SUPPLEMENTARY CODEBOOK

Who gains when interest rates fall?

Sylvain Catherine, Max Miller, James D. Paron, and Natasha Sarin

This codebook documents data construction, calibration, and solution details that the paper and its appendix summarize. Equations, tables, and figures of this codebook carry the prefix S. All other numbers refer to the paper.

1 Survey of Consumer Finances

Data on household portfolios come from the Survey of Consumer Finances (SCF). We construct networth as:

$$\begin{aligned}\texttt{networth_d} = \texttt{cash_dep} + \texttt{equity} + \texttt{fixed_inc} + \texttt{real_estate} \\ + \texttt{bus} + \texttt{vehic} - \texttt{mortgage_dbt} - \texttt{vehic_dbt} - \texttt{edu_dbt} - \texttt{other_dbt},\end{aligned}$$

where the constituent variables are defined as:

- `cash_dep`: value of cash deposits defined as liquid accounts (`liq`) which are the sum of all checking, savings, and money market accounts, call accounts at brokerages, and prepaid cards, net of the part of savings and money market accounts invested in stock, added to certificates of deposit (`cds`) and the cash value of life insurance (`cashli`).
- `equity`: value of all financial assets invested in stock, which include directly held stock, stock mutual funds, and the portion of any combination mutual funds, annuities, trusts, IRA/Keogh accounts, and other retirement accounts invested in stock.
- `fixed_inc`: value of all other remaining financial assets (`fixed_inc` = `fin` − `cash_dep` − `equity`). The largest component of this asset category is bonds held outright, in mutual funds, and in retirement accounts.
- `real_estate`: value of the primary residence (`houses`) plus the value of other residential real estate (`oresre`) and net equity in nonresidential real estate (`nnresre`).
- `bus`: reported market value of private business interest.

– `vehic`: prevailing retail value for all vehicles owned by household.

– `mortgage_dbt`: housing debt from mortgages, home equity loans, and home equity lines of credit (`mrthel`) plus debt for other residential property (`resdbt`).

– `vehic_dbt`: debt from vehicle loans (`veh_inst`).

– `edu_dbt`: debt from student loans.

– `other_dbt`: other debt, including other lines of credit plus credit card balance (`ccbal`) plus installment loans less education loans and vehicle loans (`other_dbt` = `othloc` + `ccbal` + `install` − `edn_inst` − `veh_inst`).

In addition to portfolio data, we also use data on household wage income (`wageinc`) which we combine with data on the number of adults in the household and the Social Security wage index to create a per capita wage measure that is comparable over time. We count married and widowed respondents as two-adult households and all others as single-adult households.

2 Duration component calculations

2.1 Duration of equity

The duration of `equity` is obtained using annual estimates for the duration of the aggregate stock market from Greenwald et al. (2023), Figure D2 of the September 2023 working paper version. These estimates are applied uniformly to all individuals in the SCF by survey year.

To allow equity duration to vary across the wealth distribution, we apply a mean-preserving heterogeneity adjustment analogous to the one described for private business wealth in Section 2.4. We split households into six wealth groups g—bottom 50%, 50–90%, 90–99%, 99–99.9%, 99.9–99.99%, top 0.01%—and compute the price-dividend ratio for each group using directly held stocks and mutual funds and dividend income reported in the SCF. Specifically, let $E_g = \sum_{i \in g} E_i$ and $D_g = \sum_{i \in g} D_i$ denote aggregate holdings of directly held stocks and mutual funds and dividend income in wealth group g, where both are deflated by the Social Security wage index. The group-level price-dividend ratio is $m_g^{eq} = E_g / D_g$, and the equity-weighted population average is $\bar{m}^{eq} = \sum_g E_g \, m_g^{eq} / \sum_g E_g$.

The mean-preserving adjustment factor for group g is

$$\alpha_g^{eq} = \frac{m_g^{eq}}{\bar{m}^{eq}},$$

so that the adjusted equity duration for households in group g is $\text{dur}_g^{eq} = \alpha_g^{eq} \times \text{dur}^{eq,\text{agg}}$, where $\text{dur}^{eq,\text{agg}}$ is the aggregate duration estimate from Greenwald et al. (2023). By construction, $\sum_g E_g \, \alpha_g^{eq} / \sum_g E_g = 1$. The adjusted duration applies to directly held stocks and mutual funds; equity held through retirement accounts and other managed assets receives the aggregate duration.

2.2 Duration of fixed income

Data on the Macaulay duration of government bonds, municipal bonds, corporate bonds, and mortgage backed securities come from Bloomberg where the series used are:

- US gov/credit: `LUGCTRUU`
- US Treasury: `LUATTRUU`
- Government-related: `LD08TRUU`
- US aggregate: `LBUSTRUU`
- Municipal bond: `LMBTTR`
- Corporate: `LUACTRUU`
- US MBS: `LUMSTRUU`
- Global aggregate: `LEGATRUU`

For holdings of US government bonds (`govtbnd` + `gbmutf` + `savbnd`), we use the market-value weighted average Macaulay duration of the US gov/credit, US Treasury, and, from 1994, government-related bond categories. For holdings of tax-free and municipal bonds (`notxbnd`+`tfbmutf`), mortgage-backed securities (`mortbnd`), corporate bonds (`corpbnd`), and foreign bonds (`forbnd`), we use the Macaulay duration of municipal bonds, corporate bonds, US MBS, and the global aggregate, respectively. For other financial assets (`othfin`), we assign 5.64, which is the average fixed income wealth-weighted duration

in the SCF assets for which we have data. All remaining fixed income assets receive the Macaulay duration of the US aggregate.

Similar to equity and private business, we also apply a mean-preserving heterogeneity adjustment to bond durations across wealth groups. For each of the same six wealth groups g, let $B_g = \sum_{i \in g} B_i$ and $I_g = \sum_{i \in g} I_i$ denote aggregate bond holdings and interest income (both deflated by the Social Security wage index). The group-level price-to-interest ratio is $m_g^{bond} = B_g / I_g$, and the bond-weighted population average is $\bar{m}^{bond} = \sum_g B_g \, m_g^{bond} / \sum_g B_g$. The mean-preserving adjustment factor for group g is

$$\alpha_g^{bond} = \frac{m_g^{bond}}{\bar{m}^{bond}},$$

so that the adjusted bond duration for households in group g is $\text{dur}_g^{bond} = \alpha_g^{bond} \times \text{dur}^{bond,\text{agg}}$, where $\text{dur}^{bond,\text{agg}}$ is the Bloomberg duration estimate for the relevant bond type. By construction, $\sum_g B_g \, \alpha_g^{bond} / \sum_g B_g = 1$. This allows bond duration to vary across the wealth distribution in proportion to valuation multiples implied by reported interest income.

