Rising Income Risk at the Top*

J. Carter Braxton Kyle Herkenhoff Chengdai Huang
Michael Nattinger Jonathan Rothbaum Lawrence D. W. Schmidt

March 25, 2026

Abstract

We document a broad increase in U.S. income risk from 1969 to 2019 using newly digitized IRS tax returns. The aggregate consequences of rising income risk depend on its permanence, where in the income distribution it occurs, and how it affects savings behavior. We find that since the 1970s, permanent income risk increased across the income distribution, particularly among top earners. While permanent risk has always been high for top earners, it rose by nearly 70 percent among the top 5%. We further show that large negative income shocks strongly predict financial distress among top earners, including early 401(k) withdrawals and moves to lower-income neighborhoods, and that top earners have increased their savings in response to rising risk. Quantitatively, we show that rising income risk among top earners lowered the risk-free rate, contributed to the rise in wealth inequality and the "savings glut of the rich."

*Braxton: University of Wisconsin and IZA. Herkenhoff: University of Minnesota, FRB Minneapolis, IZA, and NBER. Huang: NYU. Nattinger: University of Wisconsin. Rothbuam: US Census Bureau. Schmidt: MIT. We thank Salome Baslandze (discussant), Job Boerma, Rishabh Kirplani, Maarten Meeuwis (discussant), Ludwig Straub, Nicolas Werquin (discussant) as well as numerous seminar and conference audiences for helpful comments. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the Federal Reserve System or the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product (Data Management System (DMS) number: P-7503840 (P-7529415), Disclosure Review Board (DRB) approval numbers: CBDRB-FY24-SEHSD003-055, CBDRB-FY24-0300 and CBDRB-FY26-SEHSD003-025).

Recently, there has been substantial debate about the evolution of income risk in the United States, with some studies arguing that income risk likely declined in administrative data on individual earnings (e.g., Guvenen et al. (2022), Salgado et al. (2023)), while others have found stable or rising income risk using alternate methods and data sources (e.g., Moffitt et al. (2022), Ziliak et al. (2022)). In this paper, we revisit this question using a newly available panel of household income from the universe of U.S. 1040 tax returns spanning 1969 to the present. These data allow us to measure income at the household level, including business income, and to track its evolution across the entire income distribution over a long horizon. We use these data to address two questions: *How has income risk evolved over time and across the income distribution?* and *What are the aggregate consequences of these changes?* We find that income risk has risen for all households and especially high-income households since 1969, both in near-raw data and in our estimated income process. We then study the macroeconomic implications of these changes by nesting our income process in a quantitative Bewley-Huggett-Aiyagari model. We show that rising income risk at the top of the distribution increases savings among high-income households, lowers the natural rate of interest, and contributes to the expansion of wealth inequality.

In this paper, we make four contributions. First, empirically we document the evolution of income risk across the income distribution and over time. To motivate this exercise, we start by computing the variance of income changes over 5 and 10-year horizons in the (near) raw tax data. We document three main facts. First, the variance of 5-year income changes has been largely flat since 1969, while the variance of 10-year income changes has increased substantially. Because longer-horizon income changes primarily reflect shocks to permanent income, these patterns suggest that permanent income risk has increased over the past 50 years, while transitory risk has remained flat or declined. Second, across the income distribution, the variance of both 5- and 10-year income changes is asymmetrically U-shaped: risk declines from the bottom through roughly the 80th percentile and then rises sharply at the top of the distribution.[1] Third, we move beyond existing studies and show that these increases in longer-horizon income risk are concentrated among top earners, with the variance of 10-year income changes increasing by nearly 15% at the 95th percentile of the prior income distribution since the 1970s and nearly 40% for the top 1 percent.

Motivated by these results, we then write down a permanent-transitory income process where the variance of shocks to permanent and transitory income is allowed to vary by decade as well as by lagged income rank. We estimate our income process using the approach of Braxton et al. (2025) which allows for estimating a rich income process that makes the parameters a

[1]These cross-sectional patterns align with the variance of 1-year and 5-year income differences in Guvenen et al. (2014) and Guvenen et al. (2021), which are estimated on individual W2 earnings.

function of observables (in our context, lagged income rank and time) and recovers estimates of permanent and transitory income for every person in every period via the Kalman filter. Crucially, the fact that an individual's risk profile depends on lagged income significantly fattens the right tail of the income distribution and adds higher-order income risk. Our parameter estimates reveal two main results. First, that permanent income risk is asymmetrically U-shaped with permanent income risk being the highest among top earners. Second, we show that permanent income risk has increased substantially at the top of the income distribution, for the top 5% of households it has increased, on average, by 70% and for the top 1% it has effectively doubled since the 1970s. Our baseline results utilize adjusted gross income less interest and dividend income, and we show that our results are robust to many alternative measures, including considering only wage and salary income. We also document that these patterns hold within both self-employed and non-self-employed individuals.

Our second contribution is to validate that households at the top of the income distribution do in fact face risk. For this exercise, we show that our recovered permanent income shocks among high-income households are correlated with observable signals of economic distress. In particular, we find that in response to a 1-SD decline in permanent income, households are more likely to: (1) have an early 401(k) withdrawal and (2) move to a ZIP code with a 30% lower median income. Importantly for each of these results, we show that these patterns of responding to a permanent income shock hold within the top 1% of the income distribution.[2] Thus, we view these results as validating that high-income households do in fact face a substantial amount of income risk.

Our third contribution is to empirically measure how these changes in risk impact savings behavior. We build on Smith et al. (2023) and capitalize interest and dividend income to build proxies for wealth and savings rates. We then measure how savings-to-income ratios respond to within-income rank and across-occupation variation in income risk.[3] Conditional on income rank, we find that savings-to-income ratios rise by 4.7% of prior income for a 0.1 increase in the variance of permanent income shocks (i.e., permanent income risk). We find much smaller savings responses to increases in transitory income risk. Further, we show that the increase in savings in response to permanent risk is concentrated among younger and high-income house-

[2] In a companion paper (Braxton et al. (2025)), we show large declines in permanent income are associated with large increases in the likelihood of filing for bankruptcy, entering foreclosure, and recording a debt charge-off, even among individuals in the top 2% of the income distribution.

[3] Our empirical approach of exploiting occupation level variation in permanent risk to elicit a savings response follows a long tradition in the precautionary savings literature (e.g., Carroll and Samwick (1997) and Boar (2021), and references therein). Relative to these papers, we compare differences across occupations but within income rank. We additionally examine the heterogeneity in the savings response to income risk across the income and age distribution. Our finding that younger households have a larger savings response to permanent income risk aligns with the "buffer stock" theory of savings (e.g., Carroll (1997)).

holds. We refer to these moments as "savings elasticities" and use them to discipline our quantitative model.

Finally, our fourth contribution is to quantitatively evaluate the aggregate effects of our measured changes in income risk over the income distribution and time. To do so, we embed our income process into a finite life-cycle Bewley-Huggett-Aiyagari style model, where households face idiosyncratic income shocks that are a function of their lagged income rank and can self-insure via saving and borrowing in a risk-free asset. We calibrate the quantitative model to align with our empirical savings elasticities. Further, we show that our quantitative model is consistent with several non-targeted moments including the near-zero elasticity of savings with respect to transitory income risk and the concavity of consumption with respect to permanent income from Straub (2019). The intuition for why the model easily replicates the findings in Straub (2019) is that higher earners face more risk and thus save more.

Using our quantitative model as a laboratory, we find that changes in income risk since the 1970s raise the top 5% wealth share by more than 2.4 percentage points, with over half of this increase attributable to the rise in permanent income risk. Increased saving by top earners also lowers the risk-free rate from 3.5% to 2.9%, and the resulting decline in interest rates further magnifies wealth inequality. The economic logic is straightforward. In partial equilibrium, an increase in income risk raises precautionary saving demand throughout the income distribution. In general equilibrium, however, the additional demand for assets pushes down the risk-free rate, dampening saving incentives. For most households, especially those below the 90th percentile of the earnings distribution, this decline in the return to saving more than offsets the increase in precautionary motives, causing saving rates to fall. Among top earners, in contrast, the rise in income risk is sufficiently large that precautionary motives dominate the equilibrium decline in interest rates, leading their saving rates to increase sharply. The model therefore generates a central feature of the data emphasized by Mian et al. (2020): a "savings glut of the rich," in which saving rises at the top even as it falls for the rest of the distribution.

To isolate the role of changes in income risk among high-income households, we only allow the income process parameters for the top 10% of the distribution to change between the 1970s and 2010s. For all other percentiles of the income distribution, income risk is held fixed at 1970s levels. We find that roughly 40% of our model's predicted reduction in interest rates can be attributed to changing income risk at the top, while nearly all of the change in the top 5% wealth share is attributable to changing income risk at the top. These results suggest that the income risk of top earners is an important cause of lower interest rates in the U.S. as well as rising wealth inequality.

Literature. This paper relates to several recent literatures. First, this paper contributes to the literature that has examined the evolution of income risk in the U.S. (e.g., Moffitt et al. (2022), Salgado et al. (2023), Braxton et al. (2025) and references therein) as well as the literature that has extended income processes to allow for heterogeneity in the shocks that individuals face (e.g., Arellano et al. (2017), De Nardi et al. (2020), Guvenen et al. (2021) and references therein). Relative to these papers, we examine how income risk has evolved both *over* time as well as *across* the income distribution. Further, we provide novel evidence that individuals at the top of the income distribution do in fact face risk by examining how high-income households respond to large declines in permanent income via changes in asset positions and default.

Our paper also contributes to the literature that has examined declining interest rates in the U.S. over the past 40 years. In this literature several different explanations have arisen including: (1) increasing income inequality (e.g., Favilukis (2013), Kaymak and Poschke (2016), Auclert and Rognlie (2018), Aladangady et al. (2021), Mian et al. (2021a), and Mian et al. (2021b) among others), and (2) changing financial and demographic factors (e.g., Caballero et al. (2008), Favilukis (2013), Caballero et al. (2017), Farhi and Gourio (2018) , Auclert et al. (2021)). Relative to these papers, we consider the role of changes in income risk over time and how changes in income risk have been especially pronounced among high-income households. Further, we use an "identified moment" in the tradition of Nakamura and Steinsson (2018)) to isolate the impact of exposure to greater permanent income risk and savings behavior.

Additionally, our paper contributes to the literature that has studied wealth inequality. Work by Saez and Zucman (2016) and Smith et al. (2023) has aimed to measure the degree of wealth inequality in the U.S. and how it has evolved over time, while Hubmer et al. (2021) examines alternative theories for increasing wealth inequality.[4] Relative to these papers, we contribute to the literature by proposing increasing income risk for high-income households as a proposed explanation for increases in wealth inequality. Further, we discipline the strength of this channel by leveraging the insights of Saez and Zucman (2016) and Smith et al. (2023) to capitalize interest and dividend income to obtain proxies for savings.

Finally, our paper contributes to the literature that has examined how savings behavior varies across the income distribution (e.g., Caroll (2000), Dynan et al. (2004), Straub (2019), Mian et al. (2021b) and references therein).[5] A common theme of these papers is that high-income households save more than lower income households, and recently Mian et al. (2020)

[4] A related literature has examined the modifications and channels needed to get quantitative Bewley models to be consistent with the data on top wealth inequality. Recent papers in this literature include De Nardi and Fella (2017) and Gaillard et al. (2023). We show that accounting for high-income households facing a large degree of permanent income risk can help in matching top wealth inequality.

[5] Additionally, there is a long tradition of examining how savings and wealth behavior responds to changes in risk in quantitative models, e.g., Carroll (1997) and Gourinchas and Parker (2002).

have shown that savings rates have increased substantially among the top 1% of the income distribution in the U.S. since the 1980s. Our paper provides a channel for why high-income households have traditionally saved more and why this has increased over the past 50 years, namely because they face higher and growing income risk.

The paper proceeds as follows. We begin by discussing our panel of household income from 1040 tax returns and use the data to provide a series of motivating facts in Section 1. We detail the econometric income process in Section 2. We then present our estimation results in Section 3. We use those estimates to discipline our quantitative model in Section 4 and measure the effects of rising risk on interest rates and wealth inequality in Section 5. Lastly, we conclude in Section 6.

1 Data and Motivation

Before delving income process estimation and quantitative modeling, we provide an overview of our dataset and motivating facts for our analysis.

1.1 Data

Our analysis is based on new linkages between tax data and surveys. Our primary dataset is a panel of 1040 tax returns that starts in 1969 and is available every five years through 2019, i.e., we are able to observe an individual in our panel in years that end with 4 or 9 (1969, 1974, ..., 2014, 2019). In the 1040s, we are able to observe whether an individual has a spouse (and after 1994 if they have dependents), the household level of adjusted gross income (AGI), wage and salary income (WSI) as well as interest and dividend income.[6]

Our baseline income measure is adjusted gross income (AGI) net of interest and dividend income. This definition of income includes wage and salary income as well as business income.[7] We show in Appendix C.2 that our results are robust to using alternative definitions of income such as adjusted gross income (AGI) and wage and salary income (WSI). Central to this paper is the ability to observe interest and dividend income, which we can use to create a capitalized

[6] See Alexander, Bleckley, Fisher, Genadek, Leonard, and Magganas (2024) for information on the construction of the 1040 panel.

[7] Data on interest income is available starting in 1969 in our sample, while data on dividend income is available starting in 1974. We infer the level of dividend income in 1969 by taking the ratio of dividend income to AGI in 1974 at the household level and projecting backwards to 1969 using the 1969 level of AGI. Our results are robust to these choices. Additionally, our definition of income also includes capital gains. We are able to remove capital gains from our measure of income starting in 1994.

measure of wealth. By tracking how capitalized wealth evolves over time, we can examine how households respond to changes in income risk.[8]

While we have access to the full universe of 1040s, our baseline sample comprises the 1040s for individuals who are in the Annual Social and Economic Supplement (ASEC) to the Current Population Survey (CPS).[9] The ASEC asks questions that provide additional labor market information that is not readily available on tax forms (e.g., occupation, education, self-employment status) that will allow us to further understand the drivers of income risk across the distribution as well as how households respond to changes in risk. We link individuals from the ASEC to their income information in the 1040s panel using scrambled Social Security numbers, called Protected Identification Keys (PIKs). We show in Appendix C.3 that our results for the trends in the evolution of income risk across the income distribution are robust to using a random sample of 1040s.

Our measures of income (e.g., AGI, interest and dividend income, etc.) are reported at the household level but the data is constructed so that we have a panel of individuals. To avoid double counting incomes for households with multiple adults, we give each adult a weight of one-half in a two-adult household. Note that since our baseline sample is comprised of individuals in the ASEC, we additionally use sampling weights from the ASEC. As we discuss in more detail below, we remove changes in household composition as part of our residualization process. Thus, we think of our sample construction as sampling individuals to measure household outcomes.

To study income dynamics, we focus on a sample of individuals with a minimum degree of labor force attachment. To be included in our estimation sample, an individual must: (1) satisfy a minimum income requirement in at least 2 years, (2) satisfy the minimum income criterion in at least 50% of observations (inclusive) between the first and last year that they satisfy the minimum income criterion, and (3) be between the ages of 25 and 60. For conditions (1) and (2), we impose a minimum income criterion equal to $10,000 (in 2005 PCE dollars).[10] These sampling criteria result in a sample of 2.06 million individuals. When a household's income moves below our minimum cutoff, or is missing, we classify their income observation

[8] To capitalize interest and dividend income into a measure of wealth we follow Smith, Zidar, and Zwick (2023).

[9] In particular, we are able to link individuals to their ASEC responses if they were included in the ASEC during the years 1973, 1979, 1981-1991, 1994, and 1996-2020.

[10] In results that are available upon request, we show that our results are robust to using alternative values of the minimum income cutoff. Part of the motivation for using a higher cutoff for 1040 income compared to W2 income (as in Braxton et al. (2025)) is that households do not have to file taxes if their income is below a given cutoff. In recent years these cutoffs have ranged between $12,950 for a single person under 65 to $25,900 for a married household filing jointly. See https://www.irs.gov/newsroom/who-needs-to-file-a-tax-return for additional details. *Link last accessed on May 5, 2024.*

Table 1: Summary statistics

Variable	Mean
Real Income	\$67,210
Age	41.1
Share Non-employed/Non-filing	6.1%
Share Male	50.1%
Number of adults in household	1.6
Observations	10,440,000
Individuals	2,060,000

Note: See Section 1.1 for sample selection criteria. Real income is measured in 2005 (PCE) dollars, and corresponds to adjusted gross income (AGI) less interest and dividend income. The variable "share non-employed/non-filing" is the share of individuals whose real income in a given year does not satisfy the minimum income criterion or is missing.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

as missing and are able to retain them in the sample.[11]

Finally, we discuss how we remove the predictable component of household income, i.e., residualizing. We residualize income by remove age and year fixed effects as well as dummy variables for the number of adults in the household as in De Nardi et al. (2020).[12] Additionally, we exclude negative incomes from the residualization stage and treat those observations as part of the non-employed/non-filer category in the filter. Table 1 provides summary statistics for the households in our sample.

