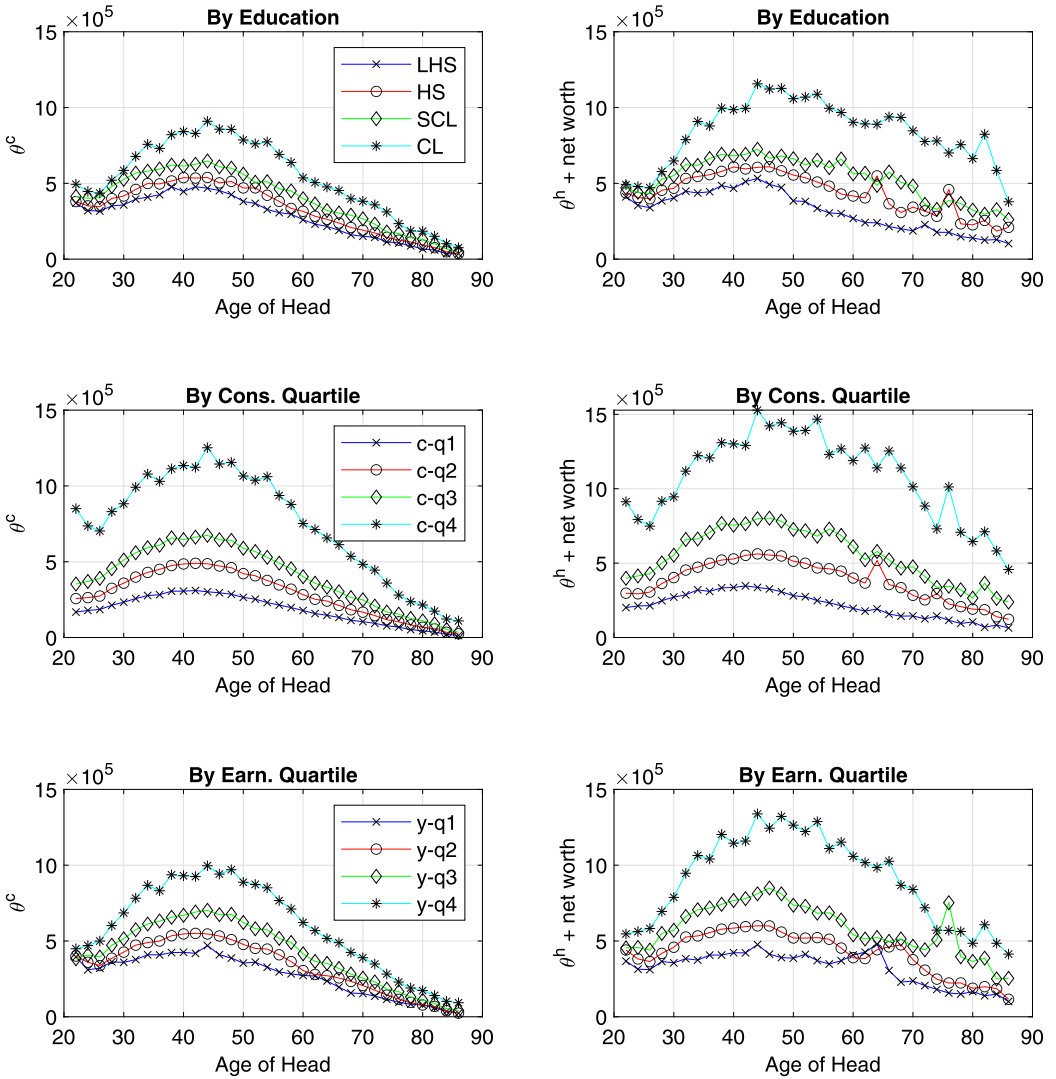


FIGURE 3. This figure shows the evolution of lifetime wealth measures, conditional on various dimensions of heterogeneity. The left panels plot the average risk-adjusted lifetime consumption (θ^c), while the right panels plot the average human-plus-financial wealth ($\theta^h + \text{net worth}$). The first row breaks down these averages by education of the household head (less than high school LHS, high school HS, some college SCL, and college CL). The second row shows average lifetime wealth measures for the four quartiles of current expenditure (household total) at each age. The third row shows average lifetime wealth measures for the four quartiles of current earnings (household totals) at each age. All values are reported in 2016 equivalent dollars.

worth (columns 3–5), and cross-sectional dispersion of lifetime welfare. We learn two notable lessons from Table 1: first, the dispersion of current consumption and earnings substantially overstates welfare inequality; second, drawing inference about the evolution of welfare inequality from the variance of net worth would result in a considerable

TABLE 1. Variances of consumption equivalents, current expenditures, current earnings, and assets, between 1983 and 2016. Standard errors in parentheses are based on 1000 nonparametric block bootstrap replications, where the raw panel data are resampled at the household level. These results illustrate how proportional dispersion of certainty-equivalent consumption has increased by a similar magnitude as that of expenditure itself. However, the level of certainty equivalent inequality is generally lower.

Year	Variance of \ln^\dagger				
	(1) \bar{c}_{ij}^c	(2) \bar{c}_{ij}^p	(3) c_{ij}	(4) y_{ij}	(5) a_{ij}^\dagger
1983	0.130 (0.008)	0.127 (0.011)	0.191 (0.005)	0.372 (0.013)	21.9 (1.30)
1989	0.153 (0.008)	0.146 (0.008)	0.206 (0.006)	0.403 (0.016)	27.8 (1.49)
1993	0.158 (0.008)	0.154 (0.007)	0.21 (0.006)	0.44 (0.017)	36.0 (1.69)
1998	0.182 (0.007)	0.188 (0.009)	0.235 (0.006)	0.415 (0.015)	38.9 (1.65)
2004	0.209 (0.007)	0.221 (0.008)	0.267 (0.007)	0.49 (0.018)	45.1 (1.66)
2010	0.226 (0.007)	0.253 (0.009)	0.283 (0.007)	0.547 (0.018)	67.6 (1.70)
2016	0.229 (0.008)	0.253 (0.009)	0.295 (0.008)	0.489 (0.015)	60.1 (1.74)

[†]All variables expressed in natural logs, except net worth a_{ij} , where the inverse hyperbolic sine transformation is used.

overestimate.¹⁵ In summary, while observable variables exhibit variation of the expected magnitude (e.g., the assets variance is much larger than its counterpart for consumption or earnings), they may provide an inaccurate portrayal of both levels and changes in welfare inequality over the sample period.

5.1 Three facts about lifetime inequality

While the steep increase in welfare inequality is striking, other conspicuous patterns shed light on the mechanics of growing inequality in lifetime wealth.

1. The magnitude of welfare dispersion. As shown above, welfare inequality is about half as large as earnings inequality in 2016. Remarkably, this is true even for the welfare measure \bar{c}^p constructed from earnings data. One possibility is that human wealth θ_{ij}^h is much more equally distributed than earnings, which could occur if a large part of earnings inequality is transitory. This is not the case, as the variance of $\ln(\theta_{ij}^h)$ is actually larger

¹⁵Transitory measurement error may arise from misreporting in the survey responses themselves, or from noise in the estimation procedure of \bar{c}_{ij}^c and \bar{c}_{ij}^p . An advantage of present value measures of welfare is that they help average out transitory measurement error in survey responses. In regards to measurement error possibly arising from estimation itself, the standard errors of estimated variances are stable over the sample period, which indicates that any such error, if present, is also stable and does not affect the precision of estimated trends. See also Gallipoli, Low, and Mitra (2020).

TABLE 2. Variances of human wealth and total lifetime wealth, with and without adjustment for age composition. All variables are in natural logs. Adjusted versions of human and lifetime wealth have been divided by the age-adjustment factor $B_{ij} = \sum_{k=j}^J \prod_{m=j}^k \beta_k$, which accounts for heterogeneity in expected remaining length of life. Standard errors in parentheses are based on 1000 nonparametric block bootstrap replications, where the raw panel data are resampled at the household level.

Year	Variance of ln			
	(1) θ_{ij}^h	(2) θ_{ij}^h/B_{ij}	(3) $\theta_{ij}^h + a_{ij}$	(4) $(\theta_{ij}^h + a_{ij})/B_{ij}$
1983	0.451 (0.024)	0.115 (0.017)	0.211 (0.009)	0.188 (0.027)
1989	0.508 (0.028)	0.152 (0.020)	0.227 (0.013)	0.209 (0.018)
1993	0.514 (0.025)	0.169 (0.015)	0.225 (0.012)	0.210 (0.015)
1998	0.529 (0.023)	0.205 (0.019)	0.249 (0.011)	0.294 (0.025)
2004	0.543 (0.021)	0.229 (0.018)	0.284 (0.011)	0.385 (0.023)
2010	0.518 (0.021)	0.222 (0.019)	0.277 (0.011)	0.368 (0.026)
2016	0.552 (0.025)	0.246 (0.020)	0.273 (0.014)	0.386 (0.022)

than the variance of $\ln(y_{ij})$, as shown in the first column of Table 2. However, the variance of $\ln(\theta_{ij}^h)$ includes age effects due to the pronounced life-cycle profile of human wealth, wherein young households are “human-wealth-rich” because they have many years of working life ahead of them. The importance of age effects in the translation of lifetime wealth into \bar{c}_{ij}^p is apparent in equation (6), where the number of elements in the sum on the left-hand side falls with age. For this reason, in the second column of Table 2 we report age-adjusted variances obtained after dividing human wealth by $B_{ij} = \sum_{k=j}^J \prod_{m=j}^k \beta_k$, which is the same annuitization factor implicitly applied when calculating \bar{c}_{ij}^p . This adjustment results in lower proportional dispersion of human wealth and in estimates that are much closer to the dispersion of certainty equivalent consumption. We conclude that the variances of current earnings and of unadjusted human wealth significantly overstate welfare inequality. In the next section, we suggest a simple procedure to explicitly account for such confounding demographic effects when breaking down welfare gains into different components.

