HCEO WORKING PAPER SERIES

Working Paper

Permanent-Income Inequality[†]

Brant Abbott
Queen's University
abbottbrant@gmail.com

Giovanni Gallipoli
University of British Columbia
gallipol@mail.ubc.ca

THIS DRAFT: FEBRUARY 17, 2019

Abstract. We characterize the distribution of permanent-income and quantify the value of assets and human capital in lifetime wealth portfolios. We estimate the distribution of human wealth using nonparametric identification results that allow for state-dependent stochastic discounting and unobserved heterogeneity. The approach imposes no restrictions on income processes or utility. Accounting for the value of human capital delivers a different view of inequality: (i) in 2016 the top 10% share of permanent-income was 1/3 lower than the corresponding share of assets; (ii) however, since 1989, the top 10% share of permanent-income has grown much faster than the corresponding share of assets. Human wealth has a mitigating influence on inequality, but this effect has waned over time due to the growing importance of assets in lifetime wealth portfolios. We find that consumption expenditures are tightly linked to permanent-income; however, liquidity constraints can lead to substantial deviations below permanent-income.

JEL Classification: E2, E21, D31, I24
Keywords: Wealth, Human Capital, Permanent Income, Consumption, Inequality

1   Introduction

A primary objective of inequality research is to understand the forces shaping differences in the economic wellbeing of individuals and households. Empirical research has made progress towards this goal by analyzing inequality of observable variables, primarily income and wealth.[1] However, a broader assessment of economic inequality would require that one also accounts for the heterogeneity associated with future earnings potential. This is apparent in the optimal redistribution branch of the literature where equalization of marginal utilities from consumption is often assumed to be the underlying policy goal, and optimal policies depend on the unobservable value of ex-ante expected future earnings.[2]

Yet-to-be realized earnings may constitute the most important determinant of economic wellbeing for many households. A young person with a steeply increasing expected earnings profile may be better off than inferred by simply measuring their current net worth or income. In this paper we estimate pecuniary measures that reflect the values of both human capital and asset wealth. We use these measures to (i) characterize the evolution of households' lifetime wealth portfolios over the life-cycle; (ii) identify trends in the concentration of lifetime wealth since 1989; and (iii) document the non-linear relationship between household consumption expenditures and permanent-income. At the heart of our analysis are nonparametric estimates of the value to individuals of their yet-to-be realized earnings, which we refer to as their human wealth.

The point that one should account for the value of lifetime earnings has long been recognized in the empirical literature.[3] We develop an empirical approach to perform present value calculations for future earning flows and recover human wealth. Our approach features state-dependent stochastic discounting to account for the ease with which consumption can be shifted across time

---

[1]See Saez and Kopczuk (2004), Piketty and Saez (2006), Saez and Zucman (2014), Bricker, Henriques, Krimmel *et al.* (2016), Kaymak and Poschke (2016) and Kuhn and Rios-Rull (2016). The literature on the distribution of wages and earnings documents widening inequality in the working population (see, among others, Levy and Murnane, 1992; Gottschalk, Moffitt, Katz *et al.*, 1994; Goldin and Katz, 2007; Autor, Katz, and Kearney, 2008). Studies of the wealth distribution focus on asset wealth held by the wider population, including the unemployed and those who do not participate in the labor market (Saez and Zucman, 2014; Bricker, Henriques, Krimmel *et al.*, 2016). Gabaix, Lasry, Lions *et al.* (2016) show that a rich model of income dynamics, featuring type-dependence, is necessary to explain the fast rise of top income inequality. Work by De Nardi, Fella, and Paz-Pardo (2016, 2019) illustrates how rich income processes (as in Guvenen, Karahan, Ozkan *et al.*, 2016) may be reflected in the equilibrium distribution of wealth. Interestingly, Athreya, Ionescu, and Neelakantan (2015; 2017) show that human capital investments affect financial portfolios in the cross-section and throughout the life-cycle.

[2]This is an extensive literature. The New Dynamic Public Finance part of the literature is surveyed by Golosov, Tsyvinski, Werning *et al.* (2006) and Kocherlakota (2010). Examples from the Ramsey planning literature include Conesa, Kitao, and Krueger (2009), Davila, Hong, Krusell *et al.* (2012), Heathcote, Storesletten, and Violante (2017) and Krueger and Ludwig (2018).

[3]See Lillard (1977). More recent contributions include Bowlus and Robin (2004) and Guvenen, Kaplan, Song *et al.* (2017). For an overview of related work, see Sanders and Taber (2012) and Abbott and Gallipoli (2018).

periods, as well as uncertainty about future earnings and consumption.[4] Being constrained by a credit limit, or facing a great deal of risk, reduces a household's valuation of their future earnings; in Section 4.2 we use this feature of the valuation problem to gauge the extent of lifetime wealth losses due to incomplete markets. Our approach allows for unobserved type heterogeneity, imposes no parametric restrictions on utility or income processes and, by design, delivers tractable estimates of the conditional distributions of future earnings.

By combining human wealth estimates with data on net worth we estimate lifetime wealth, defined as the sum of human wealth and asset wealth. From this we recover estimates of permanent-income, which is the (age-adjusted) annuity value of lifetime wealth. The latter statistic is reminiscent of permanent-income as defined by Friedman (1957), with the obvious difference that in Friedman's model human wealth is the risk-free present value of expected future earnings. We find that in 2016 the top 10% share of permanent-income was roughly 1/3 lower than the corresponding share of asset wealth. However, between 1989 and 2016, permanent-income concentration grew more than 1% per year, which is much faster than the 0.5% growth rate of asset wealth concentration. We infer that while human capital has had a mitigating influence on the level of overall inequality, this mitigating influence appears to be declining over time.

To obtain our estimates we combine data from the Panel Study of Income Dynamics (PSID) and the Survey of Consumer Finances (SCF). The PSID is useful for its panel data on earnings and consumption, which are required for identification of nonparametric human wealth valuation functions. We then apply these functions to SCF data, where the resulting estimates of human wealth can be combined with observed net worth. This allows one to obtain more accurate estimates of lifetime wealth and permanent-income. We do not make assumptions about the processes that generate risk in the labor market, and any aggregate risk present in the data is accounted for in our estimates.

As mentioned above, our estimates of human wealth account for state-dependent discount factors and for changes in marginal utility. Rather than assuming specific functional forms we estimate stochastic discount factors nonparametrically. This dispenses with several restrictions and lets data guide the choice of utility function in a flexible way.[5] Nonparametric identification of the marginal utility function is achieved by using and extending key results in Escanciano, Hoderlein, Lewbel *et al.* (2016). This involves writing the intertemporal Euler equation in such a manner that the

---

[4]Huggett and Kaplan (2016) convincingly argue that the true value of human capital is lower than what would be implied by discounting future earnings at the risk-free interest rate, an approach commonly advocated because of its simplicity (see Becker, 1975; Jorgenson and Fraumeni, 1989; R. Haveman and Schwabish, 2003). Mechanically discounting income flows does not account for state-dependent valuations of future earnings or other possible forms of heterogeneous discounting (Gabaix and Laibson, 2017).

[5]In the robustness Section 7.4 we replicate the analysis under alternative parametric utility specifications and compare results to our non-parametric benchmark.

estimated marginal utility function is the solution of a homogeneous Fredholm integral equation of the second kind, as discussed in Section 3. Given identification of the stochastic discount factor, human wealth then depends on an integral over its possible future values multiplied by the realizations of the stochastic discount factor.

Having obtained a marginal utility function, the estimated human wealth valuation equation turns out to be the solution of an *inhomogeneous* Fredholm integral equation of the second kind. Thus, we must extend the results in Escanciano, Hoderlein, Lewbel *et al.* (2016) in order to prove nonparametric identification of human wealth.

A separate issue arises from the fact that only one realization of the future state of the world is observed for each person and time-period in our sample. Hence we do not observe the entire distribution of possible future outcomes, on which an individual's human wealth depends. We address this data limitation by estimating the distribution of possible outcomes using those observed for individuals who are, in a way made clear later, ex-ante the same. This approach works under an identification assumption, which we refer to as conditional equivalence of expectations. This assumption states that individuals who are ex-ante equivalent, in terms of individual characteristics and the aggregate state, face the same distribution of ex-post outcomes. Our implementation allows for the distribution of ex-post outcomes to vary with both observable characteristics and unobservable types. Unobservable heterogeneity is potentially very important in this situation because certain forms, such as heterogeneous income profiles, could lead to differences in the distributions of ex-post outcomes even if individuals have identical ex-ante observable characteristics. To identify unobservable types we adapt the method developed by Bonhomme, Lamadon, and Manresa (2017) in such a way that the number of unobservable types is chosen to reflect the degree of ex-ante heterogeneity in the sample. Inclusion of these types in the conditioning set assuages our concern that unobserved differences in human wealth may lead to underestimates of the degree of inequality. When we test our assumption, we find no evidence of residual prior information about future outcomes after accounting for these unobserved types.

Our econometric work allows us to analyze heterogeneity in human wealth, lifetime wealth and permanent-income. In Section 4 we report findings about the life-cycle evolution of these variables, and explore the implications of these patterns for the welfare costs of incomplete markets. Section 5 overviews cross-sectional facts that can be immediately related to existing studies of inequality based on observed income or net worth. This approach allows one to ask questions about the evolution of inequality based on measures of household wealth spanning their entire life-cycle, despite the fact that both lifetime wealth and permanent-income are ex-ante magnitudes that cannot be directly observed. By their nature, these theoretical constructs are identified through a set of structural assumptions, hence the usefulness of our estimates is limited by the plausibility of those assumptions. Our use of nonparametric methods ensures that only the low level assumptions of the

theory, such as utility maximization, are used to identify the value of human wealth, rather than higher level assumptions, such as specific utility functional forms or wage generating processes. As such, we make the assumptions underlying our estimates as plausible as possible, while still maintaining comparability between our analysis and existing studies of inequality in observable variables.

In Section 6 we examine the mechanics underlying the large shifts that occurred in the cross-sectional distribution of lifetime wealth over the past four decades and we highlight the importance of the growing share of asset wealth in household portfolios. In the process, we document a sharp relationship between consumption expenditures and permanent-income. In particular, we estimate an elasticity of expenditures to permanent-income of 0.8. We also show that frictions, such as short-term liquidity constraints, can reduce expenditures to levels well below permanent-income.[6] Finally, in Section 7 we perform sensitivity and robustness checks, and consider several alternative parametric approximations of the utility function.

2 Human Wealth and its Valuation

Valuing intangible human capital is equivalent to pricing a non-traded asset that pays dividends equal to an individual's earnings. To quantify the value of this asset to an individual, we define the human wealth of individual i at time t as the idiosyncratic valuation θ_{it} of the uncertain stream of earnings that their human capital will deliver in the future. Theorem 1 in Huggett and Kaplan (2012) provides the theoretical underpinning for the use of the 'non-traded asset' approach, originally suggested by Lucas (1978), to value human capital. We posit a simple model where a general risk process leads to current earnings decisions y_{it}. The idiosyncratic valuation of the asset paying this uncertain flow of income follows a standard pricing formula:

$$\theta_{it} = \mathbb{E}_{it} \left[\beta \frac{u_c(c_{it+1}, \mathbf{v}_{it+1})}{u_c(c_{it}, \mathbf{v}_{it})} \left(y_{it+1} + \theta_{it+1} \right) \right], \tag{1}$$

where c_{it} is current consumption of agent i and u_c denotes marginal utility. The vector $\mathbf{v}_{it}$ contains variables that may affect the marginal utility from current consumption, such as leisure or past consumption. The expectation $\mathbb{E}_{it}$ is allowed to vary across individuals and time periods, capturing the fact that current realizations of aggregate and idiosyncratic shocks, as well as individual heterogeneity, may affect expectations.

Equation 1 delivers a general expression for the valuation of returns to human capital. This

---

[6]Non-linearities in expenditure decisions are relevant for the ongoing debate on the macroeconomic implications of household consumption responses to income changes (see, among recent studies, Parker, Souleles, Johnson *et al.*, 2013; Jappelli and Pistaferri, 2014; Carroll, Slacalek, and Tokuoka, 2014; Kaplan and Violante, 2014; Kaplan, Violante, and Weidner, 2014; Fuster, Kaplan, and Zafar, 2018; Straub, 2018; Bayer, Lütticke, Pham-Do *et al.*, 2019).

expression holds across a wide variety of settings and can be derived from alternative life cycle models of earnings (see Sanders and Taber, 2012). In Appendix A we explicitly derive the pricing equation 1 from a rich heterogeneous agents model that encompasses various standard models. For example, if labor supply responds to wages, then y_{it+1} in equation 1 accounts for the optimal labor supply rule, which is a function of underlying wage shocks. If wages are themselves a function of past human capital investments or labor supply, then y_{it+1} in equation 1 can be replaced by the product of optimal labor supply and endogenous wages. In this case both wages and labor supply are allowed to be functions of the general underlying risk process and of ex-ante heterogeneity: assuming that choices are optimal, the equation reverts to an expectation over state-dependent earnings realizations. In Appendix A we show that equation 1 still holds if we introduce endogenous marriage and divorce choices to the earnings model. This turns out to be useful when estimating the valuation equation using household-level data.

The Role of Borrowing Constraints: As discussed in Appendix A.3, our approach allows for the possibility that agents face binding borrowing constraints in their financial portfolio. However, even agents who are borrowing constrained in terms of different financial assets have well-defined human wealth as in equation 1. There is an important and intuitive reason for this feature of the model. Our exercise recovers a valuation at which agents would *choose* not to trade away their human capital. More specifically, our method views the quantity of human capital one possesses as a choice, but sets the price of that individual's human capital at a level that ensures they neither want to sell or buy a marginal unit. For individuals that are borrowing constrained, this price will tend to be lower than for similar unconstrained individuals. This is because constrained individuals would benefit from selling human capital as this would allow them to increase current consumption, which they value more than future consumption. Therefore, this lower price is exactly what we want to recover because borrowing constrained individuals have lower valuations of their future earnings than unconstrained individuals. Indeed, future earnings are worth less to individuals who cannot access them in advance, and the way we have structured our exercise allows us to explicitly quantify this effect.

3 Estimating Human Wealth

Our approach to the estimation of human wealth features two sequential steps. In the first step we apply the methods developed in Escanciano, Hoderlein, Lewbel *et al.* (2016) to recover non-parametric estimates of marginal utility functions, as well as an estimate of the deterministic time discount factor (β). These are then used in a second step to obtain nonparametric estimates of human wealth. We overview both steps in detail, even though only the identification results for the second step are novel. The careful description of the first step greatly aids in understanding our

identification and estimation approach in the second step.[7]

3.1 Identification

3.1.1 Nonparametric Marginal Utility Function Identification

It is helpful to use compact notation $q = (c, \mathbf{v})$ and $q' = (c', \mathbf{v}')$ to represent the current and future choices of an arbitrary individual. The consumption decision of an individual who is not at a corner solution is described by the following intertemporal Euler equation:

$$u_c(q) = \beta \mathbb{E}\left[u_c(q')R'|q\right]. \tag{2}$$

This condition is written for the return R' on an arbitrary asset; R' could be the return on any asset traded by the individual. Conditioning on current choices q is equivalent to conditioning on the entire information set because all relevant information is acted upon and reflected in these decisions.[8]

We begin by rewriting equation (2) in a form that replaces the expectation operator with the associated integral over the space of q'. In this integral the future marginal utilities are weighted by a factor corresponding to the product of (i) the conditional expectation of future rates of return and (ii) the Markov (transition) kernel describing transitions from q to q'. The notation we use for this weighting factor is $\psi(q, q') = \mathbb{E}\left[R'|q, q'\right] \times f(q'|q)$, where f is the conditional density of q'. The Euler equation (2) can be represented as

$$u_c(q) - \beta \int u_c(q')\psi(q, q')dq' = 0. \tag{3}$$

As explained by Escanciano, Hoderlein, Lewbel *et al.*, this is a homogeneous Fredholm integral equation of the second kind. The solution for $u_c(q)$ given β is well known. However, in our case both $u_c(q)$ and β must be determined, which leads to a question of identification.

Finite Support Case: Identification is easiest to understand if we restrict ourselves to the case in which the space of q_i is a finite number M of data points. Formally, the support is $q \in \left\{q^1, q^2, \ldots, q^M\right\}$. Under this assumption we can rewrite the Euler equation (3) at any current

---

[7]We also outline a new procedure to accommodate the use of biennial data in the first step.

[8]From the point of view of an econometrician, the right-hand-side of Euler equation (2) is a function of consumption and leisure choices that depend on the (yet unknown) future state of the world Ω'. That is, one could write (2) as

$$u_c(q(\Omega)) = \beta \mathbb{E}\left[u_c(q'(\Omega'))R'(\Omega')|\Omega\right].$$

For notational simplicity we omit the (Ω, Ω').

choice vector q^k as

$$u_c(q^k) - \beta \sum_{m=1}^{M} u_c(q^m)\psi_d(q^k, q^m) = 0, \tag{4}$$

where ψ_d is a discrete analogue of the transition function ψ. Rather than solving a complicated integral equation, identification in this finite example requires solving a linear system. Writing equation (4) in matrix notation, this entails solving

$$(I - \beta\Psi)U_c = 0, \tag{5}$$

where $U_c = \left(u_c(q^1), u_c(q^2), \ldots, u_c(q^M)\right)'$, and Ψ is a $M \times M$ matrix, with $\Psi_{km} = \psi_d(q^k, q^m)$.

This system has a nontrivial solution with $U_c \gg 0$ only if $\det(I - \beta\Psi) = 0$, which is true if β^{-1} is an eigenvalue of Ψ. In such cases the solution for U_c will depend on the eigenvector of Ψ associated with the eigenvalue β^{-1}. Thus, β is identified as the inverse of any real eigenvalue of Ψ such that $\beta \in (0, 1)$, and U_c is identified as the solution of the homogeneous system for the associated eigenvector. In general, Ψ may have multiple eigenvalues larger than unity, thus only set identification is achieved in the finite support case.[9]

General Case: Proof of identification in the general case where q has a continuous support requires functional analysis, but is reminiscent of the finite support case above. One first defines a linear operator A that, when applied to the unknown function $u_c(q)$, results in

$$(Au_c)(q) = \beta \int u_c(q')\psi(q, q')dq'. \tag{6}$$

This definition and equation 3 together imply that $u_c = \beta Au_c$. In the case that u_c and Au_c are positive valued (marginal utility is positive) and A is a compact operator, a solution for u_c exists only if $\beta = 1/\rho(A)$, where $\rho(A)$ is the largest real eigenvalue (spectral radius) of the operator A.[10] Therefore, if these assumptions are maintained, a unique value of β and a unique function u_c solve equation (3) and point identification is achieved.

3.1.2 Nonparametric Human Wealth Identification

We now turn to the question of nonparametric identification of θ_{it} in equation (1), taking β and the marginal utility function as given. We introduce a vector $\mathbf{z}$ containing variables that summarize an individual's information set. Unlike estimation of the marginal utility function, we now also consider individuals who may be credit constrained. Therefore, current consumption and leisure

---

[9]It is worth noting that Ψ is not simply a transition matrix (whose largest eigenvalue would be 1), but rather a transition matrix multiplied (element-wise) by expected asset returns $\mathbb{E}[R'|q, q']$.

[10]In the infinite dimensional case a linear compact positive operator has one positive eigenvector and its corresponding eigenvalue is equal to the spectral radius of the operator. Hence, we have uniqueness in this case.

may not fully summarize each individual's information set, posing an identification problem for expectations. This motivates the following assumption, which says that the general conditioning on individual and time period, an (i, t) pair, is equivalent to conditioning on a vector $\mathbf{z}_{it}$ of individual characteristics and aggregate states, as well as the individual's age j_{it}:

Definition (Conditional Equivalence of Expectations): *Expectations are conditionally equivalent with respect to the vector* $\mathbf{z}$ *if for any individual* i *and time period* t

$$\mathbb{E}_{it}\left[\beta\frac{u_c(q')}{u_c(q)}\left(y' + \theta'\right)\right] = \mathbb{E}\left[\beta\frac{u_c(q')}{u_c(q)}\left(y' + \theta'\right)\Big|\mathbf{z} = \mathbf{z}_{it}, j = j_{it}\right].$$

Conditional equivalence of expectations holds if the vector $\mathbf{z}_{it}$ is sufficient to span the current information set of any individual i at time t, given their age j_{it}. $\mathbf{z}_{it}$ includes both individual and aggregate information. In Appendix D we directly examine the empirical content of this assumption, and report evidence that $\mathbf{z}_{it}$ does a good job of capturing all relevant information, as reflected in observed choices and in their predictive power on later life outcomes.

If conditional equivalence of expectations holds, one can rewrite the human wealth equation (1) with θ_{it} replaced by a function $\theta(j, \mathbf{z})$:

$$\theta(j, \mathbf{z}) = \mathbb{E}\left[\beta\frac{u_c(q')}{u_c(q)}\left(y' + \theta(j+1, \mathbf{z}')\right)\Big|\mathbf{z}\right]. \tag{7}$$

This is a functional equation, similar to the Euler equation analyzed above.

