Permanent Income Inequality

PRELIMINARY AND INCOMPLETE

Brant Abbott
Queen's University
`abbottbrant@gmail.com`

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

THIS DRAFT: JULY 15, 2017

Abstract. We estimate the extent of inequality in lifetime wealth and permanent income, and study its evolution over time. Our definitions of these variables are related to Friedman's original concepts, but do not impose a specific utility function. We extend existing results on nonparametric identification of marginal utility functions in order to establish nonparametric identification of the human wealth component. Our method imposes no restrictions on the dynamics of income processes, features state-dependent stochastic discounting, and allows for important dimensions of unobserved heterogeneity. Our results suggest that accounting for the value of human capital significantly changes the assessment of inequality and its evolution. We find that human wealth has a mitigating influence on overall inequality, but that inequality has been increasing much faster than asset wealth alone would indicate. Specifically, (i) in 2013 the top 10% shares of lifetime wealth and permanent income were only about 2/3 as large as the corresponding share of asset wealth; (ii) however, between 1989 and 2013 the top 10% shares of lifetime wealth and permanent income grew about 25% and 40% faster, respectively, than the corresponding share of asset wealth. Finally, we find that households at the top of the asset wealth distribution have not increased their share of human wealth over time. Instead, the excess increase in the concentration of lifetime wealth is mostly due to the growing importance of asset wealth as a share of lifetime wealth portfolios.

JEL Classification: D3, G1, I24, E21, J01
Keywords: Inequality, Earnings, Wealth, Human Capital

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 that are related to wellbeing, primarily income and wealth.[1] However, a more complete assessment of economic inequality may be attainedby accounting for the heterogeneity associated to 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 ex-ante unobservable value of expected future earnings.[2]

Depending on how they are discounted, potential future earnings may constitute the most important determinant of economic wellbeing. A young person with a steeply increasing expected earnings profile may be much better off than would be indicated by simply measuring their asset wealth or current income. The extent to which these future earnings matter depends on how much they are discounted. Appropriate discounting of future earnings effectively accounts for the ease with which consumption can be shifted across time periods, and for uncertainty in future earnings and consumption.[3] Being constrained by a credit limit or facing

[1]See for example the work of Saez and Kopczuk (2004), Piketty and Saez (2006), Saez and Zucman (2014), Bricker, Henriques, Krimmel, and Sabelhaus (2016), Kaymak and Poschke (2016) and Rios-Rull, Kuhn, et al. (2016). An extensive literature on the distribution of wages and earnings documents widening inequality in the working population (see for example Levy and Murnane (1992), Gottschalk, Moffitt, Katz, and Dickens (1994), Goldin and Katz (2007) and Autor, Katz, and Kearney (2008)). Studies of the wealth distribution focus on the financial/real wealth held by the wider population, including the unemployed and those who do not participate in the labor market (see Saez and Zucman, 2014; Bricker, Henriques, Krimmel, and Sabelhaus, 2016). More recently, the work of De Nardi, Fella, and Paz-Pardo (2016) carefully illustrates how rich income processes (as those described in Guvenen, Karahan, Ozkan, and Song, 2016) may be reflected in the equilibrium distribution of wealth.

[2]This is an extensive literature. The New Dynamic Public Finance part of the literature is surveyed by Golosov, Tsyvinski, Werning, Diamond, and Judd (2006) and Kocherlakota (2010). Examples from the Ramsey planning section of the literature include Conesa, Kitao, and Krueger (2009) and Davila, Hong, Krusell, and Ríos-Rull (2012).

[3]The way future income is discounted is important. Huggett and Kaplan (2016) convincingly argue that the true value of human capital is far below the value that would be implied by discounting future net earnings at the risk-free interest rate, an approach that is commonly advocated because of its simplicity. For examples, see Becker (1975), Jorgenson and Fraumeni (1989) and R. Haveman and Schwabish (2003). Mechanically discounting income flows to approximate human capital rules out state-dependent changes in the valuation of future earnings.

a great deal of risk would reduce a household's valuation of their potential future earnings.

In this paper we estimate pecuniary statistics that reflect the values of both human capital and asset wealth, and analyze trends in inequality of these statistics over the time period 1989-2013. At the heart of our analysis are nonparametric estimates of the value to individuals of their earnings potential, which we refer to as their human wealth. These estimates differ from the simple expected present value of future earnings in several ways. Most importantly, they feature state-dependent stochastic discounting, rather than risk-free discounting. Combining human wealth estimates with observed asset wealth data allows us to estimate lifetime wealth, which is simply the sum of human wealth and asset wealth. We also construct 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 simply the risk-free present value of expected future earnings. We find that in 2013 the top 10% share of permanent income was just 2/3 as large as the top 10% share of asset wealth, and the top 10% share of lifetime wealth was even slightly lower. However, permanent income inequality grew about 40% faster than asset wealth inequality over the 1989-2013 period, while lifetime wealth inequality grew about 25% faster. Hence, we conclude that human capital has had a mitigating influence on overall inequality, but that this influence has declined 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 estimated functions to SCF data, where the resulting estimates of human wealth can be added to observed net worth. This allows us to finally compute estimates of lifetime wealth and permanent income. We do not make assumptions about the processes that generate risk in the labor market, aggregate risk is present in the data and affects estimates of human wealth. To ensure a long enough sample to identify the aggregate risk component, we impute consumption in the PSID prior to 1999 using the method suggested by Attanasio and Pistaferri (2014). Thus, all PSID data from 1968 to 2013 are utilized when estimating human wealth.

Crucially, our approach allows for the estimates of human wealth to depend on state-dependent stochastic discount factors, and therefore on marginal utility functions. Rather than assuming specific functional forms we estimate stochastic discount factors nonparametrically. This dispenses with several restrictions and enables the data to guide the choice of utility function in a flexible way. Nonparametric identification of the marginal utility function is achieved by following the method suggested by Escanciano, Hoderlein, Lewbel, Linton, and Srisuma (2015). This approach involves writing the intertemporal Euler equation in such a manner that the estimated marginal utility function is the solution of a homogeneous Fredholm equation of the second kind. Given identification of the stochastic discount factor, human wealth then depends on an integral over possible realizations of its own future value multiplied by associated realizations of the stochastic discount factor. Compared to the marginal utility function, the estimated human wealth valuation equation is the solution of an *inhomogeneous* Fredholm equation of the second kind. For this reason we need to extend existing results to prove nonparametric identification of human wealth.

A separate issue for our analysis arises from the fact that only one realization of the future state of the world is observed for each person and time-period in the sample. As such, we do not observe the entire distribution of possible future outcomes, which an individual's human wealth depends on. To address this data limitation we introduce an identification assumption, which we refer to as conditional equivalence of expectations. This assumption simply 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. We implement this by allowing 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 types of heterogeneity, such as heterogeneous income profiles, could lead to substantial differences in distributions of ex-post outcomes even if individuals are equivalent in ex-ante observable variables. 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.

The upshot of the aforementioned econometric work is an analysis of inequality of human wealth, lifetime wealth and permanent income that can be immediately related to the existing analyses of inequality based on observable variables. As an illustration, we might ask 'what is the top 1% share of lifetime wealth?' and 'how has the Gini coefficient of permanent income changed over time?' despite the fact that they are unobservable. This is important because in theory these variables are more closely related to economic wellbeing than observable income or wealth alone are. However, by their nature, these theoretical variables are identified through a set of structural assumptions, and 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 believe we have made the assumptions underlying our estimates as plausible as possible, while still maintaining the crucial comparability between our analysis and existing studies of inequality based on observable wealth and income.

2 Theory

2.1 Lifetime Wealth

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 (individual-level) risk. An individual's observable characteristics, such as education, age and gender, are contained in the vector X_{it}. An individual's unobservable type, which may be informative about their expected earnings 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. A household's wealth portfolio is 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 $\mathbf{a}_{(ij)t}$.

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. In addition to leisure and formal labor, many agents spend an exogenous fraction n_{it} of their time on housework.

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:

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

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 who 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

$$
\sum_{\kappa \in k} a_{it+1}^{\kappa} + c_{it} \leq w_{it} h_{it} + \sum_{\kappa \in k} R_t^{\kappa} a_{it}^{\kappa} - T_t\left(\mathbf{a}_{it}, w_{it}, h_{it}\right), \tag{2}
$$

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 $\ell_{it} = 1 - h_{it} - n_{it}$ 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:

$$
\begin{aligned}
&V_i^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \qquad (3)\\
&u\left(c_{it}^*, \ell_{it}^*\right)\\
&+\beta\left\{(1-\tilde{\mu}_{it})E_{\{\Omega^{t+1},\mathbf{a}_{it+1}\}}\left[V_i^S(\mathbf{a}_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1})|\mathbf{a}_{(ij)t+1}^*\right]\right.\\
&\left.+\tilde{\mu}_{it}E_{\{\Omega^{t+1}\}}\left[V_i^M(\mathbf{a}_{(ij)t+1}^*, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right]\right\} + \heartsuit_{it(j)}.
\end{aligned}
$$

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, which we describe 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 will be split and individual i will receive 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.

Married households are assumed to be unitary. Therefore, the joint optimization problem of the spouses can be viewed as that of a planner who maximizes a weighted average of the spouses' utilities, according to some Pareto weights. Above we have already 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) nonoptimized allocations, which we denote by $\tilde{V}_i^M$. The problem of the household planner is:

$$V_{(ij)}^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \max_{\mathbf{b}_{(ij)t}} \Big\{ \lambda_{(ij)} \tilde{V}_i^M(\mathbf{a}_{(ij)t}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) \tag{4}$$
$$+ (1-\lambda_{(ij)}) \tilde{V}_j^M(\mathbf{a}_{(ij)t}, X_{jt}, X_{it}, \eta_j, \eta_i, \Omega^t) \Big\},$$

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

$$\sum_{\kappa \in k} a^\kappa_{(ij)t+1} + c_{(ij)t} \le w_{it}h_{it} + w_{jt}h_{jt} + \sum_{\kappa \in k} R^\kappa_t a^\kappa_{(ij)t} - T_t\left(\mathbf{a}_{(ij)t}, w_{it}, w_{jt}, h_{it}, h_{jt}\right), \tag{5}$$

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 $\ell_{it} = 1 - h_{it} - n_{it}$ and $\ell_{jt} = 1 - h_{jt} - n_{jt}$, and the household borrowing limit $\sum_{\kappa \in k} a^\kappa_{(ij)t+1} \ge \underline{a}_{(ij)t}$ also constrain the household planner's choices.