2. The mitigating influence of human wealth on inequality. Human wealth has a mitigating effect on lifetime inequality. The top panel of Figure 1 establishes a connection between the age-dependence of human wealth and the accumulation of financial wealth over the life cycle, in the sense that human wealth is converted into financial wealth as households age. Thus, accounting for both human and financial wealth in lifetime wealth portfolios should lead to a more evenly distributed statistic. Indeed, the third column of Table 2 shows that the variance of $\ln(\theta_{ij}^h + a_{ij})$ is roughly half as large

as the variance of $\ln(\theta_{ij}^h)$. This finding is quantitatively striking because it shows that the proportional variation of total lifetime wealth, that is, the sum of θ_{ij}^h and a_{ij} , is a lot lower than the proportional variation of each individual component. To put this into context, note that the result holds even though the proportional dispersion of assets a_{ij} is more than an order of magnitude larger than that of θ_{ij}^h . This is explained, quite intuitively, by the fact that human wealth raises the value of the lifetime wealth portfolios for a large share of households who hold little or no financial wealth. Once we account for the disproportionate role that human wealth has in the lifetime value of many asset-poor households, both young and old, our view of inequality changes.

3. The rising dispersion of permanent-income. The variance of annuitized lifetime wealth has more than doubled between 1983 and 2016. When we take the step of annuitizing lifetime wealth in the fourth column of Table 2, we find that during the 1980s and early 1990s, age-adjusted lifetime wealth was more evenly distributed than its unadjusted counterpart; however, by 2016 age-adjusted lifetime wealth had grown considerably more unequal. This pattern results in a steep increase of total lifetime wealth dispersion between 1983 and 2016, reflecting a broad shift toward larger shares of financial wealth in household portfolios. Because financial wealth is more unequally distributed, the shift in portfolio composition has driven the overall dispersion higher. In aggregate, the ratio of assets to lifetime wealth increased from about 24% at the beginning of the 1980s to about 31% just before the 2008 recession, falling back after the recession to about 28% in 2016. This is largely due to increased asset holdings by older households, as is apparent in the bottom panel of Figure 2. To concisely summarize these little known patterns, we use Lorenz Curves representing the concentration of net worth (financial and tangible wealth), human wealth, and the resulting consumption-equivalent values \bar{c}_{ij}^p in 1983 and 2016. Figure 4 shows that net worth is by far the most concentrated among the three variables, while the consumption-equivalent values are the least concentrated. Moreover, all three variables exhibit growing concentration over the sample period. However, the proportional change in concentration is highest for the consumption-equivalent measure, with its Gini growing about 40% (from 0.18 in 1983

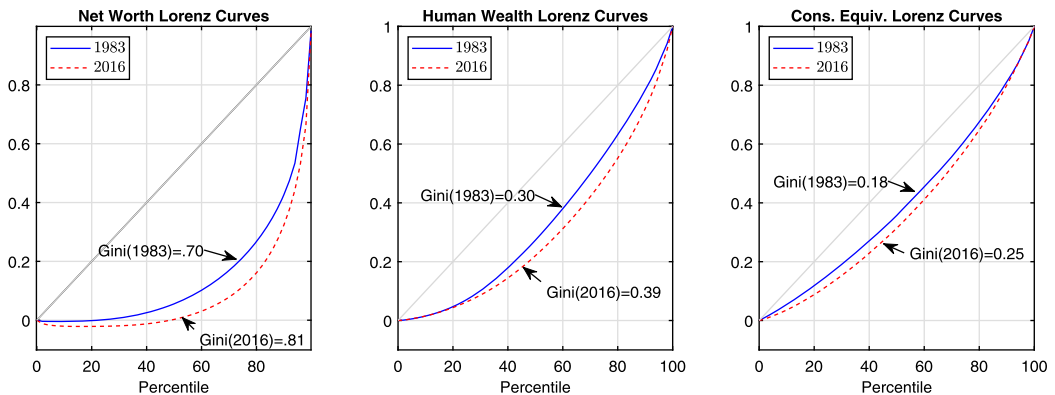


FIGURE 4. Lorenz curves of net worth, human wealth, and consumption-equivalents (permanent-income representation), 1983 versus 2016.

to 0.25 in 2016). By comparison, the Gini of net worth increased by less than 16% and the Gini of human wealth grew by 30%. The jump in consumption-equivalent concentration cannot be explained only by the increased concentration of net worth and human wealth; rather, it largely reflects the growing importance of net worth in household portfolios and, by extension, in their lifetime wealth and consumption expenditures. We conclude that, while human wealth has a mitigating effect on inequality, this effect has become less intense over time.

6. THE AGGREGATE WELFARE CONSEQUENCES OF RISING INEQUALITY

To make sense of the growing discrepancies documented in the previous section, we need a framework for assessing the contributions of different elements, such as average expenditure levels, uninsurable risk, and inequality. As originally shown in [Floden \(2001\)](#) and [Benabou \(2002\)](#), certainty equivalent measures lend themselves naturally to welfare decompositions. In what follows, we posit a utilitarian welfare function, and suggest a decomposition that accounts for the changing demographic structure.

6.1 A four-way welfare decomposition

Let $\phi_{j,t}$ denote the measure of age j individuals in period t . Utilitarian welfare at t is

$$W_t = \sum_{j=1}^J \int V_j(a_j, s^j) d\mu_j(a_j, s^j|t) \phi_{j,t},$$

where integrals deliver averages over the age-specific distribution of state variables $\mu_j(a_j, s^j|t)$. Cohorts are weighted according to $\phi_{j,t}$. We maintain $\sum_{j=1}^J \phi_{j,t} = 1$ and $\int d\mu_j(a_j, s^j|t) = 1$. Denoting expenditure rules $c_j(a_j, s^j|t) = c_{j,x|t}$ and measures $\mu_j(a_j, s^j|t) = \mu_{j,x|t}$, we define

$$\begin{aligned} \bar{C}_t &= \sum_{j=1}^J \int \bar{c}_{j,x|t} d\mu_{j,x|t} \phi_{j,t}, \\ C_t &= \sum_{j=1}^J \int c_{j,x|t} d\mu_{j,x|t} \phi_{j,t}. \end{aligned}$$

where $\bar{c}_{j,x|t}$ is the certainty-equivalent derived in equation (6) and \bar{C}_t is the average certainty-equivalent consumption across all agents alive at t ; C_t is the average consumption in the economy in the same period.

Given estimates of C_t and \bar{C}_t , one can define the “price” of inequality p_t^{ine} and the “price” of uncertainty p_t^{unc} as the implicit solutions to the following two equations:

$$\begin{aligned} \sum_{j=1}^J \phi_{j,t} \sum_{k=j}^J \beta^{k-j} u((1 - p_t^{\text{ine}}) \bar{C}_t) &= \sum_{j=1}^J \phi_{j,t} \int \sum_{k=j}^J \beta^{k-j} u(\bar{c}_{j,x}) d\mu_{j,x|t}, \\ u((1 - p_t^{\text{unc}}) C_t) &= u(\bar{C}_t). \end{aligned}$$

Price p_t^{ine} measures welfare losses due to dispersion in certainty-equivalent consumption and does not reflect uncertainty in individual consumption plans. In contrast, p_t^{unc} equates the mean certainty-equivalent consumption \bar{C}_t and the economy-wide average consumption C_t , measuring losses due to uninsurable consumption volatility.

Consumption metrics are convenient to break down welfare changes between any two periods A and B into separate components. For example, if C_A and C_B are the average consumption in the first and second periods, then $\omega^{\text{lev}} = \frac{C_B}{C_A} - 1$ measures the change in welfare due to the level shift in aggregate consumption across the two periods. The latter is just one contributor to the total consumption-equivalent welfare change from A to B , which is measured by the $\tilde{\omega}$ that satisfies the following equality:

$$\sum_{j=1}^J \int \sum_{k=j}^J \beta^{k-j} u((1 + \tilde{\omega})c_{j,x|t}) d\mu_{j,x|A} \phi_{j,A} = \sum_{j=1}^J \int \sum_{k=j}^J \beta^{k-j} u(c_{j,x|t}) d\mu_{j,x|B} \phi_{j,B}.$$

Rearranging and simplifying, we obtain

$$\begin{aligned} & u((1 + \tilde{\omega})(1 - p_A^{\text{ine}})(1 - p_A^{\text{unc}})C_A) \sum_{j=1}^J \sum_{k=j}^J \beta^{k-j} \phi_{j,A} \\ &= u((1 - p_B^{\text{ine}})(1 - p_B^{\text{unc}})C_B) \sum_{j=1}^J \sum_{k=j}^J \beta^{k-j} \phi_{j,B} \end{aligned}$$

Next, given isoelastic utility, we can derive a relationship of the form:

$$(1 + \tilde{\omega}) = (1 + \omega^{\text{ine}})(1 + \omega^{\text{unc}})(1 + \omega^{\text{lev}})(1 + \omega^{\text{dem}})$$

where the total welfare change $\tilde{\omega}$ is decomposed into four elements:

1. A change due to *increased level* of consumption: $\omega^{\text{lev}} = \frac{C_B}{C_A} - 1$.
2. A change due to *reduced uncertainty*: $\omega^{\text{unc}} = \frac{1 - p_B^{\text{unc}}}{1 - p_A^{\text{unc}}} - 1$.
3. A change due to *reduced inequality*: $\omega^{\text{ine}} = \frac{1 - p_B^{\text{ine}}}{1 - p_A^{\text{ine}}} - 1$.
4. A change due to *demographic composition*: $\omega^{\text{dem}} = \left(\frac{u^{-1}(\sum_{j=1}^J \sum_{k=j}^J \beta^{k-j} \phi_{j,B})}{u^{-1}(\sum_{j=1}^J \sum_{k=j}^J \beta^{k-j} \phi_{j,A})} \right) - 1$.