1.2 Motivating Facts: Rising Risk

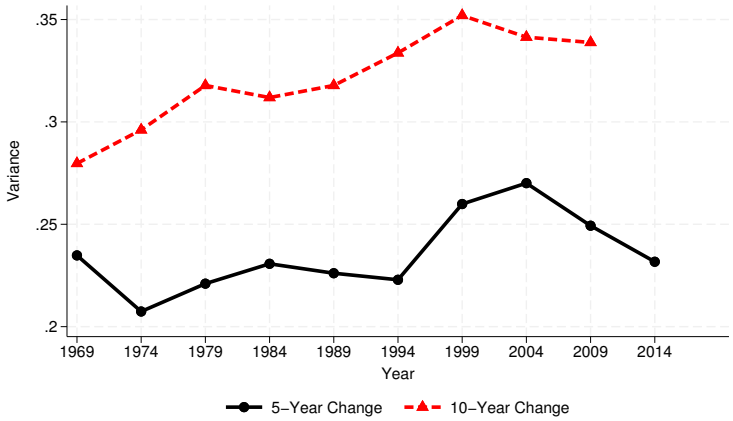
To motivate our analysis, we present a set of simple data moments designed to proxy permanent and transitory income risk. These "simple" moments provide suggestive evidence that (1) permanent income risk has increased in the U.S since the 1970s, (2) the variance of permanent income changes is U-shaped in prior income, and (3) permanent income risk has increased among high-income households.

We begin by summarizing changes in income volatility since the 1970s. In Figure 1, we plot the variance of (residual) log income changes over a 5-year (black, solid line with circle markers)

[11] Note that a household's income can move below the minimum income cutoff due to spells of non-employment, choosing to not file their taxes if their income is below the cutoff, or if due to large losses in business income their AGI is negative.

[12] We show in Appendix C.4.1 that we find similar results if we interact the number of adult fixed effects with time fixed effects.

Figure 1: Income Risk in the U.S. Since the 1970s



Note: Figure presents the variance of residual log income changes over a 5-year horizon (black, solid line with circle markers) and a 10-year horizon (red, dashed line with square markers).
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

and 10-year horizon (red, dashed line with triangle markers). The figure shows that since 1969 there has been a steady increase in the variance of income changes at a 10-year horizon, while there have been increases and decreases in the variance of 5-year changes. As discussed in Braxton et al. (2025) and references there-in, longer horizon changes in income risk are more reflective of permanent income risk as opposed to shorter horizon changes where transitory risk plays a larger role. Therefore, the steady increase in the variance of the 10-year income change likely reflects an increase in permanent income risk in the U.S. over this time period.

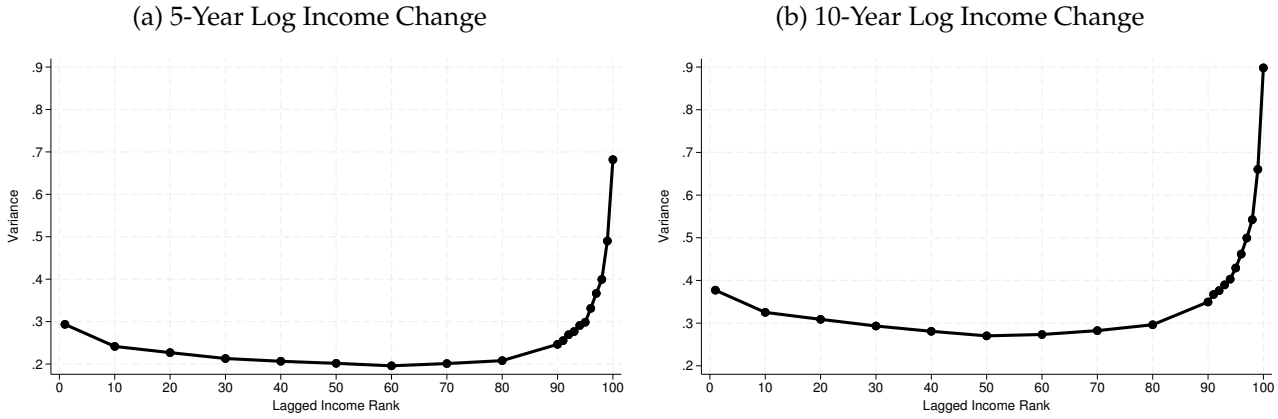
Second, we provide evidence that income risk is U-shaped in prior income. In Figure 2 we present the variance of 5-year income changes (Panel (a)) and 10-year income changes (Panel (b)) as a function of lagged income rank.[13] Both of these figures show that the variance of income changes decreases as a function of lagged income for low income households, then is relatively flat for much of the middle of the income distribution, before it increases substantially in the top of the income distribution.[14]

Lastly, we show how these simple measures of income risk have evolved over time as well as across the income distribution. In Panel (a) of Figure 3, we present the variance of 5-year

[13] As we will discuss in greater detail in Section 2.1 we rank individuals using their lagged residual log income, and we rank individuals within year and age.

[14] Recent papers showing that the variance of income changes is U-shapped include Guvenen et al. (2014) and Guvenen et al. (2021). Relative to these papers we will decompose these income changes into their permanent and transitory components and examine time trends in the evolution of income risk across the distribution. Additionally, we use household level measures of income and incorporate business income into our estimates.

Figure 2: Variance of Log Income Changes by Lagged Rank



Note: Figure presents the variance of residual log income changes over a 5-year horizon (Panel (a)) and a 10-year horizon (Panel (b)) as a function of lagged income rank.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

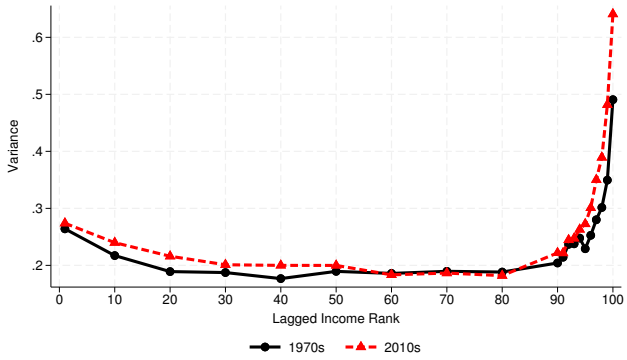
earnings changes as a function of lagged income rank for the 1970s (black, solid line with circle markers) and the 2010s (red, dashed line with triangle markers). The figure shows that at the bottom of the income distribution there have been modest increases in the variance of income shocks, and effectively no change in the upper middle of the income distribution. Conversely, as we move to the top of the income distribution, we see that the variance of income changes increases substantially. For households at the 95th percentile it has increased by nearly 20%, while for the top 1% it has increased by more than 30%. In panel (b) we show that we obtain similar results using the 10-year variance.[15]

While Figures 1, 2 and 3 provide suggestive evidence that rising income risk may be more prevalent among the highest earners, the macro implications of these trends crucially depends on (1) the relative permanence of the income shocks, and (2) whether our econometric classification of shocks truly reflect risk (i.e., were the changes unanticipated). To address these two points, we next write down a tractable but rich permanent/transitory income process, and then we estimate it by filtering (e.g., Braxton et al. (2025)). This accomplishes two goals. First, it lets us distinguish the permanent and transitory components of rising income volatility. Second, our estimation procedure yields a panel of filtered income shocks for each individual at each point in time, which we use to validate that rich households do indeed behave as if our recovered income shocks are unanticipated.

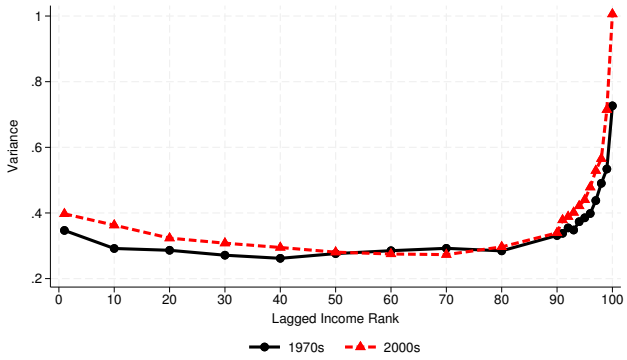
[15]Note that we cannot compute the 10-year variance using the 2010s as the base decade since our panel ends in 2019. Hence, we plot the 1970s vs. 2000s when analyzing 10-year income changes.

Figure 3: Variance of Log Income Changes by Lagged Rank and Decade

(a) 5-Year Log Income Change



(b) 10-Year Log Income Change



Note: Panel (a) presents the variance of residual log income changes over a 5-year horizon in the 1970s (black, solid line with circle markers) and 2010s (red, dashed line with triangle markers) as a function of lagged income rank. Panel (b) presents the variance of residual log income changes over a 10-year horizon in the 1970s (black, solid line with circle markers) and 2000s (red, dashed line with triangle markers) as a function of lagged income rank.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

2 Income process and estimation

We first describe the income process, which allows for shocks to permanent and transitory income to vary across the income distribution as well as over time. We then detail they way we estimate the process using the Kalman filter and an Expectations Maximization (EM) algorithm.

2.1 A tractable income process with flexible dependence on prior income

We begin with a panel of income, $Y_{i,t}$, where $i \in \{1, \ldots, N\}$ indexes households and $t \in \{1, \ldots, T\}$ indexes years. For ease of presentation, when we exposit our income process, we think of each one-unit interval in t, i.e., from $t-1$ to t as representing five years.[16] Since we are interested in income risk, we residualize log income by removing the predictable component (e.g., age and time effects) of income.[17] We denote the resulting residual log income as $y_{i,t}$.

Following our earlier work (Braxton et al. (2025)), our income process explicitly includes

[16] In Appendix B, we discuss how we translate an income process with a 1-year time step into a 5-year time step by exploiting the random walk nature of our income process, and provide evidence from the PSID that the frequency of observations does not material impact the patterns we find for permanent income risk across the income distribution. We also group time dummies such that the 1970 time dummy corresponds to 1969, 1974 and 1979 (we group the initial 3 observations), the 1980 time dummy corresponds to 1984 and 1989, the 1990 time dummy corresponds to 1994 and 1999, and so on.

[17] We discuss how we remove the predictable components of income in Section 1.1.

years with zero (and negative) income as well as missing income observations, which can occur due to non-filing of taxes. Let $l_{i,t} = [l_{E,i,t} \; l_{N,i,t}]'$ denote an household's labor market status. Element $l_{E,i,t}$ equals one when household i is employed in year t and zero otherwise. Likewise, $l_{N,i,t}$ equals one when household i is non-employed/non-filing and zero otherwise. A household is employed when they earn more than our minimum income criterion $\bar{y}$ (i.e., $Y_{i,t} > \bar{y}$) and classified as non-employed/non-filer otherwise. We observe – and condition on – employment status $l_{i,t}$ when estimating the parameters of our income process. Under assumptions often satisfied by search models, the estimation procedure is independent of the underlying process for $l_{i,t}$ (see Section 1 and Appendix A in Braxton et al. (2025) for a detailed discussion of this point). We then specify and parameterize the process for $l_{i,t}$ in Section 4, which will be governed by a separate set of parameters than those for the income process.[18]

The main innovation of this paper is to allow the shocks to permanent and transitory income to depend on where a household is in the income distribution as well as time. While our method is sufficiently tractable to characterize shocks by rich histories of prior income, our baseline approach makes income shocks a function of a household's lagged income rank, which we denote $r \in [0, 100]$. We determine an household's lagged income rank r by computing an household's percentile of lagged residual income (y_{-1}) within their age cohort a, and a given year t, i.e., $r = \text{Rank}(y_{-1}, t, a)$.[19]

For employed households, we model the process for observed income $y_{i,t}$ as the sum of permanent income $z_{i,t}$ and transitory income $\omega_{i,t}$. When a household is employed they draw their temporary income from a normal distribution with a variance $R(\cdot)$ that depends on their lagged income rank r, decade d as well as age a.[20] For households that are non-employed/non-

[18]We assume that labor market status $l_{i,t}$ is revealed before shocks to temporary and permanent income are revealed. We further assume that $l_{i,t} \perp z_i^{t-1} | (y_i^{t-1}, l_i^{t-1})$, which we refer to as augmented sequential exogeneity (ASQ) in Braxton et al. (2025). Conceptually, this assumption is that contemporary employment status is independent of lagged permanent income conditional on the history of income and employment status. We show in Braxton et al. (2025) that under ASQ the Kalman filter and smoother return the same paths of permanent income as under the assumption that employment status ($l_{i,t}$) is strictly exogenous. With the ASQ assumption, we also establish that the parameters for the employment process are separate from the parameters for the income process, and therefore can be estimated separately.

[19]For a household's first observation in the sample, we give them a lagged income rank of zero. For households with income below the minimum income cutoff, we use their value of residual log income from when they were last over the minimum income cutoff. Additionally, when an household is transitioning from being a non-employed/non-filer (N) to being employed (E), we give them a rank of zero in that period. In results that are available upon request, we show that our results are robust to the nature in which we classify lagged income ranks for new entrants and households exiting unemployment.

[20]As we discuss in more detail below, the assumption that shocks are normally distributed will buy us substantial tractability by allowing us to use an expectation maximization (EM) algorithm to estimate the parameters of our income process. However, despite the assumption of normally distributed shocks, our model still generates non-Gaussian features consistent with prior work (e.g., Guvenen et al. (2021) and references there-in). Note that shocks are drawn from a normal distribution, conditional on their rank, age and decade. Thus, in the cross-

filing, we set their residual log income to missing. We model income $y_{i,t}$ as evolving according to,

$$y_{i,t} = \begin{cases} z_{i,t} + \omega_{i,t}, & \omega_{i,t} \sim \mathcal{N}\big(0, R(r,d,a)\big) \quad \text{if } l_{E,i,t} = 1 \\ . & \text{if } l_{E,i,t} = 0 \end{cases} \tag{1}$$

As in Blundell et al. (2008), we assume that permanent income $z_{i,t}$ evolves according to a random walk. We additionally assume that permanent income is subject to a drift $B_e(r,d,a)$ that depends on the household's employment status $e \in \{E, N\}$ as well as their lagged income rank r, decade d, and age a.[21] Additionally, permanent income is subject to shocks $\nu_{i,t}$, which are drawn from a normal distribution and depend on a household's employment status, lagged income rank, decade, and age. Given this formulation permanent income evolves according to,

$$z_{i,t} = \begin{cases} z_{i,t-1} + B_E(r,d,a) + \nu_{i,t}, & \nu_{i,t} \sim \mathcal{N}\big(0, Q_E(r,d,a)\big) \quad \text{if } l_{E,i,t} = 1 \\ z_{i,t-1} + B_N(r,d,a) + \nu_{i,t}, & \nu_{i,t} \sim \mathcal{N}\big(0, Q_N(r,d,a)\big) \quad \text{if } l_{E,i,t} = 0 \end{cases} \tag{2}$$

Households draw an initial value of permanent income, which we denote $z_{i,0}$. We assume that households draw this value of initial permanent income from a normal distribution with variance $Q_0(d,a)$, which varies by decade and age,

$$z_{i,0} \sim \mathcal{N}\big(0, Q_0(a,d)\big) \tag{3}$$

Equations (1)-(3) characterize our income process. Despite contemporaneous shocks being drawn from Gaussian distributions, this income process is actually a mixture model with tail risk (i.e., fatter tails than the standard normal distribution). Conditional on income in period $t-1$ (and other observables), households draw their income shocks in period t from a Gaussian distribution. However, income in future periods is a random variable, and it determines which Gaussian distribution the household draws shocks from. This forward looking randomness over which Gaussian distribution will be drawn introduces mixing and scope for tail risk. In addition to mixing, the employment process provides an additional source of non-Gaussian risk. Non-employment/non-filing generates large, discrete changes in permanent earnings.

Functional forms. To facilitate estimation, we impose functional forms on the time, age, and rank dependence of parameters. In practice, we use B-splines to model the relationship be-

sectional of households there is mixing of normal distributions, which generates non-Gaussian features. Additionally, since a household's future rank depends on the shocks they receive in the current period, households are mixing over normal distributions when they form expectations about their future income realizations.