We rewrite equation (7) as an integral equation after making two substitutions. First, define $\delta(j, \mathbf{z}, \mathbf{z}') = \mathbb{E}[\beta(u_c'/u_c)|j, \mathbf{z}, \mathbf{z}'] \times f^j_{Z'|Z}(\mathbf{z}'|j, \mathbf{z})$, where $f^j_{Z'|Z}$ 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 given conditioning vector $\mathbf{z}$. Second, we define $g(j, \mathbf{z}) = \mathbb{E}[\beta(u_c'/u_c)y'|j, \mathbf{z}]$, which subsumes the expected discounted value of the human wealth dividend. It follows that the human wealth equation can be written as

$$\theta(j, \mathbf{z}) = g(j, \mathbf{z}) + \int \theta(j+1, \mathbf{z}')\delta(j, \mathbf{z}, \mathbf{z}')d\mathbf{z}'. \tag{8}$$

Comparing the above functional equation to the integral form of the intertemporal Euler equation, the key difference is that eq. (8) is an *inhomogeneous* Fredholm integral equation of the second kind. The lack of homogeneity is due to the presence of the term $g(j, \mathbf{z})$, which introduces the age-dependent intercept in equation (8).

One can provide conditions for a unique solution of equation (8) by exploiting the deterministic nature of age transitions. We begin by defining the vector-valued functions $G(\mathbf{z}) = (g(1, \mathbf{z}), g(2, \mathbf{z}), \ldots, g(J-1, \mathbf{z}), 0)'$ and $\Theta(\mathbf{z}) = (\theta(1, \mathbf{z}), \theta(2, \mathbf{z}), \ldots, \theta(J-1, \mathbf{z}), \theta(J, \mathbf{z}))'$, where

J is an arbitrarily old age at which earnings are zero. Furthermore, we arrange the age-specific transition functions into a $J \times J$ matrix

$$\Delta(\mathbf{z}, \mathbf{z}') = \begin{pmatrix} 0 & \delta(1, \mathbf{z}, \mathbf{z}') & 0 & \cdots & 0 \\ 0 & 0 & \delta(2, \mathbf{z}, \mathbf{z}') & & 0 \\ \vdots & \vdots & \vdots & \ddots & \delta(J-1, \mathbf{z}, \mathbf{z}') \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix}. \tag{9}$$

This matrix conforms with $\Theta(\mathbf{z}')$ in a way that permits the following representation of the integral equation (8):

$$\Theta(\mathbf{z}) = G(\mathbf{z}) + \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}')d\mathbf{z}'. \tag{10}$$

Like in Escanciano, Hoderlein, Lewbel *et al.*, we next define a linear operator B composed of a finite set of age-specific linear operators B_j. Each age-specific operator satisfies

$$(B_j\theta)(j+1, \mathbf{z}) = \int \delta(j, \mathbf{z}, \mathbf{z}')\theta(j+1, \mathbf{z}')d\mathbf{z}'. \tag{11}$$

Then, the operator B is defined as follows:

$$B = \begin{pmatrix} 0 & B_1 & 0 & \cdots & 0 \\ 0 & 0 & B_2 & & 0 \\ \vdots & \vdots & \vdots & \ddots & B_{J-1} \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix}. \tag{12}$$

This ensures that B is a linear operator such that:

$$(B\Theta)(\mathbf{z}) = \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}')d\mathbf{z}'. \tag{13}$$

Using this definition within equation (10), the function Θ is uniquely determined to be $\Theta = (I - B)^{-1}G$, provided the operator $I - B$ has a well defined inverse. The invertibility of $I - B$ follows from the assumption that, for a large enough 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 (10) recursively, starting from the last age in which human wealth has a non-zero value. At some old enough age $J-1$ the human wealth value next period is zero, which implies that human wealth in the current period is $g(J-1, z)$. The remaining human wealth functions can then be recovered recursively.

9

Finite Example: When the support of $\mathbf{z}$ is restricted to be finite, so that $\mathbf{z} \in \{\mathbf{z}^1, \mathbf{z}^2, \ldots, \mathbf{z}^M\}$, proof of a unique solution for Θ amounts to proving a unique solution for a linear system. In such a case each operator $\delta(j, \mathbf{z}, \mathbf{z}')$ becomes a sub-matrix of an upper triangular hollow block matrix Δ. Applying this to our human wealth equation, we have $\Theta = G + \Delta\Theta$, the solution of which is $\Theta = (I - \Delta)^{-1}G$, if the inverse exists. Because Δ is hollow and upper triangular, all eigenvalues of $(I - \Delta)$ are unity, and therefore the inverse exists and Θ has a unique solution.

3.2 Empirical Implementation

We consider a sample $\{q_{it}, q'_{it}, \mathbf{z}_{it}, \mathbf{z}'_{it}, R'_{it}, y'_{it}, j_{it}, j'_{it}\}$. Index $i \in N$ denotes an element within the set N of observed individuals. Index $t \in \tau(i)$ identifies the periods within the set of years $\tau(i)$ for which the variables are observed for individual i. The set $\tau(i)$ includes sample years for which i is observed in both the current and subsequent sample periods: that is, both choices q_{it} and q'_{it} must be observed. For example, if a person is observed for three subsequent waves of the data that person contributes two observations to the sample. We let $\tau_o(i) \subset \tau(i)$ be the subset of observations in which individual i is at an interior solution for assets (that is, not borrowing constrained).

3.2.1 Estimation of the Marginal Utility Function

The first step in the estimation of the marginal utility function is to replace the linear operator A in equation (6) with the estimator

$$(\hat{A}u_c)(q) = \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} u_c(q'_{it}) R'_{it} \phi_{it}(q). \tag{14}$$

The weighting functions $\phi_{it}(q)$ deliver the locally weighted average (Nadaraya-Watson) estimator of the conditional expectation in equation (6).[11] Because the estimator $\hat{A}$ has a finite dimensional range (unlike the true A), $\hat{A}$ has a finite number of eigenvalues and eigenfunctions, which can be computed by solving a linear system. Hence any eigenfunction $\hat{u}_c(q)$ of $\hat{A}$ must be a linear combination of the functions $\phi_{it}(q)$, i.e. $\hat{u}_c(q) = \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} b_{it} \phi_{it}(q)$ for some set of coefficients b_{it}. Using this result, the empirical counterpart of the intertemporal Euler equation can be re-written as

$$\sum_{i=1}^{N} \sum_{t \in \tau_o(i)} b_{it} \phi_{it}(q) = \hat{\beta} \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} \left[\sum_{m=1}^{N} \sum_{s \in \tau_o(s)} b_{ms} \phi_{ms}(q'_{it}) \right] R'_{it} \phi_{it}(q). \tag{15}$$

The left side of the equation above simply replaces $u_c(q)$ with its estimator. The right side first uses equation (14) to replace the expectation in the Euler equation (2) with its estimator, and then

---

[11]Mechanically, we construct the weighting functions as $\phi_{it}(q) = \frac{K_{it}(q)}{\hat{f}(q)}$, where $\hat{f}(q) = \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} K_{it}(q)$, and $K_{it}(q) = K^H(q - q_{it})$. The function $K^H(\cdot)$ is a multivariate Gaussian kernel with bandwidth vector H.

also replaces $u_c(q'_{it})$ with its estimator (the part in square brackets). Straightforward algebra shows that a sufficient condition for the Euler equation above to have a solution is

$$b_{it} - \hat{\beta} \sum_{m=1}^{N} \sum_{s \in \tau_o(s)} b_{ms} \phi_{ms}(q'_{it}) R'_{it} = 0, \tag{16}$$

for every $i \in N$ and $t \in \tau_o(i)$. This can be rewritten in matrix form with Φ being a square matrix with elements $\Phi_{kl} = \phi_l(q'_k) R'_k$, and $\mathbf{b}$ being a vector containing the coefficients b_{it} appropriately concatenated. Thus the restrictions in equation (16) are equivalent to $(I - \hat{\beta}\Phi)\mathbf{b} = \vec{0}$.

Letting λ^* be the largest eigenvalue of Φ, and $\mathbf{b}^*$ the associated eigenvector, the estimators of β and $u_c(q)$ are respectively

$$\hat{\beta} = \frac{1}{\lambda^*}, \quad \hat{u}_c(q) = \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} b^*_{it} \phi_{it}(q). \tag{17}$$

With no loss of generality $\hat{u}_c(q)$ can be scaled to have a unit norm. In Appendix B we show how we incorporate biennial data into this framework.

3.2.2 Estimation of Human Wealth

Point estimates of marginal utility can be recovered for each individual choice observed in our sample by evaluating the function $\hat{u}_c(q)$ at q_{it}.[12] Given these point estimates, the next step is to construct an estimator for the age-specific human wealth valuation functions. We begin by estimating the value of the expected dividend function $g(j, \mathbf{z})$ in equation (8) using the Nadaraya-Watson estimator:

$$\hat{g}(j, \mathbf{z}) = \sum_{i=1}^{N} \sum_{t \in \tau_j(i)} \hat{\beta} \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} y'_{it} \gamma_{it}(\mathbf{z}). \tag{18}$$

To do this we need to link individuals of the same age across time periods. First, we define $\tau_j(i)$ as the singleton set of years in which individual i is j years old. With a slight abuse of notation, we further let $\tau_{j(it)}(m)$ be the singleton set of years in which another individual m was the same age as individual i in period t (that is, the year when m was $j = j(it)$ years old).[13]

---

[12]This can be done even if the particular person-year observation (i, t) refers to a credit constrained individual, hence not used in the estimation procedure described above.

[13]The weighting functions $\gamma_{it}(\mathbf{z})$ are then constructed as

$$\gamma_{it}(\mathbf{z}) = \frac{K^z_{it}(\mathbf{z})}{\sum_{m=1}^{N} \sum_{t \in \tau_{j(it)}(m)} K^z_{mt}(\mathbf{z})}, \tag{19}$$

where $K^z_{it}(\mathbf{z})$ is a multivariate kernel function. Here we follow Li and Racine (2007) by defining $\mathbf{z}^c$ and $\mathbf{z}^d$ to be the sub-vectors of continuous and discrete variables contained in $\mathbf{z}$. The multivariate kernel function for a given $\mathbf{z}_{it}$ can

Next, we form estimators of the $\theta(j, \mathbf{z})$ functions. We re-write equation (8) replacing all functions by their estimators:

$$\hat{\theta}(j, \mathbf{z}) = \hat{g}(j, \mathbf{z}) + \sum_{i=1}^{N} \sum_{t \in \tau_j(i)} \hat{\theta}(j + 1, \mathbf{z}'_{it}) \hat{\beta} \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} \gamma_{it}(\mathbf{z}). \tag{20}$$

Because we have an estimate of $\hat{g}(j, \mathbf{z})$, 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 (20), 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 $\mathbf{z}'_{it}$ in order to recover the entire function $\hat{\theta}(j, \mathbf{z})$. Stacking all $\hat{\theta}(j + 1, \mathbf{z}'_{it})$ into vectors $\tilde{\Theta}_{j+1}$, and similarly stacking the $\hat{g}(j, \mathbf{z}_{it})$ into vectors $\tilde{G}_j$, we can re-write equation (20) evaluated at observed $\mathbf{z}_{it}$ values in compact form as $\tilde{\Theta}_j = \tilde{G}_j + \Gamma_j \tilde{\Theta}_{j+1}$. The 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 $\tilde{\Theta}_{j+1}$. The elements of Γ_j are

$$[\Gamma_j]_{mi} = \hat{\beta} \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} \gamma_{it}(z_{ms}). \tag{21}$$

Each column of Γ_j includes the transition kernel and stochastic discount factor of a given individual i. For each such individual i there is a corresponding age $j + 1$ human wealth estimate contained in $\tilde{\Theta}_{j+1}$. However, each row of Γ_j is evaluated at the data vector z_{ms} of each individual m who was the appropriate age in year s.[14] For each such individual 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 combine vectors $\tilde{\Theta}_j$ and $\tilde{G}_j$ into larger vectors $\tilde{\Theta}$ and $\tilde{G}$.[15] 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 9. Using this notation the set of j-specific equations $\tilde{\Theta}_j = \tilde{G}_j + \Gamma_j \tilde{\Theta}_{j+1}$ can be written compactly as $\tilde{\Theta} = \tilde{G} + \Gamma \tilde{\Theta}$. Because $(I - \Gamma)$ is invertible,[16] one can directly solve for $\tilde{\Theta} = (I - \Gamma)^{-1} \tilde{G}$.

To obtain estimators of the complete functions $\theta(j, \mathbf{z})$, rather than at the observed data points

---

then be written as $K^z_{it}(\mathbf{z}) = (\prod_{z_s \in \mathbf{z}^c} K^{h_s}(z_s - z_{s,it})) \times (\prod_{z_s \in \mathbf{z}^d} \mathbf{1}_{\{z_s = z_{s,it}\}})$. 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.

[14] We structure our data so that an observation consists of pairs $\{\mathbf{z}_{it}, \mathbf{z}'_{it}\}$ and $\{j_{it}, j'_{it}\}$. If an individual in the sample is observed over multiple periods they may contribute multiple observations to the estimation. This effectively treats the evaluation of one's own future realizations, where the individual is one year older, as a different observation.

[15] The larger vectors are defined as $\tilde{\Theta} = (\tilde{\Theta}'_1, \ldots, \tilde{\Theta}'_{J-1}, \tilde{\Theta}'_J)'$ and $\tilde{G} = (\tilde{G}'_1, \ldots, \tilde{G}'_{J-1}, \vec{0}')'$, where J is an arbitrarily old age by which all individuals have either died or retired.

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

only, we return to equation (20). Because the point estimates of $\hat{\theta}(j + 1, \mathbf{z}'_{it})$ are now available (they are the elements of $\tilde{\Theta}$), equation (20) can be evaluated at *any* point $\mathbf{z}$. Thus, the vector of estimators for the age-specific human wealth valuation functions $(\hat{\theta}(1, \mathbf{z}), \hat{\theta}(2, \mathbf{z}), \ldots, \hat{\theta}(J, \mathbf{z}))'$ has now been obtained. Appendix B shows how we apply this approach to biennial data.

3.3 Data

In this section we describe the requirements for estimating the human wealth function $\theta(j, \mathbf{z})$. In our analysis we incorporate additional data from the SCF, but we delay discussion of that data until it becomes pertinent. To obtain an empirical counterpart of the marginal utility estimator in equation (15) we need panel data on consumption and leisure (we implement $\mathbf{v}$ as leisure), as well as historical asset returns and proxies for information available to individuals when making decisions. The sample must include observations recorded over a sufficiently long time interval so as to identify the aggregate risk component of the transition kernels.

Hence the basic data requirement for the estimation of marginal utilities, discount factor and human wealth values is a sample $\{q_{it}, q'_{it}, \mathbf{z}_{it}, \mathbf{z}'_{it}, R'_{it}, y'_{it}, j_{it}, j'_{it}\}_{i \in N, t \in \tau(i)}$, where each vector q denotes a pair of consumption and leisure choices; the vector $\mathbf{z}$ includes variables that span the information sets of the decision makers; R is a historical real return from deferred consumption; and j denotes age. It turns out that the Panel Study of Income Dynamics contains much of what we need. We use panel data from the PSID covering the years 1967-2016.

Construction of q_i and q'_i involves collecting earnings and consumption data. Labour earnings is always observed. However, a more complete set of consumption expenditures is observed at the household level only after 1997. Before that date only selected categories of consumption were recorded regularly.[17] For this reason we build on the approach of Attanasio and Pistaferri (2014) to impute household consumption expenditure in periods when information is incomplete. This method relies on the ever larger availability of consumption expenditures in the PSID post-1997. The procedure estimates a demand system to impute consumption to PSID families observed in years before 1997. There are five advantages to this approach: (i) it relies on information from a single data set, making variable linkages straightforward; (ii) one can test how closely trends in consumption inequality are replicated by the imputation procedure using within-sample verification for the period during which complete expenditure data are available; (iii) because the PSID stretches all the way back to the late 1960s, this procedure delivers the longest consumption panel database currently available for the US; (iv) average consumption per household can be scaled to replicate its historical evolution; (v) last but not least, expenditure categories in the PSID appear to match national income and product account (NIPA) counterparts reasonably well. In Section 7.5

---

[17]If one goes back all the way to 1967, only food expenditures were regularly measured.

we explore the sensitivity of our results to alternative expenditure measurement approaches.[18]

For asset returns, R_i', we utilize both S&P 500 stock market returns and one-year treasury constant maturity rates, adjusted for realized annual CPI inflation. As the survey becomes biennial after 1997, we switch to the two-year return data.[19] A very important aspect of the first step in our estimation (estimating marginal utility functions) is that only households that actually hold such assets should be included as observations. We use treasury returns for households that do not appear to be credit constrained. We consider home-owners who do not experience large increases in consumption (greater than 50%) as unconstrained. In addition, for unconstrained households that indicate ownership of equities, we also use stock returns to form a second Euler equation observation. We exclude all other observations from the estimation of the marginal utility functions; however, households excluded in this step are brought back into the sample in the second stage, as the human wealth estimator does not require households to be unconstrained.

Finally, to obtain an empirical counterpart of the human wealth estimator in equation (20) we use a set of conditioning variables that spans the information set available to agents. The vector $\mathbf{z}_i$ contains (i) observable individual characteristics, including gender, education, industry, occupation, marital status, number of dependent children, (ii) an aggregate state-variable which is current average log-earnings per capita (measured in the PSID sample), and (iii) unobserved types, as we discuss in the next subsection.

3.4   Unobserved Types

The data vector $\mathbf{z}_i$ includes an unobserved type, η_i, which we allow to vary along two dimensions of heterogeneity. The first dimension captures differences in life-cycle earning profiles, identified from variation in the growth rates of earnings. The second dimension subsumes unobserved differences in consumption insurance (possibly due to family background or other informal channels), which we measure through dispersion of consumption growth rates over the life cycle.

To estimate heterogeneous types we resort to a variation of the approach originally suggested by Bonhomme, Lamadon, and Manresa (2017) that is suitable in our setting. That is, we employ a k-medians grouping algorithm to separate life-cycle moments of 'informative' variables (i.e. mean earnings growth or standard deviation of consumption growth) into clusters, where cluster membership is a type. Cluster membership is then represented through categorical variables. The idea is that variation in income and consumption growth paths conveys information about, respectively, permanent heterogeneity in income and idiosyncratic access to consumption smoothing.

To test whether our grouping procedure does a good job of estimating unobserved heterogeneity, and to establish the number of types used to model each dimension of heterogeneity, we follow

---

[18]The household consumption measures are converted to individual allocations through equivalence scales.

[19]These time series are publicly available from FRED. We also experiment with real returns for other assets.

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. More generally, 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 or consumption volatility is anticipated by agents, then their long-term choices should be consistent with later life outcomes in those dimensions.

Following Cunha, Heckman, and Navarro, we illustrate this point using the decision to attend college. Let S_i denote the college decision of individual i, taking value one if the individual completes college and zero otherwise. To the extent that heterogeneity η_i affects earnings growth, one would expect that $\text{Cov}(S_i, \eta_i) \neq 0$. Given the relationship between unobserved types and economic outcomes (such as earnings and consumption), schooling choices should be related with the (ex-post) level of earnings growth, or with the idiosyncratic dispersion of consumption growth rates. 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 or consumption. This line of reasoning offers a natural way to test whether our grouping procedure identifies the relevant "type" variation.

If the grouping algorithm successfully captures the relevant heterogeneity, the type indicator should crowd out the statistical effect of earnings profiles (and, similarly, of consumption dispersion) on college status. We find that allowing for three types to represent earnings-profile heterogeneity is sufficient to remove any direct effect of earnings growth on the expectation of college completion. In the case of consumption smoothing heterogeneity, we only need two types for the conditional expectation of college completion to be independent of consumption growth dispersion. Having established the cardinality of the type sets, we also corroborate our clustering by verifying that adding further types does not result in significant drops in within-type variances.

4   Human Wealth over the Life-Cycle

4.1   Marginal Utility and Human Wealth Estimates

Using the PSID panel sample described in the previous section we recover various objects of interest. Non-parametric estimates of the marginal utility of consumption are plotted in Figure 1. Consistent with theory, marginal utility is highly non-linear at low expenditure levels and flattens out with high expenditures.

With the estimated marginal utility function in hand, we proceed to estimate the valuations of human capital of individuals in the PSID. The left panel of Figure 2 plots the average value of human wealth by age, and contrasts it with the value estimated using a constant risk-free discount

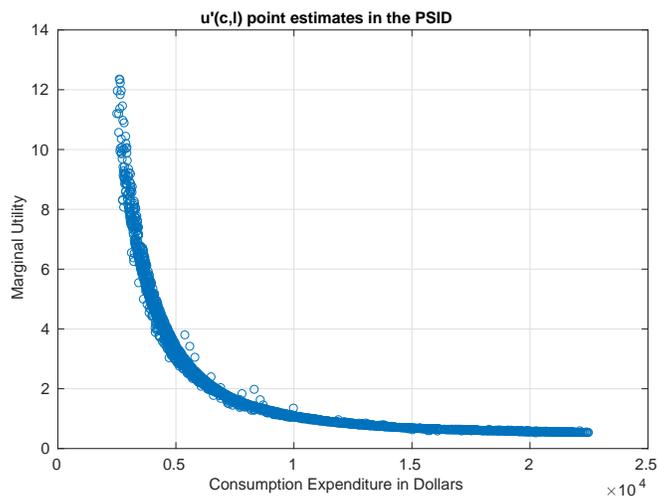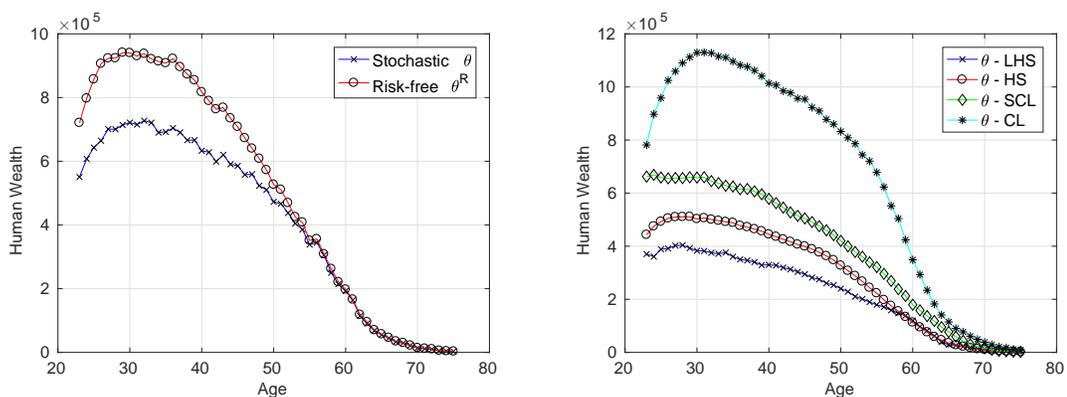Figure 1: Marginal utility as a function of consumption expenditures.