Then the demographic-corrected consumption equivalent welfare change from period A to B is the ω_{tot} that solves

$$(1 + \omega_{\text{tot}}) = \frac{(1 + \tilde{\omega})}{(1 + \omega^{\text{dem}})} = (1 + \omega^{\text{ine}})(1 + \omega^{\text{unc}})(1 + \omega^{\text{lev}}).$$

The measure ω_{tot} is the percentage change in consumption that makes agents indifferent between periods A and B , holding constant demographic composition.

6.1.1 Implementation The nonlinear nature of the calculations required to recover the different welfare components means that noise and outliers could become influential and potentially lead to inaccurate conclusions. To alleviate such concerns, we smooth data through aggregation into groups based on age and deciles of lifetime utility (value functions); the means within these groups are then used to construct the empirical counterpart of the utilitarian welfare function:

$$W_t = \sum_{k=1}^K \sum_{q_k^v=1}^{10} E[V_{ij}|q_k^v] \frac{\phi_{k,t}}{10},$$

where k indexes decade of adult life (e.g., 20s, 30s, etc.), and q_k^v indexes deciles specific to each k age group. The sums and integrals of the remaining calculations are likewise adjusted.

Importantly, the results must be interpreted with this aggregation in mind. Although averages are unaffected (the mean of equally weighted means is the overall mean), the “price” of inequality p^{ine} can be sensitive to grouping, as it reflects differences between broad segments of age-specific distributions, as opposed to finer adjustments within age-specific deciles. For example, if the top 1% get richer at the expense of households below the 90th percentile, then p^{ine} becomes larger; however, if the top 1% get richer at the expense of only those in the 98th percentile, then p^{ine} will not rise (because $E[V_{ij}|q_k^v = 10]$ would not change in this case). In the robustness section, we report on an alternative decomposition in which we leverage the larger samples of the Survey of Consumer Finances to allow separate top 1% groupings.¹⁶ The results are not appreciably different, implying that welfare costs of inequality are well captured by measuring differences between broad decile-based groups. Put differently, the simple decile-based decomposition captures the pertinent aspects of the changing inequality patterns and accurately reflects the fundamental differences between rich and poor.

6.2 The evolution of aggregate welfare and its components

Table 3 tracks changes in the welfare components described above for selected years from 1983 to 2016. The top and bottom panels refer, respectively, to consumption equivalent values for lifetime utility measures based on the permanent-income representation (\bar{c}^p) and the consumption representation (\bar{c}^c). The base year (year A in the derivations) is set to 1983, while the comparison year (year B) varies according to the first column of the table. For example, using the \bar{c}^p estimates, between 1983 and 2016 aggregate consumption increased by 87.5% ($\omega_{\text{lev}}^{2016} = 0.875$), the price of inequality increased by 14.1% ($\omega_{\text{ine}}^{2016} = -0.141$), and the price of uncertainty fell by 2.2% ($\omega_{\text{unc}}^{2016} = 0.022$), which combine to generate an overall increase in utilitarian welfare equivalent to a 64.6% of the consumption of agents alive in 1983.

¹⁶To be clear, we do not use the SCF for our main analysis because it is repeated cross-sectional data. Since panel data are required to form estimates of $\theta_j^h(\mathbf{z})$ and $\theta_j^c(\mathbf{z})$, we choose to perform all the baseline analysis using the PSID. In the robustness section, the functions $\theta_j^h(\mathbf{z})$ and $\theta_j^c(\mathbf{z})$ (estimated with the PSID) are evaluated at data points \mathbf{z} observed in the SCF.

TABLE 3. Top panel: components of utilitarian aggregate welfare based on the \bar{c}^p measure, by year. Bottom panel: same components of aggregate welfare for the \bar{c}^c measure. The base year is 1983, therefore, the values reported for each year are relative to values in 1983. Standard errors in parentheses are based on 1000 nonparametric block bootstrap replications, where the raw panel data are resampled at the household level.

Year	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}
Welfare Decomposition— \bar{c}^p				
1983	—	—	—	—
1989	0.186 (0.016)	0.213 (0.011)	-0.018 (0.003)	-0.004 (0.008)
1993	0.218 (0.018)	0.261 (0.012)	-0.026 (0.003)	-0.008 (0.009)
1998	0.337 (0.020)	0.425 (0.015)	-0.070 (0.004)	0.008 (0.009)
2004	0.491 (0.024)	0.656 (0.018)	-0.114 (0.005)	0.016 (0.010)
2010	0.502 (0.026)	0.716 (0.019)	-0.146 (0.006)	0.025 (0.011)
2016	0.646 (0.031)	0.875 (0.022)	-0.141 (0.007)	0.022 (0.011)
Welfare Decomposition— \bar{c}^c				
1983	—	—	—	—
1989	0.188 (0.016)	0.213 (0.011)	-0.024 (0.013)	0.003 (0.014)
1993	0.225 (0.017)	0.261 (0.012)	-0.035 (0.012)	0.007 (0.015)
1998	0.346 (0.019)	0.425 (0.015)	-0.064 (0.014)	0.008 (0.015)
2004	0.493 (0.023)	0.656 (0.018)	-0.096 (0.015)	-0.002 (0.015)
2010	0.514 (0.025)	0.716 (0.019)	-0.123 (0.014)	0.006 (0.015)
2016	0.668 (0.029)	0.875 (0.022)	-0.124 (0.014)	0.016 (0.015)

The welfare losses associated with inequality appear to accelerate around the year 2000, pulling down the gains due to higher expenditure levels. Indeed, in aggregate data, real consumption expenditures per capita nearly doubled from 1983 to 2016.¹⁷ However, the aggregate welfare measure ω_{tot} only grows by about 65% because of the unequal nature of consumption growth. These trends are echoed in the welfare decomposition based on \bar{c}^c , although with a somewhat smaller discounting of growth due to rising inequality, which is consistent with the differences in trends observed in Table 1. It is worth highlighting that both consumption equivalent measures indicate a small but consistent drop in the welfare costs of uncertainty, suggesting marginally improved cross-sectional insurance.

¹⁷<https://fred.stlouisfed.org/series/A794RX0Q048SBEA>

Table 3 provides a natural way to jointly discuss aggregate growth and increased inequality. It is clear that aggregate consumption growth has passed through into aggregate welfare growth: our results suggest that the progress made between 1983 and 2016 is equivalent to increasing the consumption of every agent alive in 1983 by about 65%. To give some context, this is at least an order of magnitude larger than recent estimates of the welfare gains from moving to an optimally progressive tax system.¹⁸ Yet, this gain is considerably less than the growth in aggregate consumption, primarily because of its unequal nature. If one could achieve the same growth in a manner that preserved the degree of equality that existed in 1983, aggregate welfare would have risen by the equivalent of almost 90% of the lifetime consumption of all households in 1983. Thus, about one-quarter of the potential welfare gain has been lost to inequality. Moreover, much of these losses have accrued over the two most recent decades in our sample.

7. SENSITIVITY AND ROBUSTNESS

7.1 *Comparing the PSID and the survey of consumer finances*

While the PSID is unique in that it contains long panel information about earnings, wealth, and expenditures, one concern is that the distribution of wealth is less accurate than other cross-sectional data sets specifically designed to measure wealth holdings such as the Survey of Consumer Finances (SCF). Our estimation exercise requires panel data on both consumption and income, which the SCF does not provide, to recover the human capital valuation functions in (9). Nonetheless, it is possible to leverage the wealth information in the SCF by implementing two simple steps: first, just like before, estimate the valuation functions $\theta_j^h(\mathbf{z})$ and $\theta_j^c(\mathbf{z})$ using the PSID, relying on the fact that our empirical approach delivers the whole function rather a set of point estimates; second, evaluate the estimated functions at data points within the SCF. The latter step can be performed starting from the 2004 wave of the SCF, when measures of food consumption were added and total expenditure can be imputed in precisely the same way it is done in the PSID.

Before comparing results based on the SCF to the baseline findings, it is instructive to contrast the distributions of observable variables in the two data sets. Table 4 reports trends in the variances of log consumption, income, and net worth for both SCF and PSID since 2004. The table illustrates several ways in which the distributions from the two data sets are consistent with each other. Notably, proportional dispersion in consumption is very similar across the data sets, which results from the distributions of the underlying imputation variables being close (e.g., food consumption is similarly distributed in the two data sets). The variance of log earnings is 2.5–3.5 times as large as that of log consumption; perhaps more importantly, proportional dispersion in net worth is two orders of magnitude larger than in consumption/earnings in both data sets. At a finer level of detail, it is apparent that proportional dispersion of earnings in the PSID is

¹⁸For example, Heathcote, Storesletten, and Violante (2020) estimate the gains to be about 1.8% of lifetime consumption, which Conesa and Krueger (2006) estimate a gain equivalent to 1.7%.

TABLE 4. Variances of current expenditures, current earnings, and assets in the SCF and PSID samples. For both data sets, the variance of log earnings is based on household earnings when the head is less than 66 years old and household earnings exceed \$1000 in 2016 dollars.

Year	SCF Variance of \ln^\dagger			PSID Variance of \ln^\dagger		
	c_{ij}	y_{ij}	a_{ij}^\dagger	c_{ij}	y_{ij}	a_{ij}^\dagger
2004	0.226	0.613	33.6	0.215	0.490	45.1
2010	0.241	0.668	52.8	0.218	0.547	67.6
2016	0.244	0.663	50.9	0.225	0.489	60.1

[†]All variables in natural logs, except net worth a_{ij} , where the inverse hyperbolic sine transformation is used.

smaller than in the SCF, while proportional variation in net worth is actually larger in the PSID. One salient observation regarding net worth is that, under alternative measures of inequality, the SCF data may imply more inequality than the PSID. For example, the top 10% share of wealth in 2016 in the SCF is 77% (consistent with the findings by Bricker et al., 2016), while in the PSID it is only 66%.