[21] Note the drift to permanent income can equivalently be thought of as the mean of the shock to permanent income.

tween lagged income rank r and income shocks, allowing this relationship to vary by decade d. We denote these B-spline components with a superscript s. To account for changes in the age distribution of the U.S. population over time, we also allow the functions $Q(\cdot)$, $R(\cdot)$, and $B(\cdot)$ to include a quadratic component in age, denoted by a superscript q. For example, we specify the variance of transitory income shocks as,

$$R(r,d,a) = \exp\left(\widehat{R}^s(r,d) + \widehat{R}^q(a)\right),$$

where $\widehat{R}^s(r,d)$ is the B-spline in rank and decade, and $\widehat{R}^q(a) = \lambda^q_{R,1}a + \lambda^q_{R,2}a^2$ is the quadratic in age.[22] Similarly, the variance, initial variance, and drift of permanent income are respectively given by,

$$Q_e(r,d,a) = \exp\left(\widehat{Q}^s_e(r,d) + \widehat{Q}^q_e(a)\right),\ Q_0(d,a) = \exp\left(\widehat{Q}_0(d) + \widehat{Q}^q_0(a,d)\right),\ B_e(r,d,a) = B^s_e(r,d) + B^q_e(a)$$

where $e \in \{E,N\}$ indexes employment status. We assume that $\widehat{Q}_0(d)$ captures decade-specific effects (i.e., decade fixed effects) and $\widehat{Q}^q_0(a,d)$ represents an age-dependent quadratic that is allowed to vary over time. Specifically, we estimate one set of quadratic parameters for the 1970s (including 1969), and a separate set for all subsequent decades.[23]

While richer income processes have been estimated in stationary (non-time dependent) environments (e.g. Arellano et al. (2017), De Nardi et al. (2020), and Guvenen et al. (2021)), our assumptions of linear additivity, conditional Gaussianity, and log-linear variances provide us with enough tractability to accomplish our goals of (1) allowing future income risk to depend on the path of prior income, (2) allowing past-income-dependence to evolve over time, and (3) being able to estimate the income process (and variants of the process with more flexible dependence on age, demographics, occupation etc.) on large administrative datasets. We next discuss how we tractably estimate the income process specified above.

2.2 Estimation

To estimate our income process, we adapt and apply the estimation routine from Braxton et al. (2025), which combines the Kalman filter with an Expectation Maximization (EM) algorithm. Below we discuss the role of each of these components of the estimation, and in Appendix A, we provide additional details about each step. We conclude this section by providing a brief "cookbook" of how to utilize the estimation procedure.

[22] Following Braxton et al. (2025), we model log variances to ensure that the estimated variances are positive.

[23] Allowing the age profile to shift between the initial and later periods is consistent with the approach in Braxton et al. (2025).

Kalman Filter. The first step of our estimation routine is to use the Kalman filter, which allows us to recover the path of permanent and transitory income for each household in our sample. The Kalman filter naturally applies to our model, since it is in state space form (equation (1) is the observation equation and equation (2) is the state equation). However, when applying the EM algorithm, it is convenient to modify the state space representation to include lagged permanent income in equation (2) (see Appendix A.1). Moreover, the EM algorithm requires smoothed estimates of permanent income, which we recover by applying the Kalman smoother (see Appendix A.1).[24]

Expectation Maximization (EM) Algorithm. To run the Kalman filter and smoother, we need an estimate of the parameters of the income process. We obtain parameters of our income process via maximum likelihood estimation (MLE). While there are many different ways to maximize the likelihood (simple grid search, gradient descent, etc.), we adopt the approach of Dempster et al. (1977) and apply the Expectations Maximization (EM) algorithm. The benefit of utilizing an EM algorithm that is it yields closed-form updating formulas that resemble generalized least-squares (GLS) regressions, where the weights used in the regression adjust for filtering uncertainty. As we discuss in Appendix A.2, these closed from expressions come from taking FOCs of the expected log-likelihood with respect to parameters. These closed-form expressions allow us to tractably update the parameters of our income process despite them being rich functions of observables.

Cookbook of Estimation Procedure. Finally, we conclude with a brief overview of our estimation routine. Our estimation routine proceeds as follows:

i. Guess an initial set of parameters for the income process, denoted θ_0.

ii. Using the current estimate of parameters θ_0, run the Kalman filter and smoother.

iii. Using the estimates from the Kalman filter and smoother steps, update the parameters using the GLS style formulas from the EM algorithm. Denote the updated parameters as θ_1.

iv. Repeat steps (ii.)-(iii.) until the parameters have converged.

[24] Additionally, when the estimates of the state vector are of interest, as they are in our application, Hamilton (1994) comments that we can improve the inference about the historical values of the state vector in the middle of the sample by using the Kalman smoother.

In Appendix A.3, we use a series of Monte Carlo exercises to show that: (1) our income process is identified, and (2) our method can accurately recover the parameters of the income process as well as the household level estimates of permanent income.[25]

Next, we apply our income process to the tax panel and analyze the results.

3 Estimation Results

After estimating our income process on the tax panel, we find that permanent income risk has increased, especially among high-income households. We then validate that high-income households do in fact face income risk by showing that large declines in permanent income are associated with various indicators for economic hardship (e.g., early 401K withdrawals, moving to lower income neighborhoods, etc.). Finally, we examine how savings responds to changes in income risk and find that in response to increases in permanent income risk, households save more, especially high-income and younger households.

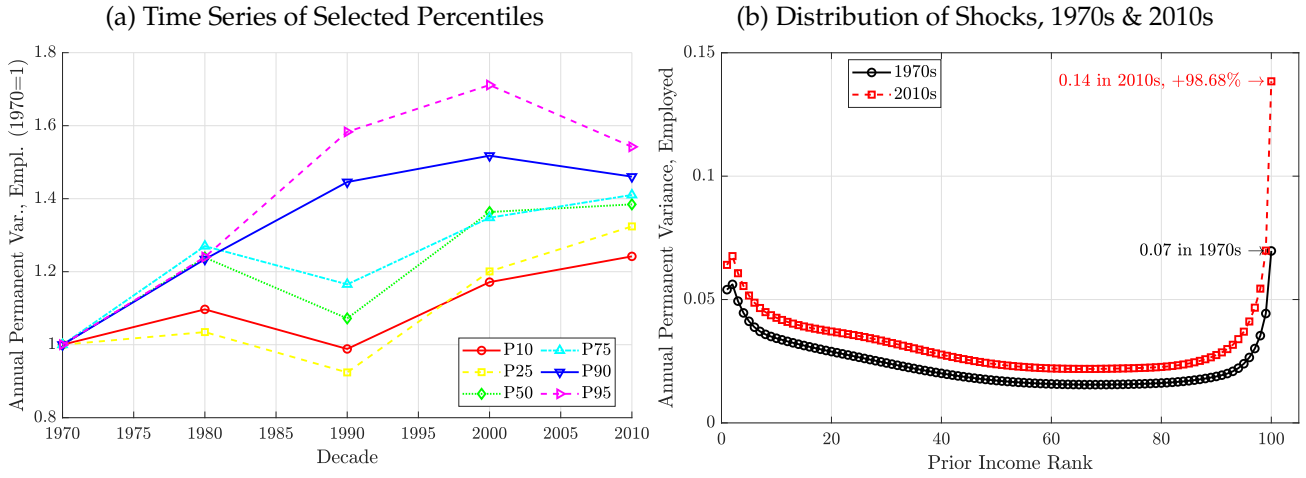
3.1 Permanent and Transitory Income Risk Across Time and People

We start by showing the evolution of permanent income risk among the employed since the 1970s for selected percentiles of the (lagged) income distribution in panel (a) of Figure 4. The pink, dashed line with triangle markers, shows the evolution of permanent income risk for households at the 95th percentile of the lagged income distribution. The figure shows that since the 1970s, permanent income risk for these households steadily increases between the 1970s and the 2000s. Over this time period, the variance of shocks to permanent income increased by over 70% for these households before declining in the 2010s. Overall these high-income households have seen their permanent income risk increase by approximately 55% since the 1970s. The blue solid line with triangle markers shows the evolution of permanent income risk among households at the 90th percentile of the income distribution and shows that these households have faced a similar pattern of rising income risk between the 1970s and the 2010s. For these households, income risk has increased by over 45% since the 1970s.

The red and yellow lines in Panel (a) of Figure 4 present the evolution of permanent income risk among lower income households, corresponding to the 10th and 25th percentiles, respectively. Lower income households have experienced a different evolution of permanent income risk relative to high-income households. Permanent income risk was effectively unchanged for these households between the 1970s and 1990s, before increasing between the 1990s and 2010s.

[25] We also show in Appendix A.3 that an estimation approach in the GMM tradition is not able to accurately recover the parameters of an income process where the parameters depend upon lagged income ranks.

Figure 4: Permanent Income Risk Among Employed: Over Time and By Lagged Income



Note: Figure presents the results of estimating the income process in Section 2.1. Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income rank for the 1970s (black, solid line with circle markers) and the 2010s (red, dashed line, with square markers). Note in panel (b) we present an annualized variance.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

Despite the increases over the past 20 years, the increases in permanent income risk have been more muted for these lower income households, increasing by between 20% and 30% since the 1970s.

While the results presented in Panel (a) of Figure 4 show the evolution of permanent income risk for each decade by selected percentile, in Panel (b) of Figure 4 we show the full profile of shocks to permanent income by lagged income rank for the 1970s (black, solid line with circle markers) and 2010s (red, dashed line with square markers).[26] The figure shows that permanent income risk is U-shaped across the income distribution. Permanent income risk decreases from the bottom of the lagged income distribution until approximately the 75th percentile of the income distribution, at which point permanent income risk starts to increase again. We find that the increases in permanent income risk as a function of lagged rank are especially steep as we move towards the very top of the income distribution. Comparing the 1970s to the 2010s, we

[26]In panel (b), we present an annualized estimate of the variance of shocks to permanent income among the employed. We annualize our estimates by dividing by 5. Additionally, our estimation routine produces an estimates of the mean shock to permanent income among the employed by lagged income rank for each decade. We present these estimates in Appendix C.1. These estimates reveal substantial mean reversion in permanent income; however we do not find large variation in these profiles between the 1970s and 2010s.

find that there has been an increase in permanent income risk across the entire income distribution (as shown in Panel (a)); however, the increase is especially large towards the top of the income distribution. Averaging over the households in the top 5% of the income distribution, we find that the variance of shocks to permanent income has increased by nearly 70%. If we look at the very top of the income distribution, we find that permanent income risk has effectively doubled. Thus, high earners have seen a substantial acceleration of permanent income risk over the past 50 years.

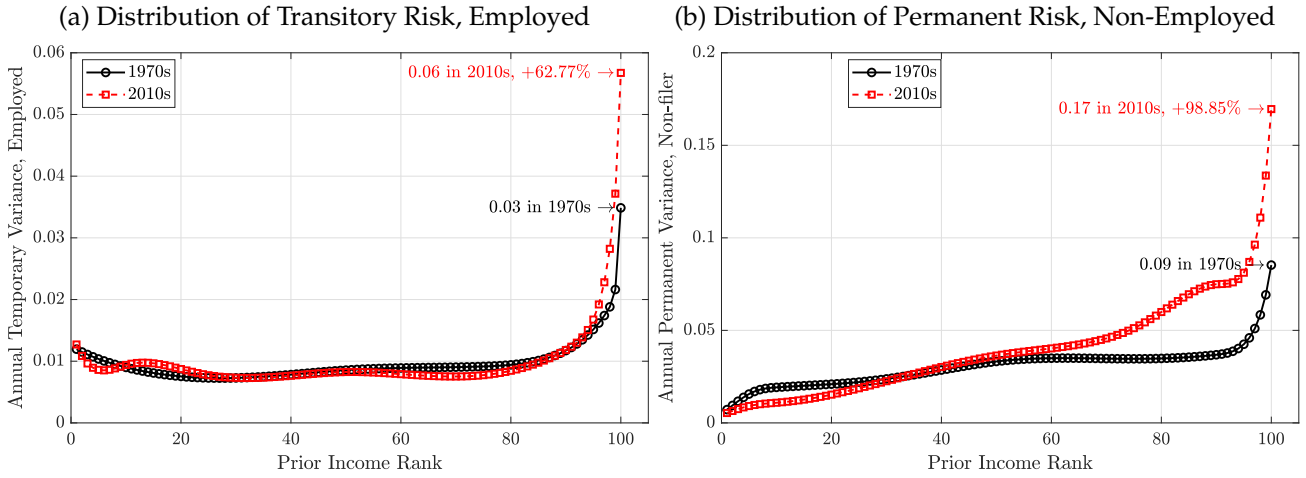
We next examine the evolution of transitory income risk over time and across the income distribution. Panel (a) of Figure 5 presents the variance of shocks to transitory income (i.e., transitory income risk) as a function of lagged income rank for the 1970s (black, solid line with circle markers) and 2010s (red, dashed line with square markers). The figure shows that transitory income risk is effectively flat for the bottom 80 percent of the income distribution and then starts to increase as we move out into the right tail of the income distribution. Comparing the 1970s and 2010s, we find that the level of transitory income risk is effectively unchanged for the bottom 95% of the income distribution. However, for households in the top 5% of the income distribution, transitory income risk is higher in the 2010s relative to the 1970s. In particular, for households at the very top of the income distribution the variance of shocks to transitory income has increased by over 60% since the 1970s. Combining these results with those presented in Figure 4, we find that high-income households are exposed to the largest amount of permanent and transitory income risk while employed and that this risk has been increasing since the 1970s.

A feature of our income process and estimation routine is that it allows us to estimate the risk that households face during periods of non-employment or when they are not filing their taxes. In Panel (b) of Figure 5, we present the variance of shocks to permanent income during periods of non-employment/non-filing for the 1970s and 2010s.[27] The figure shows that the variance of shocks to permanent income during periods of non-employment is increasing in lagged income rank (from when a household was last employed). Comparing the 1970s to the 2010s, we find that the variance of shocks to permanent income has increased during periods of non-employment and non-filing for households above the 60th percentile of the income distribution. As we move out into the right tail of the income distribution, we find larger increases in permanent income risk over time, and at the very top of the distribution find that it has nearly doubled since the 1970s.

Im summary, high-income households face the highest risk to permanent income both while

[27] In Appendix C.1, we present the mean of shocks to permanent income in periods of non-employment/non-filing as a function of lagged income rank and over time. We find that the decline in permanent income from periods of non-employment/non-filing is larger for high-income households than for lower income households.

Figure 5: Transitory Risk and Permanent Risk Among Non-Employed/Non-Filers



Note: Figure presents the results of estimating the income process in Section 2.1. Panel (a) plots the full profile of shocks to transitory income among the employed as a function of lagged income rank for the 1970s and 2010s. Panel (b) plots the full profile of shocks to permanent income among the non-employed/non-filing as a function of lagged income rank for the 1970s and the 2010s. In both figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

employed and non-employed/non-filing as well as the highest transitory income risk.[28] We additionally find that these measures of income risk have increased over time for high-income households. We next discuss a series of additional results about the evolution of income risk over time and across the income distribution.

Risk among self-employed. The results presented thus far have pooled together all individuals and examined the evolution of income risk across the income distribution and over time. Nest, we separately analyze the evolution of income risk for individuals who we designate as "self-employed" versus individuals who we designated as "non self-employed." To do so, we leverage our ability to link the sample of 1040 tax returns to the CPS-ASEC, which asks respondents about the *class of worker* they were in the prior year, i.e., were they self-employed, working for wage and salary, etc.[29] Based on the responses to this question, we split our sample into the self-employed and the non self-employed and estimate our income process separately for each

[28]Finally, our income process also allows the variance of the initial draw of permanent income to vary by decade. We show in Appendix C.1 that the variance of the initial draw of permanent income has been steadily increasing since the 1970s.

[29]See Bowler (2006) for detailed discussion of who is classified as self-employed in the CPS.

group.[30]

Figure 6 presents the profiles of permanent income risk for the self-employed (Panel (a)) and non self-employed (Panel (b)) from estimating our income process separately by self-employment status.[31] These figures highlight several points about the differences in risk faced by the self-employed compared to the non-self employed and how the risk each of these groups has evolved over time. First, comparing the profiles of risk for the self-employed to the non self-employed, we observe that self-employed individuals face higher *levels* of permanent income risk at almost every point in the income distribution in both the 1970s and the 2010s. Second, while we find that permanent income risk has increased across the income distribution for both the self-employed and non-self employed, the increase at the top of the income distribution for the self-employed is especially large. On average, for self-employed individuals in the top 5% of the income distribution permanent income risk has increased by over 130% and at the very top of the income distribution permanent income risk has increased by over 175%. While the high-income non self-employed have also seen substantial increases in permanent income risk, the magnitude of the increase for these households is smaller.

Robustness of income risk patterns We conclude this subsection by summarizing a series of robustness results.

Alternative income definitions. In Appendix C.2, we show that our results are robust to a series of alternative definitions of income. We show that our results are robust to using adjusted gross income (AGI) (Appendix C.2.1), wage and salary income (WSI) (Appendix C.2.2), as well as removing capital gains from our baseline measure of income (Appendix C.2.3).