2.2 Valuation of Human Wealth

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 a hypothetical asset that pays dividends per share equal to individual i's yearly labor income, but also requires the individual to commit to their state-contingent labor supply plan. [4] Because of this commitment we replace $w_{it}h_{it}$ from the

[4] 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

problems described above with y_{it}, with the understanding that y_{it} is state-contingent earnings under the optimal labor supply plans of problems (1) and (4) above.[5] Letting θ_{it} be the shadow price of this hypothetical asset, we can show that the valuation has the familiar Lucas-style asset pricing form:

$$\theta_{it} = \mathbb{E}_{it}\left[\beta \frac{u_c(c_{it+1}, \ell_{it+1})}{u_c(c_{it}, \ell_{it})} \left(y_{it+1} + \theta_{it+1}\right)\right]. \tag{6}$$

Thus, human wealth is the expected value of stochastically discounted future earnings (the next period dividend of the hypothetical asset) plus its future value. We next turn to deriving this relationship in the context of our model.

A complication that arises in this setting is that the current valuation of an asset depends on its effect on a person's marital bargaining power, if and when that person gets married in the future. If buying an asset does not increase the individual's utility once married, that asset would be worth less to them than otherwise. We do not attempt to explicitly estimate the effect on marital bargaining power of owning more shares of a hypothetical asset.[6] Rather, we make the simplifying assumption that bargaining between newly married couples can be represented by the symmetric Nash Bargaining solution. In other words, the ex-post Pareto weights of spouses do adjust in response to pre-marital investments, and they do so through the effect of pre-marital investments on the outside options of spouses and on the marital surplus. As we explain carefully below, the only effect of this assumption on human capital valuations is through a single person's continuation value in marriage.

Another complication that arises in this environment relates to how the hypothetical asset is allocated upon divorce. We assume that, in such circumstances, sole ownership of the asset based on individual i's labor income would go to person i, and that other assets, possibly

written against future labor services and ownership of human capital is not transferable (that is, human capital is a non-traded asset).

[5] As noted by Huggett and Kaplan (2016), this approach to valuing non-traded assets was first introduced by Lucas Jr (1978). Huggett and Kaplan (2016) also adopt this approach.

[6] In fact, this is a very interesting question in its own right but would require a much more sophisticated approach to modeling household interactions.

including a claim on alimony, would be allocated to the ex-spouse as compensation. The reason we assume that i takes ownership of the hypothetical asset is that we are valuing i's human capital, which they would own upon divorce as well.[7] This assumption, along with the one described in the previous paragraph, allows us to derive tractable formulas for valuing one's own human capital, and we exploit this tractability in the empirical analysis.

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:

$$\sum_{\kappa \in k} a^{\kappa}_{(ij)t+1} + c_{(ij)t} + \theta_{it} e_{it+1} \leq \theta_{it} e_{it} + (1 + e_{it}) y_{it} + y_{jt} \qquad (7)$$

$$+ \sum_{\kappa \in k} R^{\kappa}_{t} a^{\kappa}_{(ij)t} - T_t \left(\mathbf{a}_{(ij)t}, y_{it}, y_{jt} \right).$$

Furthermore, we include e_{it} as an additional state variable in the household planner's problem in equation (4), as well as in the definition of an individual's utility from marriage in (3). Given

[7] One can, of course, be ordered to pay alimony out of their returns to human capital in the real world. However, alimony is usually a fixed amount amount of money, so changes in earnings affect the earners' net-income, not their spouses. Thus alimony is better represented as an extra allocation of financial assets to the ex-spouse than an allocation of human capital, which is how we model it.

these adjustments we can rewrite the household planner's problem in a recursive manner as:

$$
\begin{aligned}
&V^M_{(ij)}(\mathbf{a}_{(ij)t}, e_{it}, X_{it}, X_{jt}, \eta_i, \eta_j, \Omega^t) = \qquad (8)\\
&\quad \max_{\mathbf{b}_{(ij)t}} \Big\{ \lambda_{(ij)} u\left(c_{it}, \ell_{it}\right) + (1-\lambda_{(ij)}) u\left(c_{jt}, \ell_{jt}\right) \\
&\quad + \lambda_{(ij)} \beta (1-\tilde{\mu}_{(ij)t}) E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}} \left[V^S_i(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) | \mathbf{a}_{(ij)t+1} \right] \\
&\quad + (1-\lambda_{(ij)}) \beta (1-\tilde{\mu}_{(ij)t}) E_{\{\Omega^{t+1}, \mathbf{a}_{jt+1}\}} \left[V^S_j(\mathbf{a}_{jt+1}, X_{jt+1}, \eta_j, \Omega^{t+1}) | \mathbf{a}_{(ij)t+1} \right] \\
&\quad + \beta \tilde{\mu}_{(ij)t} E_{\{\Omega^{t+1}\}} \left[V^M_{(ij)}(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) \right] \Big\},
\end{aligned}
$$

where the decision vector $\mathbf{b}_{(ij)t}$ now includes e_{it+1}. After using the budget constraint in (7) to substitute c_{it} out of the problem in (8), we can easily derive the following first-order condition for the optimal choice of e_{it+1}:

$$
\begin{aligned}
&u_c(c_{it}, \ell_{it}) \vartheta \theta_{it} = \qquad (9)\\
&\quad \beta (1-\tilde{\mu}_{(ij)t}) \frac{\partial}{\partial e_{it+1}} E_{\{\Omega^{t+1}, \mathbf{a}_{it+1}\}} \left[V^S_i(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) | \mathbf{a}_{(ij)t+1} \right] \\
&\quad + \frac{1}{\lambda_{(ij)}} \beta \tilde{\mu}_{(ij)t} \frac{\partial}{\partial e_{it+1}} E_{\{\Omega^{t+1}\}} \left[V^M_{(ij)}(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) \right].
\end{aligned}
$$

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 (8) with respect to e_{it}, noting that the c_{it} has been replaced by the budget constraint. The result is,

$$
\begin{aligned}
&\frac{\partial}{\partial e_{it+1}} V^M_{(ij)}(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1}) = \qquad (10)\\
&\quad \lambda_{(ij)} u_c(c^M_{it+1}, \ell^M_{it+1}) \vartheta \left(\theta^M_{it+1} + y^M_{it+1} \right),
\end{aligned}
$$

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

$$
\begin{aligned}
&V_i^S(\mathbf{a}_{it}, e_{it}, X_{it}, \eta_i, \Omega^t) = \qquad (11)\\
&\max_{c_{it}, \ell_{it}, h_{it}, \mathbf{a}_{it+1}} \Big\{ u\left(c_{it}, \ell_{it}\right) + \beta\left(1-\mu_{it}\right) E_{\{\Omega^{t+1}\}}\left[V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1})\right] \\
&+ \beta\mu_{it} E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}}\left[V_i^M(\mathbf{a}_{(ij)t+1}, e_{it+1}, X_{it+1}, X_{jt+1}, \eta_i, \eta_j, \Omega^{t+1})\right] \Big\}.
\end{aligned}
$$

The maximization in (11) is subject to the usual time allocation and borrowing constraints, as well the extended budget constraint,

$$
\begin{aligned}
\sum_{\kappa \in k} a_{it+1}^{\kappa} + c_{it} + \theta_{it} e_{it+1} \leq \theta_{it} e_{it} + (1+e_{it}) y_{it} \qquad (12)\\
+ \sum_{\kappa \in k} R_t^{\kappa} a_{it}^{\kappa} - T_t\left(\mathbf{a}_{it}, w_{it}, h_{it}\right).
\end{aligned}
$$

The derivative of the value function in (11) can thus be derived by replacing c_{it} with the extended budget constraint, resulting in:

$$
\frac{\partial}{\partial e_{it+1}} V_i^S(\mathbf{a}_{it+1}, e_{it+1}, X_{it+1}, \eta_i, \Omega^{t+1}) = u_c(c_{it+1}^S, \ell_{it+1}^S)\left(\theta_{it+1}^S + y_{it+1}^S\right). \qquad (13)
$$

Finally, using equations (10) and (13), one can re-arrange the first order condition for optimal e_{it+1} chosen by a married household (equation 9) into an expression describing the valuation of i's human capital θ_{it}^M (the purchase price per share of e_{it+1}):

$$\begin{aligned}\theta_{it}^M =& \beta(1-\tilde{\mu}_{(ij)t})\frac{1}{\vartheta}E_{\{\Omega^{t+1},\mathbf{a}_{it+1}\}}\left[\frac{u_c(c_{it+1}^S,\ell_{it+1}^S)}{u_c(c_{it},\ell_{it})}\left(y_{it+1}^S+\theta_{it+1}^S\right)\right] \\ &+\beta\tilde{\mu}_{(ij)t}E_{\{\Omega^{t+1}\}}\left[\frac{u_c(c_{it+1}^M,\ell_{it+1}^M)}{u_c(c_{it},\ell_{it})}\left(y_{it+1}^M+\theta_{it+1}^M\right)\right].\end{aligned} \tag{14}$$