7.2 Welfare measures based on SCF wealth data

Table 5 presents the welfare decomposition based on the permanent-income representation (\bar{c}^P) using SCF data, as well as proportional variation for both measures of certainty equivalent consumption. For ease of comparison, the PSID results are reproduced in the bottom panel of the table using 2004 as the base year for the welfare decomposition. The results are remarkably similar across data sets in both magnitudes and, notably, in trends. One interesting difference that emerges is the larger jump in the variance of $\ln(\bar{c}^P)$ between 2010 and 2016 in the SCF sample, which results from the significant increase in wealth inequality observed in the SCF over that period. This change is passed through to the welfare analysis, in the sense that the increase in the price of in-

TABLE 5. Robustness analysis, SCF versus PSID. The left-hand side of the table shows welfare decompositions in the two data sets based on the permanent-income representation between 2004 and 2016. All values are relative to the base year 2004. The right-hand side of the table reports the variances of consumption equivalents for both welfare representations in 2004, 2010, and 2016.

Year	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	\bar{c}_{ij}^c	\bar{c}_{ij}^P
		SCF Welfare Decomp.— \bar{c}^P			Var of log	
2004	—	—	—	—	0.139	0.187
2010	-0.012	0.036	-0.035	-0.012	0.160	0.220
2016	0.107	0.138	-0.079	0.057	0.161	0.252
		PSID Welfare Decomp.— \bar{c}^P			Var of log	
2004	—	—	—	—	0.209	0.221
2010	0.020	0.036	-0.043	0.029	0.226	0.246
2016	0.121	0.133	-0.041	0.032	0.229	0.253

equality is larger in the SCF sample between 2010 and 2016, resulting in smaller welfare gains in the aggregate. Nonetheless, results from both data sources reinforce the view that large losses are associated with rising certainty equivalent inequality. For example, the PSID estimates suggest that a quarter of the welfare gains from higher consumption ($\omega_{lev} = +0.13$) and better cross-sectional insurance ($\omega_{unc} = +0.03$) have been lost due to higher inequality ($\omega_{ine} = -0.04$ between 2004 and 2016). The SCF estimates, however, show welfare gains from higher consumption levels and insurance of almost 20% over the 2004–2016 period, but more than 1/3 of this growth being lost to inequality in certainty equivalent consumption. These patterns imply almost identical total changes in welfare and paint a consistent picture of the relative contribution of levels, inequality and uncertainty to economic prosperity. We do not report comparisons of welfare decompositions based on the \bar{c}^c metric because differences across the data samples are even smaller in that case.

7.3 Accounting for the top 1%

The Survey of Consumer Finances has larger samples and it oversamples rich households; these features allow us to carry out a welfare decomposition exercise in which we separately group the top 1% of households. For this sensitivity exercise, we split the top decile of lifetime utility into the 91–99 percentiles and the top 1% before implementing the welfare decomposition. As shown in Table 6, the finer percentile grouping results in almost no difference, which is evidence that comparing the top 10% to the bottom 90% adequately captures the welfare costs of rising inequality even if gains in the top 10% mostly accrue to the top 1%. This robustness follows from the fact that the mean within the top 10% is highly responsive to changes in the top 1%.

7.4 Unobserved heterogeneity

The presence of unobservable characteristics, such as prior knowledge of one’s life-cycle earnings trajectory, could imply that our estimates of θ_{ij}^h and θ_{ij}^c underestimate inequal-

TABLE 6. SCF robustness analysis. The top panel shows the welfare decomposition in the SCF sample based on the permanent-income representation, between 2004 and 2016, when the top 1% is separately accounted for. The bottom panel is the baseline welfare decomposition without separately accounting for the top 1%. All values are relative to the base year 2004.

Year	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}
SCF Welfare Decomp.— \bar{c}^p (with top 1%)				
2004	–	–	–	–
2010	–0.012	0.036	–0.035	–0.011
2016	0.107	0.138	–0.079	0.057
SCF Welfare Decomp.— \bar{c}^p (deciles)				
2004	–	–	–	–
2010	–0.012	0.036	–0.035	–0.012
2016	0.107	0.138	–0.079	0.057

ity to some degree. One reason we include current consumption c_{ij} in the conditioning set \mathbf{z}_{ij} is to account for unobserved information embedded in measured expenditures. Figure 3 shows that, indeed, consumption captures variation beyond what is contained in income and education; however, one can devise scenarios where current consumption would not adequately capture prior information.

To account for such contingencies, our empirical framework can flexibly accommodate unobserved types, denoted as η_i . We do so by focusing on a well-known dimension of income heterogeneity, namely idiosyncratic earnings profiles. To estimate unobserved types, we adopt the approach suggested by Bonhomme, Lamadon, and Manresa (2021) and make it suitable to our setting; specifically, we employ a k-medians grouping algorithm to separate mean life-cycle earnings growth (our “informative variable”) into clusters. To establish the number of clusters, we follow the reasoning of Cunha, Heckman, and Navarro (2005); these authors suggest that, if agents know their own type, they should act upon such information and make choices that are consistent with their type. It follows that it should be possible to identify heterogeneity due to ex ante types because individuals respond to this information and act on it. If any part of the variation in life-cycle earnings growth is anticipated by agents, then their long-term choices should reflect this prior information.

Following Cunha, Heckman, and Navarro, we consider the long-term decision to attend college and let S_i denote the college choice of individual i , taking value one if the individual completes college and zero otherwise. To the extent that a person's type η_i affects their long term choices, we would expect that $\text{Cov}(S_i, \eta_i) \neq 0$. Moreover, given the relationship between types and economic outcomes such as earnings, schooling choices should be related to the (ex post) level of earnings growth. By the same token, if one could control directly for the underlying type η_i , the expectation of college completion should no longer respond to these observable measures of ex post earnings. This line of reasoning offers a natural way to test whether our grouping procedure identifies the relevant “type” variation based on projecting college indicators on average earnings growth. If the grouping algorithm successfully captures the relevant heterogeneity, the type indicator should crowd out the statistical effect of earnings profiles on college status. We find that allowing for three types is sufficient to remove any direct effect of mean earnings growth on college completion.

Having established the cardinality of the type set, we replicate the welfare analysis after including the unobserved type indicator η_i in the conditioning vector \mathbf{z}_{ij} . As with comparisons to the SCF, we present welfare decomposition results (Table 7) based on the \bar{c}^p metric, where differences are largest, and we report the evolution in proportional variation of certainty equivalent consumption for both welfare measures. The results suggest that controlling for unobserved heterogeneity results in slightly flatter trends in estimated welfare inequality. This occurs, partly, because finer conditioning based on unobserved types results in larger initial dispersion in the base year of 1983. These discrepancies do not alter the conclusions from our baseline analysis in any meaningful way.

TABLE 7. Welfare decomposition and inequality when accounting for unobserved types. The left-hand side of the table shows the welfare decomposition based on the permanent-income representation between 1983 and 2016 (PSID sample). All values are relative to the base year 1983. The right-hand side of the table reports the variances of consumption equivalents for both welfare representations between 1983 and 2016.

Year	Welfare Decomp.— \bar{c}^p				Var of log	
	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	\bar{c}_{ij}^c	\bar{c}_{ij}^p
1983	—	—	—	—	0.185	0.147
1989	0.184	0.213	−0.019	−0.006	0.149	0.141
1993	0.217	0.261	−0.026	−0.009	0.154	0.149
1998	0.342	0.425	−0.066	0.008	0.177	0.181
2004	0.499	0.656	−0.108	0.015	0.204	0.212
2010	0.520	0.716	−0.136	0.025	0.217	0.240
2016	0.692	0.875	−0.120	0.025	0.217	0.234

7.5 Utility curvature

The utility curvature parameter γ is set to 2 in the baseline analysis. However, it is informative to verify how sensitive results are to this assumption and we investigate it by alternately assuming $\gamma = 2.25$ or $\gamma = 1.75$. Comparing Tables 8 and 9, which show results for each of these cases, to the baseline Table 3, it is apparent that a larger γ leads to heavier losses from rising inequality, which show up in the ω_{ine} column. The definition of the price of inequality p_t^{ine} in Section 6 illustrates why a larger value of γ implies higher inequality losses in the utilitarian welfare function. The tables also show that a larger value of γ results in a lower estimate of certainty equivalent dispersion. However, and most importantly, the trends of inequality remain consistently close to those estimated in the baseline exercise.

7.6 Accounting for labor supply inequality

To account for the disutility of labor supply in the welfare analysis, we extend the preference specification to be $U(c_j, \ell_j) = c_j^{1-\gamma}/(1-\gamma) - \nu \ell_j^{1+\chi}/(1+\chi)$, where χ is the in-

TABLE 8. Welfare decomposition and inequality trends with $\gamma = 2.25$.

Year	Welfare Decomp.— \bar{c}^p				Var of log	
	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	\bar{c}_{ij}^c	\bar{c}_{ij}^p
1983	—	—	—	—	0.103	0.108
1989	0.178	0.213	−0.020	−0.009	0.128	0.129
1993	0.212	0.261	−0.029	−0.010	0.134	0.138
1998	0.328	0.425	−0.073	0.006	0.160	0.174
2004	0.478	0.656	−0.122	0.016	0.197	0.211
2010	0.481	0.716	−0.150	0.015	0.210	0.238
2016	0.616	0.875	−0.145	0.008	0.210	0.237

TABLE 9. Welfare decomposition and inequality trends with $\gamma = 1.75$.

Year	Welfare Decomp.— \bar{c}^p				Var of log	
	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	\bar{c}_{ij}^c	\bar{c}_{ij}^p
1983	–	–	–	–	0.163	0.145
1989	0.196	0.213	–0.015	0.001	0.183	0.161
1993	0.228	0.261	–0.022	–0.005	0.187	0.168
1998	0.351	0.425	–0.066	0.015	0.211	0.201
2004	0.514	0.656	–0.104	0.020	0.231	0.230
2010	0.537	0.716	–0.138	0.039	0.247	0.265
2016	0.685	0.876	–0.136	0.039	0.254	0.268

verse of the Frisch elasticity of labor supply ℓ_j . Separable labor disutility, together with the isoelasticity assumption, guarantees that the lifetime welfare representation can be generalized in a natural way.