2.3 Duration of real estate

The duration of `real_estate` is obtained using the annual estimates of the duration of aggregate real estate from Greenwald et al. (2023), Figure D2 of the September 2023 working paper version. These estimates are applied uniformly to all individuals in the SCF by survey year.

2.4 Duration of private business wealth

The duration of private business wealth is computed for each household as the value of household businesses, `bus`, divided by the annual cashflows from those equity holdings. However, the annual cashflows from those equity holdings are not reported in the SCF, the major issue being that cashflows from private businesses partially contain implicit or explicit labor income for the entrepreneur. As such, we must estimate or difference out this labor income. We do this in four ways, depending on the household's role in the business and what is reported.

1. For households whose main respondent has an active management role in either of the household's potential actively managed businesses, reports being self-employed, and reports not receiving a salary, we estimate their predicted wage.

– The predicted wage is estimated via weighted least squares on SCF adults aged 18 to 62 with positive wages, where the adult's wage income, scaled by the Social Security wage index, is the dependent variable, and the independent variables are a third-degree polynomial in age interacted with dummies for each Race $\times$ Education $\times$ Gender group.

2. For households whose main respondent has an active management role in either of the household's potential actively managed businesses and reports being self-employed and receiving a salary or reports being employed by someone else, we add their reported wage to `bussefarminc` and subtract the maximum of their predicted wage and reported wage.

3. We repeat steps 1) and 2) for spouses who have an active management role in either of the household's potential actively managed businesses.

4. All other households with positive private business wealth who do not meet the criteria for a wage subtraction are given cashflows equal to `bussefarminc`.

We then aggregate `bus` and the estimated annual cashflows within each survey year and divide them to obtain an annual time series of valuation ratios.

Next, to allow our aggregate estimates of private business duration to vary over the wealth distribution, we perform a mean-preserving adjustment to these aggregate duration estimates. First, we split the population into different networth groups defined by whether they are in the bottom 50%, 50–90%, 90–99%, 99–99.9%, 99.9–99.99% or the top 0.01% of the wealth distribution. We then divide total business wealth (`bus`) by total income from businesses (`bussefarminc`) within each networth group. Provided that cashflows from equity are proportional to labor income, this provides a proxy for duration for each group. These price-total income ratios are then divided by the business wealth-weighted average for the aggregate population to obtain a mean-preserving adjustment which is applied to the annual aggregate private business duration estimates. The duration is then given by

$$\text{dur}(\text{Private Business}_{ct}) = \underbrace{\left(\frac{\sum_{t,i\in c}\text{Bus. Value}_{it}}{\sum_{t,i\in c}\text{Bus. Income}_{it}}\right)}_{\text{Wealth-group-specific valuation multiple}} \times \underbrace{\left(\frac{\sum_{i}\text{Bus. Income}_{it}}{\sum_{i}(\text{Bus. Income}_{it} - \text{Wages}_{it})}\right)}_{\text{Business-income-to-dividend conversion ratio}}$$

2.5 Duration of vehicles

The `vehic` category in the SCF contains detailed information on up to 4 automobiles, up to 2 non-automobile vehicles, and an aggregation of additional automobiles and non-automobile vehicles owned by the household. For the primary automobiles of the household, we attribute an expected lifetime of 8 years for 1989 and 12 years for 2019, linearly interpolating in intermediate years. We calculate the time left on an automobile's life as the model year plus the expected age minus the survey year. We assume a fixed depreciation rate to 0 over the car's remaining years, and calculate the duration using (6). We attribute a duration of one to vehicles whose age exceeds their expected lifetime.

For the aggregation of additional automobiles owned, we attribute a duration equal to the average duration of the household's automobiles of known age, or 6 years if it has none. Automobiles among the first four with no reported model year are treated the same way. For all non-automobile vehicles owned by the household, we ascribe a duration of 6 years.

2.6 Defined benefit pension wealth

We add defined benefit (DB) pensions to the analysis in three steps: (1) computing the funded portion of aggregate DB pension entitlements, (2) allocating these entitlements to individual SCF households, and (3) assigning the resulting pension wealth to asset categories with corresponding interest-rate sensitivities.

First, we take aggregate DB pension entitlements from the Federal Reserve's Distributional Financial Accounts (DFA), which reports total DB pension net worth by wealth group: bottom 50%, 50–90%, 90–99%, 99–99.9%, and top 0.1%. To isolate the funded portion, we compute a funding ratio from the Flow of Funds (Table L.117) as the ratio of pension fund assets (series FL592000075.A) to total pension entitlements (series FL594190045.A) for each year from 1945 to 2024. We apply this funding ratio to the DFA entitlements within each wealth group:

$$\text{Funded DB}_{ct} = \text{DFA DB}_{ct} \times \frac{\text{Pension fund assets}_t}{\text{Pension entitlements}_t},$$

where c denotes the wealth group. Including only the funded portion is consistent with Saez and Zucman (2016), who also restrict attention to funded DB pension wealth in their analysis.

Second, we allocate these aggregate funded DB pension amounts to individual house-

holds in the Survey of Consumer Finances (SCF) using a weighting scheme based on four factors. For each household, we construct an allocation weight:

$$w_{it} = \text{accrual}_{it} \times \text{mortality}_{it} \times \text{income}_{it} \times \mathbf{1}[\text{has DB}_{it}],$$

where:

- $\mathbf{1}[\text{has DB}_{it}]$ equals one if the respondent or spouse reports having a DB pension plan from a current or previous job (`dbplancj==1` or `dbplant==1`), and zero otherwise;
- accrual_{it} is a predicted fraction of lifetime pension benefits already accrued, obtained from a regression of the ratio of accrued Social Security benefits to total expected benefits on income centile, net worth centile, age, and sex indicators, estimated on the subsample with age ≤ 65. For households over age 65 we set $\text{accrual}_{it} = 1$; for those under 20 we set it to 0. All other predicted values are capped between 0 and 1;
- $\text{mortality}_{it} = \min\left(1,\ e_{x,t}/e_{65,t}\right)$, where $e_{x,t}$ is remaining life expectancy at age x in year t from the Human Mortality Database, computed separately by sex. This adjusts for the fact that younger households have longer expected payout horizons;
- $\text{income}_{it} = \min\left(\max\left(\text{wageinc}_{it}/\overline{\text{wageinc}}_{ct},\ 0\right),\ 5\right)$, where $\overline{\text{wageinc}}_{ct}$ is the mean wage income within wealth group c in year t. This captures the correlation between pension generosity and career earnings.