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 seminal work of Lucas (1978). 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 (11):

$$\begin{aligned}&u_c(c_{it},\ell_{it})\theta_{it} = \\ &\quad\beta(1-\mu_{it})\frac{\partial}{\partial e_{it+1}}E_{\{\Omega^{t+1},\mathbf{a}_{it+1}\}}\left[V_i^S(\mathbf{a}_{it+1},e_{it+1},X_{it+1},\eta_i,\Omega^{t+1})|\mathbf{a}_{(ij)t+1}\right] \\ &\quad+\beta\mu_{it}\frac{\partial}{\partial e_{it+1}}E_{\{\Omega^{t+1},X_{jt+1},\eta_j,\mathbf{a}_{jt+1}\}}\left[V_i^M(\mathbf{a}_{(ij)t+1},e_{it+1},X_{it+1},X_{jt+1},\eta_i,\eta_j,\Omega^{t+1})\right].\end{aligned} \tag{15}$$

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 (13) 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

$$\max_{\{V_i^M, V_j^M\}} \left(V_i^M - V_i^S\right)\left(V_j^M - V_j^S\right), \tag{16}$$

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

$$\left(V_i^M - V_i^S\right) = \frac{G_2}{G_1}\left(V_j^M - V_j^S\right). \tag{17}$$

To translate this condition into something empirically useful, note that an equivalent formulation of the household planning problem in equation (8) is:

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

Combining the first-order conditions from this problem with those from the underlying Nash

Bargaining problem results in:

$$\left(V_i^M - V_i^S\right) = \frac{1-\lambda_{(ij)}}{\lambda_{(ij)}}\left(V_j^M - V_j^S\right). \tag{18}$$

The equivalence of equations (17) and (18) is due to the fact that $\lambda_{(ij)}$ is the Pareto weight that implicitly solves the Nash Bargaining problem in equation (16).

Next, we examine equation (17) 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

$$\frac{\partial V_i^M}{\partial e_{it}} - \frac{\partial V_i^S}{\partial e_{it}} = \left(\frac{G_2}{G_1}\right)\frac{\partial V_j^M}{\partial e_{it}} + \frac{1}{G_1}\left(\frac{\partial G2}{\partial e_{it}}\left(V_j^M - V_j^S\right) - \frac{\partial G1}{\partial e_{it}}\left(V_i^M - V_i^S\right)\right). \tag{19}$$

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

$$\left(\frac{\partial G2}{\partial e_{it}} / \frac{\partial G1}{\partial e_{it}}\right) = \frac{u_c(c_it, \ell_{it})}{u_c(c_jt, \ell_{jt})} = \frac{\lambda_{(ij)}}{1-\lambda_{(ij)}}. \tag{20}$$

Therefore, the last term of equation (19) 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

$$\frac{\partial V_i^M}{\partial e_{it}} - \frac{\partial V_i^S}{\partial e_{it}} = \frac{1-\lambda_{(ij)}}{\lambda_{(ij)}}\frac{\partial V_j^M}{\partial e_{it}}. \tag{21}$$

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 (21), first note that the definition of the household planner's optimization objective in (4) implies that the envelope condition in (10) can be re-written as:

$$\lambda_{(ij)} \frac{\partial V^M_{it+1}}{\partial e_{it+1}} + (1-\lambda_{(ij)}) \frac{\partial V^M_{jt+1}}{\partial e_{it+1}} = \lambda_{(ij)} u_c(c^M_{it+1}, \ell^M_{it+1}) \vartheta \left(\theta^M_{it+1} + y^M_{it+1}\right). \tag{22}$$

Combining this with the Nash Bargaining implication in (21), we obtain an extremely useful result characterizing the effect of pre-marital investments on the utility within marriage:

$$\frac{\partial V^M_{it+1}(\cdot)}{\partial e_{it+1}} = \frac{1}{2} u_c(c^M_{it+1}, \ell^M_{it+1}) \frac{1}{\vartheta} \left(\theta^M_{it+1} + y^M_{it+1}\right) + \frac{1}{2} \frac{\partial V^S_{it+1}(\cdot)}{\partial e_{it+1}}. \tag{23}$$

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 $\partial V^S_{it+1}/\partial e_{it+1}$. An upper bound is the marginal change in their utility if the entire return on the asset, including surplus due to economies of scale, is allocated to i. With symmetric bargaining exactly one half of the component pertaining to the return that exceeds the effect on i's outside option is paid to i.

Equation (23) is useful because we now have an expression to substitute into equation (9), 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 (13), allows us to derive the following valuation formula for the human capital of a currently unmarried person i:

$$\begin{aligned} \theta^S_{it} = & \beta(1-\frac{\mu_{it}}{2}) E_{\{\Omega^{t+1}\}} \left[\frac{u_c(c^S_{it+1}, \ell^S_{it+1})}{u'(c_{it}, \ell_{it})} \left(y^S_{it+1} + \theta^S_{it+1}\right) \right] \\ & + \beta \frac{\mu_{it}}{2} E_{\{\Omega^{t+1}, X_{jt+1}, \eta_j, \mathbf{a}_{jt+1}\}} \left[\frac{u_c(c^M_{it+1}, \ell^M_{it+1})}{u'(c_{it}, \ell_{it})} \frac{1}{\vartheta} \left(y^M_{it+1} + \theta^M_{it+1}\right) \right]. \end{aligned} \tag{24}$$

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}). Having done this we can summarize the value of human capital for any individual as

$$\theta_{it} = \mathbb{E}_{it} \left[\beta \frac{u_c(c_{it+1}, \ell_{it+1})}{u_c(c_{it}, \ell_{it})} (y_{it+1} + \theta_{it+1}) \right], \tag{25}$$

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 (25) 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.

The role of borrowing constraints. An implicit assumption in the derivation of our expression for θ_{it} is that agents are unrestricted in their choice of e_{it}. They can trade short or long, and face no constraint in how much debt they can accumulate by trading short. This does not mean that these agents do not face borrowing constraints in their portfolio of real assets $\mathbf{a}_{it}$ (or $\mathbf{a}_{(ij)t}$ if married). An individual may be severely borrowing constrained in assets, but our hypothetical exercise still allows them to contemplate short selling e_{it}. There is an important and intuitive reason for the model to feature this. The reason is that our exercise aims to recover a price for the hypothetical asset at which agents would *choose* not to trade away from $e_{it} = 0$. For individuals that are borrowing constrained this price will clearly tend to be lower than for similar unconstrained individuals, entirely because they would like to use the newly introduced asset to move away from their borrowing constraint. But this lower price is exactly what we want to recover, as 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 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, Linton, and Srisuma (2015) to recover nonparametric estimates of marginal utility functions, and of the deterministic component of the time discount factor (β). Then we use these estimates in a second step to obtain nonparametric estimates of human wealth. While only the identification results for the second step are novel, we overview both steps in detail. The advantage of carefully describing the first step is that it greatly aids in understanding the second step.

3.1 Identification

3.1.1 Nonparametric Marginal Utility Function Identification

In what follows it is helpful to use compact notation such as $q = (c_{it}, \ell_{it})$ and $q' = (c_{it+1}, \ell_{it+1})$ to represent the choices of an arbitrary individual. The consumption decision of an individual who is not at a corner solution is described by an intertemporal optimality (Euler) equation, which can be written as follows:

$$u_c(q) = \beta \mathbb{E}\left[u_c(q')R'_o|q\right]. \tag{26}$$

We write this condition using the return on an arbitrary asset R_o', but it can be extended to any asset this set of agents trades in. Conditioning on q (current consumption and leisure choices) is equivalent to conditioning on the entire information set because all relevant information is acted upon in these decisions.

We begin by rewriting equation (26) 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 the conditional expectation of future rates of return and the Markov (transition) kernel estimator describing transitions from q to q'. The notation we use for this weighting factor is $\psi(q, q') = \mathbb{E}\left[R_o'|q, q'\right] \times f_{Q'|Q}(q'|q)$, where $f_{Q'|Q}$ is the conditional density of q'. Then the Euler equation (26) can be expressed as

$$u_c(q) - \beta \int u_c(q')\psi(q, q')dq' = 0. \tag{27}$$

As pointed out in Escanciano et al., this is a homogeneous Fredholm integral equation of the second kind. The solution for $u_c(q)$ given β is well known, but 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 consumption leisure pairs (although M could be a very large number, it cannot grow with the sample size). Formally, the support is $q \in \left\{q^1, q^2, \ldots, q^M\right\}$. Under this assumption we can rewrite the Euler equation (27) as

$$u_c(q^k) - \beta \sum_{m=1}^{M} u_c(q^m)\psi(q^k, q^m) = 0, \tag{28}$$

for any current choice vector q_i^k. Rather than solving a complicated integral equation, identification in this finite example requires solving a linear system. Writing equation (28) in matrix notation, this entails solving

$$(I - \beta\Psi)U_c = 0, \tag{29}$$

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

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 be the eigenvector of Ψ associated with the eigenvalue β^{-1}. Thus, β is identified as the inverse of any eigenvalue of Ψ such that $\beta \in (0, 1)$, and U_c is identified as the associated eigenvector. In general, Ψ may have multiple eigenvalues larger than unity, thus only set identification is achieved in the finite support case. 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 $\mathbb{E}\left[R_o'|q, q'\right]$. This means that if returns were a constant risk-free rate where $R_o' = R^f$ always, the largest eigenvalue of Ψ would be R^f, and we would have $\beta = 1/R^f$ as a potential solution.

General Case. Proof of identification in the general case where q_i has a continuous support requires functional analysis, but is similar to the finite support case above (we refer readers to Escanciano et al. for greater detail). One first has to define a linear operator A such that, when it is applied to the unknown function $u_c(q)$, the result is

$$(Au_c)(q) = \beta \int u_c(q')\psi(q, q')dq'. \tag{30}$$

By definition this implies that $u_c = \beta A u_c$. 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 eigenvalue (spectral radius) of the operator A.[8] Therefore, if these assumptions are maintained a unique value of β, and a unique function u_c, solve equation

[8] 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.