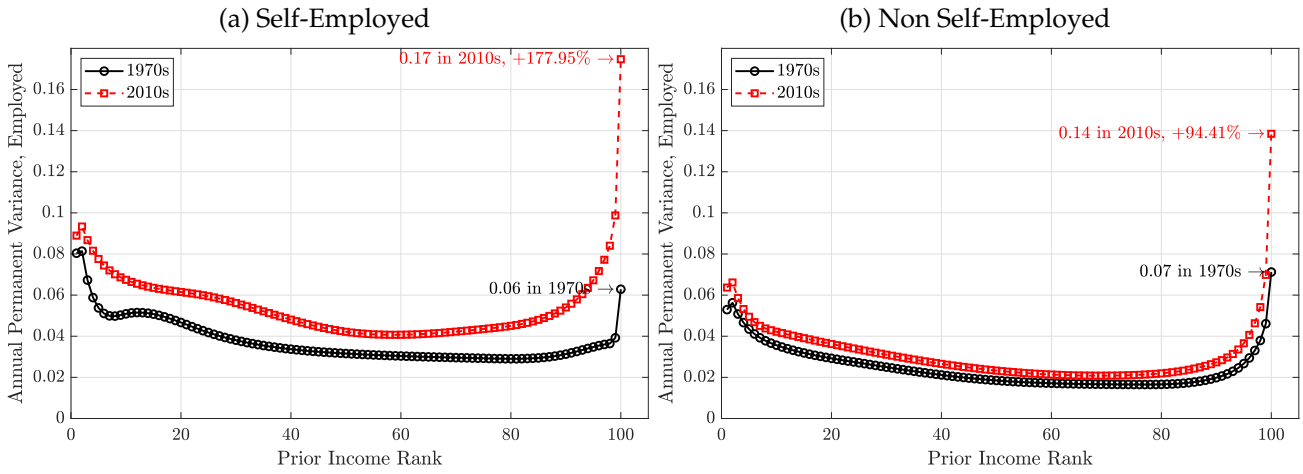
Alternative sampling. We show in Appendix C.3, that our results are robust to using a 2% random sample of all 1040s.

Demographic composition. In Appendix C.4.1, we show that our results are robust to interacting the fixed effects for year and number of adults in the household as part of our residualization process, which we interpret as evidence that our results are not driven by changes in labor force composition over time. In Appendix C.4.2, we show that we obtain similar results when we rank individuals within their 2-digit occupation, which highlights that our results are not driven by composition changes. Finally, in Appendix C.4.3, we estimate a version of our

[30] In estimating our income process separately for each group, we residualize earnings separately for each group. Additionally, note that we rank individuals within the full sample, so that being a given rank r indicates have the same level of income on average for both the self-employed and non self-employed. Finally, note our definition of self-employed combines both incorporated and non-incorporated self-employed. Our results are robust to including only incorporated self-employed.

[31] In Appendix C.1.1, we present the profiles of transitory income risk and permanent income among the non-employed risk for the self-employed and non-self-employed. We find relatively minor differences in the self-employed and non self-employed along these dimensions of income risk.

Figure 6: Permanent Income Risk Among Employed by Self-Employment Status



Note: Figure presents the results of estimating the income process in Section 2.1 where the parameters are allowed to differ by self-employment status as reported in the ASEC. Panel (a) plots the full profile of shocks to permanent income among the self-employed as a function of lagged income rank for the 1970s and the 2010s. Panel (b) plots the full profile of shocks to permanent income among the non self-employed as a function of lagged income rank for the 1970s and the 2010s. In both figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

income process where the profile of shocks across the income distribution and over time differs by whether or not an individual is above/below age 40. We find that permanent income risk has increased among both younger and older high-income households.

3.2 Do high-income households really face income risk?

In Section 3.1, we argued that high-income households are exposed to large and rising permanent income risk. One concern with this interpretation is that the income changes among these high-income households, which underlie our estimates, may not reflect risk. In this section, we examine how observable indicators of financial distress align with large declines in permanent income, and how these relationships evolve across the income distribution.

Let $Y_{i,t}$ be an indicator variable that is informative about an individual undergoing financial distress (e.g., an early 401(k) withdrawal, moving to a lower income ZIP code, etc.). Let $D^P_{i,t}$ indicate whether individual i experienced a 1 standard deviation (or larger) decline in permanent income from $t-5$ to t.[32] Finally, let $X_{i,t}$ be a vector of controls that includes fixed effects

[32]Let $z_{i,t}$ denote our estimate of permanent income for individual i in period t from the Kalman step of our estimation routine. Define $\Delta z_{i,t} = z_{i,t} - z_{i,t-5}$ as the five-year change in permanent income between periods $t-5$

for lagged income rank, year, age, ventiles of transitory shocks, and the number of adults in the household. The specification we estimate is,

$$Y_{i,t} = \alpha + \beta D^P_{i,t} + \Gamma X_{i,t} + \epsilon_{i,t} \tag{4}$$

The coefficient of interest is β, which summarizes the impact of a 1-SD decline in permanent income on the outcome variable. We additionally estimate equation (4) across the income distribution to examine how households in different spots of the income distribution respond to income shocks.

Early 401(k) Withdrawals. Our first validation exercise examines the link between large permanent income declines and early 401(k) withdrawals. Using tax form 1099-R, we observe when an individual makes a 401(k) withdrawal as well as the size of the withdrawal. To highlight the response to risk, we focus on "early withdrawals," which we define to be withdrawals that occur before age 55. To line up with the five-year timing of our income data from the 1040s, we analyze withdrawals of at least \$1,000 in any year between $t-4$ and t. Further, we require that an individual made at least \$4,000 worth of contributions into their 401(k) plan between the years $t-8$ and $t-5$.[33] This final condition is to ensure that we are considering individuals who have a 401(k) plan with contributions that can be withdrawn.[34]

Panel (a) of Figure 7 presents the results of estimating equation (4) where the outcome variable of interest is having an early 401(k) withdrawal. The blue bar reports the regression coefficient for a 1-SD decline in permanent income among our full sample of households. The coefficient indicates that in response to a 1-SD decline in permanent income, the likelihood that a household has an early 401(k) withdrawal increases by approximately 6.5 percentage points.[35] In our sample, approximately 31.5% of households will make an early 401(k) withdrawal over a 5-year period. Thus, a large decline in permanent income increases the likelihood of an early withdrawal by roughly 20% relative to the sample average, consistent with our estimates capturing meaningful income risk

We next examine how early withdrawals vary over the income distribution by estimating

and t. Let $\sigma_{r,t}$ denote the standard deviation of shocks to permanent income for individuals with lagged income rank r in period t. We define the indicator $D^P_{i,t}$ to equal one if $\Delta z_{i,t} < -\sigma_{r,t}$ for an individual with lagged rank r, and zero otherwise.

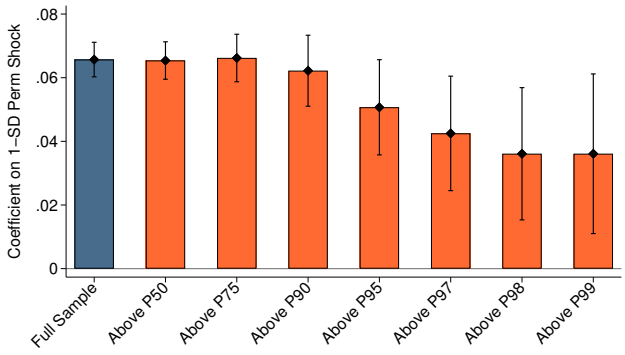
[33] We utilize Box 12 on an individuals W2 to measure the amount of 401K contributions that an individual makes in a given year.

[34] These conditions for considering early 401K withdrawals mirror those considered in Choukhmane et al. (2024), but adapted to the 5-year reporting of income in our data.

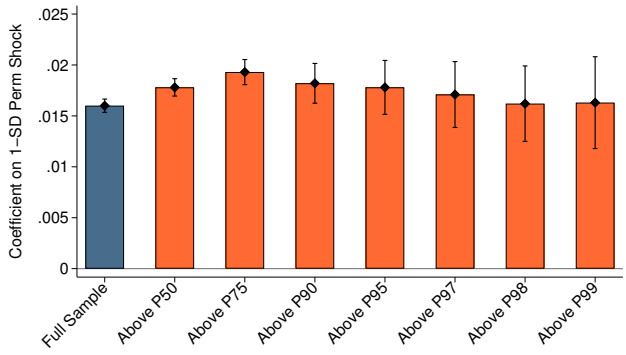
[35] For households that make an early 401(k) withdrawal, the average size of the withdrawal is approximately \$32k.

Figure 7: Response to a 1-SD Decline in Permanent Income

(a) Early 401(k) Withdrawals



(b) Moving to Lower Income Zip Codes



Note: Figure presents the results of estimating equation (4) where the outcome variable of interest is: (1) having an early 401(k) withdrawal (Panel (a), and (2) moving to a lower income ZIP code (Panel (b)). The bars correspond to the coefficient for the impact of having a 1-SD decline in permanent income. The black bars correspond to a 95% confidence interval using robust standard errors.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

equation (4) separately for households above a given lagged earnings centile. This exercise is particularly informative because higher-income households are likely to have greater access to liquid assets and credit, and therefore may be less reliant on early withdrawals in response to income declines. The orange bars in Figure 7 show that even among individuals in the top 5% of the income distribution, a 1-SD shock to permanent income increases their likelihood of an early 401(k) withdrawal by over 5 percentage points. Even at the very top of the income distribution, a 1-SD decline in permanent income still increases the likelihood of having an early 401(k) withdrawal by over 3.5 percentage points, indicating that these responses remain economically meaningful even among the highest-income households.[36]

Moving to lower income ZIP codes. Our second validation exercise examines an alternative adjustment margin through which households respond to large income declines: moving to a lower income neighborhood. For all individuals in our sample, we observe a panel of residential ZIP codes as part of their 1040 tax form. Using a 5% random sample of 1040s, we measure the annual median adjusted gross income (AGI) of each ZIP code in our panel. We then define a move to a lower income neighborhood whenever an individual changes ZIP codes from $t-5$

[36] Among households in the top 5% of the income distribution, 27.2% make an early 401(k) withdrawal over a five-year period, compared to 26.8% among those in the top 1%. Conditional on making a withdrawal, the average amounts are approximately $46,000 and $49,000, respectively.

to t and the median income at the destination ZIP code is at least 30% lower.[37]

Panel (b) of Figure 7 presents the results of estimating equation (4) where the outcome of interest is moving to a lower income ZIP code. The blue bar shows the impact of having a 1-SD decline in permanent income on the likelihood of such a move in our full sample of households. The coefficient indicates that these large declines in permanent income are associated with an increase in the likelihood of moving to a lower income ZIP code by approximately 1.6 percentage points. This increase is economically meaningful relative to a sample average of 5.9% over a five-year period.

As above, we next examine how this relationship varies across the income distribution by estimating equation (4) separately for households above a given lagged earnings centile. The orange bars in Figure 7 show that the effect of large declines in permanent income on the likelihood of moving to a lower income ZIP code is quite stable across the income distribution. Even at the very top of the income distribution, a 1-SD decline in permanent income increases the likelihood of moving to a lower income ZIP code by approximately 1.6 percentage points, indicating that these responses remain economically meaningful even among the highest-income households.[38]

Taking Stock. In this section, we've shown that in response to a large decline in permanent income, households respond by making early 401(k) withdrawals and moving to lower income neighborhoods. These results show that our econometric classification of shocks aligns with observable measures of financial hardship, even among the top 1% of the income distribution. These results suggest that high-income households do in fact face income risk. While these results examine how households respond to realized shocks, we next examine how changes in risk (i.e., changes to the variance of shocks) affect savings behavior.

3.3 Does greater income risk impact savings?

A central prediction of quantitative models with incomplete markets (e.g., Bewley, Huggett, and Aiyagari) is that households facing greater income risk accumulate more assets in order to self-insure against adverse shocks. In this section, we empirically examine the relationship between income risk and savings by exploiting variation in permanent and transitory income risk across occupations and over time, while holding constant households' income and other

[37] In results that are available upon request, our results are robust to using mean income within a ZIP code as well as alternative cutoffs for cross-ZIP income changes.

[38] Among households in the top 1% of the income distribution, 7.9% move to a lower income ZIP code over a five-year period.

observable characteristics.[39] We then use the estimated savings response to income risk to discipline the quantitative model introduced in Section 4.

For this analysis we measure income risk by occupation and over time. Specifically, we separately estimate permanent and transitory income risk over time for each of the 334 time-consistent occupation codes developed by Autor and Dorn (2013).[40] Individuals are assigned to an occupation based on their first reported occupation in the ASEC. We summarize income risk using occupation level measures of permanent and transitory income variability. Let $Q_{o,E,t}$ denote the variance of permanent income shocks for employed individuals in occupation o at time t, and let $R_{o,t}$ denote the corresponding variance of transitory income shocks.[41]

To examine how income risk affects savings behavior, we use capitalized measures of wealth and estimate how changes in wealth respond to variation in occupation level income risk. We build asset positions by capitalizing interest and dividend income from tax returns, following Smith et al. (2023). This procedure yields estimates of wealth held in savings accounts, stocks and bonds for every household between 1974-2014.[42] Let $b_{i,t}$ denote the estimated asset levels, where i indexes households and t indexes time. Define $\Delta_h b_{i,t} = (b_{i,t+h} - b_{i,t})/Y_{i,t}$ as the change in wealth relative to income for individual i between periods t and $t+h$.[43] Let $X_{i,t}$ denote a vector of controls, which includes the mean permanent shock (for each employment status and occupation o in year t), as well as fixed effects for lagged income rank, centiles of lagged wealth by year, age, number of adults in the household and 2-digit occupation. We estimate specifications of the form:

$$\Delta_h b_{i,t} = \alpha + \beta_{Q_E} Q_{o,E,t} + \beta_{Q_N} Q_{o,N,t} + \beta_R R_{o,t} + \Gamma X_{i,t} + \epsilon_{i,t}. \tag{5}$$

Equation (5) relates changes in household wealth to variation in occupation-level income risk, while flexibly controlling for income, wealth, and demographic characteristics. The coefficients of interest in equation (5) are β_j for $j \in \{Q_E, Q_N, R\}$, which capture how differences in income risk across occupations translate into differences in savings behavior, holding constant

[39] This approach follows a long tradition in the precautionary savings literature that exploits occupation-level variation in income risk to study savings behavior (e.g., Carroll and Samwick (1997), Boar (2021)).

[40] Additionally, in this estimation, we allow both the mean and variance of income shocks to vary with age within each occupation (via an age quadratic), and we residualize income separately by occupation. We do not condition on income rank, as many occupations do not span the full income distribution.

[41] We additionally define $Q_{o,N,t}$ as the permanent income risk for non-employed or non-filing individuals in occupation o at time t. We also construct occupation-level measures of mean permanent income shocks, denoted $B_{o,E,t}$ and $B_{o,N,t}$, which we include as controls to separate changes in risk from changes in mean income.

[42] See Appendix D for details. We thank Eric Zwick for sharing data on dividend capitalization factors. Our 1040 sample does not include dividend income in the year 1969. Additionally, the capitalization index for dividends, which we obtain form Smith et al. (2023), ends in the year 2016.

[43] We restrict the sample to individuals with income exceeding \$1,000 (in 2005 PCE dollars).

income rank. To isolate the effect of risk from changes in expected income, we include controls for the mean occupation-level permanent income shock. We also condition on lagged income rank, which limits comparisons to households with similar income levels, and we include fixed effects for 2-digit occupations to ensure that identification comes from variation in income risk within broadly similar types of work.

Table 2 presents the results of estimating equation (5), where we measure the change in wealth to income over a 5-year horizon.[44] For ease of exposition, we introduce the measures of income risk sequentially. In column (1) of Table 2, we show how changes in wealth to income responds to variation in occupation level permanent income risk. The positive and statistically significant coefficient on permanent income risk (Q_E) indicates that within a given income rank, individuals who are in occupations that face greater permanent income risk have a larger change in wealth to income. In particular, in response to a 0.1 increase in permanent income variance, wealth increases, on average, by 4.7% of prior income, ceteris paribus. Thus, in responses to greater permanent income risk, households respond by saving more.

We next examine the response to other measures of income risk. In column (2) of Table 2, we examine how permanent income risk when non-employed/non-filing impacts savings behavior. The coefficient on permanent income risk when non-employed/non-filing (Q_U) indicates that greater exposure to permanent risk while non-employed/non-filing *does not* impact changes in wealth to income in an economically or statistically significant manner. In column (3) of Table 2, we examine how transitory income risk impacts savings. The positive and statistically significant coefficient on transitory income (R) indicates that in response to greater transitory income risk households respond by increasing their wealth to income. In response to a 0.1 increase in an occupation's transitory income variance, individuals respond by increasing their wealth by 1.94% of prior income, ceteris paribus.