Within each wealth group c and year t, we allocate the aggregate funded DB pension amount proportionally based on survey-weighted individual weights:

$$\text{DB}_{it} = \frac{w_{it}}{\sum_{j \in c} w_{jt} \cdot \omega_j} \times \text{Funded DB}_{ct},$$

where ω_j denotes the SCF survey weight for household j.

Third, we allocate each household's DB pension wealth to three asset categories—cash deposits, fixed income, and equity—using the asset composition of private pension funds from Table L.117 of the Flow of Funds, which we compute for each year. Specifically, cash deposits include checkable and time/savings deposits; fixed income includes Treasury

securities, agency securities, corporate and foreign bonds, and mortgages; and equity includes corporate equities and mutual fund shares. Miscellaneous assets are excluded from the denominator.

These DB pension amounts are added to the corresponding categories in the household's balance sheet. For the purpose of computing interest-rate sensitivities, DB pension equity receives the aggregate stock market interest-rate sensitivity and DB pension fixed income receives the aggregate US bond interest-rate sensitivity, reflecting the diversified nature of pension fund portfolios. The interest-rate sensitivity of the combined category (e.g., public equity) is then a value-weighted average of the household's direct holdings and the DB pension component.

2.7 Duration of debts

For the debt categories, `mortgage_dbt`, `vehic_dbt`, and `other_dbt`, we break each up into their component loans as described in the SCF extract and calculate the duration of each loan separately. For each loan, we assume a fixed payment schedule, and thus its duration can be calculated using equation (6), where N is the number of years remaining on the loan and P_{nt} is the price of an n-year zero-coupon bond in year t.[1]

Under our fixed payment assumption, the only metric we need for each loan is its time remaining. Since different loan component variables contain different amounts of information in the raw SCF, we calculate the time remaining differently depending on the available information for each component loan group: primary component loans, aggregated additional loans, and lines of credit. The primary component loans of each debt category contain information on loan origination, balance, payments, and interest rates. For these loans, we calculate the number of years remaining on the loan payments using the reported origination year, length of loan at origination, and survey year. For respondents with a positive loan balance who have missing responses for loan length or a negative calculated time remaining, we impute time remaining with balance (B), initial amount (L), interest rate (R), and year of origination (p) using the equation

$$T = \frac{\log(R^p - B/L) - \log(1 - B/L)}{\log R} - p.$$

The aggregated additional loans group contains loan variables that capture an aggregation

[1] For credit card balances—which do not have this information—we assign a duration of 1 year.

of loans that the respondents hold in addition to the primary ones in each debt category. These loans include data on only loan balance and payments (X). Using the average interest rates for primary loans in the same debt category, we calculate time remaining as

$$T = -\frac{\log(1 - B(R-1)/X)}{\log R}.$$

The third group of component loans is the lines of credit. The line of credit variables contain information on loan balance, typical payments, and interest rates. With these data points, we calculate time remaining according to the same formula used for the aggregated additional loans group. Finally, there is an aggregated additional lines of credit variable, which we assign a duration equal to the average duration of the household's other lines of credit with a balance, or the median duration of first lines of credit if it has none.

We replace the duration of loans with a predicted time remaining under one year with a duration of one and give the median duration to respondents with a positive loan amount but insufficient information to calculate time remaining on the loan.

We also adjust student loan balances and duration to account for income-driven repayment (IDR) following Catherine and Yannelis (2023). Under IDR, borrowers pay 10% of earnings above 1.5 times the Social Security wage index (scaled by 0.225) annually for up to 25 years. We discount these payments at the riskfree rate to obtain the present value and corresponding cashflow duration.

Finally, since adjustable rate mortgages (ARMs) have no interest-rate exposure after they float, we treat them differently. In particular, we cap their duration at 4—the average number of years until the interest rate begins to float in the SCF. Approximately 12% of mortgages and 19% of investment property loans are adjustable rate in the data.

2.8 Interest-rate sensitivity of Social Security wealth

The interest-rate sensitivity of Social Security wealth comes from the methodology used in Catherine et al. (2025). We generate their baseline risk-adjusted Social Security wealth under the net present value wealth concept and the Treasury yield curve. We then generate Social Security wealth under an identical specification where the log forward rate at horizon h in survey year t is given by $\tilde{f}_{h,t} = f_{h,t} + \varphi^h 0.01$ where $f_{h,t}$ is the unshocked log forward rate and φ is our calibrated persistence for the interest rate process shown in Table 1. This is the same thing as applying a one-period shock of 0.01 to the log riskfree rate under the

process in equation (3).

3 Calibration and simulation inputs

3.1 Macroeconomic variables

3.1.1 Time series of riskfree interest rates

To obtain a time series of the short-term real interest rate, we use a methodology similar to that of Beeler and Campbell (2012). Using the yield on the 10-year nominal Treasury bond y_{10} and annual inflation rate π from Global Financial Data, we estimate the annual regression

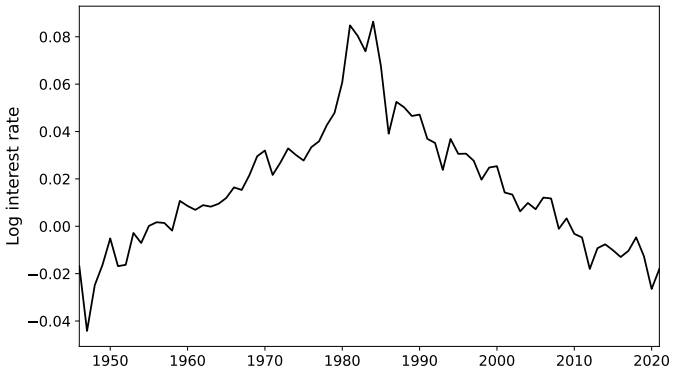
$$y_{10,t} - \pi_{t,t+1} = \beta_0 + \beta_1 y_{10,t} + \beta_2 \pi_{t-1,t} + \epsilon_{t+1} \tag{S.1}$$

on the post-war period. The fitted values are then taken as our estimate of the expected riskfree rate 10 years from time t, $\hat{f}_{10,t}$. From this, equation (3) yields the time-t riskfree rate:

$$r_{ft} = \varphi^{-10}(\hat{f}_{10,t} - (1 - \varphi^{10})\bar{r}_f). \tag{S.2}$$

We use this methodology for two main reasons. First, by using long-term rates to back out short-term rates, we smooth much of the short-term variation in measured short-term real rates that are potentially outside of our model. Second, this methodology allows us to extend our real rate series further into the past, enabling a longer simulation prior to our period of interest. This procedure yields a time series of annual realizations of real rates $\{r_{ft}\}$ and shocks $\{\epsilon_{rt}\}$ from 1791 to 2021. Figure S.1 shows the post-war time series.