Figure 2: Average human wealth over the life cycle. Values in 2016 dollars. LHS denotes less than high school education; HS is high school degree only; SCL is some college; and CL is college degree or higher.

factor. The risk-free discount factor is set equal to the average of all realizations of the stochastic discount factors in our sample, and can be interpreted as a proxy of the theoretical price of a long-term risk free bond.[20] We denote risk-free discounted human wealth by θ_{it}^R, which is computed exactly like θ_{it} with the exception that stochastic discount factors are replaced by the estimated long-run risk free factor. Risk-free discounting results in a significant overestimation of human wealth, with the largest discrepancy around the time when human wealth peaks, which is consistent with simulation-based results in Huggett and Kaplan (2016). We return to the differences between risk-free and stochastically discounted human wealth in Section 4.2 below.

Human wealth exhibits a hump-shape, with a steep drop after age 50 as retirement approaches.

---

[20]The estimated risk free discount factor is the average stochastic discount factor across all sample members that do not appear to be liquidity constrained. Such observations are those employed to estimate the marginal utility function, which requires a sample of unconstrained households.

Under stochastic discounting the average value peaks at just over $700K. This average is based on a large sample of individuals, including some who do not work. Non-employment risk is explicitly accounted for in our estimation, which considers periods of null earnings as one of the possible outcomes of a worker's employment history. The right panel of Figure 2 plots the value of human wealth by education group. As one might expect, there are large differences in both scale and shape. At their peak, college graduates hold twice as much human wealth as high-school graduates, and three times as much human wealth as high-school drop-outs. Human wealth differentials become progressively smaller as retirement approaches because all pecuniary valuations of human capital converge to very low levels.

Two interesting observations can be made at this stage. First, younger households hold substantial amounts of wealth in illiquid human capital, which exposes them to significant risks. For example, health shocks might affect their labor supply and permanently reduce the present value of their lifetime wealth. Second, the early peak in human wealth suggests that the direct (that is, unmediated by assets) contribution of human wealth to overall inequality must occur at relatively younger ages, when human wealth still accounts for a large proportion of lifetime wealth portfolios. We revisit some of these issues below, in the context of our SCF sample.

We also document significant and persistent differences in human wealth across latent type clusters, as defined in Section 3.4. In Appendix C, Figure 16 plots the life-cycle evolution of average human wealth by unobserved type cluster (namely, the three earnings growth types and the two clusters identified using consumption dispersion). Unobserved heterogeneity in earning profiles is associated with a near doubling of human wealth at the peak. This gap suggests the presence of latent sources of within group heterogeneity among observationally similar households.

4.2 Human Wealth and Market Incompleteness

As shown in Figure 2, state-dependent discounting delivers lower valuations of future earning streams relative to constant discounting. The latter risk-free discounting would arise if markets were complete. The discrepancy between human wealth values obtained under alternative discounting approaches can therefore be used as a gauge of the implicit losses in expected human wealth due to market incompleteness. The right panels of Figure 3 are scatter plots of human wealth values under stochastic discounting versus their counterparts obtained using a common and constant discount rate (complete markets discounting). The left panels show the distribution of differences between human wealth estimates under the two discounting approaches.

The plots document that stochastic discounting results in significant downward deviations in human wealth values. The discrepancy is especially large for younger individuals, who are more likely to face short term credit constraints. As shown in the left panels of Figure 3, for individuals below age 35 the present discounted value of human wealth is on average $200K below what

Figure 3: Comparing values of human wealth under stochastic and risk-free discounting. Left-hand side: densities of the difference between human wealth under risk-free (θ^R) and stochastic (θ) discounting. Right-hand side: the horizontal axes measure human wealth under risk-free discounting (complete markets, w/CM), while the vertical axes measure the values of human wealth under stochastic discounting.
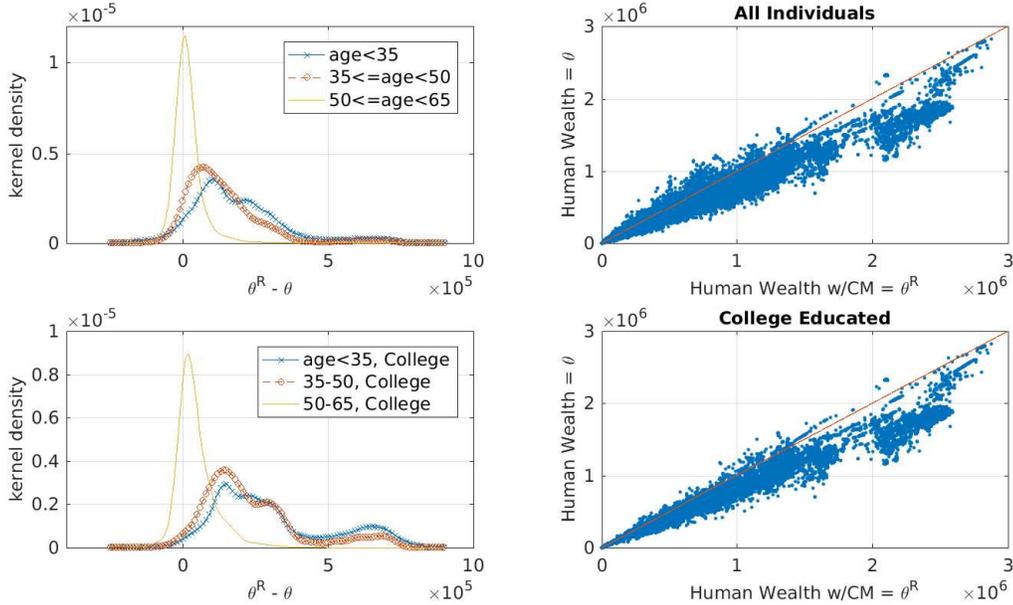
would be the case under risk-free discounting, suggesting that uninsurable life-cycle income risk and imperfect credit markets drive down the valuation of future income streams. Comparing the bottom and top right panels of Figure 3 indicates that the high human wealth individuals whose valuations are most affected by incomplete markets are highly educated and young.

As we show below, these differences in expected lifetime earnings have first-order effects on later life net worth and consumption outcomes. The average difference between risk-free and stochastically discounted human wealth in our PSID sample is about $104K. This value can be viewed as an average willingness to pay to eliminate future consumption uncertainty and thus provides an estimate of the welfare cost of market incompleteness. Figure 3 suggests that the distribution of differences in human wealth valuations is skewed, with the median willingness to pay to eliminate consumption risk being only about $54K. The observation that market incompleteness affects human capital values invites questions about the sources of discrepancy between stochastic and risk-free valuations. We explore these questions by carrying out simple counterfactuals to decompose the differences. There are several reasons why stochastic valuations tend to be smaller, as illustrated in the left panel of Figure 2, including credit constraints, age variation in consumption growth (and therefore discounting), and cross-sectional risk. In our first counterfactual we recalculate the risk-free discount factor by including the realized stochastic discount factors of individuals

Figure 4: Decomposition of Stochastic vs. Risk-free differences.
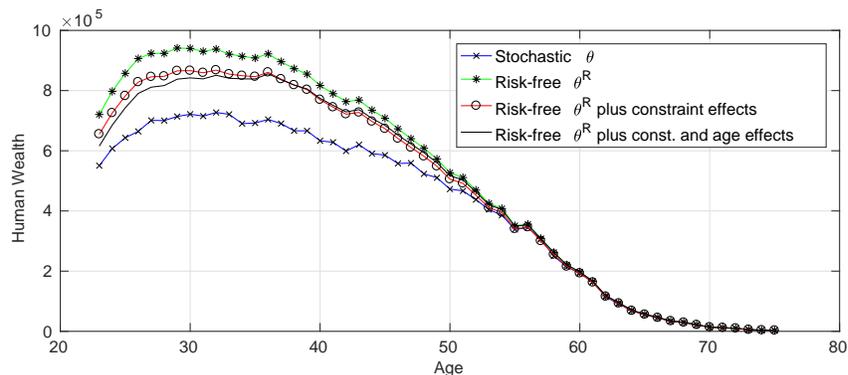
who appear to be credit constrained.[21] The resulting lifecycle mean human wealth is denoted in Figure 4 as 'θ^R plus constraint effects'. At age 30 roughly 1/3 of the difference between stochastic and risk-free valuations can be attributed to credit constraints. In our second counterfactual we additionally allow for the risk-free rate to vary with age. Younger households tend to exhibit growing consumption, which pushes down their discount factors; this effect is apparent in the profile denoted as 'θ^R plus const. and age effects', plotted in Figure 4. Counterfactual age-dependent life-cycle discounting is enough to account for nearly half of the difference between stochastic and risk-free valuations at the peak. The remainder of the difference between risk-free valuations and state-dependent valuations can be attributed to residual cross-sectional heterogeneity in earnings and consumption patterns.

4.3 Wealth Portfolios over the Life-Cycle

The PSID provides a long panel data set with sufficient information to carry out our human wealth estimation exercise. For a wider analysis of inequality, however, we rely on the Survey of Consumer Finances (SCF), which delivers much more detail on both value and composition of household wealth going back to 1989 (see Pfeffer, Schoeni, Kennickell *et al.*, 2016).

By design, the SCF captures the upper tail of the wealth distribution far better than the PSID, and contains accurate information about net worth and asset portfolios. Moreover, despite its lack of a panel dimension, we are able to use the human capital valuation function $\Theta(\mathbf{z})$ to directly evaluate equation (20) at *any* data point $\mathbf{z}$. That is, we can recover point estimates of human wealth in any data set where an appropriate counterpart of the vector $\mathbf{z}$ is available. Unfortunately, not every variable in the data vector $\mathbf{z}$ is observed in the SCF. In particular, unobserved types η cannot be estimated from repeated cross-sectional data and some variables are only observed

---

[21]In our baseline calculation, the risk-free rate is estimated as the average stochastic discount factor of those individuals who do not appear to be credit constrained and whose inter-temporal optimality conditions are likely to hold with equality. This is the same sample used to estimate the marginal utility function.

Figure 5: Average human wealth, net worth and lifetime wealth (the sum of human wealth and net worth) over the lifecycle. Values in 2016 dollars.
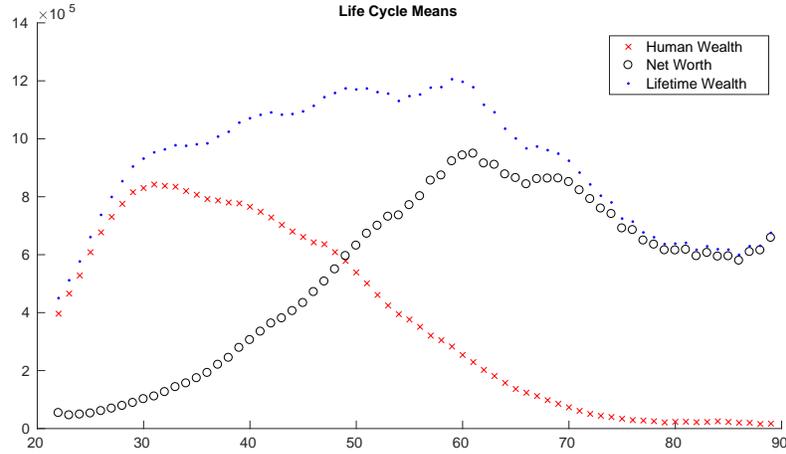
for the household head (for example education attainment and age). To deal with this problem we use a flexible non-parametric approach to impute the *full distribution* of the missing variables estimated from the PSID. It is important to impute the distributions of missing variables, rather than use their conditional expectations, because the latter would average out heterogeneity and lead to underestimation of human wealth dispersion. Details about the imputation method used to expand the cross-sectional SCF samples are presented in Appendix Section E. In the Appendix we document that both shape and scale of average human wealth estimates from SCF data closely resemble those plotted in Figure 2 using estimates from the PSID.

The Smoothness of Lifetime Wealth. Combining our estimates of human wealth with direct information about asset holdings, Figure 5 reports the average value of human and non-human wealth components, as well as lifetime wealth, for all households in our SCF sample.[22] It is apparent that lifetime wealth is remarkably stable over the life cycle, and certainly more so than its individual components. Yet, there is a subtle hump-shape in lifetime wealth, much like that of expenditures (see Fernández-Villaverde and Krueger, 2007, or Aguiar and Hurst, 2013 who decompose expenditure into its constituent categories). The portfolios of young households are heavily skewed towards human wealth, which makes shocks impacting their labor supply or health quite costly. In fact, any shocks are likely to be poorly insured among young adults because their net worth (non-human wealth) tends to be low.

The value of human wealth peaks early in life, around age 30. This is well before the peak age for earnings and draws attention to two key aspects: first, the expected length of remaining working life is important when putting a price tag on a stream of labor earnings; second, earlier

---

[22]One could easily disaggregate non human wealth in its different components.

20

investments in human capital carry a higher return while its depreciation may become more severe with age. These observations suggest that using current earnings as a measure of cross-sectional inequality is problematic, something that we revisit below.

The Changing Composition of Wealth over the Life-Cycle. The contrast between human and non-human wealth is striking. Net worth peaks around age 60 and effectively accounts for all wealth after age 70. Yet, net worth accounts for a relatively small fraction of lifetime wealth until age 40. Given these patterns, lifetime wealth is relatively stable between age 30 and 65 and, while declining to roughly 1/2 of its peak value by age 80, it exhibits less extreme proportional variation than each of its components over the course of the life cycle. The relative 'smoothness' of lifetime wealth across ages is consistent with the finding that a large chunk of lifetime wealth is determined early in life in the form of human wealth. Then, over time, lifetime wealth changes shape, shifting from illiquid human wealth to more liquid net worth. In this sense, the process of aging mostly changes the composition of wealth, while its total value varies less.

An alternative way to assess the evolution of portfolio composition is to report the share of assets out of lifetime wealth at different stages of the household's lifecycle. This share is a valuable measure of 'partial insurance' through precautionary savings (for example, see Blundell, Pistaferri, and Saporta-Eksten, 2016) and provides an estimate of the ability of households to smooth consumption in the face of earnings shocks. Figure 6 plots the portfolio asset share for all households, as well as for households in the top 1% and in the 90-99th percentiles of the lifetime wealth distribution. The shares, separately computed for each age group, confirm the growing role of asset wealth over the life cycle. The figure also illustrates how richer families tend to hold relatively more non-human wealth at earlier ages, with differences in portfolio composition being the largest around age 40. Wealth portfolios all converge to 100% assets as households age.

The peak in lifetime wealth occurs later in life for people who already own some assets when young. This becomes apparent when we examine, as we do below, the evolution of life cycle wealth for different percentiles of the lifetime wealth distribution by age.

Wealth over the Life Cycle: the Rich and the Poor. There exists significant dispersion in the distribution of human wealth at any given age (see, for example, Figure 18 in Appendix E). To document the extent of differences in the evolution of wealth holdings over the life cycle we contrast wealth patterns for households at the 25^{th}, 50^{th}, 90^{th} and 95^{th} percentiles of lifetime wealth at each given age. The results, plotted in Figure 7, confirm the presence of enormous differences in the composition of wealth portfolios between rich and poor households.

Households at the lower end of the wealth distribution hold little or no assets at any age, while richer households exhibit significant net worth at relatively early ages. Interestingly, we observe that human wealth plays a quantitatively large role even at the top end of the wealth distribution,

Figure 6: Asset share of household porfolios, by age and percentile of lifetime wealth.
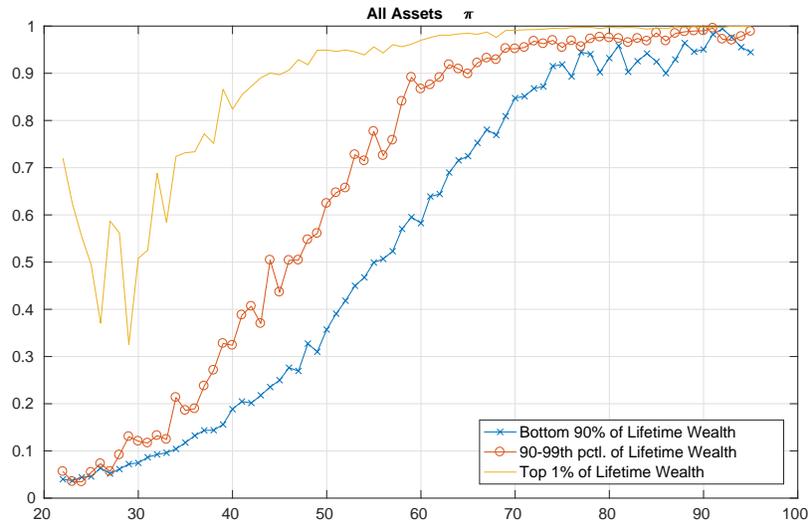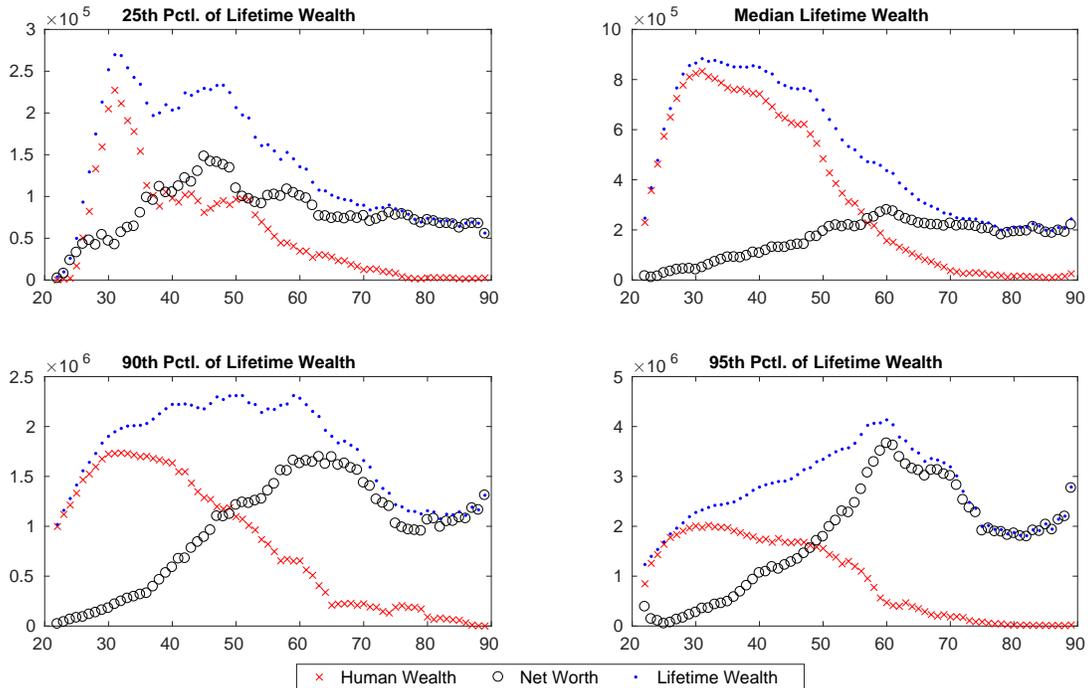

Figure 7: Average human wealth, net worth and lifetime wealth over the lifecycle, by percentile of lifetime wealth. Note: vertical axis scales vary by subplot. Values in 2016 dollars.

representing a large share of the aggregate at early ages. As we anticipated, lifetime wealth peaks early among poorer households; moreover, the human wealth of these households does not translate into equivalent amounts of net worth later in life.[23] The patterns illustrated in Figure 7 suggest that luck might have a non-trivial role in determining where one ends up in the distribution of lifetime wealth at older ages. Among older households, the richest tend to have a net worth that far exceeds their lifetime wealth at early ages, and the opposite is true of the poorest. Panel data are required to examine this conjecture in more detail, which we consider in the next subsection.

4.4 Early Human Wealth versus Net Worth at Retirement

One might wonder to what extent human wealth differences early in life are reflected in inequality at older ages. By returning to our PSID sample we are able to utilize the panel dimension of the data to answer such questions. In what follows we juxtapose early life human wealth with observed net worth and consumption outcomes later in life. Figure 8 contrasts net worth at age 65 with two different measures of human wealth at age 35: the left panel computes human wealth using expected earnings (ex-ante measure) while the right panel uses a measure of human wealth based on realized earnings over the sample period (actual lifetime earnings). The plots show that, no matter whether we use expected or realized earnings, net worth at retirement increases with human wealth valuations at age 35. That is, people who have higher expected labor earnings at age 35 also end up with bigger asset holdings at age 65. This suggests that human wealth heterogeneity, as realized early in life, may play a key role for later life net worth inequality. One aspect worth highlighting is that the ex-ante measure does a good job of predicting net worth ranks, but the ex-post measure does somewhat better at the very top of the human wealth distribution (contrast the top retirement wealth deciles in the left and right panel of Figure 8). This suggests that unanticipated shocks may play a role for the ex-post right tail of the wealth distribution in later life, while having only a moderate effect on other parts of the net worth distribution.