Finally, Column (4) considers our measures of income risk jointly and controls for the mean of shocks to permanent income. The coefficient estimates in column (4) show that we obtain very similar estimates when using all three measures of risk jointly and when controlling for the mean shocks to permanent income. We continue to find that greater exposure to permanent and transitory income risk is associated with increases in savings. We next examine the heterogeneity in the response of savings to permanent and transitory income risk.

Heterogeneity. We examine the heterogeneity in how savings responds to income shocks by estimating equation (5) for different sub-samples of our population. We focus on the response to permanent and transitory income risk as those were the parameters that we found impacted

[44] In results that are available upon request, we find that our results are robust to using a 10-year horizon to measure changes in wealth to income.

Table 2: Income Risk and Changes in Savings Behavior

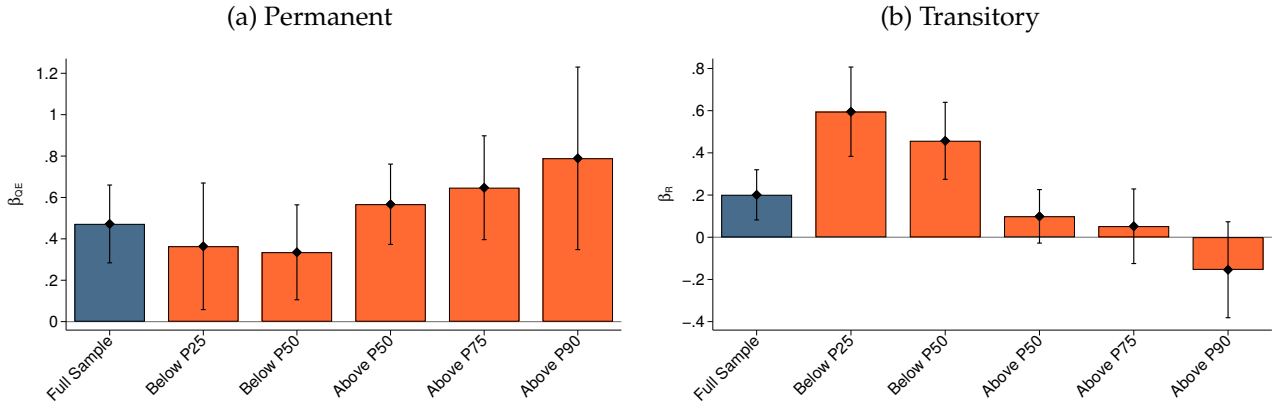
Dependent variable: 5-year change in wealth to income				
	(1)	(2)	(3)	(4)
Q_E	0.471***			0.468***
	(0.0951)			(0.0961)
Q_N		0.00319		0.00287
		(0.00228)		(0.00181)
R			0.194***	0.189***
			(0.0598)	(0.0592)
Controls	Y	Y	Y	Y
Lagged Wealth X Year FE	Y	Y	Y	Y
Lagged Income Rank FE	Y	Y	Y	Y
2-Digit Occ. FE	Y	Y	Y	Y
R-squared	0.048	0.048	0.048	0.048
No. Obs.	4,651,000	4,651,000	4,651,000	4,651,000

*Note: Table presents the results of estimating equation (5), where the dependent variable is the 5-year change in wealth to income. We measure wealth using the capitalization method of Smith et al. (2023) on interest and dividend income. Income risk is measured at the occupation level as described in Section 3.3. Controls include the mean shock to permanent income during periods of employment as well as non-employment non-filing in an occupation and year as well as fixed effects for age and the number of adults in the household. Clustered standard errors are in parentheses, where the clustering is performed at the occupation level, where occupations are classified using Autor and Dorn (2013) occupation codes. *** p<0.01, ** p<0.05, * p<0.1*
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

savings above, and consider these measures of risk sequentially. In Figure 8, we examine how savings responds to permanent (panel (a)) and transitory (panel (b)) income risk for different regions of the income distribution. The navy bar in panel (a) of Figure 8 presents our baseline estimate of how individuals respond to permanent income risk. The orange bars then present the response of savings for different regions of the income distribution. Clearly, the response to permanent income risk is increasing across the income distribution. In particular, for individuals below the median, in response to a 0.1 increase in permanent income variance, wealth increases by 3.4% of prior income. For individuals, above the median they increase their wealth by 5.7% of prior income. Thus, we find that higher income households increase their savings more in response to increasing permanent income risk compared to lower income households.

In panel (b) of Figure 8, we show that the response of savings to transitory income risk is decreasing across the income distribution. For individuals below the median, wealth increases by 4.6% of prior income in response to a 0.1 increase in transitory variance. Conversely, for

Figure 8: Income Risk and Changes in Savings Behavior by Prior Income Rank



Note: Figure presents the results of estimating equation (5) in the full sample (navy bar) as well as across the income distribution (orange bars). Panel (a) present the response of wealth to income to permanent income risk, while panel (b) presents the response of wealth to income to transitory risk. The bars represent the coefficient estimates and the whiskers represent a 95% confidence interval. To compute the 95% confidence interval, we cluster the standard errors by occupation where occupations are measured using the classification from Autor and Dorn (2013).
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

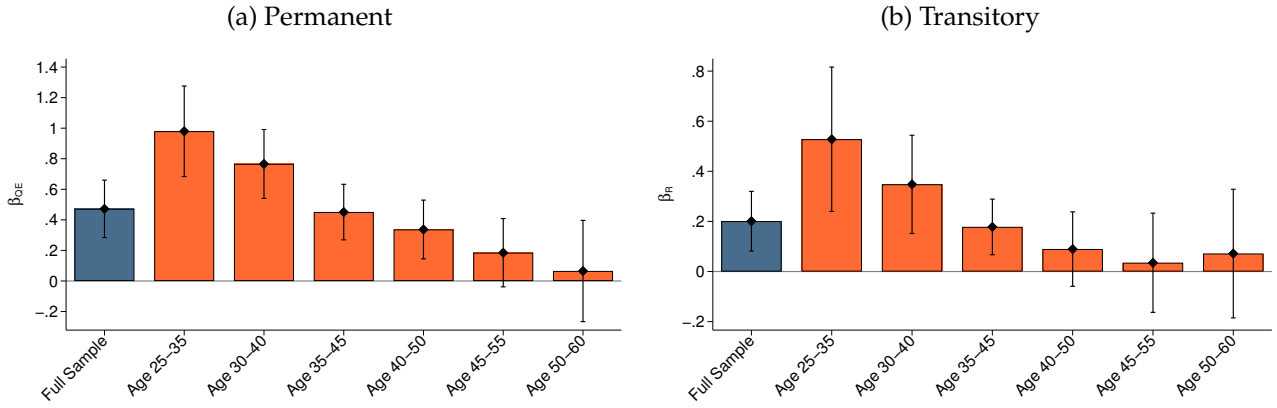
individuals above the median, wealth to income increases by less than 1% in response to a 0.1 increase in transitory income variance, and this increase is not statistically different from zero (t-stat = 1.53).

We next examine the response of savings to income risk for individuals of different ages. In panel (a) of Figure 9, we show that the response of savings to permanent income risk is decreasing in the age of an individual. For individuals between the ages of 25 and 34, in response to a 0.1 increase in permanent income variance, wealth increases by 9.8% of prior income. Conversely, for individuals between the ages of 50 and 60, the response is not statistically different from zero. Similarly, panel (b) shows that the response of savings to transitory income risk is decreasing in age. For individuals between the ages of 25 and 34, in response to a 0.1 increase in transitory variance, wealth increases by 5.3% of prior income, while for individuals between the ages of 50 and 60 there is effectively no change in wealth.[45]

In this section, we've shown that savings increases in response to permanent and transitory income risk. The response to permanent risk is driven by high-income and younger individuals, while the response to transitory income risk comes from lower income and younger individuals. We will use these estimates to calibrate and validation our quantitative model in Section

[45]The larger savings response of young households aligns with the "buffer stock" theory of savings (e.g., Carroll (1997) and Gourinchas and Parker (2002).)

Figure 9: Income Risk and Changes in Savings Behavior by Age



Note: Figure presents the results of estimating equation (5) for workers of different ages. Panel (a) present the response of wealth to income to permanent income risk, while panel (b) presents the response of wealth to income to transitory risk. The bars represent the coefficient estimates and the whiskers represent a 95% confidence interval. To compute the 95% confidence interval, we cluster the standard errors by occupation where occupations are measured using the classification from Autor and Dorn (2013).
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

4.

Taking stock. In this section, we've shown that permanent and transitory income risk has increased for top-earners in the US since the 1970s. We've additionally shown that in response to increasing permanent income risk, individuals respond by increasing their savings, especially among high earners and younger individuals. In the next section, we introduce a life-cycle Bewley-Huggett-Aiyagari with our income process. We then use the empirical estimates from this section to discipline the quantitative model. Using the quantitative model, we then examine the implication of risking permanent and transitory income risk for top earners on the risk-free interest rate and wealth inequality.

4 Model

To evaluate the effects of rising permanent income risk on the risk-free rate and wealth inequality, we integrate our income process into a finite lifecycle Bewley-Huggett-Aiyagari model. There is a unit measure of $T \geq 2$ overlapping generations of households, where a household's age is denoted $t \in \{1, \ldots, T\}$. Households save $(b > 0)$ and borrow $(b < 0)$ at interest rate r_f using a risk-free asset $b \in [\underline{b}, \infty]$, subject to the natural borrowing limit $\underline{b} \leq 0$. Their asset choices

maximize the present value of their flow utility over consumption $u(c)$, discounted at rate β. In their terminal period of life, households have warm-glow bequest motives, $\phi(b)$. Household preferences are thus given by,

$$\mathbb{E}\Big[\sum_{t=1}^{T}\beta^{t-1}u(c_t)+\beta^{T}\phi(b_{T+1})\Big].$$

Households face risky streams of permanent ($z\in\mathbb{R}$) and transitory income ($\omega\in\mathbb{R}$). The magnitude of income risk depends on their age t, employment/filing status $e\in\{E,N\}$ (where E denotes employed, and N denotes non-employed/non-filer), and their age-dependent rank in the income distribution $r\in\{0,1,\ldots,100\}$ (where $r=Rank_t\left(z,\omega,e\right)$). The transition rate across employment/non-filer status is given by $p_t\left(e'|z,\omega,r,e\right)$. Lastly, we assume households have a deterministic lifecycle profile of income ($\kappa_t\in\mathbb{R}$), and we model taxes following Heathcote et al. (2017).

We summarize pre-tax income using the function $w_t(z,\omega,e)$,

$$w_t(z,\omega,e)=\begin{cases}\exp(\kappa_t+z+\omega) & \text{if } e=E\\ \gamma\exp(\kappa_t) & \text{if } e=N.\end{cases}$$

After-tax income $\widetilde{w}_t(z,\omega,e)$ is governed by progressivity parameter α and location parameter λ,

$$\widetilde{w}_t(z,\omega,e)=\lambda w_t(z,\omega,e)^{1-\alpha}.$$

Putting this notation together, the value function for an age t household with net worth b, permanent income z, transitory income ω in employment status e is

$$V_t\left(b,z,\omega,e\right)=\max_{b'\geq\underline{b}}u\left(c\right)+\sum_{e'}p\left(e'|z,\omega,r,e\right)\beta\mathbb{E}\left[V_{t+1}\left(b',z',\omega',e'\right)\right]\forall t\leq T$$

$$V_{T+1}\left(b,z,\omega,e\right)=\phi\left(b\right)$$

subject to the budget constraint and mapping from income to ranks,

$$c+b'\leq b\left(1+r_f\right)+\lambda w_t\left(z,\omega,e\right)^{1-\nu},\qquad r=Rank_t\left(z,\omega,e\right)$$

and the laws of motion for permanent and transitory income,

$$z' = \begin{cases} z + \nu_{E,r,t+1} & \text{if } e' = E, \\ z + \nu_{U,r,t+1} & \text{if } e' = N, \end{cases} \qquad \omega' = \begin{cases} \omega_{r,t+1} & \text{if } e' = E, \\ 0 & \text{if } e' = N, \end{cases}$$

$$\nu_{e,r,t+1} \sim N(B_{e,r,t+1},\, Q_{e,r,t+1})\,,\ \omega_{r,t+1} \sim N(0,\, R_{r,t+1})\,.$$

Equilibrium. A recursive competitive equilibrium is defined by a policy function for assets $b_t'(b, z, \epsilon, e)$, an equilibrium interest rate r, a tax level parameter λ and a distribution of agents across states $\Omega : \{E, N\} \times t \times b \times z \times \epsilon \to [0, 1]$ such that: (1) Given the interest rate, the household's policy functions are optimal; (2) the equilibrium interest rate clears the asset market; (3) Ω is consistent with household policy functions; and (4) the government budget constraint is balanced.

4.1 Calibration

We next discuss the calibration of the model.[46] Some parameters are assigned using estimates from the literature, while others are calibrated to be consistent with the U.S. labor market in the 1970s.

Demographics and preferences. To align with the sample in Section 1.1, agents enter the model at age 25 ($t = 1$) , and work until age 60 ($T = 8$), with a model frequency of 5 years. When agents enter the model, they are employed and draw a level of initial assets correlated with their initial income draws, matching regression estimates of log wealth on residual income for 20-25 year olds in the 1970s from the SCF+.[47]

Agents receive utility from consumption, with preferences given by,

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma},\ \phi(b) = \frac{b^{1-\sigma_b}}{1-\sigma_b}$$

We calibrate risk aversion σ, bequest curvature σ_b, and discount factor within the model, targeting the savings regressions as in Table 2 for all households as well as age 50-60 households,as well as the top 1% wealth share from the 1970s.[48]

[46] We solve the model using VFI on discrete grids.

[47] The regression estimates are $\log(\text{wealth}_{i,t}) = 9.44 + 0.729 \times \text{resid}_{i,t} + \varepsilon_{i,t}$, $\mathbb{V}[\varepsilon_{i,t} = 1.541^2]$. When a new individual enters, we draw their permanent and transitory income as well as a residual for this regression, assuming normality, and then assign the appropriate asset level.

[48] To compute the regression coefficients within the model, we simulate agents in the baseline model as well as

Income process. In the model, agents receive a wage that is a function of their age, permanent and transitory income. The age component of income is estimated as part of the residualization process in Section 2.

When agents are non-employed/non-filing they receive (pre-tax and transfers) a fraction $\gamma \in [0, 1]$ of their age profile income. The parameter γ can be thought as the replacement rate of unemployment insurance. As in Shimer (2005), we set $\gamma = 0.4$. We next discuss the estimation of the stochastic process that governs how income evolve in the model.

Shocks to income. Within the model, we implement our estimates of the income process described in Section 2. In this model, conditional distributions of shocks for agents in a given period depend upon their within-age income rank in the prior period for all periods other than the first period of their life. The sequential exogeneity timing, in addition to the fact that ranks are within-age, make the steady-state ranks of income as a function of permanent and transitory income easy to compute outside of all equilibrium conditions.[49] Hence, before VFI and market clearing are computed, the rank(z, e, age) function can be computed, and hence conditional expectations within VFI are straightforwardly taken.

Probability of unemployment. Similar to Braxton et al. (2025), we model the likelihood that an employed individuals becomes non-employed/non-filing as a non-linear function of their prior income rank r. In particular, we use the parameter estimates from estimating equation (6) to specify the following non-linear probability of being non-employed/non-filing based on prior income rank for those who are currently employed,

$$p(N|r, E) = \alpha_0 + \alpha_1 r + \alpha_2 r^2 + \alpha_3 r^3 \tag{6}$$

We include the parameters $\{\alpha_k\}_{k=0}^3$ in Table 3. The parameter estimates imply that the probability of transitioning from employed to non-employed/non-filer is a decreasing function of prior income rank. We model the probability of remaining non-employed/non-filing if you are currently remaining non-employed/non-filing to be a constant, i.e. $p(N|r, N) = \alpha_U$. We find that the share of individuals remaining in this state across 5-year observations is 27.5%.

under a perterbation to all Q_E's of 0.04 to approximately match the variation of Q_E from 1970s to present. Then, we estimate regression (5) with rank fixed effects, controlling for an age quadratic.

[49]Specifically, a large number of agents can be initialized at age 25. Then, the income can be ranked within age 25 agents, and distributions of age 30 permanent and transitory income computed for each agent given their age-25 ranks. The agents can then be simulated forward to age 30, ranked within age-30 agents, distributions of age-35 permanent and transitory income computed, and simulated forwards again to age 35. This process is repeated until all agents reach age 60. Then, for any permanent and transitory income level at any age, the level of income can be computed and the rank determined.