Figure S.1: Time series of riskfree rates, Post-war sample



Note: This figure presents the time series of short-term riskfree rates as estimated by equation (S.1) and transformed by equation (S.2).

3.1.2 Economic growth and stock-market returns

We jointly estimate the dynamics of aggregate growth (7) and stock returns (12) via the simulated method of moments (SMM). The parameter vector to be estimated is[2]

$$\Theta \equiv \begin{bmatrix} \bar{g} & \varphi_g & \mu_g^- & \sigma_g & p_R & \lambda_{gr} & \mu_s^- & \mu_s^+ & \sigma_s \end{bmatrix}. \tag{S.3}$$

We estimate the probability of a recession p_R directly from the data as the frequency of NBER recessions. To identify the remaining 8 parameters, we use the first four moments (i.e., mean, variance, skewness, and kurtosis) of log stock returns and of log GDP growth, the autocorrelation of GDP growth, the correlation of stock returns and GDP growth, and the correlation of riskfree rates with five-year moving average growth. We use the correlation between lagged returns and GDP growth as a moment because stock markets are forward looking. This shows up in the data—as noted in Catherine et al. (2024); Fama (1990); Schwert (1990)—by stock-market returns predicting future real activity more strongly than they predict current real activity. The higher-order moments (skewness and kurtosis) are particularly useful for identifying parameters of the normal mixtures. The correlation of rates with a moving average of the growth rate, as opposed to the growth rate itself, provides identification for λ_{gr} by isolating the low-frequency covariance.

The SMM procedure then estimates Θ as follows. Letting $m(\Theta)$ denote the vector of simulated moments implied by the model and $\hat{m}$ the corresponding vector of data moments, our estimated parameters minimize the squared (percentage) errors:

$$\hat{\Theta} = \arg\min_{\Theta} \left\{ \left(\frac{\hat{m} - m(\Theta)}{\hat{m}} \right)^{\top} \left(\frac{\hat{m} - m(\Theta)}{\hat{m}} \right) \right\}, \tag{S.4}$$

where division of vectors is an element-wise operation. For the data moments, we use historical data from a number of sources. We construct the growth series using per capita real GDP growth as a proxy for aggregate wage dynamics. The primary data source is the Global Macro Database (GMD), from which we extract US real GDP per capita from 1900 to 2021. Stock-market return data cover 1900–2021 and come from Shiller. We use the log total return on the S&P 500. The stock-return moments use the full sample; the

[2]Recall that λ_{sr} is given by (C.20) and that μ_g^+ sets $p_R\mu_g^- + (1 - p_R)\mu_g^+ = 0$, so these are implicitly estimated here.

growth moments and the correlations use data from 1960. Table 3 reports the estimated parameters. Table S.1 reports the moments in the data and model.

Table S.1: Moments for estimation of growth and stock returns

Moment	Description	Data	Model
$\mathbb{E}[r_{st}]$	Mean stock return (1900-2021)	0.066	0.0659
$\sigma(r_{st})$	Standard deviation of stock returns	0.187	0.188
Skew(r_{st})	Skewness of stock returns	-0.831	-0.899
Kurt(r_{st})	Kurtosis of stock returns	3.61	3.51
$\mathbb{E}[g_t]$	Mean wage growth (1960-2021)	0.0196	0.0196
$\sigma(g_t)$	Standard deviation of wage growth	0.0205	0.0209
Skew(g_t)	Skewness of wage growth	-0.629	-0.689
Kurt(g_t)	Kurtosis of wage growth	3.42	3.34
Corr(g_t, g_{t-1})	Autocorrelation of wage growth	0.164	0.166
Corr(g_t, r_{st})	Contemporaneous correlation	0.541	0.564
Corr$(r_{f,t}, \bar{g}_{5,t})$	Risk-free rate and 5-year average growth	0.211	0.223
p_R	Recession probability	0.164	0.164

3.2 Idiosyncratic income process

Our calibration of the idiosyncratic income component of earnings follows Catherine (2022) closely, with two differences. In that study, the parameters of the income process are estimated by targeting moments of the distribution of log earnings growth from 1978 to 2011, feeding the historical mean log earnings growth of each of these years into the simulation. The targeted moments include the standard deviation of log earnings growth at the 1- and 5-year horizons, and the skewness at the 1-, 3-, and 5-year horizons. These moments are computed using Social Security administrative data and reported in Guvenen et al. (2014).

We estimate the process using the same procedure as in Appendix B.1 of Catherine (2022), with two differences. First, for years prior to 1978, we feed into the simulation the historical time series of log changes in real per capita GDP, which we use to calibrate the process for economic growth as in Section 3.1. For the period 1978–2011, we feed into the simulation the historical mean log earnings growth rate reported in Guvenen et al. (2014), but scale this variable such that it exhibits the same volatility over that period as per capita GDP growth, which is roughly 1.5 times lower. Our goal is to obtain comparable variability in idiosyncratic income risk over the business cycle in the model and data, while leveraging the longer time span of the per capita GDP growth series relative to the moments reported in Guvenen et al. (2014).