Also, it is worth relating the difference between the left and right panels of Figure 8 back to our conditional equivalence of expectations assumption, as well as our tests of that assumption in Appendix D. Ex-post human wealth realizations are somewhat more predictive of late life net worth inequality as they naturally include realized shocks, and our tests in Appendix D help assure us that these shocks are indeed unpredictable. While a sizeable proportion of late life net worth inequality is predictable based on early life human capital, there is an important component of luck as well. As shown in Figure 9, we perform a similar analysis for consumption expenditure deciles

---

[23]The message does not change if we condition on percentiles of each individual variable separately, instead of conditioning on percentiles of lifetime wealth, as shown in the Appendix Figure 17. This indicates that the ranking of lifetime wealth broadly lines up with the ranking of human wealth early in life and with the ranking of asset wealth at later ages. Importantly, pension and social security entitlements are not yet included in Figure 7, but we do observe that they are important for understanding lower wealth households when we impute them into lifetime wealth below.

Figure 8: Deciles of human wealth at age 35 vs deciles of net worth at age 65. Left panel (ex-ante human wealth) uses expected earnings; right panel (ex-post human wealth) uses realized earnings.
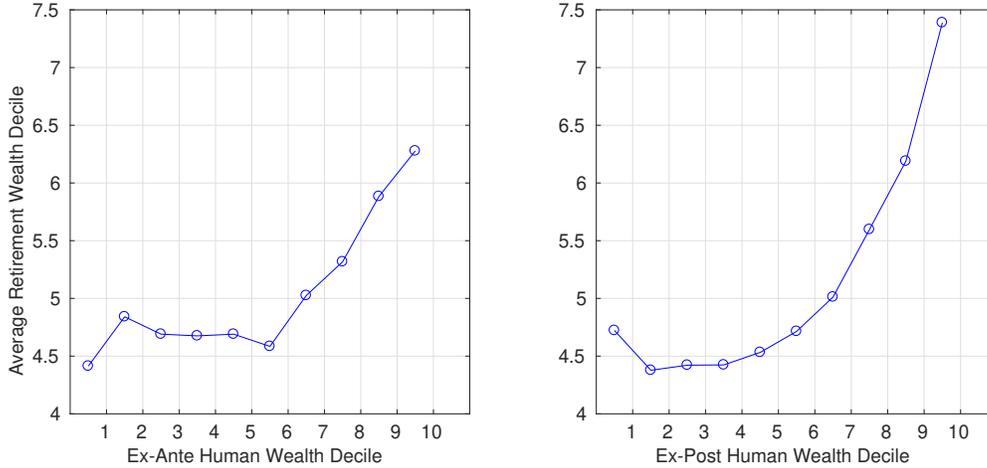
at retirement age and find very similar relationships with early life human wealth.

5   The Evolution of Cross-Sectional Inequality

Our expanded SCF data samples deliver multiple snapshots of the distribution of both lifetime wealth and permanent-income. We denote permanent-income by π and, in Appendix F, we describe how age-dependent annuity factors are used to compute it from lifetime wealth. These pieces of information are ideal to gauge the level, as well as the evolution, of cross-sectional inequality over the past three decades. In what follows we perform two simple exercises: first, we measure concentration of net worth, as in Bricker, Henriques, Krimmel *et al.* (2016); then, we contrast these measures with those obtained for human wealth, lifetime wealth and permanent-income.

5.1   Permanent-Income vs Assets: Top Shares

To facilitate comparison with existing studies, we plot measures of inequality across years. In the two panels of Figure 10 we report the shares of net worth, permanent-income and human wealth held by the top 1% (left panel) and the top 10% (right panel) of households in the distributions of the respective variables. Accounting for human wealth fundamentally changes our view of inequality and its evolution. Permanent-income is much less concentrated than asset wealth. For example, in 1989 the share of permanent-income held by the top 10% was almost half the corresponding share for net worth; this indicates that both lifetime wealth and permanent-income are significantly *less concentrated* than net worth. Similar patterns can be observed when looking at top 1% concentration measures. Figure 10 also shows that growth in concentration of permanent-

Figure 9: Deciles of human wealth at age 35 vs deciles of consumption expenditures at age 65. Left panel (ex-ante human wealth) uses expected earnings, right panel (ex-post human wealth) uses realized earnings.
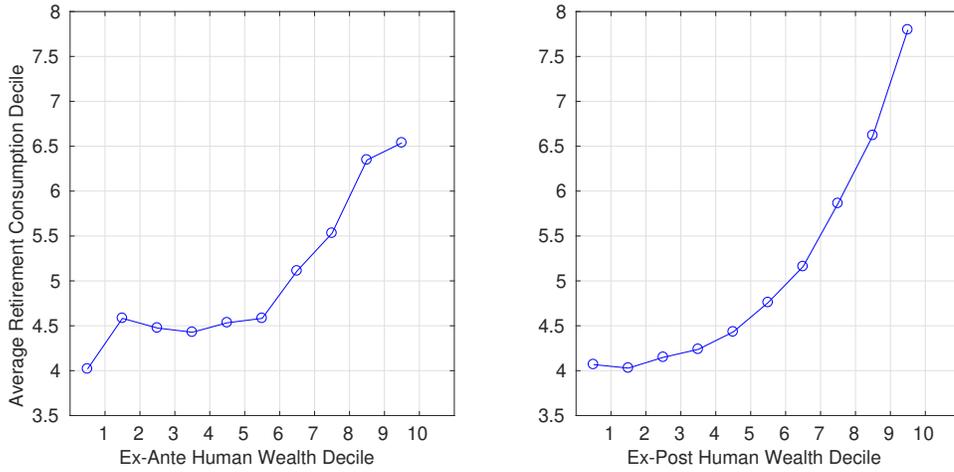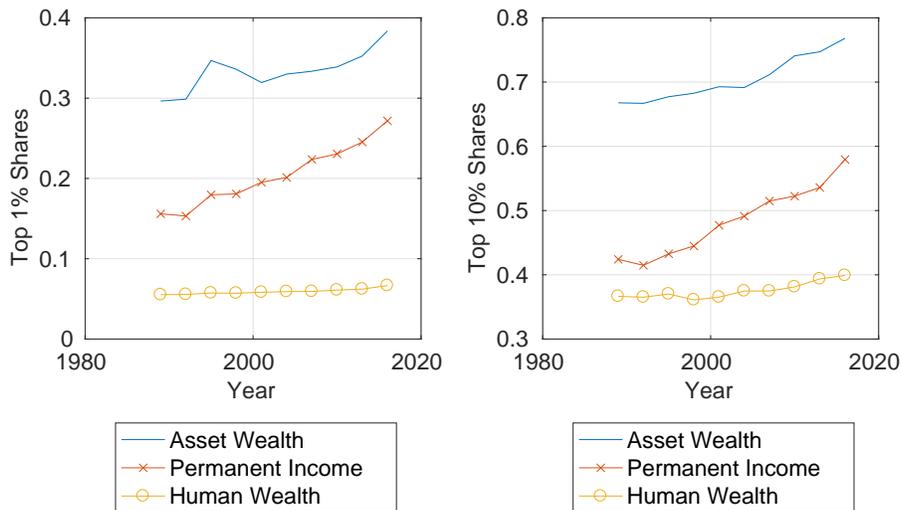

Figure 10: Concentration of net worth, human wealth and permanent-income by year (1989 to 2016. Left panel top 1%; right panel top 10%). Each plot reports the share in the hands of households at the top of the respective distribution, e.g. share of human wealth held by the top 10% of the human wealth distribution.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	Permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.366	0.405	0.372	0.424
1998	0.683	0.361	0.430	0.363	0.445
2007	0.712	0.375	0.488	0.416	0.515
2016	0.768	0.399	0.543	0.472	0.579

Table 1: This table reports the share of variable "X" in the hands of the households in the top 10% of the distribution of that same variable "X". For example, the share of permanent-income held by the households in the top 10% of the distribution of permanent-income. In Appendix C we report results for all years.

income has been strong over the past 35 years. The increase in permanent-income concentration has far outpaced that of its individual components, despite both human wealth and net worth becoming more concentrated between 1989 and 2016. This is remarkable for three reasons: first, the top 10% share of permanent-income is roughly 50% larger in 2016 than it was in 1989, while the top 1% permanent-income share almost doubled over the same interval; second, in the face of a heated debate on wealth inequality, since 1989 the pace of increase in concentration of permanent-income in the hands of the richest households is twice as fast as the well-known increase in the share of asset wealth; third, changes in human wealth concentration cannot account for the speedy growth in permanent-income concentration. As we show below, this uneven pace of concentration at the top is the by-product of the growing importance of asset wealth in household portfolios.

Comparing Measures of Concentration. To gain a more nuanced view of the changing concentration of economic resources, Table 1 reports, side by side, the shares held by the top 10% of households in the distribution of several variables for years spanning our sample length. This comparison highlights significant variation in the patterns of concentration for different variables.

Human wealth is less concentrated than current earnings, especially in the latter period of our sample. This finding is consistent with simulation results in Lucas and Moll (2014).[24] In fact, human wealth exhibits the lowest top concentration among all variables in any year. Arguably, forces that push for more concentration, such as permanent heterogeneity, are mitigated by relatively short working lives and by the accrual of shocks that depreciate human capital and limit the extent of excess-returns. On the other hand, factors possibly driving faster concentration of current earnings at the top include the growing recourse to performance-related pay among high earners, as well as changes in the relative importance of transitory shocks, especially at the higher-end of the earnings distribution.[25] Regardless of the source of this discrepancy, our findings suggest that

---

[24]See Figure 3 of Lucas and Moll (2014), who compare Lorenz curves for income and value functions in a growth model with endogenous time allocation. See also Castañeda, Díaz-Giménez, and Ríos-Rull (2003).

[25]For a discussion of human capital depreciation see Gallipoli and Turner (2011). For discussions of performance pay see Lemieux, Macleod, and Parent (2009), Makridis (2014), or Fonseca Galindo, Gallipoli, and Yedid-Levi (2018).

Figure 11: Lorenz curves for net worth, human wealth and permanent-income in 1989 (left) and 2016 (right).
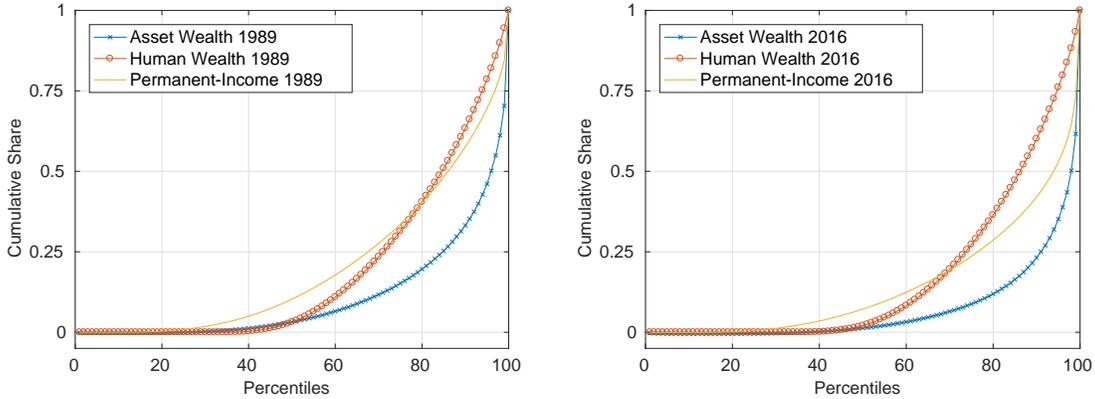
one should exercise caution when using only current earnings to draw inference about changes in long-term inequality. An overview of the distribution of different wealth measures can be obtained by plotting Lorenz curves, as we do in Figure 11 for assets, human wealth and permanent-income. The plots show that human wealth is much less concentrated than asset wealth; moreover, its share distribution did not change dramatically over the sample period. In contrast, the distribution of permanent income appears much more similar to that of asset wealth in 2016 than it did in 1989. Figure 11 and Table 1 document that, among all variables, permanent-income and lifetime wealth are the ones experiencing the fastest shift towards higher concentration. This pattern cannot be explained by swings in the concentration of human wealth and assets alone, as their distributions did not change sufficiently fast. Rather, as we discuss in Section 6 below, a pronounced growth in the share of asset wealth in household portfolios is the key driver of the sustained increases in permanent-income concentration. This appears to be the main reason why the Lorenz curve of permanent-income has become so much more similar to the that of assets.

6 The Mechanics of Increasing Inequality

The analysis in the previous section documents that households at the top of the net worth distribution have steadily increased their share of assets. An offsetting effect might come from rising concentration of human wealth in the hands of a different set of households at the top of the human wealth distribution. In contrast, having the same subset of households sit at the top of both distributions would compound and exacerbate the concentration of permanent-income. Hence, a key

---

For a discussion of the role of transitory earnings shocks for earnings dispersion see, among others, Gottschalk, Moffitt, Katz *et al.* (1994), Heathcote, Perri, and Violante (2010) or Bloom, Guvenen, Pistaferri *et al.* (2018).

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	Permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.125	0.351	0.203	0.385
1998	0.683	0.129	0.381	0.217	0.409
2007	0.712	0.122	0.457	0.267	0.490
2016	0.768	0.125	0.521	0.311	0.564

Table 2: This table reports the share of variable "X" in the hands of the households in the top 10% of the distribution of net worth. For example, the share of permanent-income held by the households in the top 10% of the distribution of net worth. In Appendix C we report results for all sample years.

question is whether the joint probability of being near the top of both the human and asset wealth distributions has changed over time.

To examine these issues we perform several checks: (i) we evaluate total and human wealth concentration among the group of households sitting at the top of the net worth distribution; (ii) we measure how much of the stock of lifetime wealth in different years is accounted for by asset wealth and document the role of portfolio composition for growing wealth concentration; (iii) we characterize the role of a changing age composition.

6.1 Which Households Have Grown Richer?

Restricting attention only to households at the top of the distribution of asset wealth, as we do in Table 2, is instructive. Human wealth concentration is effectively unchanged among the highest net worth households. This is despite the fact they seem to retain a growing share of current earnings over time. The fact that the share of human wealth held by asset-rich households has not risen indicates that growth in their earnings share has possibly been due to non-systematic and low-persistence income shocks, perhaps reflecting more performance-related pay or transitory bonuses. Some of these richer households are also relatively older so that increases in their earnings share has little or no effect on their share of the stock of human wealth. Nonetheless, the significant rise in the net worth share of these select households is clearly reflected in the strong growth of both their lifetime wealth and permanent-income shares. These observations illustrate two aspects of the run-up in economic inequality of the past 35 years: (i) looking at earnings flows (rather than human wealth stocks) may be misleading, whether studying levels or trends; (ii) while human wealth has become more concentrated (Table 1), the share of human wealth belonging to households at the top of the net worth distribution has not changed (Table 2). Thus, large increases in the top share of permanent-income cannot be due to a selected group of households holding an increasing share of both human and asset wealth. In summary, asset-rich households have not become more likely to be at the top of the human wealth distribution. Rather, human wealth has, over time, become

Figure 12: Asset wealth as a share of household lifetime wealth (aggregate by year).
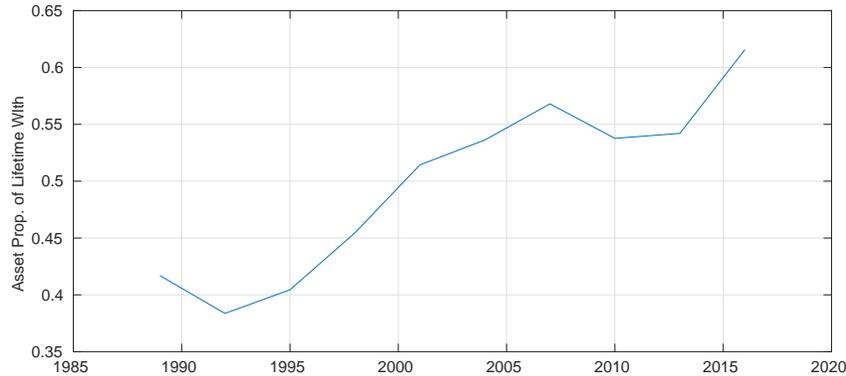
a less important determinant of inequality in permanent-income and this shift has resulted in a distribution of permanent-income that more closely resembles that of asset wealth. Put simply, being rich in human wealth was less important for permanent-income in 2016 than it was in 1989.

6.2 The Changing Importance of Assets in Household Portfolios

The previous analysis suggests that the higher concentration of permanent-income is related to the observation that assets, as a share of lifetime wealth, have on average become more prominent since 1989. This increase has wide ranging implications for the nature and extent of wealth inequality. A larger buffer of liquid assets implies better ability by households to respond to shocks like disability, unemployment and displacement; however, a diminished role for human wealth may also indicate that asset accumulation is increasingly driven by factors other than high early-life earnings and hard work.

To assess the relative importance of asset wealth over time, we use our extended SCF data and calculate the average assets-to-lifetime-wealth ratio in the cross-section of households in each sample year. The time series of this ratio is plotted in Figure 12 and clearly documents how asset wealth has become progressively more important within lifetime wealth portfolios. Starting from a value of around 40% in the early 1990s, asset wealth accounted for over 60% of lifetime wealth in 2016. The main deviation from this systematic growth pattern occurred during the recession of the late 2000s, when asset prices and valuations dipped temporarily.

Whether this run-up in the asset share of lifetime wealth portfolios is due to increased savings or, rather, price changes that pushed up the value of existing assets is the object of some debate. In Figure 13 we plot both the historical asset shares and the counterfactual proportions of lifetime wealth that the richest 1%, and the top 90th-99th percentiles, would have held in assets if the underlying prices had stayed constant. Specifically, we approximate the value of household assets

Figure 13: Actual and counterfactual portfolio shares of assets. Counterfactual shares are computed using constant 1989 prices. Left panel is for 90th to 99th percentiles of lifetime wealth, right panel for top 1%.
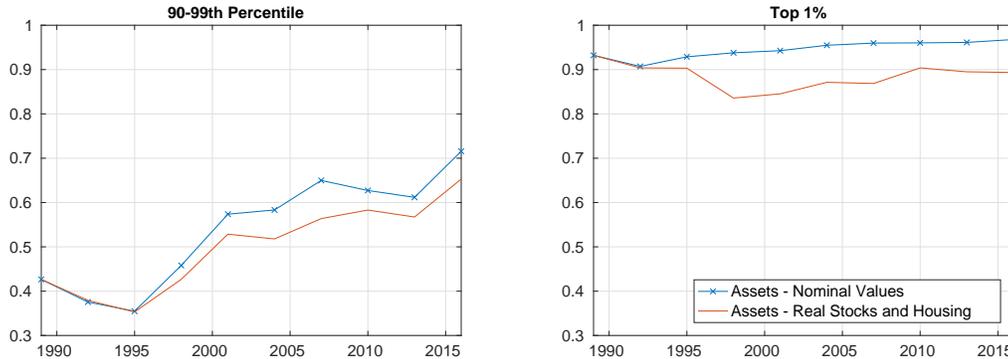
in each year at constant 1989 prices, using price indices for stocks and housing.[26] From these plots we draw two observations: (i) the run-up in the portfolio share of assets (blue line) is much starker for the top 90-99%, with the share growing from below 45% to above 70%, as opposed to the growth from 92% to 96% for the top 1% of households; (ii) our constant-price exercise (red line) suggests that incremental savings explain most of the run-up in the portfolio share of assets of the 90-to-99 group (left panel of Figure 13) but they do not account for much of the changes that occurred among the top 1%.

The lack of accurate indices to control for household-specific prices is an issue in the counterfactual exercises above. Nonetheless, it is safe to say that the portfolio composition of the top 90-99% of households in 2016 looks much more similar to that of the richest 1% than it did in 1989: remarkably, their share of tangible assets in lifetime wealth portfolios has grown by 25 percentage points over that period. Much of this convergence is accounted for by increased savings (i.e. quantities of assets) among the top 90-99 wealth percentiles, as can be observed by the growth of their asset wealth relative to their lifetime wealth when we hold prices constant.

The much smaller run-up in the portfolio share of assets among the top 1% is not surprising because tangible assets account for most of their portfolios, which means that price changes are almost equally reflected in the numerator and denominator of the portfolio share. This results in a more muted response of the asset share of the portfolio to price changes. Despite the smaller change in the asset wealth share of portfolios among the top 1%, prices do have an effect among these very rich households: if prices had stayed at their 1989 values, the portfolio share of asset wealth would be marginally down over the sample period. Moreover, this selected group of rich households appears to carry much of the aggregate asset price risk, because it owns roughly $1/3$

---

[26]We use the U.S. All-Transactions House Price Index (https://fred.stlouisfed.org/series/USSTHPI) for residential wealth and the Wilshire 5000 Total Market Full Cap Index (https://fred.stlouisfed.org/series/WILL5000INDFC) for equity prices. All types of wealth other than equities and housing are valued at current (real) prices.

of the asset wealth in the US economy, as we document in Table 6 below. Part of the convergence of asset shares in lifetime wealth portfolios might be due to the changing US age distribution. We explicitly consider this possibility later.