(27) and point identification is achieved.

3.1.2 Nonparametric Human Wealth Identification

We now turn to the second step and the question of nonparametric identification of θ_{it} in equation (6). Relying on the derivations above, we take as given that β and the marginal utility function are identified.

We now introduce the vector $\mathbf{z}$ containing variables that summarize an individual's information set. Unlike the estimation of the marginal utility function, some individuals we now consider may be credit constrained, and therefore current consumption and leisure may not fully summarize their information sets. Hence, we make the following assumption:

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

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

Conditional equivalence of expectations will hold if $\mathbf{z}_{it}$ spans the current information set of any individual i. Assuming this is the case, we can rewrite the human wealth valuation equation (6) with θ_{it} replaced by a function $\theta(j,\mathbf{z}_{it})$, where j is the age of the individual:

$$\theta(j,\mathbf{z}_{it})=\mathbb{E}\left[\beta\frac{u_c(c_{it+1},\ell_{it+1})}{u_c(c_{it},\ell_{it})}\left(y_{it+1}+\theta(j+1,\mathbf{z}_{it+1})\right)|\mathbf{z}_{it}\right]. \tag{31}$$

This is now a functional equation, somewhat similar to the Euler equation described above.

Dropping individualized subscripts (i's) and replacing time subscripts with 'prime' notation, we can rewrite equation (31) as an integral equation after defining two substitutions. First, define $\delta(j,\mathbf{z},\mathbf{z}')=\mathbb{E}[\beta(u_c'/u_c)|j,\mathbf{z},\mathbf{z}']\times f_{Z'|Z}(\mathbf{z}'|j,\mathbf{z})$, where $f_{Z'|Z}$ is the conditional density of $\mathbf{z}'$. Each $\delta(j,\mathbf{z},\mathbf{z}')$ can be described as a an appropriately discounted density function for $\mathbf{z}'$ at age j, given conditioning set $\mathbf{z}$. 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 is written as

$$\theta(j, \mathbf{z}) = g(j, \mathbf{z}) + \int \theta(j+1, \mathbf{z}')\delta(j, \mathbf{z}, \mathbf{z}')d\mathbf{z}. \tag{32}$$

Comparing the above functional equation to the integral form of the intertemporal Euler equation, the main difference is that eq. (32) 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 subsumes the current 'dividend' associated with human capital and introduces an age-dependent intercept in the functional equation (32).

One can provide conditions for a unique solution of equation (32) by exploiting the deterministic nature of the age transitions. We begin by defining the vector-valued functions

$$\begin{aligned} \Theta(\mathbf{z}) &= (\theta(1, \mathbf{z}), \theta(2, \mathbf{z}), \ldots, \theta(J-1, \mathbf{z}), 0)' \\ G(\mathbf{z}) &= (g(1, \mathbf{z}), g(2, \mathbf{z}), \ldots, g(J-1, \mathbf{z}), 0)', \end{aligned}$$

where J is an arbitrarily old age. Furthermore, combine the age specific transition kernels into a $J \times J$ matrix

$$\Delta(\mathbf{z}, \mathbf{z}') = \begin{pmatrix} 0 & \delta(1, \mathbf{z}, \mathbf{z}') & 0 & \ldots & 0 \\ 0 & 0 & \delta(2, \mathbf{z}, \mathbf{z}') & & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ \vdots & \vdots & \vdots & & \delta(J-1, \mathbf{z}, \mathbf{z}') \\ 0 & 0 & 0 & \ldots & 0 \end{pmatrix}. \tag{33}$$

This is a hollow upper triangular matrix so that it conforms with $\Theta(\mathbf{z}')$ in the following representation of the integral equation (32):

$$\Theta(\mathbf{z}) = G(\mathbf{z}) + \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}')d\mathbf{z}'. \tag{34}$$

Mirroring the approach of Escanciano et al., we now define a linear operator B composed of a finite set of age-specific linear operators B_j. Each age-specific operator is defined such that

$$(B_j\theta)(j+1, \mathbf{z}) = \int \delta(j, \mathbf{z}, \mathbf{z}')\theta(j+1, \mathbf{z}')d\mathbf{z}'. \tag{35}$$

The operator B is arranged as follows:

$$B = \begin{pmatrix} 0 & B_1 & 0 & \dots & 0 \\ 0 & 0 & B_2 & & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ \vdots & \vdots & \vdots & & B_{J-1} \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}, \tag{36}$$

which ensures that B is a linear operator such that:

$$(B\Theta)(\mathbf{z}) = \int \Delta(\mathbf{z}, \mathbf{z}')\Theta(\mathbf{z}')d\mathbf{z}'. \tag{37}$$

Applying this definition in equation (34) the function Θ is uniquely determined to be $\Theta = (I - B)^{-1}G$, provided the operator $I - B$ has a bounded inverse. In turn, this is true provided that the operator B is compact. This result clearly relies on the assumption that, for a large enough J, the value of human wealth is zero, which leads to B being upper triangular and hollow. The simple intuition for this identification result becomes apparent if one thinks about solving the pricing equation (34) recursively, starting from the last age in which human wealth has a non-zero value. In the last period, one can use the fact that the human wealth value next period is zero to solve for the value of human wealth in the current period. The solution can then be stored and used to solve for the value of human wealth one period prior, allowing for a backward recursion up to the initial age. The upper triangular and hollow shape of matrix (36) effectively uses the implicit recursiveness of the human wealth pricing equation to jointly identify the value of human wealth for different age groups.

Finite example. When the support of $\mathbf{z}$ is restricted to be finite, formally $\mathbf{z} \in \{\mathbf{z}^1, \mathbf{z}^2, \ldots, \mathbf{z}^M\}$, proof of a unique solution for Θ simplifies to proving a unique solution to a linear system. In such a case each $\delta(j, \mathbf{z}, \mathbf{z}')$ reduces to a Markov transition matrix. Each transition matrix is an element of the block matrix Δ, which is hollow and upper triangular. Applying this to the 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 now consider a sample with N observations of q_i, q_i', $\mathbf{z}_i$, $\mathbf{z}_i'$, R_i' and j (age). Observations in this sample may be weakly dependent. Furthermore, a subset $N_0 \leq N$ of these N observations corresponds to interior (not borrowing constrained) allocations. In what follows we proceed sequentially: first, we describe our implementation of Escanciano et al.'s method of estimating the marginal utility function. Second, we overview in some detail our own estimator of the human wealth values.

3.2.1 Marginal Utility Function Estimation

The first step in the estimation of the marginal utility function is to replace the linear operator A in equation (30) with the estimator

$$(\hat{A}u_c)(q) = \frac{1}{N_0}\sum_{i=1}^{N_0} u_c(q_i')R_i'\phi_i(q), \tag{38}$$

The density function $\phi_i(q)$ is estimated using the Gaussian kernel function K_h with bandwidth h as follows

$$\phi_i(q) = \frac{K_{hi}(q)}{\hat{f}(q)}, \tag{39}$$

where

$$\hat{f}(q) = \frac{1}{N_0}\sum_{i=1}^{N_0} K_{hi}(q), \tag{40}$$

and

$$K_{hi}(q) = K_h\left(c - c_i\right)K_h\left(\ell - \ell_i\right). \tag{41}$$

Because the estimator $\hat{A}$ has a finite dimensional range, unlike A itself, Escanciano et al. show that $\hat{A}$ has at most N_0 eigenvalues and eigenfunctions, which can be computed by solving a linear system. As such, any eigenfunction $\hat{u}_c(q)$ of $\hat{A}$ must be a linear combination of the functions $\phi_i(q)$, i.e. $\hat{u}_c(q) = N_0^{-1}\sum_{i=1}^{N_0} b_i\phi_i(q)$ for some set of coefficients b_i. We use this result and rewrite the Euler equation with its arguments replaced by estimators, obtaining

$$\hat{u}_c(q) - \hat{\beta}(\hat{A}\hat{u}_c)(q) = \frac{1}{N_0}\sum_{i=1}^{N_0} b_i\phi_i(q) - \hat{\beta}\frac{1}{N_0}\sum_{i=1}^{N_0}\frac{1}{N_0}\sum_{s=1}^{N_0} b_s\phi_s(q_i')R_i'\phi_i(q) = 0. \tag{42}$$

Some straightforward algebra shows that the Euler equation above has a solution only if

$$b_i - \hat{\beta}\frac{1}{N_0}\sum_{s=1}^{N_0} b_s\phi_s(q_i')R_i' = 0, \tag{43}$$

for every $i = 1, \ldots, n$. This can be rewritten in matrix form with Φ being a $N_0 \times N_0$ matrix with elements $\Phi_{is} = N_0^{-1}\phi_s(q_i')R_i'$. Each element Φ_{is} can be interpreted as an appropriately weighted future return value. If $b = (b_1, \ldots, b_{N_0})'$ is a coefficient vector, then the above restrictions can be expressed as

$$(I - \hat{\beta}\Phi)b = 0. \tag{44}$$

Letting λ^* be the largest eigenvalue of Φ in absolute value, and b^* be the associate eigenvector, then the estimators of β and $u_c(q)$ are respectively

$$\hat{\beta} = \frac{1}{|\lambda^*|} \tag{45}$$

$$\hat{u}_c(q) = \frac{1}{||b^*||}\frac{1}{N_0}\sum_{i=1}^{N_0} b_i^*\phi_i(q), \tag{46}$$

where $\hat{u}_c(q)$ is scaled to have a unit norm.