The second difference is that we account for the diversification arising from two earners

in the typical US household. As Guvenen et al. (2014) only report moments for individual earners, we proceed in two steps to estimate the income process for a two-earner household. First, we estimate all parameters for a single-earner household by targeting the moments reported in Guvenen et al. (2014). The estimated parameters are reported in the 'One earner' row of Table S.2 and are very close to those in Catherine (2022), except for the parameter λ_{zg}, which governs the relationship between GDP growth and the skewness of idiosyncratic income shocks. This parameter is roughly 1.5 times larger than in Catherine (2022), reflecting the fact that the GDP growth time series is itself about 1.5 times less volatile than the mean log earnings growth series used in that study. Second, we use the estimated parameters to generate new moments for two-earner households. To do so, we simulate individuals using our estimated income process and randomly match them to create two-earner households. We then compute log income growth moments from these households, assuming that the highest earner in each household receives 70% of the combined income. Finally, we rerun the estimation procedure targeting the new moments. The estimated parameters are reported in the 'Two earners' row of Table S.2.

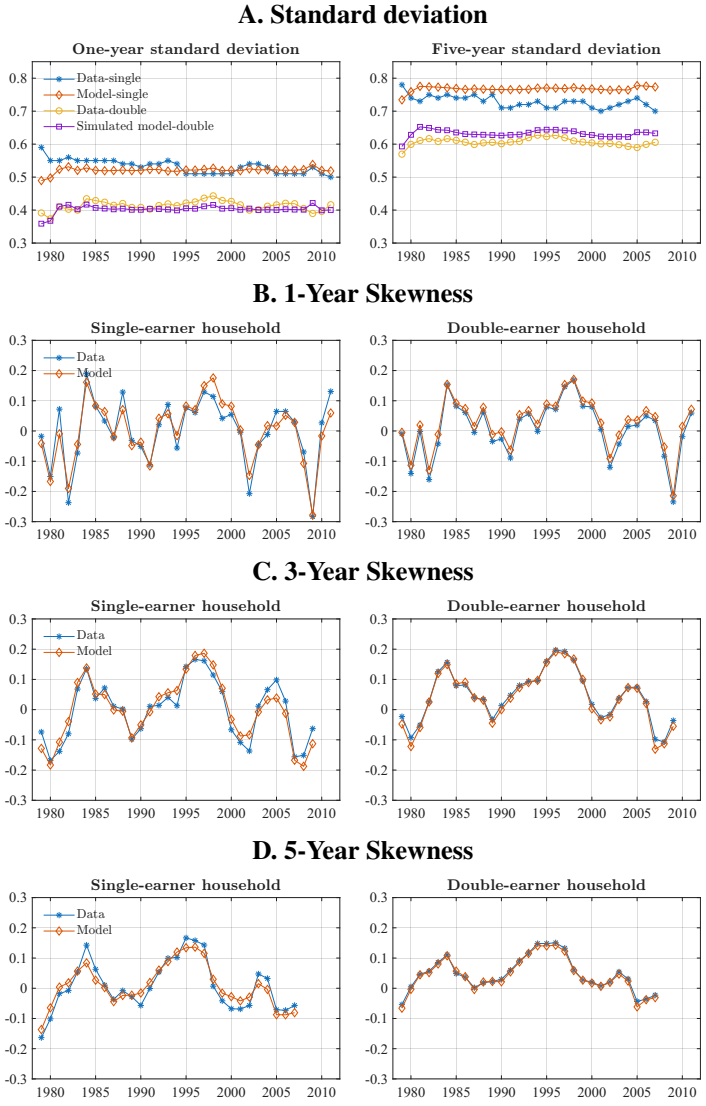
Table S.2: Estimated parameters: Idiosyncratic labor-income shocks

		Persistent shocks					Transitory shocks	
	p_z	ρ_z	$\overline{\mu_z^-}$	λ_{zg}	σ_z^-	σ_z^+	σ_η^-	σ_η^+
One earner	.145	.971	-.085	6.542	.644	.012	.758	.097
Two earners	.224	.993	-.026	5.219	.352	.059	.414	.124

This table reports parameter estimates for the idiosyncratic income process, conditional on the group-level aggregate shock. The single-earner process is estimated as in Catherine (2022) by targeting the time-series of cross-sectional moments of individual income log growth rates in the SSA administrative data from Guvenen et al. (2014). To estimate the double-earner process, we first simulate data for couples using the single-earner process and recompute the target moments using the sum of two incomes. We assume a 30% gap between the highest and lowest earner and no assortative matching.

Figure S.2 presents the values and fit of targeted and simulated moments from the estimation of the idiosyncratic labor income shocks. Panels on the left report targeted and simulated moments for single-earner households, where targeted moments are from Guvenen et al. (2014). Panels on the right report targeted and simulated moments for households with two earners.

Figure S.2: Fit of the idiosyncratic labor-income risk model



3.3 Income tax rates

Households pay income tax on labor income and Social Security benefits. In the model, households face the following marginal tax rates:

$$\text{Marginal Tax Rate}_{it} = \begin{cases} 0.10 & \text{if } \tilde{L}_{it} < 0.18, \\ 0.12 & \text{if } 0.18 \leq \tilde{L}_{it} < 0.72, \\ 0.22 & \text{if } 0.72 \leq \tilde{L}_{it} < 1.54, \\ 0.24 & \text{if } 1.54 \leq \tilde{L}_{it} < 2.94, \\ 0.32 & \text{if } 2.94 \leq \tilde{L}_{it} < 3.73, \\ 0.35 & \text{if } 3.73 \leq \tilde{L}_{it} < 9.32, \\ 0.37 & \text{if } \tilde{L}_{it} \geq 9.32. \end{cases} \tag{S.5}$$

The bendpoints in this formula are the limits of the 2020 tax brackets divided by the wage index.

3.4 Entrepreneurs

We calibrate the private business income (30) and wealth (32) processes to data from the SCF. We need to calibrate the deterministic component of income $\ell_E(a_{it})$, the initial volatility σ_{Ez0} and time-series volatility σ_{Ez} of the permanent component z_{it}, the mean valuation ratio υ, and the probability of becoming an entrepreneur by age $p_E(a_{it})$. In the SCF, we consider a household to be an entrepreneur if the main respondent actively manages a private business and the household's business and farm income exceeds its wage income. We restrict our sample to households before retirement (i.e., aged 18 to 60).