Why richer households choose to save a high share of their resources is an open question (Dynan, Skinner, and Zeldes, 2004). One possibility is that the opportunity cost of consumption is higher for them; evidence from Norwegian data in Fagereng, Guiso, Malacrino *et al.* (2016) suggests that returns on investments during the late 1990s and early 2000s were generally increasing in wealth. Other explanations relate to the presence of non-homotheticity in preferences (Straub, 2018), motives to transfer wealth to offspring (De Nardi, 2004; Boar, 2017) or save it for late life medical expenses (De Nardi, French, and Jones, 2010).[27]

6.3 Permanent-Income and Consumption Expenditures

If saving patterns play a role for concentration at the top, their effect should be detectable in consumption expenditure data (Attanasio and Pistaferri, 2016). Since 1999 the PSID reports an extensive range of such expenditure data, covering most of non-durable outlays for sample households. Using these expenditure records one can estimate each household's yearly outlays, and then link them to measures of permanent-income. For the purposes of this analysis we use measures of human wealth adjusted to reflect post-tax income and include imputed values of pension and social security entitlements. The idea is that consumption decisions depend on net, rather than gross, earnings, and that illiquid pension entitlements may have an effect on expenditures alongside other savings. We also adjust reported consumption expenditures so that yearly averages are consistent with NIPA data.[28] We discuss these adjustments in Appendix F.

Figure 14 shows a scatter plot juxtaposing permanent-income and consumption expenditures for households in our PSID sample (years 1999 and onwards). One immediate observation is that a large proportion of consumption expenditures lie below the 45 degree line, as many households spend less than what is implied by their permanent-income. This is consistent with results in Straub (2018), who estimates an average propensity to consume out of permanent-income below one. In contrast, the theoretical prediction of classic versions of the permanent-income hypothesis is that consumption responses to changes in permanent income should be one for one (see Zeldes, 1989; Carroll and Kimball, 1996). The average gap in Figure 14 is roughly 25%, and the discrepancy is present among both wealthy and poorer households. Expenditures are systematically below annuitized wealth among the richest, consistent with the observation that these households annuitize

---

[27]Heterogeneity in saving rates across age groups seems to play a role (Huggett and Ventura, 2000). One view, discussed in Carroll (1998), is that some households regard the accumulation of wealth as an end to itself. This heterogeneity may be related to long-term parental investments and intrafamily transmission of patience and risk-aversion (Doepke and Zilibotti, 2017) or to the fact that wealth yields a flow of services in the guise of status.

[28]Unadjusted average expenditures in the PSID tend to be lower than yearly estimates based on aggregate data.

Figure 14: Permanent-income adjusted for taxes and pension entitlements versus consumption expenditures (constant dollars, yearly frequency, in logarithms).
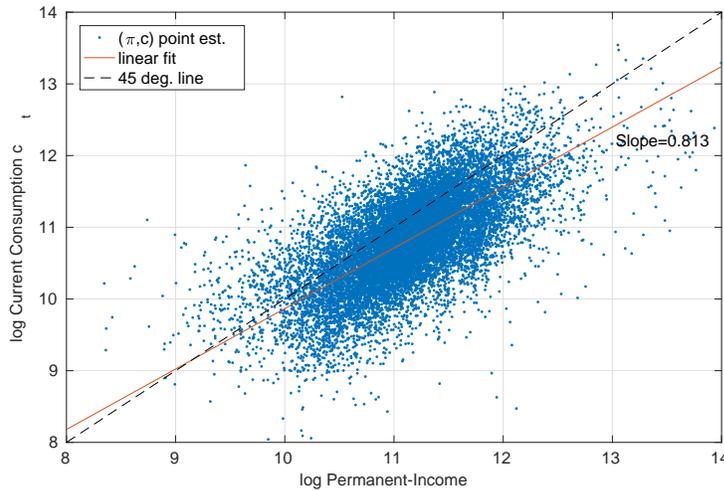
a very small fraction of their wealth (see Carroll, 1998). Figure 14 also indicates a tight connection between expenditures and permanent-income. Despite consumption being generally less than permanent-income, we observe a strong propensity to consume out of incremental permanent-income, with the slope of the fitted line in Figure 14 just above 0.81. Table 3 reports corroborating evidence from regressions of consumption expenditures on permanent-income.[29] Estimates of the marginal propensity to consume out of permanent-income suggest, again, values around 0.8. Slopes are precisely estimated and do not change much across alternative specifications or samples. Our measures of the consumption sensitivity account for responses to permanent-income from all wealth sources (human and non-human) and, while larger than the baseline value of 0.7 reported in Straub (2018), they lie well below one.[30]

Table 3 also documents significant heterogeneity in the propensity to consume out of permanent-income. This heterogeneity stems from cross-sectional variation in net worth. Households with a larger asset share exhibit a stronger propensity to consume out of lifetime wealth: holding permanent-income constant, a 10% increase in the value of net worth is associated with a consumption expenditure pass-through from permanent-income that is 0.8% higher. This heterogeneity is consistent with the view that net worth is more liquid than human wealth, as well as with the observation that households with significant shares of asset wealth in their portfolios tend to be much older.[31] However, even for households with higher expenditures out of permanent-income,

---

[29]As in Blundell, Pistaferri, and Saporta-Eksten (2016); Arellano, Blundell, and Bonhomme (2017), we examine a sample of households with a working head. We also restrict the age to be 65 or less.

[30]Both the definition and measurement of permanent income are different in Straub (2018). In particular, his baseline proxies of permanent income do not account for income from asset wealth or state-dependent discounting.

[31]There is growing recognition that portfolio composition and liquidity play a key role for consumption hetero-

	Dependent Variable: $log\ c_{it}$			
	Married	All	Married	All
	(1)	(2)	(3)	(4)
Panel A: OLS				
log P.I.	0.828	0.797	0.659	0.633
	(0.013)	(0.012)	(0.017)	(0.014)
log Net Worth			0.080	0.079
			(0.005)	(0.004)
Panel B: MAD (outlier robust) regressions				
log P.I.	0.844	0.813	0.673	0.648
	(0.007)	(0.006)	(0.010)	(0.008)
log Net Worth			0.080	0.079
			(0.003)	(0.003)
N of observations	11,660	13,333	11,660	13,333

Table 3: Propensity to consume out of permanent-income. This table presents results for regressions of (household-level) log consumption on the log of permanent-income and, in some specifications, on the log of net worth (asset wealth). The sample includes households with a working head. Panel A refers to OLS regressions; Panel B refers to outlier-robust regressions. Standard errors are in parenthesis.

we estimate a pass-through coefficient well below one. These findings buttress the view that aggregate consumption dynamics are sensitive to changes in the distribution of resources over time, an observation that has important implications for policy.

Consumption Deviations from Permanent-Income. To further explore the role of liquidity in the relationship between consumption and permanent-income, we identify households whose behaviour in the data is consistent with short-term liquidity constraints.

In Figure 15 we report plots of consumption vs permanent-income where we highlight in red the households with expenditure patterns consistent with short-term credit constraints. Specifically, the left-panel identifies households with large consumption growth between t and $t + 1$, and the right panel identifies households with large consumption growth between $t + 1$ and $t + 2$.[32] Rapid consumption growth in the near future (within the next 4 years in this case) is consistent with households being unable to smooth consumption by borrowing at time t. The plots show that, for either definition of large future consumption growth, the red dots are more likely to be below the 45-degree line and to exhibit a larger gap.[33] Arguably, households corresponding to red dots may face

---

geneity and its aggregate effects (see for example Kaplan and Violante, 2014; Kaplan, Violante, and Weidner, 2014; Bayer, Lütticke, Pham-Do *et al.*, 2019).

[32]Large consumption growth is defined as $\frac{c_{t+1}}{c_t}$ larger than 1.25 ($\frac{c_{t+2}}{c_{t+1}} > 1.25$ in the right panel). We also experiment with other thresholds and obtain very similar results.

[33]A concern here is classical measurement error in consumption: that is, a large c_{t+1}/c_t could be due to a downward error in c_t, which is also on the plot's vertical axis. However, the right-panel is not subject to this criticism as errors

Figure 15: Tax and pension adjusted permanent-income versus consumption expenditures (in constant dollars). Red bubbles refer to households experiencing large consumption growth (.5 log points or more) in the following period, which is consistent with being short-term credit constrained.
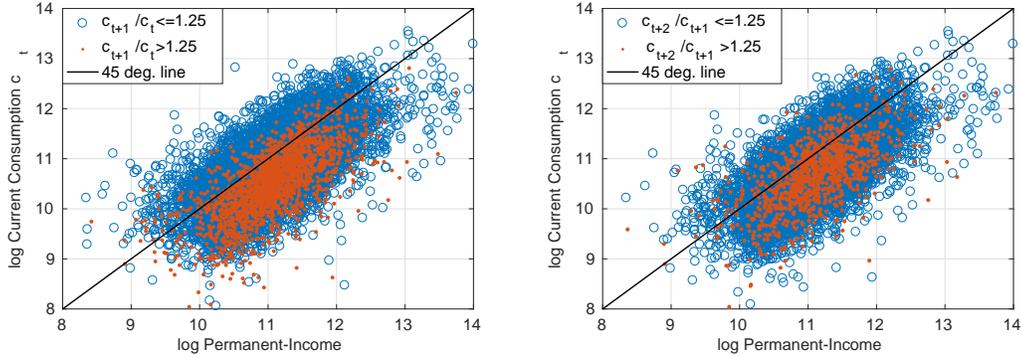
short-term credit constraints that limit their ability to spend according to their permanent-income (Zeldes, 1989). We identify a significant share of households that seem short-term constrained in their expenditure behavior and are prone to consuming below permanent-income, which implies that looking only at consumption expenditures is not ideal to draw inference about the distribution of welfare outcomes over the life-cycle.

Further analysis of the discrepancies between permanent-income and expenditures confirms the importance of credit-constraints in consuming out of permanent-income. In Table 4 we report results from two regressions that quantify both the likelihood and extent of consumption deviations: (i) column 1 in Table 4 reports estimates for a linear probability model of the likelihood that current household consumption expenditures are below permanent-income; (ii) column 2 in the same table reports estimates of the proportional consumption deviation from permanent-income. The dependent variables are projected on dummies that take unit value when consumption growth between any two consecutive periods after t is larger than 25%, which we interpret as evidence of short-term liquidity constraints. We consider four such dummies, spanning the time periods between t and $t+4$ (which spans 10 years in our biennial data).[34]

After controlling for jumps in consumption up to four periods ahead, we estimate that the baseline probability that current expenditures lie below π_{it} (the intercept in column 1) is roughly $2/3$; as one would expect, this probability grows further when one conditions on large jumps in household consumption growth after period t, which suggests that consumption levels are more likely to be below permanent-income for households that appear liquidity constrained. When looking at

---

in c_t do not affect c_{t+2}/c_{t+1}. Furthermore, in the regressions in Table 4 we show that large c_{t+3}/c_{t+2} and c_{t+4}/c_{t+3} are also associated with c_{it} being unusually low compared to π_{it}, which cannot be due to errors in c_{it} measurement.

[34]In Appendix C we report results when we set dummies equal to one when consumption growth exceeds 50%. Moreover, to account for the possibility of spurious correlation between consumption values in period t and consumption growth in $t+1$ due to classical measurement error, we also replicate the analysis controlling only for episodes of large consumption growth occurring after $t+1$. Results are almost identical and available upon request.

	Dependent Variables	
	$\mathbb{I}\{c_{it} < \pi_{it}\}$	$\ln(\pi_{it}) - \ln(c_{it})$
	(1)	(2)
constant	0.6768	0.1969
	(0.0083)	(0.0089)
$\mathbb{I}\left\{\frac{c_{it+1}}{c_{it}} > 1.25\right\}$	0.1691	0.2821
	(0.0091)	(0.0123)
$\mathbb{I}\left\{\frac{c_{it+2}}{c_{it+1}} > 1.25\right\}$	0.1126	0.1841
	(0.0110)	(0.0143)
$\mathbb{I}\left\{\frac{c_{it+3}}{c_{it+2}} > 1.25\right\}$	0.0827	0.1402
	(0.0106)	(0.0132)
$\mathbb{I}\left\{\frac{c_{it+4}}{c_{it+3}} > 1.25\right\}$	0.0532	0.0960
	(0.0110)	(0.0136)
N-observations	7,386	7,386

Table 4: Deviations of consumption flows c_{it} from permanent-income π_{it}. This table reports results for two linear regressions: (1) a linear probability model with dependent variable equal to 1 when $\{c_{it} < \pi_{it}\}$ and zero otherwise; (2) a continuous model in which the dependent variable is the log difference $\ln(\pi_{it}) - \ln(c_{it})$. Both regressions include dummies taking unit value if consumption growth exceeds 25% between any consecutive periods between t and $t+4$. Standard errors in parenthesis.

the size of consumption deviations from permanent-income, column 2 in Table 4 shows that baseline consumption is, on average, lower than permanent-income (intercept is estimated to be 0.2); however, among households experiencing large consumption jumps in the near future, this gap becomes much larger; for example, it doubles if c_{it+2}/c_{it+1} exceeds 25%. This evidence suggests that the majority of households in our sample exhibit expenditures significantly below their permanent-income, and this is especially true for short-term credit-constrained households. These patterns reinforce the view that consumption expenditures alone might not deliver a complete picture of the prevailing degree of permanent-income inequality.

6.4 The Role of Demographic Change

It is possible that an aging population has affected the dynamics of inequality. To examine this hypothesis, we carry out a counterfactual re-weighting exercise. This allows us to assess how inequality in the variables we observe would have changed had the age distribution stayed the same as in 1989. Following DiNardo, Fortin, and Lemieux (1996), we use probit regressions with a full set of age dummies to construct counterfactual sample weights. Then, using these age-corrected weights, we generate counterfactual versions of our concentration measures, which we report in Figure 5. Under the counterfactual, there is almost no change in human wealth concentration, unlike our baseline results. This discrepancy illustrates that the rising concentration of human

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.366	0.405	0.372	0.424
1998	0.687	0.360	0.427	0.360	0.440
2007	0.723	0.362	0.473	0.407	0.497
2016	0.783	0.367	0.507	0.450	0.542

Table 5: Shares of variable "X" in the hands of the households in the top 10% of the distribution of that same variable "X". For example, the share of permanent-income held by the households in the top 10% of the distribution of permanent-income. In Appendix C we report results for all sample years.

wealth might largely follow from the fact that in 2016 a smaller segment of the population was at their peak of human wealth relative to 1989, when the baby boom cohort was close to their highest human wealth years. It follows that fewer households outside the top 10% have large human wealth stocks in 2016 than in 1989 and, mechanically, this accounts for some of the run-up in human wealth concentration. In contrast, under the counterfactual the concentration of net worth would have increased marginally more over the sample period. Again, this is due to the fact that in 2016 there is a relatively larger cohort of older and richer people, which results in the top 10% being a more selected group of high-net worth households. Put differently, in 1989 less wealth was needed to be in the top 10% because the set of people at peak net worth ages was proportionally smaller.

7 Sensitivity and Robustness

In this section we assess the sensitivity of our findings to alternative ways of examining the data. First, we replicate our concentration analysis for the top 1% of each variable. Second, we present results for alternative measures of lifetime wealth inequality, making allowances for the role of government and value of leisure. Next, we consider a range of parametric specifications of the utility function and gauge the sensitivity of human wealth to these alternatives. Lastly, we consider sensitivity to different consumption measurement approaches.

7.1 Concentration of Wealth: the Top 1%

Table 6 reports top 1% shares of different variables. The patterns over time are similar to those for the top 10%. The top shares of all variables are rising, with permanent-income concentration rising faster than its individual components, although the discrepancy between the growth of asset wealth and permanent-income is smaller than what we find for the top 10%. Table 7 replicates the analysis only for the set of households at the top of the distribution of net worth, providing

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	Permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.296	0.055	0.131	0.101	0.156
1998	0.336	0.057	0.162	0.102	0.181
2007	0.333	0.059	0.197	0.144	0.224
2016	0.384	0.067	0.244	0.194	0.272

Table 6: This table reports the share of variable "X" in the hands of the households in the top 1% of the distribution of that same variable "X". For example, the share of permanent-income held by the households in the top 1% of the distribution of permanent-income. In Appendix C we report results for all sample years.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	Permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.296	0.011	0.130	0.039	0.153
1998	0.336	0.016	0.161	0.053	0.176
2007	0.333	0.016	0.196	0.066	0.212
2016	0.384	0.019	0.243	0.107	0.265

Table 7: This table reports the share of variable "X" in the hands of the households in the top 1% of the distribution of net worth. For example, the share of permanent-income held by the households in the top 1% of the distribution of net worth. In Appendix C we report results for all sample years.

additional evidence that asset-rich households are not holding increasingly large shares of human wealth. The fact that their share of human wealth does not rise, while that of permanent-income grows significantly, indicates again that asset wealth has become increasingly important over time.

7.2 Concentration of Wealth between the 50th and 90th Percentiles

So far we have mostly focused on the top concentration of different wealth measures. In what follows we document the evolution of the wealth share held by households ranked between the 50th and 90th percentiles of each variable's distribution. This exercise is helpful to reconcile the shifts at the top with the adjustments occurring among less wealthy households.

Table 8 shows that the net worth share of this group fell from 30% to 22%, consistent with the view that a major shift in the asset distribution has been taking place over the last three decades. The share of human wealth, on the other hand, has only marginally gone down for the 50-90 group over this period. Just as in the case of the richest households, the share of permanent income has become more similar to that of assets since 1989. The large drop in the share of permanent income cannot be explained by changes in the concentration of human wealth and assets alone; this suggests that changes in the share of permanent income are partly due to the growing importance

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	Permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.300	0.602	0.499	0.536	0.473
1998	0.285	0.601	0.468	0.524	0.445
2007	0.262	0.591	0.417	0.495	0.388
2016	0.219	0.580	0.385	0.464	0.349

Table 8: This table reports the share of variable "X" in the hands of the households ranked between the 50th and 90th percentiles of the distribution of that same variable "X". For example, the share of permanent-income held by households in the 50-90 percentiles of the distribution of permanent-income.

of assets in lifetime wealth portfolios, just like for richer households. These results, together with the evidence presented in Table 1, indicate that households in the bottom 50% of the wealth distribution did not experience large changes in their shares of asset and human wealth. Most of the reallocation of lifetime resources can be accounted for by flows from the 50-90th percentiles to the top 10%.

7.3 Alternative Measures of Permanent-Income

In Table 9 we explore three alternative ways to gauge changes in the concentration of resources and contrast them with previous results. For comparison, column (1) in Table 9 reproduces results from Table 1. Our baseline estimation does not value the opportunity cost of leisure time. To account for this, we re-estimate human wealth under the assumption that every available hour is valued at its market value, up to 35 hours per week. This approach assumes out the effect of differences in lifetime labor supply. Column (3) in Table 9 shows that the resulting measure, denoted 'full-time equivalent', exhibits slightly higher concentration than our baseline measure. This is not unexpected: part of the equalizing influence of accounting for human wealth derives from age-differences in life-cycle labor supply; however, we now include the value of the leisure of retired households in their human wealth. Nonetheless, the change in concentration over time closely tracks baseline estimates, suggesting that wages, rather than hours worked, are responsible for the growing concentration of human capital. In column (4) we report estimates of the concentration of post-tax permanent-income. These estimates are obtained from the distribution of earnings after taxes and transfers.[35] This adjusted measure of permanent-income appears only slightly less concentrated than our baseline estimates. Moreover, it exhibits an almost identical pattern over time. In column (5) we further account for the role of taxes and transfers by re-estimating with imputed social security and pension benefits included. This further reduces the concentration of permanent-income as the benefits are a larger part of the lifetime wealth of poorer households.

---

[35]We approximate taxes using the power function approach described in Guner, Kaygusuz, and Ventura (2014).

| Year | Permanent-income | | | | Lifetime cons. |
| | baseline | 'Full-Time' Eqv | post tax | post tax + pens | Expected |
	(1)	(2)	(3)	(4)	(5)
1989	0.424	0.430	0.411	0.361	0.379
1998	0.445	0.452	0.432	0.381	0.384
2007	0.515	0.525	0.502	0.450	0.392
2016	0.579	0.588	0.566	0.502	0.410

Table 9: Shares of variable "X" owned by households in the top 10% of the distribution of that same variable "X". E.g., the share of earnings held by the households in the top 10% of the distribution of earnings.

In the last column of the same table we also report a proxy of permanent-income concentration based on lifetime consumption expenditures. This measure is obtained from the expected present value of stochastically discounted life-cycle consumption (calculated similarly to θ_{it}). Lifetime consumption appears significantly less concentrated than permanent-income, particularly in later years; moreover, it exhibits only a small increase between 1989 and 2016. The cross-sectional distribution of lifetime consumption across households is also less concentrated than that of post-tax permanent-income. However, after pension entitlements are imputed the concentration of expected lifetime consumption is only smaller than permanent-income in the second half of the sample period. This suggests that richer households may have grown their expenditures much less than their permanent-income, an observation reminiscent of the results presented in Section 6. A caveat in the construction of this variable is that we include net worth in the information vector $\mathbf{z}_{it}$, which limits data available for estimation to years when net worth is observable. This may be responsible for some of the discrepancies we document.[36] Moreover, non-classical measurement error in expenditures (see Aguiar and Bils, 2015) may imply that consumption concentration has risen by more than these results suggest.