3.2.2 Human Wealth Estimation

Despite the fact that $\hat{u}_c(q)$ was estimated using a subsample $N_0 \leq N$ of observations,[9] point estimates of marginal utility can be constructed for any of the N observations. Furthermore, we can transform $\hat{u}_c(q)$ into $\hat{u}_c^z(\mathbf{z})$ by employing the Nadaraya-Watson kernel as follows:

$$\hat{u}_c^z(\mathbf{z}) = \frac{\sum_{i=1}^N K_{hi}(\mathbf{z})\hat{u}_c(q_i)}{\sum_{i=1}^N K_{hi}(\mathbf{z})}. \tag{47}$$

This step provides an estimate of marginal utility as a function of the conditioning variables $\mathbf{z}$, and is needed in order to estimate the entire $\Theta(\mathbf{z})$ function, although it is not necessary for estimation at each of the observed $\mathbf{z}_i$ points, i.e. to estimate $\Theta(\mathbf{z}_i)$.

[9]The subset of observations at interior solutions.

The next step is to replace each $(B_j\theta)(j+1,\mathbf{z})$ in equation (35) with its estimator

$$(\hat{B}_j\Theta)(j+1,\mathbf{z}) = \frac{1}{N_j}\sum_{i=1}^{N_j}\theta(j+1,\mathbf{z}'_i)\gamma_{ji}(\mathbf{z}), \tag{48}$$

where N_j is the sample size for those of age j, and $\gamma_{ji}(\mathbf{z})$ is constructed as

$$\gamma_{ji}(\mathbf{z}) = \xi_{ji}(\mathbf{z})\hat{\beta}\frac{\hat{u}_c(q'_i)}{\hat{u}^z_c(\mathbf{z})}, \tag{49}$$

where

$$\xi_{ji}(\mathbf{z}) = \frac{K_{hji}(\mathbf{z})}{\frac{1}{N_j}\sum_{s=1}^{N_j}K_{hjs}(\mathbf{z})}, \tag{50}$$

and

$$K_{hji}(\mathbf{z}) = \prod_{d=1}^{D}K_h\left(z_d - z_{di}\right). \tag{51}$$

Then, we proceed by constructing a set of age-specific matrices $\mathit{\Gamma}_j$. Each row of this matrix is the observed vector of $\gamma_{ji}(\mathbf{z})$'s evaluated at a particular value of $\mathbf{z}$. Each individual row of $\mathit{\Gamma}_j$ contains the vector $\gamma_{ji}(\mathbf{z})$ evaluated at an observed $\mathbf{z}_i$. Hence the row-i column-s element is $[\mathit{\Gamma}_j]_{is} = \gamma_{js}(\mathbf{z}_i)$. Arranging the $\mathit{\Gamma}_j$'s into a block matrix, we obtain

$$\mathit{\Gamma} = \begin{pmatrix} 0 & \mathit{\Gamma}_1 & 0 & \dots & 0 \\ 0 & 0 & \mathit{\Gamma}_2 & \dots & 0 \\ 0 & 0 & 0 & \ddots & 0 \\ \vdots & \vdots & \vdots & & \mathit{\Gamma}_{J-1} \\ 0 & 0 & 0 & \dots & 0 \end{pmatrix}. \tag{52}$$

Therefore, we can rewrite equation (34) in terms of its estimator at the points where $\mathbf{z}$ is observed:

$$\tilde{\Theta} = \tilde{G} + \Gamma\tilde{\Theta}, \tag{53}$$

where $\tilde{G}$ is the Nadaraya-Watson kernel estimator of the conditional expectation $G(\mathbf{z})$ evaluated at the observed $\mathbf{z}_i$ points. Because $(I - \Gamma)$ is invertible, the estimator for $\tilde{\Theta}$ is $\hat{\tilde{\Theta}} = (I - \Gamma)^{-1}\tilde{G}$.

Lastly, we would like to estimate the entire function $\Theta(\mathbf{z})$, rather than a collection of points along it. This is where the restriction that $\theta(J, z) = 0$ for large enough age J is used in the implementation. Specifically, we now use the estimated $\tilde{\Theta}(\mathbf{z}_i)$ to form an estimator $\hat{\Theta}(\mathbf{z})$ age-wise as follows:

$$\hat{\theta}(\mathbf{z}) = \hat{G}(\mathbf{z}) + \frac{1}{N_j}\sum_{i=1}^{N_j}\tilde{\theta}(j+1, z_i')\gamma_{ji}(\mathbf{z}). \tag{54}$$

Note that this step is where the kernel function $\hat{u}_c^z(\mathbf{z})$ has finally been used (in previous steps we evaluated $\hat{u}_c$ at known points $\mathbf{z}_i$).

With an estimate $\hat{\Theta}(\mathbf{z})$ available, several other statistics of interest can easily be computed. First, human wealth can be estimated at the household level by taking $\hat{\Theta}^h(\mathbf{Z}_{ij}) = \hat{\Theta}(\mathbf{z}_i) + \hat{\Theta}(\mathbf{z}_j)$ for any couple i and j, where $\mathbf{Z}_{ij} = (\mathbf{z}_i, \mathbf{z}_j)$. The lifetime wealth of such a household can be estimated by adding observed net worth to $\hat{\Theta}^h(\mathbf{Z}_{ij})$, and their permanent income can be estimated by computing the annuity value of this, according to an age-dependent annuitization factor.

3.3 Data

The basic data requirement for the estimation of marginal utilities, discount factor and human wealth values is a sample $\{q_i, q_i', \mathbf{z}_i, \mathbf{z}_i', R_i', j_i\}_{i=1}^N$. Each vector q denotes a pair of consumption and leisure choices; vector $\mathbf{z}$ includes observables that approximate the information set of the decision makers; R is an historical real return from deferred consumption; and j denotes age. The sample must include observations recorded over a sufficiently enough time interval so to identify the aggregate risk component of the transition kernels.

To obtain an empirical counterpart of the estimator in equation (), and to recover the marginal utilities, we only need panel information about consumption and leisure choices, as well as data on historical asset returns. It turns out that the Panel Study of Income Dynamics contains much of what we need, hence we use panel data from the PSID spanning the years 1967-2013. We assume the repeated observations on the same individuals in this data set satisfy the required weak dependence assumption.

Construction of q_i and q_i' involves collecting earnings and consumption data. Labour earnings is always observed. However, a fairly complete set of consumption expenditures is observed at the household level only after 1997. Before that date only few categories of consumption were recorded regularly. One complication is that consumption expenditure data are limited to relatively few categories until 1997. For this reason we build on the approach of Attanasio and Pistaferri (2014) to approximate household consumption expenditure in periods when information is incomplete. This method relies on the ever larger availability of consumption expenditures in the PSID since post-1997. The procedure effectively estimates a demand system and uses it to impute consumption to PSID families observed in the years before 1997, when only food expenditures was regularly measured. There are four clear advantages to this approach: (i) the procedure relies on information from a unique data set, making variable linkages more obvious; (ii) one can test how closely trends in consumption inequality are replicated by the imputation procedure using the period during which complete expenditure data are available (in-sample verification); (iii) since the PSID stretches all the way back to the late 1960s, this procedure delivers the longest longitudinal consumption data base currently available for the US; (iv) last but not least, expenditure categories in the PSID appear to match NIPA counterparts reasonably well.

As a proxy for real asset returns we set R_i' to be the one-year treasury constant maturity rate minus realized annual CPI inflation, when using annual data prior to 1997. As the survey becomes biannual after 1997, we switch to the two-year treasury constant maturity rate minus realized CPI inflation.[10]

After estimating the marginal utilities, one can proceed to obtain the empirical counterpart of the human wealth estimator in equation (54). To this purpose we need to use a set of conditioning variables that approximate the information set available to the agents. The vector $\mathbf{z}_i$ contains information about individual characteristics, such as gender, education, industry, occupation, marital status, and number of dependent children.

3.4 Estimating Unobserved Types

The vector $\mathbf{z}_i$ includes also the unobserved type η_i. We recover unobserved type variation using a procedure in the spirit of recent work by Bonhomme, Lamadon, and Manresa (2017). First, we assume that η_i can vary along two dimensions of heterogeneity. The first dimension captures any heterogeneity in expected (life-cycle) earnings profiles, and we use information about heterogeneity in the growth rates of earnings to identify it. The second dimension subsumes unobserved differences in wealth, which we measure by gauging the dispersion of consumption growth rates over the life cycle of different sample members. In both cases, the idea is that measures of earnings growth and consumption dispersion convey information about, respectively, heterogeneity in earnings profiles and access to wealth used to smooth consumption.

We use a grouping algorithm (k-medians) to assign unobserved types to individuals in the sample. To establish 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 the reasoning of Cunha, Heckman, and Navarro (2005). The basic idea is that, if agents know about their own wealth and earnings type, they should act upon such information and make choices that are consistent with their type. More generally, it

[10]These time series are publicly available from FRED.

should be possible to identify any ex-ante types that are observable by agents because such agents would respond to this information and act on it.

In this sense, idiosyncratic earnings growth rates and the dispersion of life-cycle consumption growth are the by-product of different underlying "types". Hence, they should be useful in predicting long term choices. Following Cunha et al., 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 over the life-cycle, one would expect that $E(S_i|\eta_i) \neq 0$. Moreover, given the relationship between unobserved types and outcomes such as earnings and consumption, one would expect that 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 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 eliminate any direct effect of earnings growth on the conditional expectation of college completion. In the case of wealth heterogeneity, we only need two types for the conditional expectation of college completion to be independent of consumption growth dispersion measures.[11] Therefore, in the estimation of marginal utilities and human wealth values, we use three earnings growth types and two wealth types to account for unobserved heterogeneity.

[11]Having established the cardinality of the different types of heterogeneity, we also corroborate our classification by verifying that this type of grouping is associated to large drops in within-group variances. Adding further types does not result in large decreases in within-group variances.

3.5 Imputing Human Wealth in the Survey of Consumer Finances

Although we carry out our main human wealth estimation exercise using the PSID, we carry out our analysis of inequality using the SCF. We do this because the SCF captures the upper tail of the wealth distribution far better than the PSID, and provides much more detail on assets. We cannot estimate human wealth in the SCF directly because it is not panel data. However, because we estimate the entire function $\Theta(\mathbf{z})$ in our main exercise, we can recover point estimates of human wealth in any data set where $\mathbf{z}$ is observed.