Letting w_j be the net wage rate, we can substitute $y(s^j) = w_j \ell_j$ in the lifetime budget constraint (3). Isoelasticity is useful insofar it allows the value function to be expressed as (all derivations are in Appendix C)

$$V_j(a_j, s^j) = \frac{u'(c_j)}{1 - \gamma} E_j \left[\sum_{k=j}^J \beta^k \frac{u'(c_k)}{u'(c_j)} \left(c_k + \frac{\gamma - 1}{1 + \chi} w_k \ell_k \right) \right] \tag{11}$$

$$= \frac{u'(c_j)}{1 - \gamma} \theta_j^{c,adj}, \tag{12}$$

where we denote the labor-adjusted present value $E_j[\sum_{k=j}^J \beta^k \frac{u'(c_k)}{u'(c_j)} (c_k + \frac{\gamma-1}{1+\chi} w_k \ell_k)]$ as $\theta_j^{c,adj}$. Intuitively, as the labor supply elasticity $\frac{1}{\chi}$ tends to zero, lifetime utility collapses back to the baseline consumption representation in (4). For nonzero labor elasticity, the adjusted value of lifetime consumption $\theta_j^{c,adj}$ is estimated in the same manner as θ_j^c and θ_j^h . Since $w_k \ell_k = 0$ whenever $\ell_k = 0$, there are no selection issues due to unobserved earnings for those out of employment.

The certainty-equivalent (CE) consumption equation can also be generalized to include labor supply in its LHS. We define the CE as the steady flow of consumption that makes the household indifferent, holding labor supply constant. That is, we solve

$$\sum_{k=j}^J \beta^{k-j} \frac{(\bar{c}(a_j, s^j))^{1-\gamma}}{1 - \gamma} - E_j \left[\sum_{k=j}^J \beta^k v \frac{\ell_k^{1+\chi}}{1 + \chi} \right] = V_j(a_j, s^j), \tag{13}$$

where ℓ_k is the optimal labor supply policy associated with $V_j(a_j, s^j)$. Using the intratemporal optimality condition, and replacing the marginal utility $u'(c)$ with its analytical counterpart $c^{-\gamma}$, the term $-E_j[\sum_{k=j}^J \beta^k v \frac{\ell_k^{1+\chi}}{1+\chi}]$ in (13) can be conveniently rewritten

TABLE 10. Welfare decomposition and inequality trends after adjusting for labor supply ($\chi = 2$).

Year	Welfare Decomp.—Labor Adjusted				Var. of log \bar{c}_{ij}
	$\chi = 2$				
	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	
1983	–	–	–	–	0.125
1989	0.188	0.213	–0.022	0.002	0.146
1993	0.227	0.261	–0.032	0.006	0.150
1998	0.349	0.425	–0.060	0.007	0.174
2004	0.513	0.656	–0.084	–0.002	0.194
2010	0.529	0.716	–0.111	0.003	0.210
2016	0.675	0.875	–0.117	0.012	0.216

as

$$\begin{aligned}
 -E_j \left[\sum_{k=j}^J \beta^k v \frac{\ell_k^{1+\chi}}{1+\chi} \right] &= \frac{c_j^{-\gamma}}{1-\gamma} \times E_j \left[\sum_{k=j}^J \beta^k \frac{c_k^{-\gamma}}{c_j^{-\gamma}} \left(\frac{\gamma-1}{1+\chi} \right) w_k \ell_k \right] \\
 &= \frac{c_j^{-\gamma}}{1-\gamma} \left(\frac{\gamma-1}{1+\chi} \right) \theta_j^h.
 \end{aligned} \tag{14}$$

One can then recover the CE consumption $\bar{c}(a_j, s^j)$ by solving

$$\sum_{k=j}^J \beta^{k-j} \frac{(\bar{c}(a_j, s^j))^{1-\gamma}}{1-\gamma} = \frac{c_j^{-\gamma}}{1-\gamma} \left(\theta_j^{c,adj} - \frac{\gamma-1}{1+\chi} \theta_j^h \right). \tag{15}$$

Both $\theta_j^{c,adj}$ and θ_j^h are straightforward to compute. Table 10 reports the proportional variance of the CE measure \bar{c} in (15) alongside the welfare decomposition results based on a Frisch elasticity of 1/2. While CE dispersion grows slightly less between 1983 and 2016, and the welfare costs of inequality ω_{ine} are marginally smaller, accounting for variation in labor supply has no material effects on the baseline findings. Doubling the Frisch elasticity to 1, as shown in Table 11, further confirms the robustness of the benchmark analysis.

8. CONCLUSIONS

We suggest two alternative approaches to estimate certainty equivalent consumption (CE) measures of lifetime welfare for individuals and households. The resulting estimates rest on different representations of lifetime welfare, which we label the consumption representation and the permanent-income representation.

Given estimates of the distribution of CE consumption for 1983 through 2016, we quantify (i) the evolution of inequality in lifetime welfare, and (ii) the offsetting effects of higher levels and dispersion of consumption expenditures over the sample period. By either CE measure, inequality of welfare is substantially lower than that of in-

TABLE 11. Welfare decomposition and inequality trends after adjusting for labor supply ($\chi = 1$).

Year	Welfare Decomp.—Labor Adjusted				Var. of log \bar{c}_{ij}
	$\chi = 1$				
	ω_{tot}	ω_{lev}	ω_{ine}	ω_{unc}	
1983	–	–	–	–	0.118
1989	0.188	0.213	–0.022	0.001	0.139
1993	0.230	0.261	–0.029	0.005	0.141
1998	0.354	0.425	–0.057	0.007	0.164
2004	0.534	0.656	–0.073	–0.001	0.179
2010	0.542	0.716	–0.103	0.003	0.196
2016	0.685	0.875	–0.111	0.011	0.204

come or consumption; however, welfare dispersion grows faster than inequality in income or consumption. The primary source of these discrepancies is the value of future earnings, that we call human wealth. The latter is accounted for in our CE measures, but is not fully reflected in the cross-sectional distributions of current earnings and consumption. We document how human and financial wealth shape the evolution of welfare inequality over time, with net worth playing a more influential role in recent decades.

CE consumption has the benefit of providing a simple statistic to jointly evaluate the aggregate impact of the rising average and dispersion of household expenditures. Under a utilitarian welfare criterion, we assess both aggregate and distributional changes observed over the past few decades. Between 1983 and 2016 overall welfare increased by the equivalent of increasing the expenditure of all agents alive in 1983 by about 65%. This measure holds demographic composition constant as it was in 1983. By comparison, average consumption increased by 88% over the same period, implying that CE inequality reduced aggregate welfare by roughly a quarter. These welfare losses occurred mostly after the year 2000, offsetting the gains from much higher consumption and marginally better cross-sectional insurance.

APPENDIX A: LIFETIME BUDGET CONSTRAINT

To derive the risk-adjusted lifetime budget constraint in (3), we define state-dependent stochastic discount factors $m(s_{j+1}, s^j)$ (SDFs). The notation implies that any such factor applies to a state occurring at $j + 1$, conditional on the history up to age j . Transition probabilities $\pi(s_{j+1}|s^j)$ are subsumed into the SDFs, which can be written as

$$m(s_{j+1}, s^j) = \pi(s_{j+1}|s^j)\beta \frac{u'(c_{j+1})}{u'(c_j)}. \quad (16)$$

Next, we multiply each possible realization of the age $j + 1$ budget constraint from (1) by the corresponding stochastic discount factor $m(s_{j+1}, s^j)$. Summing all of these terms together with the age j intertemporal budget constraint results in the following iden-

tity:

$$\begin{aligned}
 & \sum_{k=j}^{j+1} E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} c_k \right] \\
 &= \sum_{k=j}^{j+1} E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} y(s^k) \right] + a_j \\
 & \quad - E_j \left[\beta \frac{u'(c_{j+1})}{u'(c_j)} a_{j+1} (1+r) \left(1 - (1+r) \sum_{s_{k+1}} m(s_{k+1}, s^k) \right) \right] \\
 & \quad - E_j \left[\beta \frac{u'(c_{j+1})}{u'(c_j)} \frac{a_{j+2}}{1+r} \right].
 \end{aligned} \tag{17}$$

This summation can be extended to span the entire life cycle, resulting in

$$\begin{aligned}
 & \sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} c_k \right] \\
 &= \sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} y(s^k) \right] + a_j \\
 & \quad - \sum_{k=j}^{J-1} E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} a_{k+1} (1+r) \left(1 - (1+r) \sum_{s_{k+1}} m(s_{k+1}, s^k) \right) \right] \\
 & \quad - E_j \left[\beta \frac{u'(c_{j+1})}{u'(c_j)} \frac{a_{j+1}}{1+r} \right].
 \end{aligned} \tag{18}$$

Optimality implies that the term in the second line of this expression is

$$\left(1 - (1+r) \sum_{s_{k+1}} \pi(s_{k+1}|s^k) \beta \frac{u'(c_{k+1})}{u'(c_k)} \right) = \frac{\lambda(s^k)}{u'(c_k)},$$

where $\lambda(s^k) \geq 0$ is the Lagrange multiplier from the age k borrowing constraint, that is,

$$u'(c_k) = \beta(1+r)E_k[u'(c_{k+1})] + \lambda(s^k).$$

Optimality also implies $a_{J+1} = 0$. The risk-adjusted lifetime budget constraint can thus be written as

$$\sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} c_j \right] = \sum_{k=j}^J E_j \left[\beta^{k-j} \frac{u'(c_k)}{u'(c_j)} y(s^k) \right] + a_j - \sum_{k=j}^{J-1} E_j \left[\beta^{k-j} \frac{\lambda(s^k)}{u'(c_j)} \underline{a} \right].$$

The last term on the right-hand side is zero when households are not credit constrained. Under the working assumption that \underline{a} is the natural borrowing limit, which

is never binding with CRRA preferences, the term is also zero. Similarly, if one imposed the tight borrowing constraint $\underline{a} = 0$, the term would drop out from all calculations. Finally, if liquidity constraints did bind for $\underline{a} < 0$, the last term would be positive. However, in such cases, the actual amount of unsecured debt would be scaled down by the probability of hitting the constraint, reflected in the expectation operator $E_j[\cdot]$. This probability would be higher for values of \underline{a} close to zero and would go to zero as \underline{a} approaches the natural borrowing limit. In either case, the term would remain very small, if present at all, and the lifetime budget constraint would continue to hold as an approximation. Incidentally, it is worth emphasizing that the last term on the right-hand side, when positive, would affect only the permanent-income representation of lifetime welfare but would have no bearing on welfare measures based on the consumption representation. Positive realizations of that term would imply that the wealth representation delivers a lower bound of the true welfare and lies, on average, below welfare estimates based on the consumption representation. As we show in the empirical analysis, this is not true and the opposite tends to occur.