7.4 Parametric Utility Functions

We repeat the analysis of top share concentration in lifetime wealth using various parametric utility function specifications. This exercise clarifies how sensitive the valuation of human wealth is to parametric restrictions, while also showing which restrictions deliver results similar to those obtained using non-parametric estimation. We consider two sets of utility functions: standard CRRA and Epstein-Zin recursive utility. Table 10 reports results for different values of the intertemporal elasticity of substitution (IES) and relative risk aversion (RRA). It is apparent that higher curvature of the utility function implies higher estimated concentration at the top. Relatively poorer and less insured households, with more variable consumption, end up discounting expected earnings more

---

[36]When we do not include net worth, we find that $\mathbf{z}_{it}$ does not fully capture information about future consumption.

Year	Baseline	CRRA			Epstein-Zin			
	(non parametric)				IES=1.5	IES=1.5	IES=1.75	IES=1.75
		RRA=1.6	RRA=2	RRA=3	RRA=2	RRA=3	RRA=2	RRA=3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
1989	0.405	0.416	0.428	0.479	0.439	0.502	0.430	0.462
1998	0.430	0.440	0.451	0.496	0.464	0.525	0.454	0.486
2007	0.488	0.501	0.513	0.554	0.519	0.564	0.512	0.540
2016	0.543	0.557	0.570	0.610	0.574	0.612	0.567	0.594

Table 10: Share of lifetime wealth in the hands of the households in the top 10% of the distribution of lifetime wealth. Each column corresponds to a different parametric specification of the utility function.

heavily. A CRRA specification with RRA of 1.6, or a EZ utility with IES of 1.75 and RRA of 2, deliver results that are close to our non-parametric estimates. The latter specification is consistent with preference for early resolution of uncertainty over the life-cycle (Epstein and Zin, 1989).

7.5 Alternative Measures of Consumption Expenditures

To assuage concerns that our baseline results may be sensitive to changes in the way we measure consumption, we re-estimate human wealth and permanent income using alternative consumption measures. Table 11 summarizes results for the top 10% concentration of human wealth and permanent income under these alternatives. Column (1) reports our baseline results for comparison. First, in columns (2) and (5) we report estimates based on actual, rather than imputed, expenditure for the years 1998-2014, when the PSID expanded the set of consumption categories.[37] Results are only slightly sensitive to this alternative approach, with neither human wealth nor permanent-income exhibiting significant changes in top concentration. In columns (3) and (6) we also replicate the analysis using only food expenditures, which are consistently measured over our sample period. We find slightly larger concentration of human wealth and permanent income, with no evidence of changes in trends over time. These robustness checks indicate that our baseline results are not very sensitive to the use of alternative measures of consumption.

8 Conclusions

This paper proposes a new approach to quantify the value of human capital and asset wealth held by households. Our analysis brings together different data sources and delivers estimates of household lifetime wealth and permanent-income, over the life-cycle and in the cross-section. The measurement approach does not entail strong assumptions about preferences or income processes as it relies on a non-parametric identification procedure.

---

[37]Before 1998 we continue to use imputed expenditures, so to be able to perform the human wealth calculation.

Year	Human Wealth			Permanent-Income		
	Baseline	Obs. 1998-2014	Food Cons.	Baseline	Obs. 1998-2014	Food Cons.
	(1)	(2)	(3)	(4)	(5)	(6)
1989	0.366	0.356	0.391	0.424	0.437	0.458
1998	0.361	0.352	0.389	0.445	0.461	0.480
2007	0.375	0.363	0.398	0.515	0.528	0.551
2016	0.399	0.386	0.421	0.579	0.592	0.615

Table 11: Shares of human wealth and permanent income for (i) our baseline results; (ii) when preferences are estimated in step one using only reported expenditure between 1998 and 2014; and (iii) using food consumption in lieu of imputed consumption in all estimation steps. In approach (ii) we estimate preferences using only reported data, but employ imputed consumption for human wealth estimation in prior years.

Our estimates deliver new insights on household-level heterogeneity along a variety of wealth and income measures. This information is especially instructive when examining the changing patterns of wealth inequality over the past three decades. Accounting for heterogeneity in wealth and lifetime resources is key to provide a broader assessment of cross-sectional inequality and its evolution. We document that human wealth is significantly less concentrated than net worth, and that inequality in permanent-income is actually lower than inferred from dispersion measures that focus exclusively on asset wealth. It is also apparent that richer households have accrued a considerably larger share of permanent-income over the past decades. In fact, concentration of permanent-income has grown much faster than concentration of net worth. As a consequence, effective inequality has increased more than previously thought, albeit from a lower initial level.

We document that changes in the marginal distributions of net worth and human wealth only account for a small part of the significant increase in permanent-income concentration. The share of households that sit at the top of both net worth and human wealth distributions has not changed much between 1989 and 2016. This suggests that increased concentration of permanent-income is not explained by a small set of households holding increasing shares of all types of wealth. Instead, a crucial driver behind the run-up in wealth inequality is the growing value of assets as a share of the lifetime wealth portfolios of rich households. High net worth households account for a larger share of total permanent-income in 2016 than they did in 1989.

We use our estimates of permanent-income to gauge the relationship between consumption expenditures and lifetime wealth. Our analysis suggests that the average propensity to consume out of permanent income is around 0.8. We also find evidence of significant cross-sectional dispersion in the sensitivity of consumption to lifetime wealth; this heterogeneity is due to differences in both the value and composition of lifetime wealth portfolios. Our findings support the view that aggregate consumption responds to changes in the distribution of lifetime resources, an observation that has considerable implications for policy.

References

Abbott, B. and Gallipoli, G. (2018). 'Human Capital Inequality: Empirical Evidence.' Tech. rep., Vancouver School of Economics, UBC. Working Paper.

Aguiar, M. and Bils, M. (2015). 'Has consumption inequality mirrored income inequality?' *American Economic Review*, vol. 105(9):pages 2725–56.

Aguiar, M. and Hurst, E. (2013). 'Deconstructing life cycle expenditure.' *Journal of Political Economy*, vol. 121(3):pages 437–492.

Arellano, M., Blundell, R., and Bonhomme, S. (2017). 'Earnings and consumption dynamics: a nonlinear panel data framework.' *Econometrica*, vol. 85(3):pages 693–734.

Athreya, K., Ionescu, F., and Neelakantan, U. (2015). 'Stock market participation: The Role of Human Capital.' Tech. rep., Working Paper, Federal Reserve Bank of Richmond.

Athreya, K. B., Ionescu, F., and Neelakantan, U. (2017). 'College or the Stock Market, or College and the Stock Market?' Tech. rep., Board of Governors of the Federal Reserve System (US).

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

Attanasio, O. P. and Pistaferri, L. (2016). 'Consumption inequality.' *Journal of Economic Perspectives*, vol. 30(2):pages 3–28.

Autor, D. H., Katz, L. F., and Kearney, M. S. (2008). 'Trends in US Wage Inequality: Revising the Revisionists.' *The Review of Economics and Statistics*, vol. 90(2):pages 300–323.

Bayer, C., Lütticke, R., Pham-Do, L., and Tjaden, V. (2019). 'Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk.' *Econometrica*, vol. 87(1):pages 255–290.

Becker, G. S. (1975). *Human Capital*. Second Edition, National Bureau of Economic Research, New York, NY.

Benzoni, L. and Chyruk, O. (2015). 'The Value and Risk of Human Capital.' Working paper, Chicago Fed.

Bloom, N., Guvenen, F., Pistaferri, L., Sabelhaus, J., Song, J., and Salgado, S. (2018). 'Why Has U.S. Earnings Volatility Been Declining for Four Decades?' Tech. rep., mimeo, University of Minnesota and Stanford.

Blundell, R., Pistaferri, L., and Saporta-Eksten, I. (2016). 'Consumption inequality and family labor supply.' *The American Economic Review*, vol. 106(2):pages 387–435.

Boar, C. (2017). 'Dynastic Precautionary Savings.' *Manuscript*.

Bonhomme, S., Lamadon, T., and Manresa, E. (2017). 'Discretizing unobserved heterogeneity.' Tech. rep., University of Chicago.

Bowlus, A. J. and Robin, J.-M. (2004). 'Twenty years of rising inequality in US lifetime labour income values.' *The Review of Economic Studies*, vol. 71(3):pages 709–742.

Bricker, J., Henriques, A., Krimmel, J., and Sabelhaus, J. (2016). 'Measuring Income and Wealth at the Top using Administrative and Survey Data.' *Brookings Papers on economic Activity*.

Carroll, C. D. (1998). 'Why do the rich save so much?' Tech. rep., National Bureau of Economic Research.

Carroll, C. D. and Kimball, M. S. (1996). 'On the concavity of the consumption function.' *Econometrica:*

Journal of the Econometric Society, pages 981–992.

Carroll, C. D., Slacalek, J., and Tokuoka, K. (2014). 'The Distribution of Wealth and the MPC: Implications of New European Data.' *American Economic Review*, vol. 104(5):pages 107–11.

Castañeda, A., Díaz-Giménez, J., and Ríos-Rull, J.-V. (2003). 'Accounting for U.S. Earnings and Wealth Inequaltiy.' *Journal of Political Economy*, vol. 111(4):pages 818–857.

Conesa, J. C., Kitao, S., and Krueger, D. (2009). 'Taxing capital? Not a bad idea after all!' *The American Economic Review*, vol. 99(1):pages 25–48.

Cunha, F., Heckman, J., and Navarro, S. (2005). 'Separating Uncertainty from Heterogeneity in Life Cycle Earnings.' *Oxford Economic Papers*, vol. 57(2):pages 191–261.

Davila, J., Hong, J. H., Krusell, P., and Ríos-Rull, J.-V. (2012). 'Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks.' *Econometrica*, vol. 80(6):pages 2431–2467.

De Nardi, M. (2004). 'Wealth inequality and intergenerational links.' *The Review of Economic Studies*, vol. 71(3):pages 743–768.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2016). 'The Implications of Richer Earnings Dynamics for Consumption and Wealth.' NBER working paper no. 21917.

De Nardi, M., Fella, G., and Paz-Pardo, G. (2019). 'Nonlinear Household Earnings Dyanimics, Self-Insurance, and Welfare.' The Journal of the European Economic Association, forthcoming.

De Nardi, M., French, E., and Jones, J. (2010). 'Why Do the Elderly Save? The Role of Medical Expenses.' *Journal of Political Economy*, vol. 118(1):pages 39–75.

DiNardo, J., Fortin, N. M., and Lemieux, T. (1996). 'Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach.' *Econometrica*, vol. 64:pages 1001–1044.

Doepke, M. and Zilibotti, F. (2017). 'Parenting with style: Altruism and paternalism in intergenerational preference transmission.' *Econometrica*, vol. 85(5):pages 1331–1371.

Dynan, K. E., Skinner, J., and Zeldes, S. P. (2004). 'Do the rich save more?' *Journal of Political Economy*, vol. 112(2):pages 397–444.

Epstein, L. G. and Zin, S. (1989). 'Substitution, Risk Aversion and the Temporal Behavior of Asset Returns: a Theoretical Framework.' *Econometrica*, vol. 57:pages 937–969.

Escanciano, J. C., Hoderlein, S., Lewbel, A., Linton, O., and Srisuma, S. (2016). 'Nonparametric Euler Equation Identification and Estimation.' Tech. rep., mimeo, Boston College.

Fagereng, A., Guiso, L., Malacrino, D., and Pistaferri, L. (2016). 'Heterogeneity and persistence in returns to wealth.' Tech. rep., National Bureau of Economic Research.

Fernández-Villaverde, J. and Krueger, D. (2007). 'Consumption over the life cycle: Facts from consumer expenditure survey data.' *The Review of Economics and Statistics*, vol. 89(3):pages 552–565.

Fonseca Galindo, J., Gallipoli, G., and Yedid-Levi, Y. (2018). 'Match Quality, Contractual Sorting and Wage Cyclicality.' Tech. rep., UBC. Vancouver School of Economics.

Friedman, M. (1957). 'The Permanent Income Hypothesis.' In *A theory of the Consumption Function*, pages 20–37. Princeton University Press.

Fuster, A., Kaplan, G., and Zafar, B. (2018). 'What Would You Do With $500? Spending Responses to Gains, Losses, News and Loans.' Tech. rep., National Bureau of Economic Research.

Gabaix, X. and Laibson, D. (2017). 'Myopia and Discounting.' Tech. rep., Harvard.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). 'The dynamics of inequality.' *Econometrica*, vol. 84(6):pages 2071–2111.

Gallipoli, G. and Turner, L. (2011). 'Household Responses to Individual Shocks: Disability and Labour Supply.' Mimeo.

Goldin, C. and Katz, L. F. (2007). 'The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005.' NBER Working Paper 12984.

Golosov, M., Tsyvinski, A., Werning, I., Diamond, P., and Judd, K. L. (2006). 'New Dynamic Public Finance: A User's Guide [with Comments and Discussion].' *NBER macroeconomics annual*, vol. 21:pages 317–387.

Gottschalk, P., Moffitt, R., Katz, L. F., and Dickens, W. T. (1994). 'The growth of earnings instability in the US labor market.' *Brookings Papers on Economic Activity*, vol. 1994(2):pages 217–272.

Guner, N., Kaygusuz, R., and Ventura, G. (2014). 'Income taxation of U.S. households: Facts and parametric estimates.' *Review of Economic Dynamics*, vol. 17(4):pages 559 – 581. ISSN 1094-2025.

Guvenen, F., Kaplan, G., Song, J., and Weidner, J. (2017). 'Lifetime incomes in the United States over six decades.' Tech. rep., National Bureau of Economic Research.

Guvenen, F., Karahan, F., Ozkan, S., and Song, J. (2016). 'What Do Data on Millions of U.S. Workers Reveal about Life-Cycle Earnings Dynamics?' Working Paper, University of Minnesota.

Heathcote, J., Perri, F., and Violante, G. L. (2010). 'Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006.' *Review of Economic Dynamics*, vol. 13(1):pages 15–51.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). 'Optimal Tax Progressivity: An Analytical Framework.' *The Quarterly Journal of Economics*, vol. 132(4):pages 1693–1754.

Huggett, M. and Kaplan, G. (2011). 'Human capital values and returns: Bounds implied by earnings and asset returns data.' *Journal of Economic Theory*, vol. 146:page 897919.

Huggett, M. and Kaplan, G. (2012). 'The Money Value of a Man.' NBER WP 18066.

Huggett, M. and Kaplan, G. (2016). 'How large is the Stock Component of Human Capital?' *Review of Economic Dynamics*, vol. 22:page 2151.

Huggett, M. and Ventura, G. (2000). 'Understanding why high income households save more than low income households.' *Journal of Monetary Economics*, vol. 45(2):pages 361–397.

Jappelli, T. and Pistaferri, L. (2014). 'Fiscal policy and MPC heterogeneity.' *American Economic Journal: Macroeconomics*, vol. 6(4):pages 107–36.

Jorgenson, D. and Fraumeni, B. (1989). 'The Accumulation of Human and Non-Human Capital, 1948-1984.' In R. Lipsey and H. Tice, eds., *The Measurement of Saving, Investment and Wealth*. NBER Studies in income and Wealth.

Kaplan, G. and Violante, G. L. (2014). 'A model of the consumption response to fiscal stimulus payments.' *Econometrica*, vol. 82(4):pages 1199–1239.

Kaplan, G., Violante, G. L., and Weidner, J. (2014). 'The Wealthy Hand-to-Mouth.' *Brookings Papers on Economic Activity*, vol. 2014(1):pages 77–138.

Kaymak, B. and Poschke, M. (2016). 'The evolution of wealth inequality over half a century: The role of taxes, transfers and technology.' *Journal of Monetary Economics*, vol. 77:pages 1–25.

Kocherlakota, N. R. (2010). *The new dynamic public finance*. Princeton University Press.

Krueger, D. and Ludwig, A. (2018). 'Optimal Taxes on Capital in the OLG Model with Uninsurable Idiosyncratic Income Risk.' CEPR Working Paper 12717.

Kuhn, M. and Rios-Rull, J.-V. (2016). '2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics.' *Quarterly Review*, pages 1–75. April.

Lemieux, T., Macleod, W. B., and Parent, D. (2009). 'Performance Pay and Wage Inequality.' *Quarterly Journal of Economics*, vol. 124(1).

Levy, F. and Murnane, R. J. (1992). 'US Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations.' *Journal of Economic Literature*, vol. 30(3):pages 1333–1381.

Li, Q. and Racine, J. S. (2007). *Nonparametric econometrics: theory and practice*. Princeton University Press.

Lillard, L. A. (1977). 'Inequality: Earnings vs. Human Wealth.' *The American Economic Review*, pages 42–53.

Lucas, R. E. (1978). 'Asset prices in an exchange economy.' *Econometrica: Journal of the Econometric Society*, pages 1429–1445.

Lucas, R. E. and Moll, B. (2014). 'Knowledge Growth and the Allocation of Time.' *Journal of Political Economy*, vol. 122(1):pages 1–51.

Makridis, C. (2014). 'The Performance Pay Premium, Human Capital and Inequality: Evidence from Over Forty Years of Microdata.' Tech. rep., Working paper.

Parker, J. A., Souleles, N. S., Johnson, D. S., and McClelland, R. (2013). 'Consumer spending and the economic stimulus payments of 2008.' *American Economic Review*, vol. 103(6):pages 2530–53.

Pfeffer, F. T., Schoeni, R. F., Kennickell, A., and Andreski, P. (2016). 'Measuring wealth and wealth inequality: Comparing two US surveys.' *Journal of Economic and Social Measurement*, vol. 41(2):pages 103–120.

Piketty, T. and Saez, E. (2006). 'The evolution of top incomes: a historical and international perspectives.' In *American Economic Review Papers and Proceedings*, pages 200–205. Volume 96, number 2.

R. Haveman, R. Bershadker, A. and Schwabish, J. (2003). 'Human Capital in the United States from 1975 to 2000: Patterns of Growth and Utilization.' Upjohn Institute.

Saez, E. and Kopczuk, W. (2004). 'Top Wealth Shares in the United States: 1916-2000: Evidence from Estate Tax Returns.' NBER Working Paper 10399.

Saez, E. and Zucman, G. (2014). 'Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data.' NBER working paper 20625.

Sanders, C. and Taber, C. (2012). 'Life-cycle wage growth and heterogeneous human capital.' *Annu. Rev. Econ.*, vol. 4(1):pages 399–425.

Straub, L. (2018). 'Consumption, Savings, and the Distribution of Permanent Income.' Unpublished manuscript, Harvard University.

Zeldes, S. P. (1989). 'Consumption and liquidity constraints: an empirical investigation.' *Journal of Political Economy*, vol. 97(2):pages 305–346.

A The Lifetime Wealth of Households

A.1 Preliminaries

In what follows we provide an example of a rich household model that gives rise to the human capital valuation equation 1. For this illustration we focus on a partial equilibrium setting and treat prices and asset returns as exogenous. The model features both single individuals and married couples; agents face both idiosyncratic and aggregate sources of risk, and realized shocks can lead to marriage and divorce.

We consider a partial equilibrium model because (i) we are only interested in using it to derive an agent's human capital valuation equation, and (ii) this means that we can keep the model flexible in a number of dimensions without losing tractability. We allow for a general vector of time-varying individual characteristics, many of which may affect wages. For example, such individual characteristics can include variables related to human capital; current labor supply decisions can affect future wages through the dynamics of these characteristics. By choice, we do not restrict the shape of the effects of these variables; this preserves the generality of the model and illustrates that we do not need to take a stand on the mechanics of every individual or household process in order to derive the human capital valuation equation. Rather, we only need to explicitly write out the Bellman equations of households and be precise about their expectations over future state variables. In that regard, we develop a notation where subscripts on expectation operators make clear the variables with respect to which expectations are being taken.

In our notation, the state of the economy at time t is represented by Ω_t. The history of states of the world is then $\Omega^t = \{\Omega_0, \Omega_1, \ldots, \Omega_t\}$. Ω_t includes realizations of all aggregate and idiosyncratic risk. An individual's observable characteristics, such as education, age, gender, etc., are contained in the vector X_{it}. An individual's unobservable type, which may be informative about their expected earnings or consumption profile, is denoted by η_i. If an individual is married they will have a spouse with observable characteristics X_{jt} and unobservable type η_j.

In what follows we define a household's portfolio of (non human) wealth as a vector containing various assets and liabilities. For an unmarried household this vector is $\mathbf{a}_{it} = \{a_{it}^\kappa\}_{\kappa \in k}$, where a_{it}^κ is the individual's position in asset κ. For a married household consisting of an individual i and their spouse j, the wealth portfolio is $\mathbf{a}_{(ij)t}$. Individuals receive wage offers w_{it} every period; implicitly these offers may depend on all household state variables, for example X_{it}. Households pay taxes and receive transfers according to a function $T_t(\cdot)$.

A.2 Human Wealth: Household's Problem

To obtain the pricing equation for this general model specification, we proceed in two steps. First, we posit the life-cycle problem of individuals who choose their savings, how much labor they supply in the market, as well as their marital status. Second, we use the value functions of these individuals to derive a general human capital pricing equation.

Individual Value Functions. An individual enjoys utility from consumption and leisure, denoted $u(c_{it}, \ell_{it})$, and (possibly) from being married to their spouse, denoted $\heartsuit_{it(j)}$. An individual's value function when single, V_i^S, depends on their own state variables and their beliefs about marital prospects. The value function when married, V_i^M, depends on both own and spousal state variables, and beliefs about the prospect of remaining married. An individual may supply a fraction h_{it} of their time in the labour market, for which they earn a wage w_{it}. Wages vary with X_{it} and Ω_t.