We choose the mean log valuation ratio υ to match the average log aggregate valuation ratio:

$$\upsilon = \mathbb{E}\left[\log \frac{\sum_{i=1}^{N_t} E_{it}}{\sum_{i=1}^{N_t}(L_{it} - L_{it}^{\text{wage}})}\right] = 3.461, \tag{S.6}$$

where L_{it}^{wage} is the predicted wage paid to the entrepreneur, as computed in the empirics. Recall that this aggregate ratio was how we computed the duration and interest-rate sensitivity of private business in Section 2.4, so this calibration is consistent with those moments as well. To estimate the deterministic life-cycle profile of business income, we fit a cubic polynomial in age to the moments $\ell_E(a_{it}) = \mathbb{E}\big[\log \frac{L_{it}}{\bar{L}_t} \mid a_{it}\big]$, which implies an estimate

$$\ell_E(a_{it}) = -5.749 + 0.318 \times a_{it} - 0.0526 \times \frac{a_{it}^2}{10} + 0.00250 \times \frac{a_{it}^3}{100}. \tag{S.7}$$

To estimate the variances of z_{it}, we use the fact that $\text{var}(z_{it}|a_{it}) = \text{var}\big(\log \frac{L_{it}}{\bar{L}_t} \mid a_{it}\big)$. Consequently, the time-series volatility can be computed from the change in variance of

business income over the life cycle,

$$\sigma^2_{Ez} = \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}\right) - \text{var}\left(\log \frac{L_{it}}{\bar{L}_t}\middle| a_{it}-1\right), \tag{S.8}$$

and the initial volatility can be computed as the level at the initial age, 20: $\sigma^2_{Ez0} = \text{var}\big(\log \frac{L_{it}}{\bar{L}_t} \mid a_{it} = 20\big)$. We estimate these by running a linear regression of the variances $\text{var}(\log(E_{it}/\bar{L}_t)|a_{it})$ on age a_{it}, and find a constant $\sigma^2_{Ez0} = 2.535$ and a slope $\sigma^2_{Ez} = 0.0563$. Finally, we compute the entry probability $p_E(a_{it})$ as the change in the share of households who are entrepreneurs from one age to the next.[3] Specifically, we first compute the share of households at each age a_{it} who are entrepreneurs (call this $s_E(a_{it})$), then fit a cubic polynomial to this share, and then take the change in this fitted line as the entry probability, from the fact that $p_E(a_{it}) = \frac{s_E(a_{it}) - s_E(a_{it}-1)}{1-s_E(a_{it}-1)}$.

3.5 Aggregate shocks in the overlapping-generations simulation

The overlapping-generations (OLG) model has 4 aggregate shocks that must be mapped to data: these are shocks to (1) riskfree rates ϵ_r, (2) wage growth ϵ_g, (3) stock returns ϵ_s, and (4) house prices ϵ_h. Shocks to the riskfree rate come from the procedure outlined in Section 3.1.1. This procedure also provides us with the level of real interest rates r_f. The series of wage income growth rates g are proxied by the rate of growth of per capita real GDP, just as in Section 3.1. We obtain g_1 using the relation $g_{1t} = g_t - \lambda_{gr} r_{ft}$.

Stock returns come from the S&P500 and are added directly to the simulation after subtracting out the portion driven by shocks to interest rates—which is akin to subtracting $\sigma_{n_s}\epsilon_r$ from the log stock return. This means that we add both ϵ_g and ϵ_s into the simulation simultaneously. Similarly, aggregate house prices come from the Case-Shiller index and are added after differencing out the portion of the capital gain attributable to interest rate shocks $\sigma_{n_h}\epsilon_r$. For the idiosyncratic portion of house price shocks, we assign each simulated individual a shock drawn from a normal distribution with variance $\tilde{\sigma}^2_h$.

When not using data in the simulation, we randomly assign agents aggregate shocks from 50 randomly generated time series. When computing top wealth shares, we do so within a group of shocks and then take the average top 10% share among the 50 groups.

[3] We do not model exit, so the model attributes all changes in the share of entrepreneurs to entry.

4 Solution method

This section details how we solve the model numerically. To render the model solvable, we exploit two key properties of the model that allow us to rescale the problem, eliminating state variables; and to transform the wealth return, simplifying the portfolio choice problem.

4.1 Rescaling the model

We eliminate two state variables, the wage index $\bar{L}_t$ and the house price P_{ht}, by rescaling the problem as follows. Henceforth, let $\hat{Z}_t$ denote wage-indexed quantities $Z_t/\bar{L}_t$ for any variable Z_t. The utility function equals $u(C_{it}, H_{it}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} u(\hat{C}_{it}, \hat{P}_{ht} H_{it})$. From this, it follows (and we will prove) that there exists a function $\hat{V}$ such that the value function equals[4]

$$\begin{aligned} V_{it} &= V(\bar{L}_t, P_{ht}, W_{it}, r_{ft}, g_{1t}, z_{it}, a_{it}, \text{AIYE}_{it}) \\ &= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}(\hat{W}_{it}, r_{ft}, g_{1t}, z_{it}, a_{it}, \hat{\text{AIYE}}_{it}), \end{aligned} \tag{S.9}$$

provided that the agent chooses the rescaled policies $\{\hat{C}_{it}, \hat{P}_{ht} H_{it}\}$ instead of $\{C_{it}, H_{it}\}$. Henceforth, we will use the notation $\hat{H}_{it} \equiv \hat{P}_{ht} H_{it}$ to emphasize that this is a single choice variable.

Before proving the claim, note that all constraints can be rescaled accordingly. Wealth follows

$$\hat{W}_{i,t+1} = \frac{\bar{L}_t}{\bar{L}_{t+1}} (\hat{W}_{it} + \tilde{L}_{it} + \hat{B}_{it} + \hat{I}_{it} - \hat{T}_{it} - \hat{C}_{it} - \chi P_{n_h t} \hat{H}_{it} \mathbb{I}_{\{\pi_{h,it}=0\}} - c_s) R_{Wi,t+1}, \tag{S.10}$$

which is independent of $\bar{L}_t$ and P_{ht}. The minimum and maximum house size constraints are $P_{n_h t} \hat{H}_{it} \geq \kappa_{\min}$ and $\pi_{h,it} \leq \kappa_{\max}$. The scaled value of a private business $\hat{E}_{it}$ is also independent of $\bar{L}_t$. Thus, the constraints can all be written in terms of controls $\{\hat{C}_{it}, \hat{H}_{it}\}$ and without any dependence on $\bar{L}_t$ and P_{ht}.