Table 3: Model parameters

Variable	Value	Description
r_f	$(1+3.5\%)^5-1$	Risk-free interest rate
α	0.181	Progessivity of tax function
λ	5.620	Location parameter of tax function
γ	0.4	Replacement Rate UI
α_U	0.275	Unemployed/nonfiler rate for unemployed/nonfiler
α_0	0.193	Added constant in unemployment probability for employed
α_1	$-.00484$	Linear term in unemployment probability for employed
α_2	$5.38e-05$	Square term in unemployment probability for employed
α_3	$-1.80e-07$	Cube term in unemployment probability for employed

Note: Table presents model parameters for the baseline estimation of the quantitative model.

Taxes. We use a progressive tax function as in Benabou (2002) and Heathcote et al. (2017). We set the tax progressivity (α) parameters to be equal to 0.181 based on the estimates in Heathcote et al. (2017), and $\lambda = 5.620$ to balance the government's budget constraint. Based on our parameterized tax function, agents with pre-tax incomes below approximately $\$10K$ receive transfers from the government, while individuals with pre-tax incomes greater than $\$10K$ pay labor income taxes.

Asset Markets. Agents are able to save and borrow at the annualized risk-free rate of 3.5%, which is the average value of the risk-free rate in the 1970s as measured by Laubach and Williams (2003). This interest rates results in assets being in positive net supply, and in the interest rate experiment in Section 5, we will hold this level of assets fixed and adjust the equilibrium interest rate to clear the asset market. We set the borrowing limit $\underline{b}$ to the natural borrowing limit.[50]

Table 3 presents the parameters that govern the model economy, while Table 4 shows the model fit of the calibrated parameters.[51]

4.2 Non-targeted moments

In this section, we present a series of non-targeted moments that serve as a model validation exercise.

[50] In our model, implementing the natural borrowing limit is equivalent to requiring that individuals die with zero debt.

[51] We calibrate an annual discount factor of 0.901. This estimate is within the range of estimates in recent consumer credit papers (e.g., Bornstein and Indarte (2023), Chatterjee et al. (2023), and Aguiar et al. (2024)).

Table 4: Parameters calibrated within the model

Variable	Value	Moment	Data Value	Model Value
σ	1.41	Savings Regression Coefficient	0.47	0.43
σ_b	1.15	Savings Regression Coefficient, 50-60	0.07	0.17
$\beta^{1/5}$	0.901	Top 1% Wealth Share, 1970s	23.0	18.0

Note: Table presents fitted parameters and targeted data moments. Parameters are curvature of utility and bequest functions as well as discount factor, and targets are the regression coefficients from estimating equation (5) via OLS for all households and for 50-60 year old households, as well as the top 1% wealth share in the 1970s. Moments are jointly determined within the model, related parameters are for intuition only.

Savings response to transitory risk. We start by examining the savings response of agents in the model to changes in transitory income risk. As in Section 4.1, we perform an experiment where we vary the transitory risk that households face, and replicate the empirical specification in equation (5) using model simulated data. Estimating this regression with variation in transitory risk, we find a savings elasticity of transitory risk that is equal to 0.163, which is 95% confidence interval of our empirical estimate presented in Table 2.

Wealth concentration. We next examine concentration of wealth predicted by the quantitative model and compare it to estimates from the SCF+ in 1970.[52] In the 1970s, the top 5% of the wealth distribution contained 49.4% of all wealth in the economy. In our quantitative model, the top 5% of the wealth distribution contains 44.4% of all wealth. In Table 5, we show that that the model is also able to produce estimates of the wealth share held by the top 20% of the distribution as well as the top 40%.

Table 5: Non-targeted moments: wealth concentration

Variable	(1) Model	(2) Data
Top 5 Percent Wealth Share, 1970s	44.4	49.4
Top 20 Percent Wealth Share, 1970s	79.2	76.9
Top 40 Percent Wealth Share, 1970s	97.7	93.3

Note: Table presents estimates of the share of wealth held across the wealth distribution in the quantitative model (column (1)) and in the 1970s SCF+ (column (2)).

[52]Note that while the model is calibrated to match the wealth share of the top 1% of the distribution, the moments presented in this section for the top 5%, 20% and 40% are untargeted in our calibration.

Table 6: Non-targeted moments: Tails of Wealth and Income, Top 0.1% Share to Top 1% Share

	(1) Baseline Model	(2) Flat Risk Profile	(3) Data (1970s)
Wealth	0.308	0.239	0.322
Income	0.24	0.181	0.254

Sources: This table reports the ratio of the top 0.1% share of a variable (wealth or income) to its top 1% share. As discussed in Hubmer et al. (2021), a higher value of this ratio implies a thicker tail. Column (1) reports estimates from the baseline model, while column (2) reports estimates from a version of the model with a common (flat) income risk profile across the income distribution. The 1970s data estimates in column (3) are from Saez and Zucman (2016) for wealth and Piketty and Saez (2003) for income.

Savings and income. A series of influential papers show that higher-income individuals save more (e.g., Caroll (2000), Dynan et al. (2004), Straub (2019), Mian et al. (2020), among others). Straub (2019) show that consumption is concave in permanent income. In particular, he estimates that a 1 log point increase in permanent income increases log consumption by 0.7 log points (s.e. = 0.05). Using model-simulated data, we find that a 1 log point increase in permanent income is associated with a 0.74 log point increase in consumption, which lies within the confidence interval reported by Straub (2019). Our model matches this empirical estimate because top earners face more income risk and therefore save more.

An implication of our results is that the model generates substantial wealth concentration at the very top of the distribution. To assess how well it captures upper-tail inequality, we compute the ratio of the top 0.1% share to the top 1% share for both wealth and income.[53] We then compare the model to the corresponding estimates in Saez and Zucman (2016) for wealth and Piketty and Saez (2003) for income. Table 6 shows that the model matches these features of the data quite well.[54] This success reflects the concentration of income risk at the top of the distribution. Because high-income households face substantial income risk in the model, they save at high rates, which in turn helps generate a thick upper tail of wealth. To isolate this mechanism, we also solve a version of the model with a common, or "flat," profile of income risk across the income distribution. Column (2) of Table 6 shows that, under this alternative specification, the top-tail ratios decline notably for both wealth and income. Elevated risk among top earners therefore plays an important role in the model's ability to match the observed thickness of the wealth distribution.

[53] As discussed in Hubmer et al. (2021), ratios of top wealth shares provide a measure of tail thickness.

[54] These findings are also consistent with Gaillard et al. (2023) and the references therein, which show that the wealth distribution has a much thicker tail than the income distribution.

5 Implications of rising income risk

In this section, we examine how changes in income risk across the distribution have impacted the risk-free rate and wealth inequality. In this experiment, we change the parameters of the income process from the baseline 1970s estimates to the 2010s estimates, and adjust the interest rates to keep the level of assets in positive net supply fixed compared to the baseline economy.[55] We focus on the impact of risk increases for the employed, both originating from permanent Q_E and transitory R risk, by income rank.

We start by considering all changes in income risk between the 2010s and the 1970s. In response to the increase in risk, especially in the right tail of the income distribution, we find that there is a 0.6pp decline in the risk-free rate (column (2) of Table 7). Thus, the changes in income risk (primarily for high-income individuals) have led to a substantial decline in the risk-free rate over the prior 40 years. Given that Laubach and Williams (2003) measure an average real interest rate of appoximatley 1% in the 2010s, we conclude that structural changes in income process dynamics can account for approximately 25% of the fall in real interest rates over the 1970-2019 period.

We next examine, which components of the change in income risk is driving the decline in the risk-free. We first consider how the increase in permanent risk among the employed impacts savings behavior and the risk-free rate. Panel (b) of Figure 4 presents the profile of shocks to permanent income among the employed across the income distribution in the 1970s economy (black, solid line) and the 2010s economy (red, dashed line). In response to the increase in permanent risk at the top of the distribution, savings increases and the risk-free rate declines substantially. In column (3) of Table 7, we show that to clear the asset market following the rise in permanent risk the risk-free rate must decline to an annualized value of 3.1%, 0.4pp lower than its 1970s value. Thus, increases in permanent income risk at the top of the distribution play a substantial role in decreasing the risk-free rate between the 1970s and the 2010s.

In column (4) of Table 7, we show that the increase in transitory income risk for the employed has had a modest impact on the risk-free rate. From the rise in transitory income risk, the annualized risk-free rate which clears the asset market drops by just over 0.1pp to 3.4%. This is perhaps unsurprising given that the rise in transitory risk is concentrated at the very top of the income distribution, but the savings response to transitory risk is effectively zero for these individuals given the results from Figure 8.

We additionally examine the impact of changes in income risk on wealth inequality by mea-

[55] Additionally to highlight the changes in risk (and not changes in mean income due to Jensen's inequality), we adjust the age income profile to keep the level of average income fixed across the estimations of the model.

Table 7: Impact of changes in income risk

Economy	(1) Baseline (1970*s*)	(2) 2010*s*	(3) $\Delta_t Q_E$	(4) $\Delta_t R$
Counterfactual alters only...		All Parameters	Perm. Income Risk Emp.	Temp. Income Risk
Risk-free rate	3.5%	2.9%	3.1%	3.4%
Change in r_f vs. baseline	–	-0.6pp	-0.4pp	-0.1pp
Top 5% Wealth Share	44.4%	46.7%	45.3%	44.5%
Change in share vs. baseline	–	2.3pp	0.9pp	0.1pp

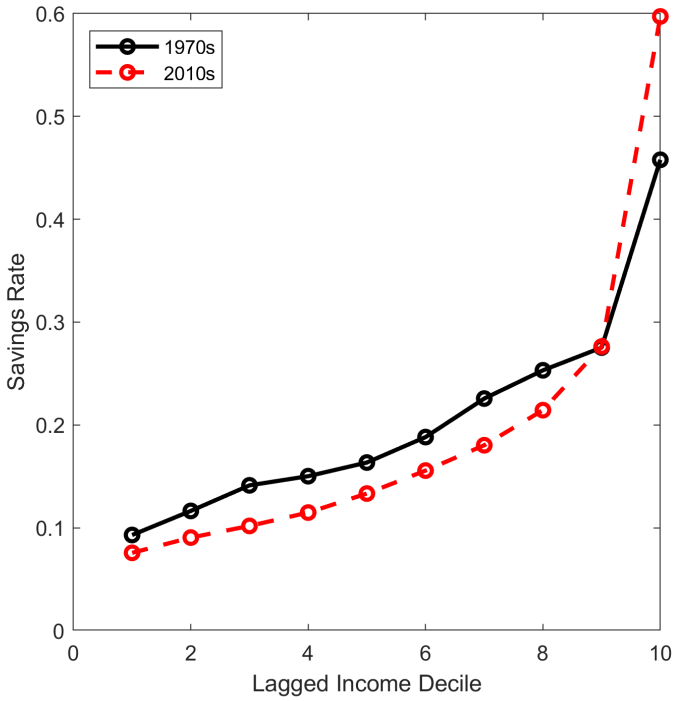
Note: Table presents the results of changing income risk across the distribution between the 1970s and 2010s. The first column corresponds to our 1970s estimation of the model. In columns (2)-(4), we adjust the profile of income risk for a given set of parameters (column (2)), or individual parameter (columns (3)-(4)), to their 2010s values, and adjust the risk-free rate to clear the asset market.

suring changes in the top 5% wealth share across steady states of the quantitative model.[56] In column (2) of Table 7, we show that the changes in income risk we document between the 1970s and 2010s have increased the share of wealth held by the top 5% of the wealth distribution by 2.3 percentage points. Decomposing the source of the increase in the top 5% wealth share, nearly half of the increase has come from changes in permanent income risk. Thus, the increase in permanent income risk we document has led to an increase in wealth inequality.

Changes in savings behavior across the income distribution. To better understand the rise in wealth inequality, we examine how savings behavior evolves across the income distribution in the quantitative model, as shown in Figure 10. The black solid line, with circle markers, shows that in the baseline version of the model, high-income households have higher saving rates, consistent with the empirical evidence in Caroll (2000), Dynan et al. (2004), and Straub (2019). As we move to the 2010s, income risk rises across the distribution. In partial equilibrium, this increase in risk would raise desired precautionary savings for all households. In general equilibrium, however, the asset market clears through a decline in the risk-free rate, which weakens incentives to save. The red dashed line in Figure 10 shows that, in equilibrium, households below the top decile reduce their saving rates relative to the 1970s, while households in the top decile increase their savings sharply. Thus, for most households the decline in the interest rate more than offsets the increase in precautionary saving motives, whereas for top-income households the rise in income risk is strong enough to dominate the equilibrium price effect. Our

[56] We find similar results using different measures of wealth inequality.

Figure 10: Changes in Savings Behavior by Income



Note: Figure presents the profile of savings as a function of income decile in the baseline model with the 1970s level of income risk (black, solid line with circle markers) and the model with the 2010s level of income risk (red, dashed lien with circle markers).

model therefore predicts that rising income risk increases savings at the top of the income distribution while reducing saving rates for the rest of the distribution, a pattern that is consistent with the recent empirical evidence in Mian et al. (2020).

Isolating the role of income risk at the top. To isolate the role of risk among top earners, we repeat our previous exercise only allowing income process parameters from p90 and above to change between the 1970s and 2010s. For all other percentiles of the income distribution, income risk is held fixed at 1970s levels. Table 8 reports the results. We find that roughly 40% of the model's predicted reduction in interest rates can be attributed to changing income risk at the top. Additionally, effectively all of the increase in the top 5% wealth share has come from the changes in income risk at the top of the distribution. These results suggest that income risk among top earners is an important cause of lower interest rates in the U.S. as well as the evolution of wealth inequality.

Table 8: Changes in income risk among top 10% of the distribution

Change in:	(1) Risk Free Rate (r_f)	(2) Top 5% Wealth Share
Shift all risk p1-p100	-0.58pp	2.4pp
Shift only risk for p90-p100	-0.22pp	2.4pp
Percent due to p90+ risk	39%	100%

Note: Table presents the results of a quantitative experiment in which we only allow income process parameters to evolve for those in the top 10% of the income distribution.

6 Conclusion

We contribute to two ongoing debates over the evolution of income risk in the U.S. (Moffitt et al. (2022), Salgado et al. (2023), Braxton et al. (2025)) and the causes and consequences of rising saving rates among the richest households (Mian et al. (2020) and Mian et al. (2021a)). The measurement component of our paper uses newly digitized and linked Census-IRS tax returns from 1969 to 2019 (Alexander et al. (2024)). These historical 100-percent IRS files were not retained by the IRS, but rather rediscovered on magnetic tapes at the Census Bureau, digitized and then subsequently assigned Protected Identification Keys (PIKs) in order to facilitate linkages over time and across datasets. These data provide a unique 50-year panel of business and labor income, asset income, and demographic Census data.

Unlike existing U.S. income risk estimates that assume homoskedasticity, we flexibly filter the data in a way that allows for differential trends in permanent income risk across the income distribution. We find that permanent income risk among high-income workers (i.e., those in the 95th percentile of the income distribution and above) rose by over 67% over our sample window. Because our Kalman filter yields a panel of permanent and transitory shocks for every individual, we validate that top earners indeed face "risk" by documenting that negative permanent shocks coincide with observable proxies for adverse life events, changes in asset positions, and financial distress, including bankruptcy and foreclosure.

We then capitalize the interest and dividend income on the tax returns, and we show that high-income individuals respond to increasing permanent risk by significantly increasing their savings. We attempt to establish a causal link between savings behavior and income risk by exploiting within-income-rank, across-occupation variation in risk profiles. Conditional on income rank, we find that savings-to-income ratios rise by 6.6% of prior income for a 0.1 increase in permanent income shock variance.

Lastly, we embed our income process into a Bewley-Huggett-Aiyagari model, and we find

that rising permanent income risk at the top of the income distribution has put downward pressure on interest rates and increased wealth inequality. Moreover, rising permanent income risk at the top of the income distribution provides an alternate rationale for increased savings rates among the richest U.S. households, rather than attributing it to non-homothetic preferences.

References

Aguiar, M., M. Bils, and C. Boar (2024). Who are the hand-to-mouth? *Review of Economic Studies*, rdae056.

Aladangady, A., E. Gagnon, B. K. Johannsen, and W. Peterman (2021). Macroeconomic implications of inequality and income risk.

Alexander, J. T., D. Bleckley, J. Fisher, K. Genadek, S. H. Leonard, and A. Magganas (2024). Newly available individual-level us tax data from 1969-1994.

Altonji, J. G., A. A. Smith Jr, and I. Vidangos (2013). Modeling earnings dynamics. *Econometrica 81*(4), 1395–1454.

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

Auclert, A., H. Malmberg, F. Martenet, and M. Rognlie (2021). Demographics, wealth, and global imbalances in the twenty-first century. Technical report, National Bureau of Economic Research.