Unfortunately, not every variable in the vector $\mathbf{z}$ is observed in the SCF, in particular unobserved types η cannot be estimated from that data. In addition, many variables are only observed for the household head, for example educational attainment and age. To address this we impute the distribution of the missing variables from the PSID. It is important to impute the distributions of missing variables rather than replace by their conditional expectations, because the latter would average out heterogeneity and lead to underestimates of inequality.

We proceed with this imputation by first partitioning $\mathbf{Z}$ into observed variables $\mathbf{Z}^+$ and unobserved variables $\mathbf{Z}^-$. Next, we estimate the conditional distribution function $\Pi(\mathbf{Z}^-|\mathbf{Z}^+)$ in the PSID. Because $\mathbf{Z}^-$ takes discretely many values, this distribution can be viewed as a probability mass function $\Pi(\mathbf{Z}^-|\mathbf{Z}^+) = \{\pi_1(\mathbf{Z}^+), \pi_2(\mathbf{Z}^+), \ldots, \pi_M(\mathbf{Z}^+)\}$, where M is the number of points in the support $\mathbf{Z}^-$. In turn, each $\pi_m(\mathbf{Z}^+)$ can be estimated using the Nadaraya-Watson kernel estimator (still using PSID data).

Next, we begin working with the SCF data. We create M versions of this data set, where each imputes the value of $\mathbf{Z}^-$ to be that of the corresponding point in the support of $\mathbf{Z}^-$. The sample weight for observation i in version m of the data is reduced by $\pi_m(\mathbf{Z}_i^+)$. We then stack the data into a single data set, where each observations will appear M times, but the total weight of those observations equals the original sample weight of the observation. Human wealth can be computed for each observation and analysis can proceed using the sample weights appropriately.

4 Estimation Results

THESE ARE OLD RESULTS. TO BE UPDATED WITH MOST RECENT ESTIMATES

Using the mthods described in the previous sections we obtain a set of estimates for the overall lifetime wealth of each household in our SCF sample, as well as a detailed decomposition of the relative components of each household's wealth portfolio at any given point in time.

Figure 1: The value of human wealth, net worth and the sum of the two (total wealth) over the lifecycle. Values expressed in dollars are averages across all households.

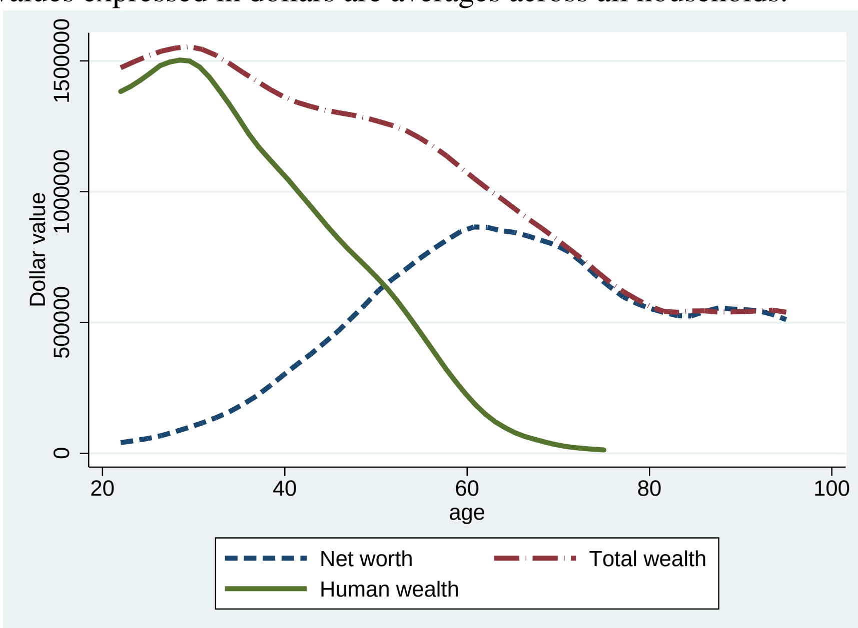

Figure 1 reports the estimated human and non-human wealth components, as well as the total wealth, averaged across households in our sample. It is immediate to see that young households tend to have most of their wealth concentrated in human wealth, which makes any shocks impacting labor supply or health extremely costly. In addition, shocks to human wealth are poorly insured because the stock of net worth (non-human wealth) is typically very low among young adults.

The value of human wealth peaks very early in the life cycle, between age 25 and 30, well before the age at which earnings typically peak. This can be explained by considering two aspects: first, the expected length of remaining working life is important when putting a price tag on a stream of labor earnings; second, earlier investments in human capital tend to carry a much higher return while effective human capital depreciation may become more severe with age.

The contrast between human and non-human wealth is quite striking: as is well known, net worth peaks after age 60 after accounting for a relatively small fraction of total wealth until age 40. However, its decline is much more gradual as it effectively accounts for all wealth after age 70. Given these patterns, total wealth appears to peak quyite early (around age 30) and, while declining to roughly 1/3 of its peak value by age 80, it exhibits less extreme proportional changes over the course of the life cycle.

In the two panels of Figure 2 we report the share of, respectively, net worth (real/financial wealth), 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. Importantly, the households in the top of the distribution of one of these variables may be different than the households in the top of the distribution of another of these variables.

Including human wealth when doing inequality accounting changes both the static and dynamic view of inequality over the past few decades. First, one can see that permanent income is much less concentrated than real/financial wealth: the total share of wealth held by the top 1% is a little over 15% and roughly half of their share of net worth which is well over 30%. This means that overall well-being, measured by permanent income, is much *less concentrated* than net worth. Given the fact that human wealth peaks very early in the life cycle, this is consistent with the fact that total wealth inequality is especially driven by later life discrepancies in net worth (late in life the distribution of human wealth becomes close to a degenerate with the first moment converging towards zero). Similar patterns can be observed for the richest top 10% of households. The share of permanent income held by the richest 10%

Figure 2: Concentration of net worth (real/financial assets), permanent income, and human wealth wealth. Each plot reports the share in the hands of the households at the top of the respective distribution.

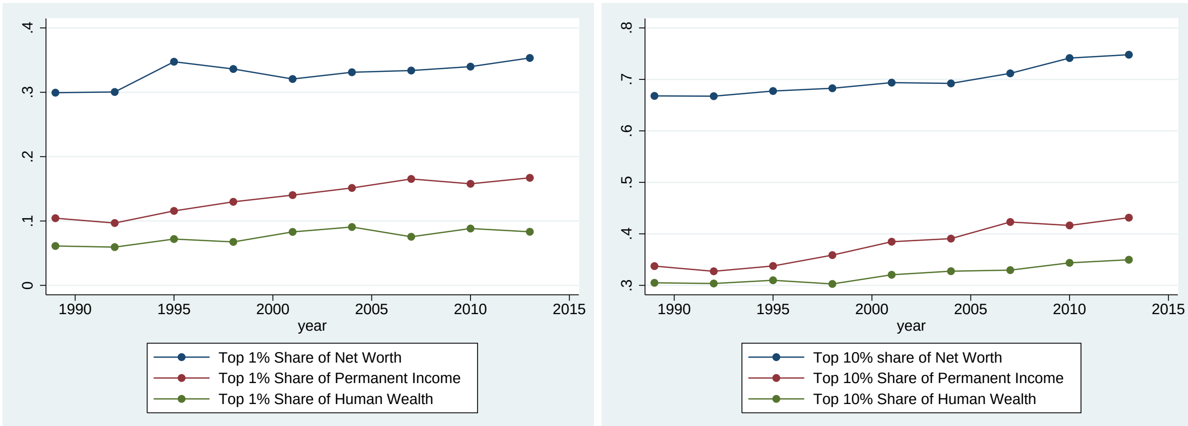

(a) Shares held by households in the top 1% of the respective distribution

(b) Shares held by households in the top 10% of the respective distribution

of households is roughly half the share of real/financial wealth held by the same households, indicating that well-being is more evenly distributed overall.

This narrative, however, becomes much less reassuring when we consider the evolution of permanent income concentration over time. While both human wealth and net worth have become significantly more concentrated between 1989 and 2013, it is apparent that the growth in permanent income concentration has far outpaced the growth in net worth concentration. This is especially striking if one considers the extreme attention and concern that has accompanied the growth of net worth. In fact, our analysis suggests that since 1989 the speed at which permanent income has concentrated in the hands of the richest households is almost twice as large as the increase in their share of real/financial wealth.

5 The Mechanics of Increasing Inequality

While the preceding analysis provides a rich portrayal of the historical patterns of US inequality over the past quarter century, it does leave some questions unanswered. How did such a steep increase in permanent income concentration come about? Have households in the top of the distribution of one variable (say, net worth) been adding to their share of other variables (say, human wealth)? How do young (and old) households in 2013 compare to their counterparts in 1989?

While it is clear that households in the top of the real/financial wealth distribution have been steadily increasing their share of real/financial wealth, this might in principle be offset by rising concentration of human wealth if a different set of households were at the top of the human wealth distribution. In contrast, if the same subset of households sit at the top of both distributions, this would compound and exacerbate the concentration of permanent income. Hence, a more subtle question is: how has the joint probability of being near the top of both the human and real/financial wealth distributions changed over time?

To answer these questions we proceed in two steps: (i) first, we characterize the changes in the empirical (marginal) distributions, and highlight the remarkable differences between the concentration of human wealth and that of the associated income/earning flows; (ii) next, we provide a way to account for the joint evolution of the distribution of real/financial wealth and human wealth.

5.1 Contrasting Measures of Concentration: Stocks versus Flows

To provide an overview of the changing concentration of economic resources, we report in tables 1 and 2 the values of the shares held by the top 10% of households. In particular, Table 1 reports the share of each variable held by the top 10% of households in the distribution *of that variable*, while Table 2 reports the share of each variable held by the top 10% of households *for real/financial wealth*.