We can now prove the claim by backward induction. First, consider the bequest (24). In the terminal period (i.e., $\bar{b}$ periods after death), the claim is true, as

$$V_{it} = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it}, \hat{P}_{ht} H_{it}\}} u(\hat{C}_{it}, \hat{P}_{ht} H_{it}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{S.11}$$

[4]For simplicity, we will consider the case of a worker who never becomes an entrepreneur, and the same logic will apply when we add entrepreneurship.

Now suppose that it is true at any time $t+1$ within the time span after death. This implies

$$V_{it} = \max_{\{C_{it}, H_{it}, \pi_{it}\}} \left\{ u(C_{it}, H_{it}) + \beta \mathbb{E}_t \left[\bar{L}_{t+1}^{1-\gamma} P_{h,t+1}^{-\nu(1-\gamma)} \hat{V}_{i,t+1} \right] \right\} \tag{S.12}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it}, \hat{H}_{it}, \pi_{it}\}} \Bigg\{ u(\hat{C}_{it}, \hat{H}_{it}) + \beta \mathbb{E}_t \left[\left(\frac{\bar{L}_{t+1}}{\bar{L}_t} \right)^{1-\gamma} \left(\frac{P_{h,t+1}}{P_{ht}} \right)^{-\nu(1-\gamma)} \hat{V}_{i,t+1} \right] \Bigg\} \tag{S.13}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{S.14}$$

The second equality changes the choice variables to $\{\hat{C}_{it}, \hat{H}_{it}, \pi_{it}\}$, which is possible because the rescaled utility function and budget constraints only feature these choice variables. The third equality follows from the fact that the maximand has no dependence on $\bar{L}_t$ and P_{ht}. It follows that the initial bequest utility equals

$$b(\bar{L}_t, P_{ht}, W_{it}, r_{ft}, g_{1t}) = \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{b}(\hat{W}_{it}, r_{ft}, g_{1t}). \tag{S.15}$$

This proves the claim for the bequest.

The rest of the proof follows the same inductive reasoning. We know, upon death, the value function rescales as claimed (since the value function is the bequest utility). Now, for the induction step, suppose the value function conditional on living to $t+1$ rescales as claimed. Then

$$V_{it} = \max_{\{C_{it}, H_{it}, \pi_{it}\}} \left\{ u(C_{it}, H_{it}) + \beta \mathbb{E}_t \left[\bar{L}_{t+1}^{1-\gamma} P_{h,t+1}^{-\nu(1-\gamma)} \left((1-m_{it}) \hat{V}_{i,t+1} + m_{it} \hat{b}_{i,t+1} \right) \right] \right\} \tag{S.16}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \max_{\{\hat{C}_{it}, \hat{H}_{it}, \pi_{it}\}} \Bigg\{ u(\hat{C}_{it}, \hat{H}_{it}) \tag{S.17}$$

$$+ \beta \mathbb{E}_t \left[\left(\frac{\bar{L}_{t+1}}{\bar{L}_t} \right)^{1-\gamma} \left(\frac{P_{h,t+1}}{P_{ht}} \right)^{-\nu(1-\gamma)} \left((1-m_{it}) \hat{V}_{i,t+1} + m_{it} \hat{b}_{i,t+1} \right) \right] \Bigg\} \tag{S.18}$$

$$= \bar{L}_t^{1-\gamma} P_{ht}^{-\nu(1-\gamma)} \hat{V}_{it}. \tag{S.19}$$

The equalities follow identical reasoning as above. Most importantly, $\hat{V}_{it}$ in the last equa-

tion is independent of $\bar{L}_t$ and P_{ht}, proving the conjecture.

4.2 Transforming the portfolio problem

In choosing a portfolio, the agent must select optimal weights $\{\pi_{n,it}, \pi_{s,it}, \pi_{h,it}\}$ on the risky assets. Numerically, this problem becomes intractable because all three assets are exposed to interest-rate shocks. We circumvent this tractability problem by re-expressing the problem in terms of exposures to separate asset-market shocks: the agent instead chooses $\{\varepsilon_{r,it}, \pi_{s,it}, \pi_{h,it}\}$, where $\varepsilon_{r,it}$ is the interest-rate sensitivity of the total portfolio. This transformation is also ideal because the long-term bond share π_n is itself not that meaningful, as it arbitrarily depends on our choice of the bond duration n; in contrast, ε_r is invariant to n.

To implement this change of choice variables, we need to rewrite the wealth return (29) in a way that replaces the long-term bond share $\pi_{n,it}$ with the interest-rate sensitivity of the total portfolio $\varepsilon_{r,it}$. To do this, start from the result of Appendix G.5 that, to a second order, the log wealth return (suppressing i) equals

$$r_{w,t+1} = r_{ft} + \pi_t^\top (r_{t+1} - r_{ft}\iota) + \frac{1}{2}\left(\pi_t^\top \text{diag}(\Sigma\Sigma^\top) - \pi_t^\top \Sigma\Sigma^\top \pi_t\right) - (\pi_{h,t} + \pi_{s,t} - 1)^+\theta, \tag{S.20}$$

where π_t is the vector of portfolio weights, r_{t+1} is the vector of risky returns, ι is a vector of ones, $\Sigma\Sigma^\top$ is the covariance matrix of returns with respect to the shock vector, and θ is shorthand for the mortgage premium.

To transform the wealth return, first note that we can rewrite the log return on any asset $j \in \{n, s, h\}$ as

$$r_{j,t+1} = \log R_{j,t+1} = r_{ft} + \mu_j - \sigma_{jr}\epsilon_{r,t+1} + \sigma_{j,-r}^\top \epsilon_{-r,t+1}, \tag{S.21}$$

where ϵ_{-r} is defined as the (standardized) vector of shocks orthogonal to ϵ_r, $\epsilon_{-r,t+1} \equiv [\epsilon_{g,t+1}, \epsilon_{s,t+1}, \epsilon_{h,t+1}, \tilde{\epsilon}_{h,t+1}]^\top$, and $\sigma_{j,-r}$ the corresponding vector of loadings. In this notation, the loadings on the rate shock are $\{\sigma_{nr}, \sigma_{sr}, \sigma_{hr}\} = \{\sigma_n, \lambda_{sr}, \sigma_{n_h}\}$. Stacking these returns in vector form, we have $r_{t+1} = r_{ft}\iota + \mu - \Sigma_r \epsilon_{r,t+1} + \Sigma_{-r}\epsilon_{-r,t+1}$, meaning that the wealth return is

$$\begin{aligned} r_{w,t+1} = r_{ft} + \pi_t^\top (\mu - \Sigma_r \epsilon_{r,t+1} + \Sigma_{-r}\epsilon_{-r,t+1}) + \frac{1}{2}\big(\pi_t^\top \text{diag}(\Sigma_r\Sigma_r^\top) + \pi_t^\top \text{diag}(\Sigma_{-r}\Sigma_{-r}^\top) \\ - \pi_t^\top \Sigma_r \Sigma_r^\top \pi_t - \pi_t^\top \Sigma_{-r}\Sigma_{-r}^\top \pi_t\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+\theta. \end{aligned} \tag{S.22}$$