Auclert, A. and M. Rognlie (2018). Inequality and aggregate demand. Technical report, National Bureau of Economic Research.

Autor, D. H. and D. Dorn (2013). The growth of low-skill service jobs and the polarization of the us labor market. *American economic review 103*(5), 1553–1597.

Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica 70*(2), 481–517.

Blundell, R., L. Pistaferri, and I. Preston (2008). Consumption inequality and partial insurance. *The American Economic Review*, 1887–1921.

Boar, C. (2021). Dynastic precautionary savings. *The Review of Economic Studies 88*(6), 2735–2765.

Bornstein, G. and S. Indarte (2023). The impact of social insurance on household debt.

Bowler, M. (2006). Understanding the employment measures from the cps and ces survey. *Monthly Lab. Rev. 129*, 23.

Braxton, J. C., K. Herkenhoff, J. Rothbaum, and L. Schmidt (2025). Changing income risk across the us skill distribution: Evidence from a generalized kalman filter. *American Economic Review 115*(12), 4438–4475.

Braxton, J. C., K. F. Herkenhoff, C. Huang, M. Nattinger, and L. Schmidt (2025). Tail risk and bankruptcy in the united states. *Manuscript*.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2008). An equilibrium model of" global imbalances" and low interest rates. *American Economic Review*.

Caballero, R. J., E. Farhi, and P.-O. Gourinchas (2017). The safe assets shortage conundrum. *Journal of economic perspectives 31*(3), 29–46.

Caroll, C. (2000). Why do the rich save so much? does atlas shrug? the economic consequences of taxing the rich (pp. 465-484). ed. by joel b. slemrod.

Carroll, C. D. (1997). Buffer-stock saving and the life cycle/permanent income hypothesis. *The Quarterly journal of economics 112*(1), 1–55.

Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of monetary Economics 40*(1), 41–71.

Chatterjee, S., D. Corbae, K. Dempsey, and J.-V. Ríos-Rull (2023). A quantitative theory of the credit score. *Econometrica 91*(5), 1803–1840.

Choukhmane, T., J. Colmenares, C. O'Dea, J. L. Rothbaum, and L. D. Schmidt (2024). Who benefits from retirement saving incentives in the us? evidence on gaps in retirement wealth accumulation by race and parental income. Technical report, National Bureau of Economic Research.

De Nardi, M. and G. Fella (2017). Saving and wealth inequality. *Review of Economic Dynamics 26*, 280–300.

De Nardi, M., G. Fella, and G. Paz-Pardo (2020). Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association 18*(2), 890–926.

Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological) 39*(1), 1–22.

Dynan, K. E., J. Skinner, and S. P. Zeldes (2004). Do the rich save more? *Journal of political economy 112*(2), 397–444.

Farhi, E. and F. Gourio (2018). Accounting for macro-finance trends: Market power, intangibles, and risk premia. Technical report, National Bureau of Economic Research.

Favilukis, J. (2013). Inequality, stock market participation, and the equity premium. *Journal of Financial Economics 107*(3), 740–759.

Gaillard, A., C. Hellwig, P. Wangner, and N. Werquin (2023). Consumption, wealth, and income inequality: A tale of tails.

Gourinchas, P.-O. and J. A. Parker (2002). Consumption over the life cycle. *Econometrica 70*(1), 47–89.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Guvenen, F., L. Pistaferri, and G. L. Violante (2022). Global trends in income inequality and income dynamics: New insights from grid. *Quantitative Economics 13*(4), 1321–1360.

Hamilton, J. D. (1994). State-space models. *Handbook of econometrics 4*, 3039–3080.

Heathcote, J., K. Storesletten, and G. L. Violante (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics 132*(4), 1693–1754.

Hubmer, J., P. Krusell, and A. A. Smith Jr (2021). Sources of us wealth inequality: Past, present, and future. *Nber macroeconomics annual 35*(1), 391–455.

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.

Laubach, T. and J. C. Williams (2003). Measuring the natural rate of interest. *Review of Economics and Statistics 85*(4), 1063–1070.

Mian, A., L. Straub, and A. Sufi (2021a). Indebted demand. *The Quarterly Journal of Economics 136*(4), 2243–2307.

Mian, A. R., L. Straub, and A. Sufi (2020). The saving glut of the rich. Technical report, National Bureau of Economic Research.

Mian, A. R., L. Straub, and A. Sufi (2021b). What explains the decline in r*? rising income inequality versus demographic shifts. *Rising Income Inequality Versus Demographic Shifts (September 22, 2021). University of Chicago, Becker Friedman Institute for Economics Working Paper* (2021-104).

Moffitt, R., J. Abowd, C. Bollinger, M. Carr, C. Hokayem, K. McKinney, E. Wiemers, S. Zhang, and J. Ziliak (2022). Reconciling trends in us male earnings volatility: Results from survey and administrative data. *Journal of Business & Economic Statistics 41*(1), 1–11.

Nakamura, E. and J. Steinsson (2018). Identification in macroeconomics. *Journal of Economic Perspectives 32*(3), 59–86.

Piketty, T. and E. Saez (2003). Income inequality in the united states, 1913–1998. *The Quarterly journal of economics 118*(1), 1–41.

Saez, E. and G. Zucman (2016). Wealth inequality in the united states since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics 131*(2), 519–578.

Salgado, S., N. Bloom, F. Guvenen, L. Pistaferri, and J. Sabelhaus (2023). Evaluating the great micro moderation. *Available at SSRN 4585714*.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American economic review*, 25–49.

Shumway, R. H. and D. S. Stoffer (1982). An approach to time series smoothing and forecasting using the em algorithm. *Journal of time series analysis 3*(4), 253–264.

Smith, M., O. Zidar, and E. Zwick (2023). Top wealth in america: New estimates under heterogeneous returns. *The Quarterly Journal of Economics 138*(1), 515–573.

Straub, L. (2019). Consumption, savings, and the distribution of permanent income. *Unpublished manuscript, Harvard University*.

Ziliak, J. P., C. Hokayem, and C. R. Bollinger (2022). Trends in earnings volatility using linked administrative and survey data. *Journal of Business & Economic Statistics 41*(1), 12–19.

A Additional details: Estimation

In this appendix, we provide a series of additional details and results for our estimation routine. We first provide additional details about the Kalman Filtering Step of the estimation in Appendix A.1.

A.1 Kalman Filter

Our derivations of the Kalman filter and smoother are primarily based on Hamilton (1994). For convenience in computing parameter updates in the EM algorithm, we modify equation (2) so that the state variable includes both the current realization of permanent income and its lag. Let $\zeta_{i,t}$ denote an household i's unobserved state in period t:

$$\zeta_{i,t} = \begin{bmatrix} z_{i,t} \\ z_{i,t-1} \end{bmatrix},$$

For casting the income process from equations (1)-(3) into state-space form, we make the following adjustments to notation to facilitate running the Kalman filter for each household i. We let $B_t(l_{i,t}, r_{i,t-1})$ denote the drift for a household i in period t who has employment status $l_{i,t}$ and has prior rank $r_{i,t-1}$. We analogously define $R_t(l_{i,t}, r_{i,t-1})$ and $Q_t(l_{i,t}, r_{i,t-1})$.

State equation. From equation (2), the law of motion for the state vector $\zeta_{i,t}$ (i.e., the *state equation*) can be rewritten as,

$$\zeta_{i,t} = \begin{bmatrix} z_{i,t} \\ z_{i,t-1} \end{bmatrix} = \underbrace{\begin{bmatrix} B_t(l_{i,t}, r_{i,t-1}) \\ 0 \end{bmatrix}}_{\widehat{B}_t(l_{i,t}, r_{i,t-1})} + \underbrace{\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}}_{\widehat{F}} \zeta_{i,t-1} + \begin{bmatrix} \nu_{i,t} \\ 0 \end{bmatrix}. \tag{7}$$

Measurement equation. Using the definition of the state vector $\zeta_{i,t}$, the law of motion for income in equation (1) can be rewritten using an observation matrix $H(l_{i,t})$, where $H(l_{i,t}) = \begin{bmatrix} l_{E,i,t} & 0 \end{bmatrix}$ governs the relationship between the state vector $(\zeta_{i,t})$ and income $y_{i,t}$ among households who are employed $(l_{E,i,t} = 1)$.[57] The resulting law of motion is our *measurement equation*,

$$y_{i,t} = H(l_{i,t})\zeta_{i,t} + l_{E,i,t}\, \omega_{i,t}, \tag{8}$$

[57] When a household is non-employed/non-filing ($l_{E,i,t} = 0$), the observation $y_{i,t}$ provides no signal about latent income other than what can be inferred from other observables. In those instances, the Kalman filter does not use $y_{i,t}$ to update its guess about $z_{i,t}$. $H(l_{i,t})$ captures that intuition.

We apply the Kalman filter and smoother to the state-space system defined by equations (7) and (8). In the Kalman filtering step of the estimation, we assume the parameters of the income process are known. In the EM step of our estimation routine, which we discuss in Appendix A.2, we estimate the parameters that govern the income process. Starting with estimates $\widehat{\zeta}_{i,1|0}$ and $M_{i,1|0}$, which we will define below, we obtain an estimate of the time series of permanent income (and its lag) from the Kalman filter, which we denote $\widehat{\zeta}_{i,t|t}$, as follows:[58]

1. Estimate the Kalman Gain:

$$K_{i,t} = M_{i,t|t-1}H^{'}(l_{i,t})\left[H(l_{i,t})M_{i,t|t-1}H^{'}(l_{i,t}) + R_t(l_{i,t}, r_{i,t-1})\right]^{-1}.$$

2. Update the state vector:

$$\begin{aligned}\widehat{\zeta}_{i,t|t} &= \widehat{\zeta}_{i,t|t-1} + K_{i,t}\left(y_{i,t} - H(l_{i,t})\widehat{\zeta}_{i,t|t-1}\right)\\ \widehat{\zeta}_{i,t+1|t} &= \widehat{F}\widehat{\zeta}_{i,t|t} + \widehat{B}_{t+1}(l_{i,t+1}, r_{i,t}).\end{aligned}$$

3. Update the MSE matrix:

$$\begin{aligned}M_{i,t|t} &= M_{i,t|t-1} - K_{i,t}H(l_{i,t})M_{i,t|t-1}\\ M_{i,t+1|t} &= \widehat{F}M_{i,t|t}\widehat{F}' + Q_{t+1}(l_{i,t+1}, r_{i,t})e_1^2 e_1^{2'}.\end{aligned}$$

where $e_1^2 = [1, 0]'$. Repeat steps 1-3 for $t = 2, ...T$, and for each household $i \in \{1, .., N\}$. Finally, to run the Kalman filter we need an estimate of the mean of the state-vector ($\widehat{\zeta}_{i,1|0}$) and the MSE matrix ($M_{i,1|0}$). When an household enters into the labor market, we set their lagged rank to zero and thus we set $\widehat{\zeta}_{i,1|0} = \widehat{B}_1(l_{i,1}, 0)$, and we initialize the variance of the state vector using,[59]

$$M_{i,1|0} = \begin{bmatrix} Q_0 + Q(l_{i,1}, 0) & Q_0 \\ Q_0 & Q_0 \end{bmatrix}.$$

Kalman Smoother In order to apply the EM algorithm below, we must apply the Kalman smoother after running the Kalman filter. We denote the resulting posterior estimates from running the Kalman smoother step as $\{\{\widehat{\zeta}_{i,t|T}\}_{t=0}^{T}\}_{i=1}^{N}$ and $\{\{M_{i,t|T}\}_{t=0}^{T}\}_{i=1}^{N}$. The steps for the smoothed Kalman filter are:

[58] We use $\widehat{\zeta}_{i,t|t-1}$ to denote the estimate of permanent income for household i in period t using data through period $t-1$. Similarly, we let $M_{i,t|t-1}$ denote the uncertainty in this estimate.

[59] For a households first observation in the labor market, we need to specify a rank for drawing their permanent and transitory shock. We give them a "rank" of zero. Our results are robust to this choice.

1. Run the Kalman filter as presented above storing the sequences $\{M_{i,t|t-1}\}_{t=1}^{T}$ and $\{M_{i,t|t}\}_{t=1}^{T}$ as well as $\{\hat{\zeta}_{i,t|t-1}\}_{t=1}^{T}$ and $\{\hat{\zeta}_{i,t|t}\}_{t=1}^{T}$.

2. Store the element $\hat{\zeta}_{i,T|T}$ from $\{\hat{\zeta}_{i,t|t}\}_{t=1}^{T}$.

3. Calculate the sequence of smoothed estimations $\{\hat{\zeta}_{i,t|T}\}_{t=1}^{T-1}$ in reverse order by iterating on:

$$\hat{\zeta}_{i,t|T} = \hat{\zeta}_{i,t|t} + J_{i,t}(\hat{\zeta}_{i,t+1|T} - \hat{\zeta}_{i,t+1|t})$$

   for $t = T-1, T-2, ..., 1$, where $J_{i,t} = M_{i,t|t}\hat{F}' M_{i,t+1|t}^{-1}$.

4. Update the sequence of MSE by iterating on:

$$M_{i,t|T} = M_{i,t|t} + J_{i,t}(M_{i,t+1|T} - M_{i,t+1|t})J_{i,t}'.$$

A.2 Expectation Maximization (EM) Algorithm.

In this appendix we discuss how we use an EM algorithm to tractably update the parameters of our income process. To derive the EM updating formulas, the first step is to write down the full-information log likelihood assuming that state-variables (i.e., $z_{i,t}$) are observed:

$$\begin{aligned} LL_i(\{y_{i,t}\}_{t=1}^{T}, \{z_{i,t}\}_{t=1}^{T} \mid \{l_{i,t}\}_{t=1}^{T}, \theta_0) = & -\frac{T+1}{2}\log(2\pi) \\ & - \frac{1}{2}\log(Q_0) - \frac{1}{2}\frac{(z_{i,0})^2}{Q_0} \\ & - \frac{1}{2}\sum_{t=1}^{T}\log(Q_t(l_{i,t}, r_{i,t-1})) - \frac{1}{2}\sum_{t=1}^{T}\frac{(z_{i,t} - z_{i,t-1} - B_t(l_{i,t}, r_{i,t-1}))^2}{Q_t(l_{i,t}, r_{i,t-1})} \\ & - \frac{1}{2}\sum_{t=1}^{T}\log(R_t(l_{i,t}, r_{i,t-1})) - \frac{1}{2}\sum_{t=1}^{T}\frac{(l_{E,i,t})(y_{i,t} - z_{i,t})^2}{R_t(l_{i,t}, r_{i,t-1})}, \end{aligned} \tag{9}$$

where the parameter vector is denoted $\theta_0 = \{Q_t(\cdot), R_t(\cdot), B_t(\cdot), Q_0\}$.

E-Step. The second step is to take the expectation of the likelihood with respect to $\{z_{i,t}\}$, given $\{y_{i,t}\}_{t=1}^{T}$:

$$\mathcal{L}(\theta_0) = E_{\{z_{i,t}\}_{t=1}^{T}}\left[LL_i(\{y_{i,t}\}_{t=1}^{T}, \{z_{i,t}\}_{t=1}^{T} \mid \{l_{i,t}\}_{t=1}^{T}, \theta_0)\middle| \{y_{i,t}\}_{t=1}^{T}\right]$$

This amounts to integrating equation (9) with respect to the distribution of the latent states $F(\{z_{i,t}\} \mid \{l_{i,t}, y_{i,t}\}_{t=1}^{T}, \theta_0)$ to express the likelihood in terms of posterior means and covariances.

The Kalman filter yields $F(\{z_{i,t}\} \mid \{l_{i,t}, y_{i,t}\}_{t=1}^{T}, \theta_0)$, making this integration simple.

Why is the integration with respect to z simple? Shumway and Stoffer (1982) show, this expectation yields the likelihood function in terms of $\{\widehat{\zeta}_{i,t|T}\}_{t=1}^{T}$ and $\{M_{i,t|T}\}_{t=1}^{T}$, both of which are estimated by the Kalman filter and smoother.[60] In other words, when we take the expectation of the likelihood given observed data $\{y_{i,t}\}_{t=1}^{T}$, the unobserved latent states $z_{i,t}$ become simple posteriors from the Kalman filter, $\widehat{z}_{i,t|T} = E[z_{i,t} \mid \{l_{i,t}, y_{i,t}\}_{t=1}^{T}, \theta_0]$, and covariances are given by $\{M_{i,t|T}\}_{t=1}^{T}$.