Rate of return heterogeneity The risk-adjusted lifetime constraint can accommodate ex ante heterogeneity in rates of return across households. The model is generalized by writing the period budget constraint as

$$c_j + \frac{a_{j+1}}{1+r(s^j)} = a_j + y(s^j).$$

This extension allows for general heterogeneity in both capital returns and earnings. Specifically, the return r in the second line of equation (18) is replaced by $r(s^k)$ so that the Lagrange multiplier becomes the solution to

$$u'(c_k) = \beta(1+r(s^k))E_k[u'(c_{k+1})] + \lambda(s^k).$$

The remainder of the derivation would then proceed as in the restricted case without rate of return heterogeneity.

APPENDIX B: ESTIMATION

This section describes the steps for the identification and estimation of human wealth θ_{ij}^h . We use an identical approach to recover the present value of lifetime consumption θ_{ij}^c , with the qualification that in the latter case the dividend functions are estimated with expenditure data c_{ij+1} in place of earnings y_{ij+1} .

B.0.1 Nonparametric identification of human wealth We rewrite the unknown recursive function equation in (9) as an integral equation after making two substitutions. First, define $\delta(j, \mathbf{z}, \mathbf{z}') = \mathbb{E}[\beta(u'_c/u_c)|j, \mathbf{z}, \mathbf{z}'] \times f_{\mathbf{Z}'|Z}^j(\mathbf{z}'|j, \mathbf{z})$, where $f_{\mathbf{Z}'|Z}^j$ is the age-specific conditional density of \mathbf{z}' . Each $\delta(j, \mathbf{z}, \mathbf{z}')$ can be described as an appropriately discounted density function for \mathbf{z}' at age j , for a given conditioning vector \mathbf{z} . It follows that the human wealth equation can be written as

$$\theta(j, \mathbf{z}) = y(j, \mathbf{z}) + \int \theta(j+1, \mathbf{z}')\delta(j, \mathbf{z}, \mathbf{z}') d\mathbf{z}'. \quad (19)$$

Next, we define two sets of functions for earnings, $Y(\mathbf{z}) = (y(1, \mathbf{z}), y(2, \mathbf{z}), \dots, y(J, \mathbf{z}))'$, and human wealth, $\Theta(\mathbf{z}) = (\theta(1, \mathbf{z}), \theta(2, \mathbf{z}), \dots, \theta(J, \mathbf{z}))'$. These vectors contain one function for each age, up to J which is an arbitrarily old age. Finally, we arrange the age-specific transition functions into a $J \times J$ matrix

$$\Delta(\mathbf{z}, \mathbf{z}') = \begin{pmatrix} 0 & \delta(1, \mathbf{z}, \mathbf{z}') & 0 & \dots & 0 \\ 0 & 0 & \delta(2, \mathbf{z}, \mathbf{z}') & & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & \delta(J-1, \mathbf{z}, \mathbf{z}') \\ & & & & 0 \end{pmatrix}. \tag{20}$$

This matrix conforms with $\Theta(\mathbf{z}')$ so that the following representation of the integral equation (19) holds:

$$\Theta(\mathbf{z}) = Y(\mathbf{z}) + \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}') d\mathbf{z}'. \tag{21}$$

We next define a linear operator B composed of a finite set of age-specific linear operators B_j . Each age-specific operator satisfies

$$(B_j\theta)(j+1, \mathbf{z}) = \int \delta(j, \mathbf{z}, \mathbf{z}')\theta(j+1, \mathbf{z}') d\mathbf{z}'. \tag{22}$$

Then the operator B is constructed from the age-specific B_j operators as follows:

$$B = \begin{pmatrix} 0 & B_1 & 0 & \dots & 0 \\ 0 & 0 & B_2 & & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & B_{J-1} \\ & & & & 0 \end{pmatrix}. \tag{23}$$

This ensures that B is a linear operator such that

$$(B\Theta)(\mathbf{z}) = \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}') d\mathbf{z}'. \tag{24}$$

Using this definition within equation (21), the function Θ is uniquely determined as $\Theta = (I - B)^{-1}Y$, provided the operator $I - B$ has a well-defined inverse. The invertibility of $I - B$ follows from the assumption that after age J the value of human wealth is zero, which leads to B being upper triangular with all zeros on the diagonal. The simple intuition for this identification result becomes apparent if one solves the pricing equation (21) recursively, starting from the last age in which human wealth has a nonzero value. At some old enough age J , the human wealth value next period is zero, which implies that human wealth in the current period is $y(J, z)$. The remaining human wealth functions can then be recovered by backward recursion.

B.1 Empirical implementation

We consider a sample $\{c_{ij}, c_{ij+1}, \mathbf{z}_{ij}, \mathbf{z}_{ij+1}, y_{ij+1}\}$ consisting of observations for consumption, individual characteristics and earnings. Index $i \in N_j$ denotes an element within the

set N_j of individuals who are observed at both age j and $j + 1$. If a person is observed for three subsequent waves of the data, that person contributes two observations to the sample, etc.

B.1.1 Estimation of human wealth We estimate the value of the expected future human wealth in (19) using the Nadaraya–Watson estimator:¹⁹

$$\hat{\theta}(j, \mathbf{z}) = y(j, \mathbf{z}) + \sum_{i=1}^{N_j} \hat{\theta}(j+1, \mathbf{z}'_{ij}) \hat{\beta} \frac{\hat{u}'(c_{ij+1})}{u'(c_{ij})} \xi_{ij}(\mathbf{z}). \quad (25)$$

Because we observe $y(j, \mathbf{z})$, the only obstacle to obtain an estimate of the current human wealth function $\hat{\theta}(j, \mathbf{z})$ is that the future function $\hat{\theta}(j+1, \mathbf{z}')$ is so far unknown. However, as it is clear from equation (28), the entire function $\hat{\theta}(j+1, \mathbf{z}')$ need not be known. Rather, one only needs to have estimates of its value at the subset of observed points \mathbf{z}_{ij+1} in order to recover the entire function $\hat{\theta}(j, \mathbf{z})$. Stacking all $\hat{\theta}(j+1, \mathbf{z}_{ij+1})$ into vectors $\tilde{\Theta}_{j+1}$, and similarly stacking the $y(j, \mathbf{z}_{ij})$ into vectors \tilde{Y}_j , we can rewrite (28) in compact form, evaluated at observed \mathbf{z}_{ij} values, as $\tilde{\Theta}_j = \tilde{Y}_j + \Gamma_j \tilde{\Theta}_{j+1}$. The elements of Γ_j are

$$[\Gamma_j]_{mi} = \hat{\beta} \frac{\hat{u}'(c_{ij+1})}{u'(c_{ij})} \xi_{ij}(\mathbf{z}_{mj}). \quad (26)$$

Each column of Γ_j includes the weighting function $\xi_{ij}(\mathbf{z})$ and stochastic discount factor of a given individual i . Moreover, each row of Γ_j is evaluated at the data vector \mathbf{z}_{mj} for each individual m who was of the appropriate age j .²⁰ Put differently, each weighting function is evaluated at many data points, each associated with an age j observation. Thus, the m in \mathbf{z}_{mj} could be any individual in the sample when they are of the correct age.

¹⁹The weighting functions $\xi_{ij}(\mathbf{z})$ are then constructed as

$$\xi_{ij}(\mathbf{z}) = \frac{K_{ij}^z(\mathbf{z})}{N_j \sum_{m=1}^{N_j} K_{mj}^z(\mathbf{z})},$$

where $K_{ij}^z(\mathbf{z})$ is a multivariate kernel function. Here, we follow Li and Racine (2007) by defining \mathbf{z}^c and \mathbf{z}^d to be the subvectors of continuous and discrete variables contained in \mathbf{z} . The multivariate kernel function for a given \mathbf{z}_{ij} can then be written as $K_{ij}^z(\mathbf{z}) = (\prod_{z_s \in \mathbf{z}^c} K^{h_s}(z_s - z_{s,ij})) \times (\prod_{z_s \in \mathbf{z}^d} \mathbf{1}_{\{z_s = z_{s,ij}\}})$. The first product includes univariate Gaussian kernel functions with bandwidth h_s . The second product includes indicator functions, which ensure the kernel has positive value for an observation with the corresponding values of all discrete variables, and zero otherwise. For example, this means that female data will have no influence on the conditional expectation for a male observation, and vice versa.