One immediately noticeable finding from this analysis is that human wealth clearly exhibits the lowest concentration at the top among all variables. In particular, when comparing earnings inequality and human wealth inequality we find that, while human wealth is a func-

Year	Asset Wealth	Human Wealth	Earnings	Lifetime Wealth	Perm. Income
1989	0.668	0.305	0.408	0.337	0.360
1992	0.667	0.304	0.400	0.327	0.344
1995	0.677	0.310	0.407	0.338	0.359
1998	0.683	0.303	0.389	0.359	0.379
2001	0.694	0.321	0.434	0.385	0.411
2004	0.692	0.328	0.419	0.391	0.423
2007	0.712	0.330	0.443	0.423	0.457
2010	0.741	0.344	0.464	0.416	0.456
2013	0.748	0.350	0.462	0.431	0.471

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 earnings held by the households in the top 10% of the distribution of earnings.

Year	Asset Wealth	Human Wealth	Earnings	Lifetime Wealth	Perm. Income
1989	0.668	0.101	0.224	0.283	0.330
1992	0.667	0.093	0.222	0.262	0.307
1995	0.677	0.098	0.239	0.276	0.325
1998	0.683	0.104	0.237	0.311	0.355
2001	0.694	0.107	0.280	0.348	0.396
2004	0.692	0.099	0.273	0.356	0.409
2007	0.712	0.101	0.291	0.389	0.441
2010	0.741	0.100	0.304	0.385	0.445
2013	0.748	0.105	0.306	0.395	0.455

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 earnings held by the households in the top 10% of the distribution of net worth.

tion of potential future earnings, it tends to exhibit significantly lower concentration shares than earnings at the top. This suggests that (i) a non-trivial component of earnings concentration is due to transitory shocks; (ii) human wealth concentration may be mitigated by the fact that working lives can be relatively short and may suffer from setbacks that depreciate human capital and limit the excess-returns associated to skills.

It is worth stressing that, although the overall top 10% shares of both variables in Table 1 have risen, the top 10% share of earnings appears to have risen by more. Moreover, when looking at the top 10% of households in the financial/real wealth distribution (Table 2) only the share of earnings exhibits some growth. This discrepancy indicates that either the age of those in the top 10% of net worth has increased (hence higher earnings), or that there is a growing correlation between transitory shocks to earnings and real/financial wealth.

A simple comparison also reveals a great deal of information about the composition of households at the top of different wealth distributions. Even though human wealth has become more concentrated, as shown in Table 1, the share of human wealth attributable to the top 10% of the net worth distribution has not changed, as seen in Table 2. Thus, the large increases in the top 10% share of permanent income imply that the importance of net worth in household portfolios must have risen over time. Put simply, it appears that households who are in the top 10% or finanacial/real wealth have not become more likely to also be at the top of the human wealth distribution. Rather, this evidence suggests that human wealth has become a less important determinant of inequality in permanent income. Taken together, these observations illustrate how considering only earning flows (rather than human wealth stocks) would be misleading, whether studying levels or time trends.

Next, in tables 3 and 4, we overview the changes in concentration among the top 1% of households. This analysis broadly confirms the findings for the top 10% of households. The top shares of all variables are rising; however, when we focus only on households in the top 1% of net worth things look different. The top 1% share of human wealth does not rise, but their

year	Net Worth	Human Wealth	Earnings	Income	Permanent Income
1989	0.299	0.061	0.108	0.252	0.104
1992	0.300	0.059	0.109	0.156	0.097
1995	0.347	0.072	0.121	0.190	0.116
1998	0.336	0.068	0.114	0.230	0.130
2001	0.321	0.083	0.157	0.287	0.140
2004	0.331	0.091	0.137	0.245	0.151
2007	0.334	0.076	0.159	0.333	0.165
2010	0.340	0.088	0.160	0.256	0.158
2013	0.353	0.083	0.152	0.316	0.167

Table 3: 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	Earnings	Income	Permanent Income
1989	0.299	0.011	0.045	0.205	0.104
1992	0.300	0.010	0.048	0.113	0.096
1995	0.347	0.011	0.061	0.153	0.115
1998	0.336	0.013	0.061	0.169	0.129
2001	0.321	0.012	0.072	0.203	0.139
2004	0.331	0.013	0.078	0.197	0.151
2007	0.334	0.013	0.077	0.256	0.164
2010	0.340	0.011	0.068	0.184	0.157
2013	0.353	0.012	0.069	0.241	0.166

Table 4: 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.

share of permanent income increases more than their share of real/financial wealth. Again, this can only occur because real/financial wealth has become a more important determinant of permanent income (as opposed to a larger probability of being in the top of both human wealth and real/financial wealth distributions).

Finally, we stress that all these results are static, and one must consider the possibility that rising human wealth concentration early in life will generate rising real/financial wealth inequality later in life. In the following section we revisit some of these issues more formally.

5.2 The Role of Demographic Change

One striking feature of Human Wealth is how much it varies with age, as illustrated in Figure 1. Given this, one might hypothesize that the well-known trend towards an older population could be responsible for some of the observed changes in Human Wealth Inequality, and therefore also Permanent Income Inequality. To this end we carry out a re-weighting exercise in the spirit of Denardo, Fortin and Lemieux (1986) that allows us to consider how inequality in the variables we observe would have changed had the age distribution not changed since 1989. In particular, we fit probit regressions with a full set of age dummy variables as the independent variables, and then use the predicted probabilities to transform the SCF sample weights into a new set of weights that forces the age distribution to be constant. Figure 3 illustrates how the age distribution actually changed between 1989 and 2013, and how our counterfactual weights reshape the 2013 age distribution.

Given the above re-weightings we can produce counterfactual versions of the tables in the subsection above. In particular, Table 5 reproduces the top 10% shares of each variable under the counterfactual re-weighting. The first interesting thing to notice in Table 5 is that there is almost no change in Human Wealth Inequality, as opposed to the five percentage point

Figure 3: Actual age distributions in 1989 and 2013 (left) and re-weighted counterfactual age distribution in 2013 with actual age distribution in 1989 (right).

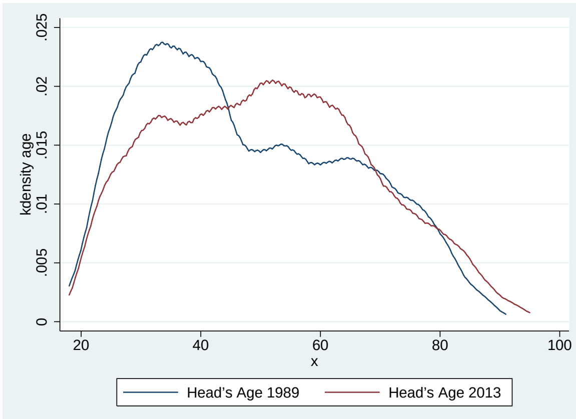

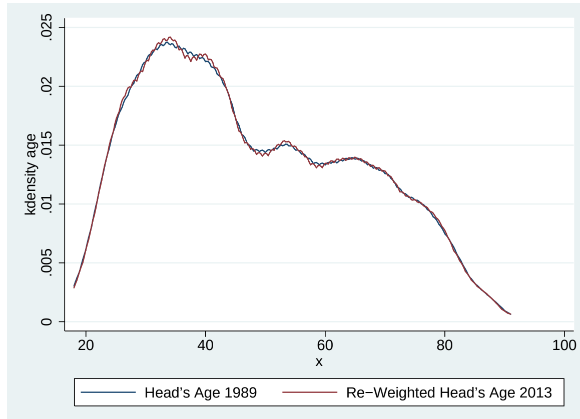

increase we observe in the actual data. This tells us that the rising Human Wealth top 10% share is due to the fact that a much smaller segment of the population is currently at their peak human wealth age, and so fewer individuals outside the top 10% have large Human Wealth stocks in 2013 than in 1989. Next we observe the effect that this has on Lifetime Wealth and Permanent Income. Lifetime Wealth and Permanent Income Inequality do not rise as much under the counterfactual weights, but they do still rise a great deal. This indicates that asset wealth is increasingly important as a component of Lifetime Wealth inequality, and possibly more important than rising Human Wealth inequality.

5.3 Who Got Richer? Inequality Accounting

While describing the changes is the marginal distributions of net worth and human wealth is instructive, it is not sufficient to understand observed changes in permanent income. As mentioned above, increasing concentration of permanent income may be due to an increasing overlap of the households sitting at the top of the marginal distributions of net worth and human wealth, but it might also be due to a progressive shift towards real/financial wealth in the composition of permanent income. In this section we address the question of which features

Year	Asset Wealth	Human Wealth	Earnings	Lifetime Wealth	Perm. Income
1989	0.668	0.305	0.408	0.337	0.360
1992	0.669	0.298	0.397	0.322	0.339
1995	0.683	0.303	0.404	0.335	0.355
1998	0.690	0.299	0.405	0.353	0.375
2001	0.706	0.313	0.424	0.378	0.407
2004	0.703	0.314	0.415	0.378	0.411
2007	0.720	0.316	0.437	0.404	0.441
2010	0.759	0.306	0.444	0.379	0.419
2013	0.764	0.307	0.441	0.393	0.433

Table 5: This table reports the counterfactual 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. These counterfactuals hold that age distribution constant as it was in 1989.

of the data are most important in accounting for the increasing concentration of permanent income (PI). To answer this question we explicitly quantify the extent to which the same set of households may be responsible for the increased concentration of PI.