Next, let us substitute out $\pi_{n,t}$ using the fact that the interest-rate sensitivity of the portfolio equals $\varepsilon_{r,t} = \pi_t^\top \Sigma_r$. Substituting this into the wealth return, we have

$$r_{w,t+1} = r_{ft} + \pi_t^\top \mu - \varepsilon_{r,t}\epsilon_{r,t+1} + \pi_t^\top \Sigma_{-r}\epsilon_{-r,t+1} + \frac{1}{2}\big(\pi_t^\top \text{diag}(\Sigma_r \Sigma_r^\top) + \pi_t^\top \text{diag}(\Sigma_{-r}\Sigma_{-r}^\top) - \varepsilon_{r,t}^2 - \pi_t^\top \Sigma_{-r}\Sigma_{-r}^\top \pi_t\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+ \theta. \quad \text{(S.23)}$$

Because the first row of Σ_{-r} is zero (i.e., $r_{n,t+1}$ has no loadings on $\epsilon_{-r,t+1}$), and because the first element of $\mu + \text{diag}(\Sigma_r \Sigma_r^\top)/2$ equals zero (i.e., $\mu_n = -\sigma_n^2/2$), this rewritten wealth return has no dependence on $\pi_{n,t}$. We have completely replaced it with $\varepsilon_{r,t}$. Indeed, (S.23) can be written as[5]

$$r_{w,t+1} = r_{ft} + \pi_t^{*\top}(\mu_t^* - \Sigma_t^* \epsilon_{t+1}) + \frac{1}{2}\big(\pi_t^{*\top} \text{diag}(\Sigma_t^* \Sigma_t^{*\top}) - \pi_t^{*\top} \Sigma_t^* \Sigma_t^{*\top} \pi_t^*\big) - (\pi_{h,t} + \pi_{s,t} - 1)^+ \theta, \quad \text{(S.24)}$$

$$\pi_t^* \equiv \begin{bmatrix} 1 \\ \pi_s \\ \pi_h \end{bmatrix}, \quad \mu_t^* \equiv \begin{bmatrix} -\varepsilon_{r,t}^2/2 \\ \mu_s + \lambda_{sr}^2/2 \\ \mu_h \end{bmatrix}, \quad \text{and} \quad \Sigma_t^* \equiv \begin{bmatrix} -\varepsilon_{r,t} & 0 & 0 & 0 & 0 \\ 0 & \lambda_{sg} & 1 & 0 & 0 \\ 0 & 0 & 0 & \sigma_h & \tilde{\sigma}_h \end{bmatrix}. \quad \text{(S.25)}$$

This transformed expression then implies an equivalent expression for the wealth return:

$$R_{Wi,t+1} = R_{ft} + (e^{-\varepsilon_{r,it}^2/2 - \varepsilon_{r,it}\epsilon_{r,t+1}} - 1)R_{ft} + \pi_{s,it}(e^{r_{ft} + \mu_s + \lambda_{sr}^2/2 + \lambda_{sg}\epsilon_{g,t+1} + \epsilon_{s,t+1}} - R_{ft}) + \pi_{h,it}(e^{r_{ft} + \mu_h + \sigma_h \epsilon_{h,t+1} + \tilde{\sigma}_h \tilde{\epsilon}_{hi,t+1}} - R_{ft}) - (\pi_{h,it} + \pi_{s,it} - 1)^+ (R_{Mit} - R_{ft}) \quad \text{(S.26)}$$

(To see this, simply apply the same second-order approximation to (S.26) as we did to get (S.20), and the result will be (S.23)). Thus, the optimal allocation can equivalently be stated as a choice over $\{\varepsilon_{r,t}, \pi_{s,t}, \pi_{h,t}\}$, and these choices map to an optimal long-term bond allocation equal to $\pi_{n,t} = \frac{\varepsilon_{r,t} - \pi_{s,t}\lambda_{sr} - \pi_{h,t}\sigma_{n_h}}{\sigma_n}$. This allocation is unique given the maturity n of the bond.

References

Beeler, Jason and John Y. Campbell, "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment," *Critical Finance Review*, 2012, *1* (1).

Catherine, Sylvain, "Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle," *The Review of Financial Studies*, 2022, *35* (9), 4016–4054.

[5] The third element of μ_t^* is $\mu_h + \mu_{n_h} + \sigma_{n_h}^2/2$ and we have substituted in $\mu_{n_h} = -\sigma_{n_h}^2/2$.

Catherine, Sylvain and Constantine Yannelis, "The Distributional Effects of Student Loan Forgiveness," *Journal of Financial Economics*, 2023, *147* (2), 297–316.

Catherine, Sylvain, Max Miller, and Natasha Sarin, "Social Security and Trends in Wealth Inequality," *The Journal of Finance*, 2025, *80* (3), 1497–1531.

Catherine, Sylvain, Paolo Sodini, and Yapei Zhang, "Countercyclical Income Risk and Portfolio Choices: Evidence from Sweden," *Journal of Finance*, 2024, *79* (3), 1755–1788.

Fama, Eugene F., "Stock Returns, Expected Returns, and Real Activity," *The Journal of Finance*, 1990, *45* (4), 1089–1108.

Greenwald, Daniel, Matteo Leombroni, Hanno N. Lustig, and Stijn van Nieuwerburgh, "Financial and Total Wealth Inequality with Declining Interest Rates," 2023. Stanford University Graduate School of Business Research Paper.

Guvenen, Fatih, Serdar Ozkan, and Jae Song, "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 2014, *122* (3), 621 – 660.

Saez, Emmanuel and Gabriel Zucman, "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data," *The Quarterly Journal of Economics*, 2016, *131* (2), 519–578.

Schwert, G. William, "Stock Returns and Real Activity: A Century of Evidence," *The Journal of Finance*, 1990, *45* (4), 1237–1257.