²⁰The matrix Γ_j has number of rows equal to the number of observations stacked in $\tilde{\Theta}_j$ and number of columns equal to the number of observations stacked in $\tilde{\Theta}_{j+1}$. For each column i , there is a corresponding age $j+1$ human wealth estimate contained in $\tilde{\Theta}_{j+1}$. For each row m , there is a corresponding age j human wealth estimate in $\tilde{\Theta}_j$. If the data are unbalanced, one may have different numbers of observations at each age. In this case, Γ_j will not be square and the lengths of $\tilde{\Theta}_j$ and $\tilde{\Theta}_{j+1}$ will differ. We structure our data so that an observation consists of pairs $\{\mathbf{z}_{ij}, \mathbf{z}_{ij+1}\}$. If an individual in the sample is observed over only two consecutive years, they contribute one observation to the sample.

We combine vectors $\tilde{\Theta}_j$ and \tilde{Y}_j into larger vectors $\tilde{\Theta}$ and \tilde{Y} .²¹ We also arrange the matrices Γ_j into a block matrix Γ with elements arranged on the off diagonal as in the matrix Δ in equation (20). Using this notation, the set of j -specific equations $\tilde{\Theta}_j = \tilde{Y}_j + \Gamma_j \tilde{\Theta}_{j+1}$ can be written compactly as $\tilde{\Theta} = \tilde{Y} + \Gamma \tilde{\Theta}$. Because $(I - \Gamma)$ is invertible,²² one can directly solve for $\tilde{\Theta} = (I - \Gamma)^{-1} \tilde{Y}$.

To obtain estimators of the complete functions $\theta(j, \mathbf{z})$, rather than at the observed data points only, we return to equation (28). Because the point estimates of $\hat{\theta}(j + 1, \mathbf{z}_{ij+1})$ are now available (they are the elements of $\tilde{\Theta}$), equation (28) can be evaluated at *any* point \mathbf{z} . Thus, the vector of estimators for the age-specific human wealth valuation functions $(\hat{\theta}(1, \mathbf{z}), \hat{\theta}(2, \mathbf{z}), \dots, \hat{\theta}(J, \mathbf{z}))'$ has now been obtained.

Biennial data in human wealth estimation For an observation drawn during a period of biennial sampling, equation (10) can be rewritten by iterating the valuation equation one year further into the future:

$$\hat{\theta}(j, \mathbf{z}) = y(j, \mathbf{z}) + \hat{g}(j, \mathbf{z}) + \sum_{i=1}^{N_j^2} \hat{\theta}(j + 2, \mathbf{z}_{ij+2}) \left(\prod_{k=j}^{j+1} \beta_k \right) \frac{\hat{u}'(c_{ij+2})}{\hat{u}'(c_{ij})} \xi_{ij}(\mathbf{z}). \tag{27}$$

N_j^2 is the set of observations of j year-old individuals in biennial data. The function $\hat{g}(j, \mathbf{z})$ are estimates of the conditional expectation of discounted earnings one year ahead, for a j year-old individual with current state vector \mathbf{z} . These estimates are computed using the Nadaraya–Watson estimator and data from the annual sample period:

$$\hat{g}(j, \mathbf{z}) = \sum_{i=1}^{N_j} \beta_j \frac{\hat{u}'(c_{ij+1})}{u'(c_{ij})} y(j + 1, \mathbf{z}_{ij+1}) \xi_{ij}(\mathbf{z}). \tag{28}$$

As before, we form vectors $\tilde{\Theta}$ and \tilde{Y} , as well as a matrix Γ such that $\tilde{\Theta} = \tilde{Y} + \Gamma \tilde{\Theta}$. Some elements of $\tilde{\Theta}$ and \tilde{Y} are based on annual observations using equation (28), and others are based on biennial observations using equation (27). Where \tilde{Y} is based on biennial data, its elements are $y(j, \mathbf{z}) + \hat{g}(j, \mathbf{z})$. The matrix Γ is somewhat more complicated because rows corresponding to biennial observations must conform with columns of $\tilde{\Theta}$ corresponding to values two years ahead. Thus, Γ now must have the form

$$\Gamma = \begin{pmatrix} 0 & \Gamma_1^1 & \Gamma_1^2 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & & 0 & \Gamma_{J-1}^1 \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}, \tag{29}$$

where Γ_j^1 and Γ_j^2 are constructed as explained in equation (26). The reason we now have two blocks in each row of Γ is to allow rows corresponding to annual observations to

²¹The larger vectors are defined as $\tilde{\Theta} = (\tilde{\Theta}'_1, \dots, \tilde{\Theta}'_{J-1}, \tilde{\Theta}'_J)'$ and $\tilde{Y} = (\tilde{Y}'_1, \dots, \tilde{Y}'_J)'$, where J is an arbitrarily old age by which all individuals have either died or retired.

²²Note that $(I - \Gamma)$ is upper triangular with ones on the leading diagonal so $\det(I - \Gamma) = 1$.

multiply $\tilde{\Theta}_{j+1}$, and rows corresponding to biennial observations to multiply $\tilde{\Theta}_{j+2}$. Rows of Γ_j^1 corresponding to annual observations will contain elements as in equation (26), whereas rows corresponding to biennial observations will consist of zeros. Zeroes will appear in the rows of Γ_j^2 wherever Γ_j^1 is nonzero. After constructing such a matrix Γ , we can solve for $\tilde{\Theta} = (I - \Gamma)^{-1} \tilde{Y}$ as before.

The last step is to construct an estimator for the general function $\hat{\theta}(j, \mathbf{z})$, once estimates have been recovered by computing $\tilde{\Theta}$ at the observed sample points. This requires a weighting of equations (28) and (27). We define numbers of annual and biennial observations $n^1 = \sum_{j=22}^J N_j$ and $n^2 = \sum_{j=22}^J N_j^2$. Using these counts, we form the estimator as

$$\begin{aligned} \hat{\theta}(j, \mathbf{z}) = & y(j, \mathbf{z}) + \frac{n^1}{n^1 + n^2} \left(\beta_j \sum_{i=1}^{N_j} \hat{\theta}(j+1, \mathbf{z}_{ij+1}) \frac{\hat{u}'(c_{ij+1})}{\hat{u}'(c_{ij})} \xi_{ij}(\mathbf{z}) \right) \\ & + \frac{n^2}{n^1 + n^2} \left(y(j, \mathbf{z}) + \hat{g}(j, \mathbf{z}) \right. \\ & \left. + \left(\prod_{k=j}^{j+1} \beta_k \right) \sum_{i=1}^{N_j^2} \hat{\theta}(j+2, \mathbf{z}_{ij+2}) \frac{\hat{u}'(c_{ij+2})}{\hat{u}'(c_{ij})} \xi_{ij}(\mathbf{z}) \right). \end{aligned} \quad (30)$$

Weighting in this way ensures that, if there are only a small number of biennial observations, these observations have a limited influence on the estimated functions.

APPENDIX C: WELFARE WITH ELASTIC LABOR SUPPLY

Given expenditures c_j and labor ℓ_j at age j , preferences are $U(c_j, \ell_j) = c_j^{1-\gamma}/(1-\gamma) - \nu \ell_j^{1+\chi}/(1+\chi)$, where $(1+\chi)$ is the Frisch elasticity. Writing earnings $y_j = w_j \ell_j$, isoelasticity implies that the present discounted value (PDV) of expected utility can be written as

$$\begin{aligned} E_j \sum_{k=j}^J \left[\beta^{k-j} \left(\frac{c_k^{1-\gamma}}{1-\gamma} - \frac{\nu \ell_k^{1+\chi}}{1+\chi} \right) \right] \\ = c_j^{-\gamma} E_j \left[\sum_{k=j}^J \beta^k \left(\frac{1}{(1-\gamma)} \frac{c_k^{-\gamma}}{c_j^{-\gamma}} c_k - \frac{\nu \ell_k^\chi}{(1+\chi)} \frac{1}{c_j^{-\gamma}} \ell_k \right) \right] \end{aligned} \quad (31)$$

$$= \frac{c_j^{-\gamma}}{1-\gamma} E_j \left[\sum_{k=j}^J \beta^k \left(\frac{c_k^{-\gamma}}{c_j^{-\gamma}} c_k - \frac{\nu \ell_k^\chi}{(1+\chi)} \frac{(1-\gamma)}{c_j^{-\gamma}} \ell_k \right) \right]. \quad (32)$$

Using the intratemporal optimality condition $\nu \ell_k^\chi = w_k c_k^{-\gamma}$, we can express the PDV of future utility as a function of consumption and earnings,

$$E_j \sum_{k=j}^J \left[\beta^{k-j} \left(\frac{c_k^{1-\gamma}}{1-\gamma} - \frac{\nu \ell_k^{1+\chi}}{1+\chi} \right) \right]$$

$$= \frac{c_j^{-\gamma}}{1-\gamma} E_j \left[\sum_{k=j}^J \beta^k \left(\frac{c_k^{-\gamma}}{c_j^{-\gamma}} c_k + \frac{v \ell_k^\chi (\gamma-1)}{c_j^{-\gamma} (1+\chi)} \ell_k \right) \right] \tag{33}$$

$$= \frac{c_j^{-\gamma}}{1-\gamma} E_j \left[\sum_{k=j}^J \beta^k \frac{c_k^{-\gamma}}{c_j^{-\gamma}} \left(c_k + \frac{(\gamma-1)}{(1+\chi)} w_k \ell_k \right) \right]. \tag{34}$$

The expected present value $E_j[\sum_{k=j}^J \beta^k \frac{c_k^{-\gamma}}{c_j^{-\gamma}} (c_k + \frac{(\gamma-1)}{(1+\chi)} w_k \ell_k)]$ is the $\theta_j^{c,adj}$ defined in equation (11).

APPENDIX D: SUPPLEMENTARY TABLES

TABLE 12. Summary statistics in the PSID for all sample years. Averages of age, expenditures, and earnings are for individuals. Averages of net worth are for households. Expenditures, earnings, and wealth statistics are in year 2016 equivalent dollars.