For this accounting exercise we account for the joint evolution of net worth and human wealth. By definition, the share of PI held by the households in the top 10% of the PI distribution in period t can be written as

$$s_{PI}^{10}(t) = \frac{PI^{10}(t)}{PI^{10}(t) + PI^{90}(t)},$$

where $PI^{x}(t)$ is the aggregate value of permanent income held by households in the top $x\%$ of the distrbution of PI in year t. Of course, the share $s_{PI}^{10}(t)$ can be split between the net worth component and the human wealth component, as follows

$$s_{PI}^{10}(t) = \frac{NW^{10}(t) + HW^{10}(t)}{NW^{10}(t) + NW^{90}(t) + HW^{10}(t) + HW^{90}(t)}.$$

Here NW^n denotes the aggregate net worth value held by the top $n\%$ of the PI distrbution of households, while HW^n denotes the aggregate human wealth held by the same set of households.

Next, we define the wedge $\delta_x(t)$ as the value of variable x that, if redistributed from the top 10% to the bottom 90% of households, would make their relative shares of x in year t identical to those observed in year 1989. That is, we define $\delta_x(t)$ as the value such that

$$\frac{x^{10}(t)-\delta_x(t)}{x^{90}(t)+\delta_x(t)}=\frac{x^{10}(1989)}{x^{90}(1989)}.$$

The wedge $\delta_x(t)$ allows one to compute counterfactual inequality values for the distribution of a variable x, which can be used to account for changes in the concentration of x over time.[12]

For instance, if the relative distribution of real/financial net worth had not changed between 1989 and year t, the counterfactual share of permanent income (PI) held by households in the top 10% of the PI distribution in period t would be

$$\widetilde{s}_{PI}^{10}(t,\delta_{NW})=\frac{NW^{10}(t)+HW^{10}(t)-\delta_{NW}(t)}{NW^{10}(t)+NW^{90}(t)+HW^{10}(t)+HW^{90}(t)},$$

an expression that features the net worth wedge $\delta_{NW}(t)$ only at the numerator.

Similar reasoning suggests that, absent any changes in the distribution of human wealth after 1989, the counterfactual share of permanent income held by the the top 10% of househlds in the PI distribution would be

$$\widetilde{s_{PI}}^{10}(t,\delta_{HW})=\frac{NW^{10}(t)+HW^{10}(t)-\delta_{HW}(t)}{NW^{10}(t)+NW^{90}(t)+HW^{10}(t)+HW^{90}(t)}.$$

[12]It can be shown that

$$\delta_x(t)=x^{10}(t)\cdot\left[\frac{x^{90}(1989)}{x^{10}(1989)+x^{90}(1989)}\right]-x^{90}(t)\cdot\left[\frac{x^{10}(1989)}{x^{10}(1989)+x^{90}(1989)}\right].$$

Of course, any change in the observed value of the share s_{PI}^{10} (between 1989 and t) that is not accounted for by δ_{NW} and δ_{HW} must be due to changes in the relative importance of net worth and human capital in the composition of permanent income.

In practice, the difference $\varDelta_{NW} = s_{PI}^{10}(t) - \widetilde{s}_{PI}^{10}(t, \delta_{NW})$ measures how much of the increased concentration of PI is due to additional concentration of net worth in the hands of the top 10% of households. Similarly, the difference $\varDelta_{HW} = s_{PI}^{10}(t) - \widetilde{s}_{PI}^{10}(t, \delta_{HW})$ quantifies the role of more human wealth hoarding by the top households. Finally, the difference $\varDelta_{resid} = (s_{PI}^{10}(t) - s_{PI}^{10}(1989) - \varDelta_{NW} - \varDelta_{HW})$ identifies how much of the change in PI concentration is due to a shift in the composition of PI towards NW or HW, rather a change in the marginal distributions of NW or HW.

We use this decomposition to make sense of the changes in the concentration of permanent income between 1989 and 2013. As shown in Table 2 the share of permanent income in the hands of the top 10% of households (ranked according to their net worth) went up from 0.283 to .395, meaning that in 2013 the share of permanent income held by the top 10% of households was larger by more than 11 basis points (that is, they managed to lay claims on an extra 11% of total resources in the economy). Out of this 11% gain, roughly 3.6% was due to higher concentration in the marginal distribution of net worth, while only 0.2% can be attributed to more concentration in the marginal distribution of human wealth. The remaining change (roughly 7.4%) can only be attributed to (i) a change in the composition of permanent income that puts more weight on net worth and less on human wealth, and/or (ii) to an increase in the share of households that sit at the top of both the net worth and human wealth distribution. However, we verify that the share of households who belong to the top 10% of both marginal distributions (of net worth and human wealth) has actually decreased slightly from 16.6% in 1989 to 15.1% in 2013. This slightly lower share of households who are both rich in human and real/financial wealth implies that the higher concentration of permanent income is mostly due to the increasing importance of real/financial net worth as a component of permanent income.

While different families populate the top of the distributions of net worth and human wealth, the role of real/financial wealth has a driver of permanent income appears to have increased significantly between 1989 and 2013, and this largely explains the higher concentration of permanent income in the hand of high net worth households.

6 Conclusions

Accounting for heterogeneity in wealth composition is important to provide a general assessment of cross-sectional inequality, and of its evolution. In this paper we develop an approach that allows to quantify the amount of human wealth held by different households. Our analysis brings together different data sources and allows to obtain estimates of total household wealth, and of its composition. We show that these estimates contain new and valuable information about the way individuals and families hold wealth. This information is especially useful when accounting for the changing patterns of wealth inequality over the past three decades.

We show that human wealth is less concentrated than net worth. This suggests that inequality in permanent income is actually lower than inferred from popular measures of inequality based on real/financial wealth. However, it is also apparent that richer households have a growing share of permanent income. In fact, concentration of permanent income has grown much faster than concentration of real/financial net worth. As a consequence, effective wealth inequality has been growing much faster than previously thought, albeit from a lower initial level.

Through simple accounting exercises we show that the increasing concentration of permanent income is, to a large extent, due to the increasing importance of real/financial wealth as a share of total wealth. Changes in the marginal distributions of net worth and human wealth only account for a small part of the increase in permanent income inequality. Moreover, we find that the share of households who sit at the top of *both* net worth and human wealth distributions has actually decreased between 1989 and 2013, indicating that increased concentration of permanent income is not explained by a small set of households hoarding all types of wealth. Instead, we find that the key driver of permanent income concentration in the past decades seems to be the expansive growth of real/financial assets as share of the wealth portfolio of rich households. High net worth households, rather than high human wealth households, account for a larger share of total permanent income in 2013 than they did in 1989, suggesting that changes in wealth composition may be key to understand recent inequality patterns.

References

ATTANASIO, O., AND L. PISTAFERRI (2014): “Consumption Inequality over the Last Half Century: Some Evidence Using the New PSID Consumption Measure,” *The American Economic Review, P P*, 104(5), 122–126.

AUTOR, D. H., L. F. KATZ, AND M. S. KEARNEY (2008): “Trends in US Wage Inequality: Revising the Revisionists,” *The Review of Economics and Statistics*, 90(2), 300–323.

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

BENZONI, L., AND O. CHYRUK (2015): “The Value and Risk of Human Capital,” Working paper, Chicago Fed.

BONHOMME, S., T. LAMADON, AND E. MANRESA (2017): “Discretizing unobserved heterogeneity,” Discussion paper, University of Chicago.

BRICKER, J., A. HENRIQUES, J. KRIMMEL, AND J. SABELHAUS (2016): “Measuring Income and Wealth at the Top using Administrative and Survey Data,” *Brookings Papers on economic Activity*.

CONESA, J. C., S. KITAO, AND D. KRUEGER (2009): “Taxing capital? Not a bad idea after all!,” *The American economic review*, 99(1), 25–48.

CUNHA, F., J. HECKMAN, AND S. NAVARRO (2005): “Separating Uncertainty from Heterogeneity in Life Cycle Earnings,” *Oxford Economic Papers*, 57(2), 191–261.

DAVILA, J., J. H. HONG, P. KRUSELL, AND J.-V. RÍOS-RULL (2012): “Constrained efficiency in the neoclassical growth model with uninsurable idiosyncratic shocks,” *Econometrica*, 80(6), 2431–2467.

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

Escanciano, J. C., S. Hoderlein, A. Lewbel, O. Linton, and S. Srisuma (2015): “Nonparametric Euler Equation Identification and Estimation,” Discussion paper, mimeo, Boston College.

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

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

Gottschalk, P., R. Moffitt, L. F. Katz, and W. T. Dickens (1994): “The growth of earnings instability in the US labor market,” *Brookings Papers on Economic Activity*, 1994(2), 217–272.

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

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

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

KAYMAK, B., AND M. POSCHKE (2016): "The evolution of wealth inequality over half a century: The role of taxes, transfers and technology," *Journal of Monetary Economics*, 77, 1–25.

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

LEVY, F., AND R. J. MURNANE (1992): "US Earnings Levels and Earnings Inequality: A Review of Recent Trends and Proposed Explanations," *Journal of Economic Literature*, 30(3), 1333–1381.

LUCAS JR, R. E. (1978): "Asset prices in an exchange economy," *Econometrica: Journal of the Econometric Society*, pp. 1429–1445.

PIKETTY, T., AND E. SAEZ (2006): "The evolution of top incomes: a historical and international perspectives," in *American Economic Review Papers and Proceedings*, pp. 200–205. Volume 96, number 2.

R. HAVEMAN, R. BERSHADKER, A., AND J. SCHWABISH (2003): "Human Capital in the United States from 1975 to 2000: Patterns of Growth and Utilization," Upjohn Institute.

RIOS-RULL, J.-V., M. KUHN, ET AL. (2016): "2013 Update on the US Earnings, Income, and Wealth Distributional Facts: A View from Macroeconomics," *Quarterly Review*, pp. 1–75, April.

SAEZ, E., AND W. KOPCZUK (2004): "Top Wealth Shares in the United States: 1916-2000: Evidence from Estate Tax Returns," NBER Working Paper 10399.

SAEZ, E., AND G. ZUCMAN (2014): "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," NBER working paper 20625.