If individual i is single at time t their value function V_i^S will depend on a continuation value at time $t+1$ that includes the possibilities of choosing to get married or remain single in the following period:

$$V_i^S(\mathbf{a}_{it}, X_{it}, \eta_i, \Omega^t) = \max_{c_{it}, \ell_{it}, h_{it}, \mathbf{a}_{it+1}} \left\{ u\left(c_{it}, \ell_{it}\right) \right. \tag{22}$$
$$+ \beta\left(1 - \mu_{it}\right) E_{\{\Omega^{t+1}\}} \left[V_i^S(\mathbf{a}_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1})\right]$$
$$\left. + \beta\mu_{it} E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}} \left[V_i^M(\mathbf{a}_{(ij)t+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right] \right\}.$$

The probability $\mu_{it} = \mu\left(X_{it}, \eta_i, \Omega^t\right)$ is the conditional probability that i chooses to get married next period, after meeting potential partners. This probability depends on individual characteristics and the state of the world. In the event that i chooses to marry, their indirect utility will depend on the wealth and characteristics of their partner, $\mathbf{a}_{jt+1}$ and X_{jt+1}, as well as the state of the world next period. Thus, the expected value of being married is taken over the distribution of these variables among the j individuals that person i might choose to marry. The assets of a newly formed married household will be the sum of the spouses initial individual assets: $\mathbf{a}_{(ij)t+1} = \mathbf{a}_{it+1} + \mathbf{a}_{jt+1}$. The consumption choice of i is defined over their current budget set

$$\sum_{\kappa \in k} a_{it+1}^\kappa + c_{it} \leq w_{it}h_{it} + \sum_{\kappa \in k} R_t^\kappa a_{it}^\kappa - T_t\left(\mathbf{a}_{it}, w_{it}, h_{it}\right), \tag{23}$$

where R_t^κ is the one-period return on asset κ, and $T\left(\mathbf{a}_{it}, w_{it}, h_{it}\right)$ is a function summarizing all tax liabilities. The individual's time constraint $\ell_{it} = 1 - h_{it}$ and current borrowing constraint $\sum_{\kappa \in k} a_{it+1}^\kappa \geq \underline{a}_{it}$ also affect these choices.

If individual i is married to individual j at time t, then i's value function will include a continuation value that allows for the possibilities of staying married or separating in the following year:

$$V_i^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \tag{24}$$
$$u\left(c_{it}^*, \ell_{it}^*\right)$$
$$+ \beta\left\{(1 - \tilde{\mu}_{it})E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}} \left[V_i^S(\mathbf{a}_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) | \mathbf{a}_{(ij)t+1}^*\right]\right.$$
$$\left. + \tilde{\mu}_{it} E_{\{\Omega^{t+1}\}} \left[V_i^M(\mathbf{a}_{(ij)t+1}^*, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right]\right\} + \heartsuit_{it(j)}.$$

In the above equation the values $\left(\mathbf{a}_{(ij)t+1}^*, c_{it}^*, \ell_{it}^*\right)$ are the values of household savings, as well as consumption and leisure for individual i, that result from the joint household optimization problem described below. The parameter $\tilde{\mu}_{it} = \mu\left(X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t\right)$ is the conditional probability of a household choosing to stay

married. If the household divorces before next period their asset portfolio is split and individual i receives a part $\mathbf{a}_{it+1}$ of it. Because there may be uncertainty about the divorce settlement, a conditional expectation over possible asset divisions is taken when evaluating the divorce part of the continuation value. While we don't model the choice of getting married explicitly, we assume that the marriage shock $\heartsuit_{it(j)}$ captures the presence of non-pecuniary returns to being married to person j. These returns are assumed to be additively separable and drop out of all marginal calculations.

Household Planner Problem. Once married, the joint optimization problem of the spouses can be viewed as that of a planner who maximizes a weighted average of the spouses' utilities using a set of Pareto weights. Above we have denoted by V_i^M the utility of person i when they are assigned the allocations that the household planner finds optimal. Next, we need to distinguish this from person i's utility under (possibly) non-optimized allocations, which we denote by $\tilde{V}_i^M$. The problem of the household planner is:

$$
\begin{aligned}
V_{(ij)}^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \max_{\mathbf{b}_{(ij)t}} \Big\{ & \lambda_{(ij)} \tilde{V}_i^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) \\
& + (1 - \lambda_{(ij)}) \tilde{V}_j^M(\mathbf{a}_{(ij)t}, X_{jt}, X_{it}, \eta_j, \eta_i, \Omega^t) \Big\},
\end{aligned}
\tag{25}
$$

where the decision vector is $\mathbf{b}_{(ij)t} = \{c_{it}, c_{jt}, \ell_{it}, \ell_{jt}, h_{it}, h_{jt}, \mathbf{a}_{(ij)t+1}\}$, and $\lambda_{(ij)}$ is the Pareto weight on individual i in the household planning problem.

The feasible consumption set for married households is determined by the budget constraint

$$
\sum_{\kappa \in k} a_{(ij)t+1}^{\kappa} + c_{(ij)t} \leq w_{it} h_{it} + w_{jt} h_{jt} + \sum_{\kappa \in k} R_t^{\kappa} a_{(ij)t}^{\kappa} - T_t \left(\mathbf{a}_{(ij)t}, w_{it}, w_{jt}, h_{it}, h_{jt} \right),
\tag{26}
$$

where $c_{(ij)t}$ is total consumption expenditure of the household. This is related to the consumption resources allocated to each spouse by the constraint $c_{(ij)t} = \vartheta(c_{it} + c_{jt})$, where ϑ represents an adult equivalence scale. Individual time allocation constraints $\ell_{it} = 1 - h_{it}$ and $\ell_{jt} = 1 - h_{jt}$, and the household borrowing limit $\sum_{\kappa \in k} a_{(ij)t+1}^{\kappa} \geq \underline{a}_{(ij)t}$ also constrain the household planner's choices.

A.3 Derivation of the Pricing Equation

We derive an individual's valuation of his/her own human capital by determining the shadow price of an asset that exactly replicates that individual's state-contingent labor market outcomes. To accomplish this we introduce an asset that pays dividends equal to i's earnings, while viewing i as committed to their state-contingent labor supply plan.[38] Because of this commitment we replace $w_{it} h_{it}$ from the problems described above with y_{it}, with the understanding that y_{it} is state-contingent earnings under the optimal labor supply plans of problems (22) and (25) above. As noted by Huggett and Kaplan (2012), this approach to valuing

---

[38]Of course, in reality no one would be willing to buy this asset from i because of the inherent commitment problem. Hence, the valuation we derive is truly a shadow price representing what human capital is worth to its owners. As discussed at length by Benzoni and Chyruk (2015), it is not normally possible to enforce contracts written against future labor services and ownership of human capital is not transferable (that is, human capital is a non-traded asset).

non-traded assets was first introduced by Lucas (1978). Theorem 1 in Huggett and Kaplan (2012) formally states the conditions under which human capital can be valued using the 'non-traded asset' approach.

Human Capital Valuations of Married Individuals. We begin by valuing individual i's human capital when i is married. The number of shares of the hypothetical asset that i's household owns at time t is e_{it}, and the price of this asset is θ_{it}. We could also introduce an asset based on j's human capital, but that is not necessary to value i's human capital, hence we suppress that notation for now. When the hypothetical asset e_{it} is introduced, the budget constraint for a married household becomes:

$$\sum_{\kappa \in k} a_{(ij)t+1}^{\kappa} + c_{(ij)t} + \theta_{it}e_{it+1} \leq \theta_{it}e_{it} + (1 + e_{it})y_{it} + y_{jt} + \sum_{\kappa \in k} R_t^{\kappa} a_{(ij)t}^{\kappa} - T_t\left(\mathbf{a}_{(ij)t}, y_{it}, y_{jt}\right). \quad (27)$$

Furthermore, we include e_{it} as an additional state variable in the household planner's problem in equation (25), as well as in the definition of an individual's utility from marriage in (24). Given these adjustments we can rewrite the household planner's problem in a recursive manner as:

$$V_{(ij)}^M(\mathbf{a}_{(ij)t}, e_{it}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \max_{\mathbf{b}_{(ij)t}} \Big\{ \lambda_{(ij)} u\left(c_{it}, \ell_{it}\right) + (1 - \lambda_{(ij)})u\left(c_{jt}, \ell_{jt}\right) \quad (28)$$

$$+ \lambda_{(ij)}\beta(1 - \tilde{\mu}_{(ij)t})E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}}\left[V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1})|\mathbf{a}_{(ij)t+1}\right]$$

$$+ (1 - \lambda_{(ij)})\beta(1 - \tilde{\mu}_{(ij)t})E_{\{\Omega^{t+1}, \mathbf{a}_{jt+1}\}}\left[V_j^S(\mathbf{a}_{jt+1}, X_{jt+1}, \eta_j, \Omega^{t+1})|\mathbf{a}_{(ij)t+1}\right]$$

$$+ \beta\tilde{\mu}_{(ij)t}E_{\{\Omega^{t+1}\}}\left[V_{(ij)}^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right] \Big\},$$

where the decision vector $\mathbf{b}_{(ij)t}$ now includes e_{it+1}. After using the budget constraint in (27) to substitute c_{it} out of the problem in (28), we derive the following first-order condition for the optimal choice of e_{it+1}:

$$u_c(c_{it}, \ell_{it})\vartheta\theta_{it} = \beta(1 - \tilde{\mu}_{(ij)t})\frac{\partial}{\partial e_{it+1}}E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}}\left[V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1})|\mathbf{a}_{(ij)t+1}\right] \quad (29)$$

$$+ \frac{1}{\lambda_{(ij)}}\beta\tilde{\mu}_{(ij)t}\frac{\partial}{\partial e_{it+1}}E_{\{\Omega^{t+1}\}}\left[V_{(ij)}^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right].$$

To proceed we must calculate the derivatives of the married and single continuation values using envelope conditions. For the married continuation value this involves straightforward differentiation of equation (28) with respect to e_{it}, noting that the c_{it} has been replaced by the budget constraint. The result is,

$$\frac{\partial}{\partial e_{it+1}}V_{(ij)}^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) = \quad (30)$$

$$\lambda_{(ij)}u_c(c_{it+1}^M, \ell_{it+1}^M)\vartheta\left(\theta_{it+1}^M + y_{it+1}^M\right),$$

where the superscript M indicates quantities that arise during marriage. To obtain the derivative of a single person's value function we must first be explicit about the problem they solve when single. Extending

49

equation (22) to include the hypothetical asset e_{it+1} results in the following problem:

$$V_i^S(\mathbf{a}_{it}, e_{it}, X_{it}, \eta_i, \Omega^t) = \tag{31}$$

$$\max_{c_{it}, \ell_{it}, h_{it}, \mathbf{a}_{it+1}} \left\{ u\left(c_{it}, \ell_{it}\right) + \beta\left(1 - \mu_{it}\right) E_{\{\Omega^{t+1}\}} \left[V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) \right] \right.$$

$$\left. + \beta\mu_{it} E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}} \left[V_i^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) \right] \right\}.$$

The maximization in (31) is subject to the usual time allocation and borrowing constraints, as well the extended budget constraint,

$$\sum_{\kappa \in k} a_{it+1}^\kappa + c_{it} + \theta_{it} e_{it+1} \leq \theta_{it} e_{it} + (1 + e_{it}) y_{it} \tag{32}$$

$$+ \sum_{\kappa \in k} R_t^\kappa a_{it}^\kappa - T_t\left(\mathbf{a}_{it}, w_{it}, h_{it}\right).$$

The derivative of the value function in (31) can thus be derived by replacing c_{it} with the extended budget constraint, resulting in:

$$\frac{\partial}{\partial e_{it+1}} V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) = u_c(c_{it+1}^S, \ell_{it+1}^S)\left(\theta_{it+1}^S + y_{it+1}^S\right). \tag{33}$$

Finally, using equations (30) and (33), one can re-arrange the first order condition for optimal e_{it+1} chosen by a married household (equation 29) into an expression describing the valuation of i's human capital θ_{it}^M (the purchase price per share of e_{it+1}):

$$\theta_{it}^M = \beta(1 - \tilde{\mu}_{(ij)t}) \frac{1}{\vartheta} E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}} \left[\frac{u_c(c_{it+1}^S, \ell_{it+1}^S)}{u_c(c_{it}, \ell_{it})} \left(y_{it+1}^S + \theta_{it+1}^S\right) \right] \tag{34}$$

$$+ \beta\tilde{\mu}_{(ij)t} E_{\{\Omega^{t+1}\}} \left[\frac{u_c(c_{it+1}^M, \ell_{it+1}^M)}{u_c(c_{it}, \ell_{it})} \left(y_{it+1}^M + \theta_{it+1}^M\right) \right].$$

The result that stochastic discount factors are a component of the value of human capital in this model is related to general asset pricing formulations found in the literature following the work of Lucas (1978) and Huggett and Kaplan (2011, 2012). The probability of a change in marital status, and the surplus generated by marriage (through the economies of scale parameter ϑ) also factor into our valuation results.

Human Capital Valuations for Single Individuals. We derive the human capital valuation equations of an unmarried individual by considering their first-order condition for the optimal choice of e_{it+1} in

problem (31):

$$u_c(c_{it}, \ell_{it})\theta_{it} = \tag{35}$$

$$\beta(1 - \mu_{it})\frac{\partial}{\partial e_{it+1}} E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}} \left[V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) | \mathbf{a}_{(ij)t+1} \right]$$

$$+ \beta\mu_{it}\frac{\partial}{\partial e_{it+1}} E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}} \left[V_i^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) \right].$$

As was the case when deriving valuations for married individuals, we need to substitute out the derivatives of continuation values. For the derivative of $V_i^S(\cdot)$ this is straightforward, and in fact we have the expression in equation (33) already. However, the derivative of $V_i^M(\cdot)$ proves more difficult because we cannot resort to a standard envelope condition. This is the case because $V_i^M(\cdot)$ is *not* an indirect utility function, or in other words is *not* the solution to an individual optimization problem. Rather, $V_i^M(\cdot)$ is a component of the objective function maximized by the household planner. To compute the necessary derivative here we must first characterize the effect of pre-marital investments on the utility allocated to the spouse making those investments, which requires us to make assumptions about how the Pareto weight $\lambda_{(ij)}$ is determined in the event that i gets married. Indeed, valuation of pre-marital human capital investments is inextricably linked to the household bargaining process upon marriage.

As anticipated above, we assume symmetric Nash Bargaining over the surplus generated by marriage. Under this assumption we can derive a relationship pinning down how the marital utility of person i changes if they make pre-marital investments. Symmetric Nash Bargaining implies that i's utility in marriage must increase by at least as much as their outside option (utility from being single), plus half of any surplus generated by pre-marital investment.

Specifically we assume that a married household's Pareto weight solves

$$\max_{\{V_i^M, V_j^M\}} \left(V_i^M - V_i^S \right) \left(V_j^M - V_j^S \right), \tag{36}$$

where we have suppressed the state variables within the value functions for clarity. Let $G(V_i^M, V_j^M) = 0$ be the Pareto frontier of household allocations, in which case the Nash Bargaining solution must satisfy

$$\left(V_i^M - V_i^S \right) = \frac{G_2}{G_1} \left(V_j^M - V_j^S \right). \tag{37}$$

To translate this condition into something empirically useful, note that an equivalent formulation of the household planning problem in equation (28) is:

$$\max \left\{ \lambda_{(ij)} V_i^M + (1 - \lambda_{(ij)}) V_j^M \right\}$$
$$\text{subject to}$$
$$G(V_i^M, V_j^M) = 0.$$

Combining the first-order conditions from this problem with those from the underlying Nash Bargaining problem results in:

$$\left(V_i^M - V_i^S\right) = \frac{1 - \lambda_{(ij)}}{\lambda_{(ij)}} \left(V_j^M - V_j^S\right). \tag{38}$$

The equivalence of equations (37) and (38) is due to the fact that $\lambda_{(ij)}$ is the Pareto weight that implicitly solves the Nash Bargaining problem in equation (36).

Next, we examine equation (37) evaluated at the point at which person i brings exactly zero units of e_{it} to the marriage, as this is the solution we observe in the data. Computing the total differential of this equation with respect to e_{it} results in

$$\frac{\partial V_i^M}{\partial e_{it}} - \frac{\partial V_i^S}{\partial e_{it}} = \left(\frac{G_2}{G_1}\right) \frac{\partial V_j^M}{\partial e_{it}} + \frac{1}{G_1} \left(\frac{\partial G2}{\partial e_{it}} \left(V_j^M - V_j^S\right) - \frac{\partial G1}{\partial e_{it}} \left(V_i^M - V_i^S\right)\right). \tag{39}$$

While this expression may seem intractable, one can easily show that at the optimal solution to the household planner's problem

$$\left(\frac{\partial G2}{\partial e_{it}} \middle/ \frac{\partial G1}{\partial e_{it}}\right) = \frac{u_c(c_i t, \ell_{it})}{u_c(c_j t, \ell_{jt})} = \frac{\lambda_{(ij)}}{1 - \lambda_{(ij)}}. \tag{40}$$

Therefore, the last term of equation (39) equals zero when evaluated at the solution to the bargaining problem. Thus, a final simplified relationship between the derivatives of individual utilities, evaluated at the solution to the bargaining problem, is

$$\frac{\partial V_i^M}{\partial e_{it}} - \frac{\partial V_i^S}{\partial e_{it}} = \frac{1 - \lambda_{(ij)}}{\lambda_{(ij)}} \frac{\partial V_j^M}{\partial e_{it}}. \tag{41}$$

Intuitively, the extent to which i's utility in marriage will increase in excess of their outside option depends on their ex-post Pareto weight and how valuable the hypothetical asset would be to their spouse.

To utilize equation (41), first note that the definition of the household planner's optimization objective in (25) implies that the envelope condition in (30) can be re-written as:

$$\lambda_{(ij)} \frac{\partial V_{it+1}^M}{\partial e_{it+1}} + (1 - \lambda_{(ij)}) \frac{\partial V_{jt+1}^M}{\partial e_{it+1}} = \lambda_{(ij)} u_c(c_{it+1}^M, \ell_{it+1}^M) \vartheta \left(\theta_{it+1}^M + y_{it+1}^M\right). \tag{42}$$

Combining this with the Nash Bargaining implication in (41), we obtain an extremely useful result characterizing the effect of pre-marital investments on the utility within marriage:

$$\frac{\partial V_{it+1}^M(\cdot)}{\partial e_{it+1}} = \frac{1}{2} u_c(c_{it+1}^M, \ell_{it+1}^M) \frac{1}{\vartheta} \left(\theta_{it+1}^M + y_{it+1}^M\right) + \frac{1}{2} \frac{\partial V_{it+1}^S(\cdot)}{\partial e_{it+1}}. \tag{43}$$

The intuition for this equation relates to how much of the return on the hypothetical asset will be allocated to individual i by the household planner. A lower bound is the change in their utility if they exercise their outside option, which is captured by $\partial V_{it+1}^S/\partial e_{it+1}$. An upper bound is the marginal change in their utility if the entire return on the asset, including economies of scale, is allocated to i. With symmetric bargaining

52

exactly half of the component pertaining to returns that exceeds the effect on i's outside option is paid to i.

Equation (43) is useful because we now have an expression to substitute into equation (29), which was our objective when we set out to analyze the bargaining problem. Doing this, and substituting the envelope condition for single households in equation (33), allows us to derive the following valuation formula for the human capital of a currently unmarried person i:

$$
\begin{aligned}
\theta_{it}^S =& \beta(1 - \frac{\mu_{it}}{2})E_{\{\Omega^{t+1}\}}\left[\frac{u_c(c_{it+1}^S, \ell_{it+1}^S)}{u'(c_{it}, \ell_{it})}\left(y_{it+1}^S + \theta_{it+1}^S\right)\right] \\
&+ \beta\frac{\mu_{it}}{2}E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}}\left[\frac{u_c(c_{it+1}^M, \ell_{it+1}^M)}{u'(c_{it}, \ell_{it})}\frac{1}{\vartheta}\left(y_{it+1}^M + \theta_{it+1}^M\right)\right].
\end{aligned}
\tag{44}
$$

While this expression is similar to canonical asset pricing formulations, it makes clear that the correct pricing relationship involves a biased expectation of future returns to human capital, where the bias derives from the implicit extra weight single households place on outcomes in the event of remaining single. The above equation is also informative as to how one would test the robustness of the symmetric bargaining assumption: asymmetric bargaining weights would result in factors other than 1/2 (but still on the unit interval) being used to re-weight single and married outcomes.

We can subsume all sources of uncertainty into a single expectation operator $\mathbb{E}_{it}$, which also accounts for the re-weighting of unmarried future outcomes (as opposed to an unweighted expectation E_{it}). For simplicity, in our empirical work we assume bargaining weights such that $E_{it} = \mathbb{E}_{it}$. Having done this we can summarize the value of human capital as we do in equation (1), where future variables implicitly depend on marital status. Clearly, valuations of one's own human capital depend on stochastic discount factors. Thus, state-contingent realizations of individual consumption matter for valuing state-contingent human capital payoffs. The last step in our analysis is to evaluate equation (1) at the point $e_{it} = 0$ so that the equation is analogous to real-world valuations where human capital assets are not traded. Then, given some estimate of the distribution of state-contingent consumption realizations and appropriate weighting of future outcomes, human capital valuations can be estimated.