Year	Summary Statistics					
	N-individuals	N-Households	Avg. Age	Avg. Expend.	Avg. Earnings	Avg. Wealth
1983	5735	3461	43.3	20,598	20,920	242,470
1989	6110	3691	44.6	24,921	24,770	260,930
1993	7064	4311	44.8	26,015	25,886	276,640
1998	7152	4387	45.4	29,556	28,494	354,180
2004	7992	4908	45.8	34,745	31,652	428,730
2010	8237	5121	46.5	36,015	31,492	377,200
2016	8069	5026	47.3	39,706	32,867	418,910

TABLE 13. Variances of consumption equivalents, current expenditures, current earnings, and assets, between 1983 and 2016. The underlying consumption data are linearly scaled (rather than exponentially scaled as in the baseline analysis) to match aggregate expenditures. Results illustrate how proportional dispersion of certainty equivalents has increased by a similar magnitude as that of current expenditure itself, although the level of certainty equivalent inequality is generally lower.

Year	Variance of \ln^\dagger with linear expenditure scaling				
	(1) \bar{c}_{ij}^c	(2) \bar{c}_{ij}^p	(3) c_{ij}	(4) y_{ij}	(5) a_{ij}^\dagger
1983	0.1343	0.1274	0.1787	0.372	21.9
1989	0.148	0.1387	0.1869	0.403	27.8
1993	0.1551	0.1484	0.1934	0.44	36
1998	0.1714	0.1762	0.2142	0.415	38.9
2004	0.1919	0.2049	0.2429	0.49	45.1
2010	0.2055	0.2341	0.2563	0.547	67.6
2016	0.2062	0.2311	0.2651	0.489	60.1

[†]All variables in natural logarithms, except net worth a_{ij} , where the inverse hyperbolic sine transformation is used.

REFERENCES

- Abbott, Brant and Robin Brace (2020), “Has consumption inequality mirrored wealth inequality in the survey of consumer finances?” *Economics Letters*, 109–16. [1030]
- Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Gianluca Violante (2019), “Education policy and intergenerational transfers in equilibrium.” *Journal of Political Economy*, 127, 2569–2624. [1025]
- Abbott, Brant and Giovanni Gallipoli (2020), “Human capital inequality: Empirical evidence.” In *Oxford Research Encyclopedia of Economics and Finance*, Oxford University Press. [1024]
- Abbott, Brant and Giovanni Gallipoli (2022), “Supplement to ‘Permanent-income inequality.’” *Quantitative Economics Supplemental Material*, 13, <https://doi.org/10.3982/QE1851>. [1023]
- Aguiar, Mark and Mark Bilz (2015), “Has consumption inequality mirrored income inequality?” *American Economic Review*, 105 (9), 2725–2756. [1024, 1030]
- Athreya, Kartik, Felicia Ionescu, and Urvi Neelakantan (2021), “Stock market participation: The role of human capital.” *Review of Economic Dynamics*. [1024]
- Attanasio, Orazio, Erik Hurst, and Luigi Pistaferri (2014), “The evolution of income, consumption, and leisure inequality in the United States, 1980–2010.” In *Improving the Measurement of Consumer Expenditures*, 100–140, University of Chicago Press. [1024]
- Attanasio, Orazio and Luigi Pistaferri (2014), “Consumption inequality over the last half century: Some evidence using the new PSID consumption measure.” *The American Economic Review*, P P, 104 (5), 122–126. [1030]
- Attanasio, Orazio P. and Luigi Pistaferri (2016), “Consumption inequality.” *Journal of Economic Perspectives*, 30 (2), 3–28. [1024]
- Banks, James, Richard Blundell, and Sarah Tanner (1998), “Is there a retirement-savings puzzle?” *American Economic Review*, 88 (4), 769–788. [1033]
- Benabou, Roland (1996), “Heterogeneity, stratification, and growth: Macroeconomic implications of community structure and school finance.” *American Economic Review*, 86 (3), 584–609. [1025]
- Benabou, Roland (2002), “Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency?” *Econometrica*, 70 (2), 481–517. [1039]
- Blundell, Richard, Luigi Pistaferri, and Ian Preston (2008), “Consumption inequality and partial insurance.” *The American Economic Review*, 98 (5), 1887–1921. [1024]
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa (2021), “Discretizing unobserved heterogeneity.” *Econometrica*. [1046]
- Bowlus, Audra J. and Jean-Marc Robin (2004), “Twenty years of rising inequality in US lifetime labour income values.” *The Review of Economic Studies*, 71 (3), 709–742. [1024]

- Bricker, Jesse, Alice Henriques, Jacob Krimmel, and John Sabelhaus (2016), “Measuring income and wealth at the top using administrative and survey data.” *Brookings Papers on Economic Activity*. [1024, 1044]
- Conesa, Juan Carlos and Dirk Krueger (2006), “On the optimal progressivity of the income tax code.” *Journal of Monetary Economics*, 53 (7), 1425–1450. [1043]
- Cunha, Flavio, James Heckman, and Salvador Navarro (2005), “Separating uncertainty from heterogeneity in life cycle earnings.” *Oxford Economic Papers*, 57 (2), 191–261. [1046]
- Curtis, Chadwick, Julio Garrn, and Robert Lester (2021), “Working, consuming, and dying: Quantifying the diversity in the American experience.” Mimeo. [1024]
- De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo (2016), “The implications of richer earnings dynamics for consumption and wealth.” NBER working paper no. 21917. [1024]
- De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo (2020), “Nonlinear household earnings dynamics, self-insurance, and welfare.” *Journal of the European Economic Association*, 18 (2), 890–926. [1024]
- Floden, Martin (2001), “The effectiveness of government debt and transfers as insurance.” *Journal of Monetary Economics*, 48 (1), 81–108. [1025, 1039]
- Friedman, Milton (1957), “The permanent income hypothesis.” In *A Theory of the Consumption Function*, 20–37, Princeton University Press. [1024]
- Gabaix, Xavier and David Laibson (2017), *Myopia and Discounting*. National bureau of economic research. [1026]
- Gallipoli, Giovanni, Hamish Low, and Aruni Mitra (2020), *Consumption and Income Inequality Across Generations*. CEPR Discussion Paper No. DP15166. [1036]
- Geweke, John and Michael Keane (2000), “An empirical analysis of earnings dynamics among men in the PSID: 1968–1989.” *Journal of econometrics*, 96 (2), 293–356. [1024]
- Guvenen, Fatih, Greg Kaplan, Jae Song, and Justin Weidner (2017), *Lifetime Incomes in the United States Over Six Decades*. National Bureau of Economic Research. [1024]
- Guvenen, Fatih, Fatih Karahan, Serdar Ozkan, and Jae Song (2021), “What do data on millions of US workers reveal about lifecycle earnings dynamics?” *Econometrica*, 89 (5), 2303–2339. [1024]
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2010), “The macroeconomic implications of rising wage inequality in the US.” *Journal of Political Economy*, 118 (4), 681–722. [1024]
- Heathcote, Jonathan, Kjetil Storesletten, and Giovanni L. Violante (2020), “Optimal progressivity with age-dependent taxation.” *Journal of Public Economics*, 104074. [1043]
- Huggett, Mark and Greg Kaplan (2016), “How large is the stock component of human capital?” *Review of Economic Dynamics*, 22, 21–51. [1026, 1027]

- Kaymak, Barış and Markus Poschke (2016), “The evolution of wealth inequality over half a century: The role of taxes, transfers and technology.” *Journal of Monetary Economics*, 77, 1–25. [1024]
- Keane, Michael and Kenneth Wolpin (1997), “The career decisions of Young men.” *Journal of Political Economy*, 105 (3), 473–522. [1024]
- Kopczuk, Wojciech and Emmanuel Saez (2004), “Top wealth shares in the United States, 1916–2000: Evidence from estate tax returns.” *National Tax Journal*, 57 (2), 445–487. [1024]
- Krueger, Dirk and Fabrizio Perri (2006), “Does income inequality lead to consumption inequality? Evidence and theory.” *The Review of Economic Studies*, 73 (1), 163–193. [1023, 1024]
- Kuhn, Moritz and Jose-Victor Ríos-Rull (2016), “2013 update on the US earnings, income, and wealth distributional facts: A view from macroeconomics.” *Quarterly Review*, 1–75. [1024]
- Li, Qi and Jeffrey Scott Racine (2007), *Nonparametric Econometrics: Theory and Practice*. Princeton University Press. [1054]
- Lillard, Lee A. (1977), “Inequality: Earnings vs. human wealth.” *The American Economic Review*, 67 (2), 42–53. [1024]
- Lucas, Robert E. (1978), “Asset prices in an exchange economy.” *Econometrica*, 46 (6), 1429–1445. [1027]
- Parker, Jonathan, Annette Vissing-Jorgensen, and Nicholas Ziebarth (2009), “Inequality in expenditure in the twentieth century.” presentation slides, NBER Summer Institute. [1030]
- Pijoan-Mas, Josep and José-Víctor Ríos-Rull (2014), “Heterogeneity in expected longevities.” *Demography*, 51 (6), 2075–2102. [1031]
- Piketty, Thomas and Emmanuel Saez (2006), “The evolution of top incomes: A historical and international perspectives.” *American Economic Review Papers and Proceedings*, 96 (2), 200–205. [1024]
- Saez, Emmanuel and Gabriel Zucman (2016), “Wealth inequality in the United States since 1913: Evidence from capitalized income tax data.” *The Quarterly Journal of Economics*, 131 (2), 519–578. [1024]
- Sanders, Carl and Christopher Taber (2012), “Life-cycle wage growth and heterogeneous human capital.” *Annu. Rev. Econ.*, 4 (1), 399–425. [1024]
- Storesletten, Kjetil, Christopher I. Telmer, and Amir Yaron (2004), “Consumption and risk sharing over the life cycle.” *Journal of Monetary Economics*, 51 (3), 609–633. [1023]

Co-editor Kjetil Storesletten handled this manuscript.

Manuscript received 22 February, 2021; final version accepted 21 December, 2021; available online 13 January, 2022.