B Incorporating Biennial Data

Biennial Data in Marginal utility Estimation. If data are available only at two year intervals, the empirical counterpart of the Euler equation becomes

$$
\hat{u}_c(q) = \hat{\beta}^2 \sum_{i=1}^N \sum_{t \in \tau_o(i)} \hat{u}_c(q_{it}')R_{it}'\phi_{it}(q),
\tag{45}
$$

where q_{it}' and R_{it}' denote decisions and asset returns two years later. One complication that the interpretation of the largest eigenvalue of Φ is $1/\beta^2$, rather than $1/\beta$ as for the annual observations. Our solution entails the transformation $\hat{\beta}^2 = \hat{\beta}\beta_0$, where β_0 is some initial estimate of $\hat{\beta}$ (possibly based only on the annual

data sample). Then, after replacing R'_{it} with $\tilde{R}'_{it} = \beta_0 R'_{it}$, we employ biennial observations in the following moment condition:

$$\hat{u}_c(q) = \hat{\beta} \sum_{i=1}^{N} \sum_{t \in \tau_o(i)} \hat{u}_c(q'_{it}) \tilde{R}'_{it} \phi_{it}(q). \tag{46}$$

Now the largest eigenvalue of Φ can be correctly interpreted as $1/\beta$ for all observations in a combined sample of annual and biennial data. However, the estimates of $\hat{\beta}$ and $\hat{u}_c$ are conditional on β_0, hence they can be improved upon if a better estimate of β_0 becomes available. We replace β_0 by $\hat{\beta}$ and re-estimate, iterating this procedure until $\hat{\beta}$ is approximately equal to the guess β_0 and no further improvement is feasible.

Biennial Data in Human Wealth Estimation. To accommodate biennial data in this step we denote $\tau_j^1(i)$ and $\tau_j^2(i)$ as the set of annual and biennial sample years, respectively, in which i was of age j. For an observation drawn during a period of biennial sampling, equation (20) can be rewritten by iterating the valuation equation one-year further into the future:

$$\hat{\theta}(j, \mathbf{z}) = \hat{g}^1(j, \mathbf{z}) + \hat{g}^2(j, \mathbf{z}) + \sum_{i=1}^{N} \sum_{t \in \tau_j^2(i)} \hat{\theta}(j+2, \mathbf{z}'_{it}) \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} \gamma_{it}(\mathbf{z}). \tag{47}$$

Our notation is such that, for the biennial sample, q'_{it} and $\mathbf{z}'_{it}$ are data observations two years into the future. The functions $\hat{g}^1(j, \mathbf{z})$ and $\hat{g}^2(j, \mathbf{z})$ are estimates of the conditional expectation of discounted earnings, one and two years ahead, for a j year old individual with current state vector $\mathbf{z}$. These estimates are computed as in equation (18), where $\hat{g}^1$ is estimated using data from the annual sample period and $\hat{g}^2$ is estimated using data from the biennial sample period.

As before, we form vectors $\tilde{\Theta}$ and $\tilde{G}$, as well as a matrix Γ such that $\tilde{\Theta} = \tilde{G} + \Gamma \tilde{\Theta}$. Some elements of $\tilde{\Theta}$ and $\tilde{G}$ are based on annual observations using equation (20) with $\tau(i)$ replaced by $\tau^1(i)$, and others are based on biennial observations using equation (47). 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

$$\Gamma = \begin{pmatrix} 0 & \Gamma_1^1 & \Gamma_1^2 & \dots & 0 \\ 0 & 0 & \ddots & \ddots & \vdots \\ \vdots & \vdots & & 0 & \Gamma_{J-1}^1 \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}, \tag{48}$$

where Γ_j^1 and Γ_j^2 are constructed as explained in equation (21). The reason we now have two blocks in each row of Γ is to allow rows corresponding to annual observations to 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 (21), 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 non-zero. After constructing such a matrix Γ we can solve for $\tilde{\Theta} = (I = \Gamma)^{-1} \tilde{G}$ 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 (20) and (47). We define numbers of annual and biennial observations $n^1 = \sum_{i=1}^{N} \sum_{t \in \tau^1(i)} 1$ and $n^2 = \sum_{i=1}^{N} \sum_{t \in \tau^2(i)} 1$. Using these counts we form the estimator as

$$
\hat{\theta}(j, \mathbf{z}) = \hat{g}^1(j, \mathbf{z}) + \frac{n^1}{n^1 + n^2} \left(\hat{\beta} \sum_{i=1}^{N} \sum_{t \in \tau_j^1(i)} \hat{\theta}(j+1, \mathbf{z}'_{it}) \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} \gamma_{it}(\mathbf{z}) \right) \tag{49}
$$

$$
+ \frac{n^2}{n^1 + n^2} \left(\hat{g}^2(j, \mathbf{z}) + \hat{\beta}^2 \sum_{i=1}^{N} \sum_{t \in \tau_j^2(i)} \hat{\theta}(j+2, \mathbf{z}'_{it}) \frac{\hat{u}_c(q'_{it})}{\hat{u}_c(q_{it})} \gamma_{it}(\mathbf{z}) \right).
$$

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.

C   Supplementary Figures and Tables

Figure 16: Average human wealth over the life cycle by latent type. Values in 2016 dollars.

(a) Average human wealth by earnings type.

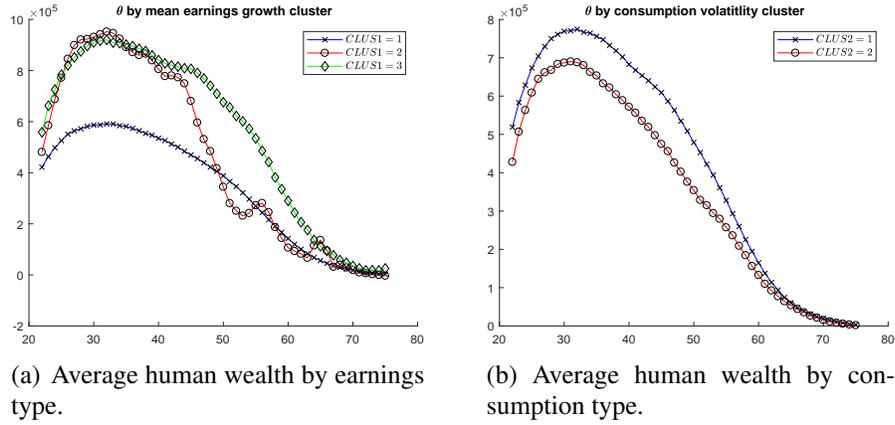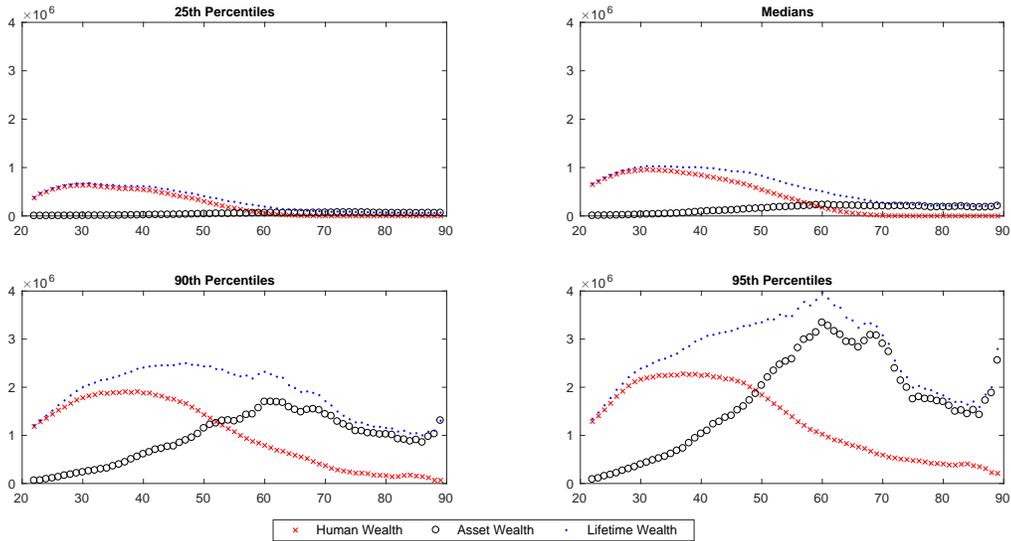(b) Average human wealth by consumption type.

Figure 17: Average human wealth, net worth and lifetime wealth over the lifecycle, across their respective percentiles. Values in 2016 dollars.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.366	0.405	0.372	0.424
1992	0.667	0.365	0.396	0.372	0.415
1995	0.677	0.370	0.412	0.376	0.433
1998	0.683	0.361	0.430	0.363	0.445
2001	0.693	0.365	0.454	0.394	0.477
2004	0.691	0.375	0.463	0.388	0.491
2007	0.712	0.375	0.488	0.416	0.515
2010	0.741	0.381	0.483	0.429	0.522
2013	0.747	0.394	0.497	0.432	0.536
2016	0.768	0.399	0.543	0.472	0.579

Table 12: All sample years: this table reports the share of variable "X" in the hands of the households in the top 10% of the distribution of that same variable "X". For example, the share of earnings held by the households in the top 10% of the distribution of earnings.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.125	0.351	0.203	0.385
1992	0.667	0.120	0.330	0.202	0.369
1995	0.677	0.126	0.349	0.215	0.385
1998	0.683	0.129	0.381	0.217	0.409
2001	0.693	0.131	0.420	0.256	0.452
2004	0.691	0.125	0.429	0.246	0.469
2007	0.712	0.122	0.457	0.267	0.490
2010	0.741	0.121	0.454	0.272	0.501
2013	0.747	0.125	0.462	0.279	0.509
2016	0.768	0.125	0.521	0.311	0.564

Table 13: All sample years: this table reports the share of variable "X" in the hands of the households in the top 10% of the distribution of Net Worth. For example, the share of earnings held by the households in the top 10% of the distribution of net worth.

	Dependent Variables	
	$\mathbb{I}\{c_{it} < \pi_{it}\}$	$\ln(\pi_{it}) - \ln(c_{it})$
	(1)	(2)
constant	0.7452	0.2976
	(0.0062)	(0.0068)
$\mathbb{I}\left\{\frac{c_{it+1}}{c_{it}} > 1.5\right\}$	0.1416	0.2960
	(0.0108)	(0.0176)
$\mathbb{I}\left\{\frac{c_{it+2}}{c_{it+1}} > 1.5\right\}$	0.0494	0.1420
	(0.0156)	(0.0213)
$\mathbb{I}\left\{\frac{c_{it+3}}{c_{it+2}} > 1.5\right\}$	0.0689	0.1124
	(0.0140)	(0.0187)
$\mathbb{I}\left\{\frac{c_{it+4}}{c_{it+3}} > 1.5\right\}$	0.0459	0.0900
	(0.0152)	(0.0194)
N-observations	7,386	7,386

Table 14: Deviations of consumption flows c_{it} from permanent-income π_{it}. This table reports results for two linear regressions: (1) a probability model with dependent variable equal to 1 when $\{c_{it} < \pi_{it}\}$ and zero otherwise; (2) a continuous model in which the dependent variable is the log difference $\ln(\pi_{it}) - \ln(c_{it})$. Both regressions include dummies taking unit value if consumption growth exceeds 50% between any consecutive periods between t and $t + 4$.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.668	0.366	0.405	0.372	0.424
1992	0.668	0.364	0.395	0.371	0.407
1995	0.679	0.370	0.410	0.378	0.425
1998	0.687	0.360	0.427	0.360	0.440
2001	0.698	0.363	0.449	0.389	0.467
2004	0.702	0.367	0.454	0.376	0.474
2007	0.723	0.362	0.473	0.407	0.497
2010	0.756	0.364	0.460	0.416	0.487
2013	0.761	0.370	0.471	0.406	0.496
2016	0.783	0.367	0.507	0.450	0.542

Table 15: All sample years: this table reports the share of variable "X" in the hands of the households in the top 10% of the distribution of that same variable "X". For example, the share of earnings held by the households in the top 10% of the distribution of earnings.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.296	0.055	0.131	0.101	0.156
1992	0.299	0.056	0.124	0.099	0.153
1995	0.347	0.057	0.149	0.109	0.180
1998	0.336	0.057	0.162	0.102	0.181
2001	0.319	0.058	0.173	0.141	0.195
2004	0.330	0.059	0.185	0.124	0.201
2007	0.333	0.059	0.197	0.144	0.224
2010	0.339	0.061	0.189	0.142	0.231
2013	0.352	0.062	0.198	0.141	0.245
2016	0.384	0.067	0.244	0.194	0.272

Table 16: All sample years: this table reports the share of variable "X" in the hands of the households in the top 1% of the distribution of that same variable "X". For example, the share of earnings held by the households in the top 1% of the distribution of earnings.

Year	Net Worth	Human Wealth	Lifetime Wealth	Earnings	permanent-income
	(1)	(2)	(3)	(4)	(5)
1989	0.296	0.011	0.130	0.039	0.153
1992	0.299	0.013	0.123	0.043	0.150
1995	0.347	0.014	0.148	0.055	0.174
1998	0.336	0.016	0.161	0.053	0.176
2001	0.319	0.016	0.172	0.065	0.185
2004	0.330	0.016	0.184	0.071	0.196
2007	0.333	0.016	0.196	0.066	0.212
2010	0.339	0.013	0.188	0.060	0.210
2013	0.352	0.014	0.198	0.063	0.240
2016	0.384	0.019	0.243	0.107	0.265

Table 17: All sample years: this table reports the share of variable "X" in the hands of the households in the top 1% of the distribution of net worth. For example, the share of earnings held by the households in the top 1% of the distribution of net worth.

D   Testing the Sufficiency of the Information Vector z

We take several steps to test our assumption that the vector $\mathbf{z}_{it}$ does span the information sets of individuals, so that all forward looking information available to households is captured. We first identify three observable decisions (consumption, labor supply and home ownership) that individuals make at age 30 based on information available at that age. Next, we show that these decisions are predictive of realized future earnings and human wealth later in life (age 50). Lastly, we show that the residual in the age 30 decision variables, computed by subtracting variation explained by the vector $\mathbf{z}$, has no additional predictive power for age 50 earnings and human wealth.

| | *Dependent Variable* | | |
| | **Earnings Age 50** | **Realized Lifetime Earn** | **Human Wealth Age 50** |
Explanatory Variable			
Age 30 Consumption	0.094	0.771	0.736
	(0.049)	(0.171)	(0.083)
Resid. Age 30 Cons.	-0.051	0.035	0.159
	(0.062)	(0.236)	(0.104)
Age 30 Home Owner	-0.071	-0.230	-0.193
	(0.031)	(0.139)	(0.058)
Resid. Age 30 Owner	-0.039	0.029	-0.028
	(0.033)	(0.148)	(0.062)
Age 30 Hours Worked	0.248	0.054	0.561
	(0.126)	(0.629)	(0.243)
Resid. Age 30 Hours	0.156	0.182	-0.174
	(0.139)	(0.708)	(0.280)

Table 18: Estimated coefficients on age 30 variables from simple regressions of age 50 dependent variables. All regressions are log-log, with the exception of home ownership, which is binary. Standard errors in parentheses. Residualized explanatory variables are the residuals from initial regressions of the corresponding explanatory variable on $\mathbf{z}_{it}$.

E   Non-Parametric Set Imputation in the SCF Sample

As mentioned in Section 4.3, a subset of the variables in the PSID sample (including unobserved type proxies) cannot be obtained from the repeated cross-section data of the SCF. To address this data limitation we impute the *full distribution* of the missing variables estimated from the PSID by performing a form of non-parametric 'set imputation'. First, using PSID data we partition the data vector $\mathbf{Z}$ into observed variables $\mathbf{Z}^+$ and unobserved variables $\mathbf{Z}^-$. We then define the conditional distribution function $P(\mathbf{Z}^-|\mathbf{Z}^+)$. Because $\mathbf{Z}^-$ takes discretely many values, this distribution can be viewed as a probability mass function

Figure 18: Mean and standard deviation of human wealth over the life cycle (SCF sample, values in 2016 dollars).

(a) Average of human wealth by age

(b) Standard Deviation of human wealth by age
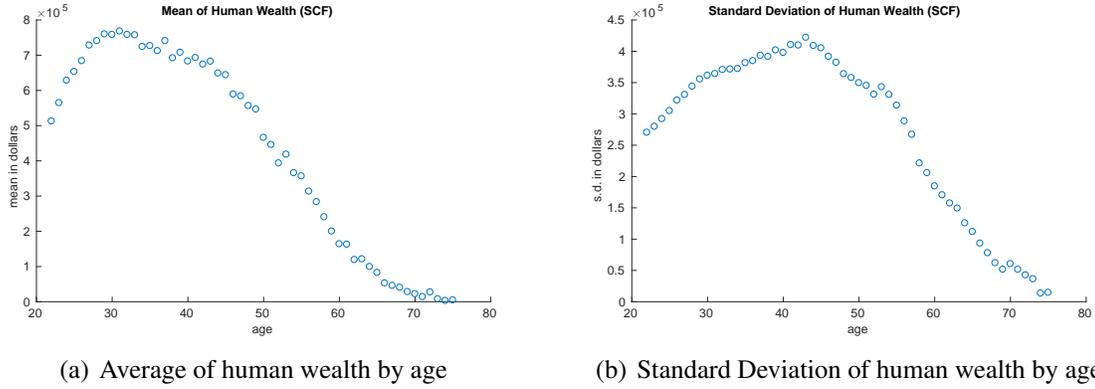
$P(\mathbf{Z}^-|\mathbf{Z}^+) = \{p_1(\mathbf{Z}^+), p_2(\mathbf{Z}^+), \ldots, p_M(\mathbf{Z}^+)\}$, where M is the number of points in the support set of $\mathbf{Z}^-$. In turn, each $p_m(\mathbf{Z}^+)$ can be estimated using the Nadaraya-Watson kernel estimator using PSID data.

Next, we expand the SCF data set so that it replicates the cross-sectional variation of $\mathbf{Z}^-$. We do this by creating M versions of the extended SCF sample, one for each of the M points in the support set of $\mathbf{Z}^-$. Hence, each such version imputes a different point in the support of $\mathbf{Z}^-$. The sample weight for observation i in data version $m \in M$ is rescaled by $p_m(\mathbf{Z}_i^+)$. Finally, we stack these subsamples into a single data set. Each original SCF observation appears M times in the expanded data set, but the total weight of these M replications is rescaled to equal the sample weight of the original observation. Human wealth can then be computed for each observation in the expanded sample, and analysis proceeds using the adjusted weights.

Human Wealth Estimates: SCF Sample. Using the valuation function in equation (49) we recover an estimate of human wealth for each household in the SCF sample.[39] This allows us to quantify the relative size of human wealth in their wealth portfolio. In the left panel of Figure 18 we report the life-cycle evolution of average human wealth. Both shape and scale of average human wealth closely track those estimated from PSID data and plotted in Figure 2. The right panel of Figure 18 also plots the standard deviation of human wealth, which is roughly half the size of the average human wealth at any given age. For example, average human wealth peaks at just below $800,000 (per household), when the standard deviation stands at roughly $350,000. This means that two standard deviations below the average corresponds to a value close to zero, while adding two standard deviations doubles the average value. Interestingly, dispersion remains fairly high until age 50. Hence, the contribution of human wealth to overall inequality is largest between ages 35 and 55. The fact that dispersion remains elevated long after average human wealth has started its decline indicates that some workers are exiting full time employment relatively young: low employment and non-employment risks are explicitly accounted for by our estimation, which considers periods of null or low earnings as possible outcomes of each worker's history.

---

[39] As we discussed, this requires linking the full distribution of unobservable types to each observation in the SCF.

F   Various Calculations and Data Adjustments

Permanent-Income Annuitization.   We use permanent-income, i.e. the annuitized value of lifetime wealth, in a number of places in our analysis, for example to produce Figure 14. We denote permanent-income by π_{it}, and estimate it based on the age of the head of household as follows:

$$\hat{\pi}_{it} = (a_{it} + \hat{\theta}_j(\mathbf{z}_{it}))\frac{r}{1 - (1+r)^{95-j}}. \tag{50}$$

In this description we take $\hat{\theta}_j(\mathbf{z}_{it}))$ as the estimated human wealth of the household, which may be the sum of both spouses human wealth. The sum of household human wealth and asset wealth, a_{it}, is household life-time wealth. Permanent-income is defined as the annuity value of lifetime wealth, where the annuitization takes variation in the remaining lifetime horizon into account. We adopt a fairly conservative approach in assuming that all households plan to live to 95 years old. The reason this should be considered conservative is that shortening the horizon would only increase permanent-income, accentuating deviations from current consumption.

Scaling Consumption Expenditures.   In Figure 14, as well as in Table 4, we emphasize how consumption expenditures tend to be lower than permanent-income. One concern in these calculations is that expenditures reported in the PSID are only a subset of the total; thus, systematic mis-measurement could be partly responsible for permanent-income exceeding consumption. To address this concern we rescale real expenditures observed in any given wave of the PSID by a factor that makes the average expenditure equal to the real aggregate expenditures per capita reported in NIPA data. There is some variation by year, but on average this rescaling significantly increases consumption expenditures and makes them closer to permanent-income.

Imputing Pension and Social Security Entitlements.   Workers generally have some entitlement to a social security pension, and many are also entitled to defined-benefit pensions through their employer. Although the value of future defined-contribution pensions are accounted for in net worth, social security and defined-benefit pensions are not. We estimate these values as follows: (i) we calculate the average social security and pension income of individuals in the SCF according to their education level; (ii) we take the smaller of this average and the maximum social security benefit as the annual entitlement. We assume a 30 year collection period up to age 95 when calculate the value of these entitlements, discounting at the risk free rate. Thus, for a retired household the value of entitlements is based on the number of years of life that remain up to age 95; for working age households the value is based on the present value of the 30 years of payments that will commence at age 66.