M-step. The third and final step is to take first order conditions of $\mathcal{L}(\cdot)$ with respect to the income process parameter vector, θ_0.[61] Consider the drift $B_t(\cdot)$. One may intuit from the quadratic functional form of the log likelihood that the resulting first order conditions will resemble ordinary least squares regressions formulas. The intuition is correct, except the updating formulas are generalized least squares regression formulas with weights that adjust for filtering uncertainty. The intuition comes from taking the expectation of equation (2) with respect to $\{z_{i,t}\}$, given $\{y_{i,t}\}_{t=1}^{T}$, which naturally suggests that B can be updated via regression:

$$\widehat{z}_{i,t|T} = \widehat{z}_{i,t-1|T} + B_t(l_{i,t}, r_{i,t-1})$$

We update the variances by taking the first order condition of $\mathcal{L}(\cdot)$ with respect to $Q_t(\cdot)$ and $R_t(\cdot)$, and then solving the first order conditions via Newton's method. We restrict the variance terms to be log-linear in parameters, yielding a well-defined concave optimization problem for $Q_t(\cdot)$ and $R_t(\cdot)$.

Algorithm. Our algorithm can be summarized in a few short steps:

1. Guess parameters $\theta^{(0)}$
2. Run the Kalman filter and smoother using $\theta^{(0)}$ to obtain $\{\widehat{\zeta}^{(0)}_{i,t|T}\}_{t=1}^{T}$ and $\{M^{(0)}_{i,t|T}\}_{t=1}^{T}$.
3. **E-step:** Write down the complete information likelihood (9). Integrate with respect to the unobservable distribution (which comes for "free" from the Kalman smoother). This yields $\mathcal{L}(\theta^{(0)})$ in terms of $\{\widehat{\zeta}^{(0)}_{i,t|T}\}_{t=1}^{T}$ and $\{M^{(0)}_{i,t|T}\}_{t=1}^{T}$.

[60] Recall that $\widehat{\zeta}_{i,t|T} = \begin{bmatrix} \widehat{z}_{i,t|T} & \widehat{z}_{i,t-1|T} \end{bmatrix}$.

[61] As we discussed in Section 2.1, we separate the income process parameters θ_0 from the parameters that govern the employment process. Section 1 and Appendix A of Braxton et al. (2025) have a detailed discussion on the statistical assumptions that are needed to separate the income process parameters and employment process parameters into two distinct groups.

4. **M-step:** Maximize $\mathcal{L}(\theta^{(0)})$ with respect to $\theta^{(0)}$. $\mathcal{L}(\theta^{(0)})$ is concave, so first order conditions are necessary and sufficient and yield simple updating formulas. This yields the new guess $\theta^{(0)} \to \theta^{(1)}$.

5. Run the Kalman filter and smoother again to obtain the new unobservable distribution $\{\widehat{\zeta}^{(1)}_{i,t|T}\}_{t=1}^{T}$ and $\{M^{(1)}_{i,t|T}\}_{t=1}^{T}$.

6. Repeat until $|\theta^{(N)} - \theta^{(N+1)}| < \epsilon$, where $\epsilon > 0$ is a prespecified tolerance.

A.3 Monte Carlo Evidence

In this appendix, we use a series of Monte Carlo exercises to establish identification and to show how simplier GMM style procedures fail to work in this setting because of the dynamic updating of an individuals lagged income rank. Finally, we show that our estimation routine is able to accurately recover individual level estimates of permanent income as well as its changes.

Identification. We start by numerically establishing identification as in Altonji et al. (2013). For this exercise, we hypothesize a set of parameters and simulate a series of income histories. We then apply our estimation routine on this simulated data, and show that the parameters which maximize the likelihood are close to the hypothesized parameters used to simulate the data. The parameters that we hypothesize are our estimated parameter from filtering our baseline income definition, AGI less interest and dividend income, on our baseline sample of 1040 income histories for individuals who have been in the CPS-ASEC. [62] Simulating data for our income process also requires that we take a stance on the statistical process for unemployment. As in our quantitative model in Section 4, we assume that the unemployment probability is a cubic function of an agent's lagged earning rank when employed, but is a constant probability when unemployed.[63] We simulate $N = 25,000$ individuals for $T = 8$ periods, the number of periods between 25 and 60, inclusively. We simulate $D = 40$ Monte Carlo draws of this experiment, and Figure 11 presents the results.

In panel (a) of Figure 11 we present the estimates of the variance of the shocks to permanent income among the employed as a function of lagged income rank from the MC exercise. The black, solid line corresponds to the true estimate that is used to simulate data in the MC exercise. The blue dashed line is our mean estimate across the 40 MC simulations, and the red dashed lines correspond to a 95 percent simulation interval. The figure shows that the true

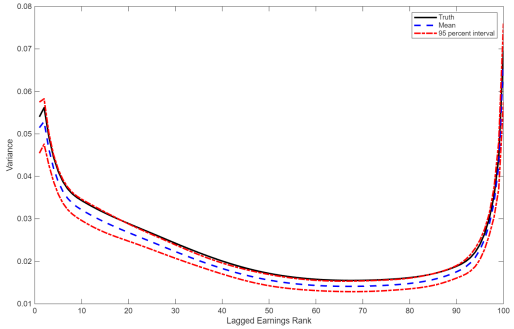
[62] Additionally, in the MC exercise, we do not consider time-variation and simulate data corresponding to the 1970s parameter estimates.

[63] In particular, we estimate equation (6) and the results are presented in Table 3.

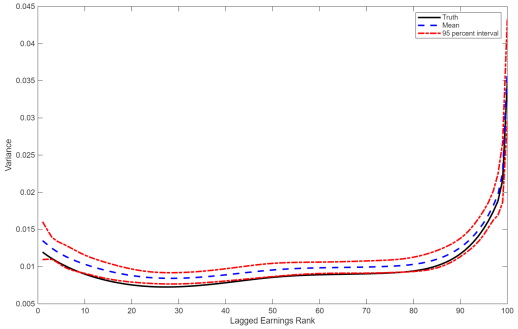
parameter values are within our simulation interval for nearly the entire distribution. The remaining panels of Figure 11 show that we obtain similar results for the variance of transitory shocks, the variance of permanent shocks to the non-employed (and non-filing) as well as the mean of permanent shocks and the mean of permanent shocks for the non-employed/non-filing. These results highlight that our estimation procedure is able to accurately recover the nature of income shocks across the income distribution.

Figure 11: Monte Carlo Results: Recovering Parameters by Lagged Income Rank

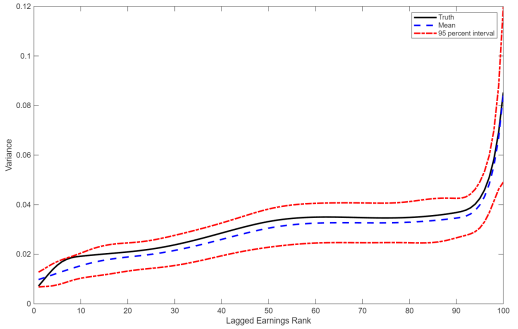
(a) Variance Permanent Shock



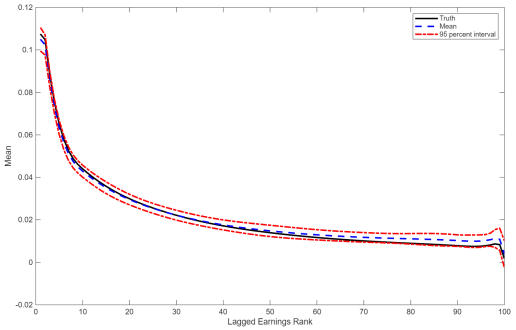
(b) Variance Transitory Shock



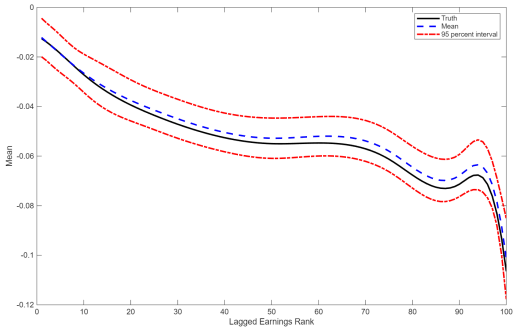
(c) Variance of Permanent Shock, Non-employed



(d) Mean Permanent Shock



(e) Mean Permanent Shock, Non-employed



Note: The figure presents the results of the Monte Carlo exercise in which we establish local identification of our model. In each panel, the black, solid line corresponds to our estimates from the 1970s for AGI less interest and dividend income, our assumed data generating process, while the blue dashed line represents the mean estimate across Monte Carlo draws from estimating our model's parameters on the simulated data. The red dash-dotted lines are the 95% simulation intervals of the point estimate across Monte Carlo draws. Panel (a) presents permanent income risk, panel (b) presents transitory income risk, panel (c) presents permanent income risk for individuals who are classified as non-employed / non-filers, panel (d) presents the mean of permanent shocks, and panel (e) presents the mean of permanent shocks for individuals who are classified as non-employed/non-filers.

GMM style estimation. We next examine how a GMM style estimation, which uses functions of means and variances for identification would perform in recovering the parameters of our income process. For intuition for how this type of GMM style procedure works, let us first consider a simpler model where the parameters do not vary by lagged income rank. Specifically, consider the following:

$$\begin{aligned} y_{i,t} &= \begin{cases} z_{i,t} + \omega_{i,t} & \text{if } l_{E,i,t} = 1 \\ . & \text{if } l_{E,i,t} = 0 \end{cases} \\ \omega_{i,t} &\sim N\left(0, R\right), \end{aligned} \tag{10}$$

where $\omega_{i,t}$ is independent of $z_{i,t}$ conditional on $l_{i,t}$ and $z_{i,t}$ evolves according to:

$$\begin{aligned} z_{i,t+1} &= z_{i,t} + B(l_{i,t+1}) + \nu_{i,t+1} \\ \nu_{i,t+1} &\sim N\left(0, Q(l_{i,t+1})\right) \\ z_{i,0} &\sim N(0, Q_0(y_0)). \end{aligned} \tag{11}$$

It is straightforward to verify that the parameters from the simple process in (10) and (11) are identified from simple moments using the following formulas:[64]

$$\begin{aligned} B(E) &= \mathbb{E}[y_{i,t+1} - y_{i,t} | l_{i,t} = l_{i,t+1} = E] \\ B(U) &= \mathbb{E}[y_{i,t+2} - y_{i,t} | l_{i,t} = l_{i,t+2} = E, l_{i,t+1} = U] - \mathbb{E}[y_{i,t+1} - y_{i,t} | l_{i,t} = l_{i,t+1} = E] \\ Q(E) &= \mathbb{V}[y_{i,t+2} - y_{i,t} | l_{i,t} = l_{i,t+1} = l_{i,t+2} = E] - \mathbb{V}[y_{i,t+1} - y_{i,t} | l_{i,t} = l_{i,t+1} = E] \\ Q(U) &= \mathbb{V}[y_{i,t+2} - y_{i,t} | l_{i,t} = l_{i,t+2} = E, l_{i,t+1} = U] - \mathbb{V}[y_{i,t+1} - y_{i,t} | l_{i,t} = l_{i,t+1} = E] \\ R &= (2\mathbb{V}[y_{i,t+1} - y_{i,t} | l_{i,t} = l_{i,t+1} = E] - \mathbb{V}[y_{i,t+2} - y_{i,t} | l_{i,t} = l_{i,t+1} = l_{i,t+2} = E]])/2 \end{aligned} \tag{12}$$

One could then condition these formulas on the date-t income rank to arrive at an estimate of the parameters of the income process as a function of an individuals lagged income. We will refer to this estimation procedure as the GMM estimation approach as well as the "simple" approach. Figure 12 presents the results of applying the "simple" approach on our simulated data from the MC exercise. In the figure, the green, solid line with circle markers, reports the mean estimate of the parameters across MC simulations while the red-dashed line reports a 95 percent simulation interval.

Panel (a) of Figure 12 shows that the "simple" approach is not able to recover the parameters of the income process for the bottom 30 percent of the distribution. While the data generating

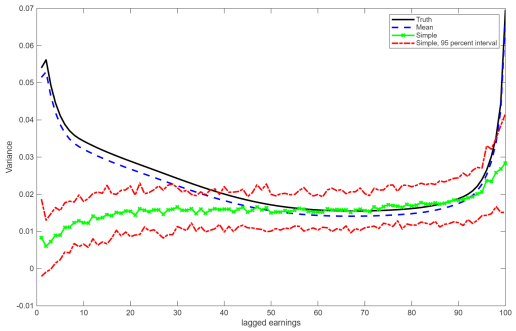
[64] See Appendix B.1 of Braxton et al. (2025).

process (DGP) exhibits declining variance of permanent shocks as an individual moves from the lowest ranks towards the 30th percentile, the simple approach fails to capture this pattern and recovers a slightly increasing pattern. Additionally, the simple approach only records a modest uptick in the variance of shocks to permanent income in the far right tail, while the DGP exhibits a much sharper increase.

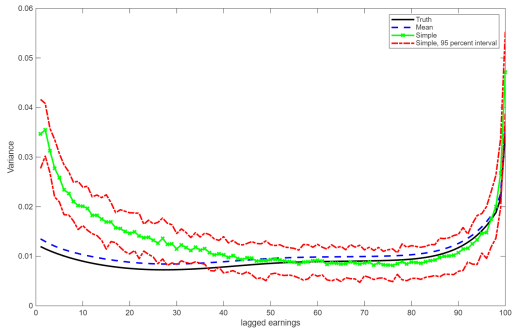
Figure (b) of Figure 12 shows the simple approach incorrectly infers a very high level of transitory risk for the lowest income households in the economy and that the profile of transition risk decreases sharply with increasing income. Thus from panels (a) and (b), we observe that the simple approach struggles to decipher permanent versus transitory risk at the bottom of the distribution. Our estimation procedure does not exhibit these problems in recovering the DGP at the bottom of the income distribution. Our procedure is able to dynamically update an individual's rank while the GMM style approach does not, resulting in the GMM style approach failing to recover the true DGP while our procedure is successful.

Figure 12: Monte Carlo Results: Recovering Parameter Using a GMM Approach

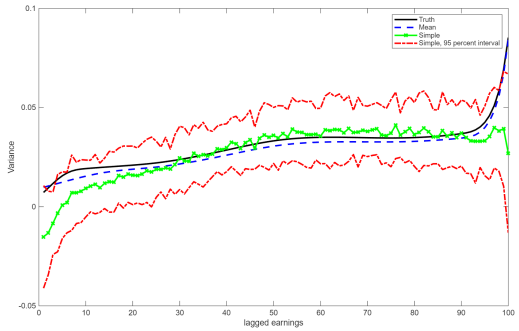
(a) Permanent Risk



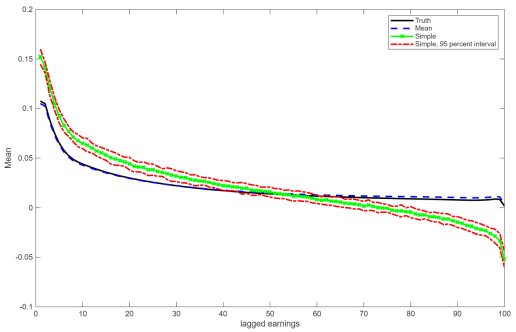
(b) Transitory Risk



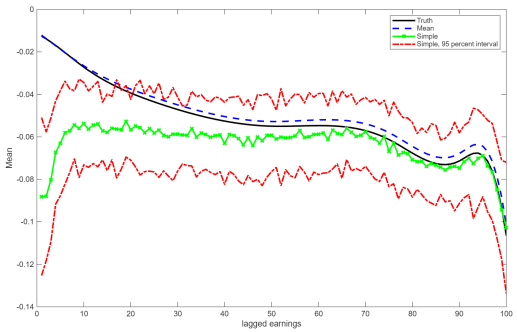
(c) Permanent Risk, Non-employed/non-filer



(d) Mean Permanent Shock



(e) Mean Permanent Shock, Non-employed/non-filer



Note: The figure presents the results of the Monte Carlo exercise in which we examine the ability of a GMM style approach to estimate our income process. In each panel, the black, solid line corresponds to our estimates from the 1970s for AGI less interest and dividend income, our assumed data generating process, while the blue dashed line represents the mean estimate across Monte Carlo draws from estimating our model's parameters on the simulated data. The green-x line is the mean estimate of the simple model computed using equations (12) while the red dash-dotted lines are the 95% simulation intervals of the point estimate of the simple model across Monte Carlo draws. Panel (a) presents permanent income risk, panel (b) presents transitory income risk, panel (c) presents permanent income risk for individuals who are classified as non-employed / non-filers, panel (d) presents the mean of permanent shocks, and panel (e) presents the mean of permanent shocks for individuals who are classified as non-employed/non-filers.

Table 9: Monte Carlo Results: Recovering Individual Level Estimates of Permanent Income

	Levels	Changes
Mean $\hat{\beta}^{MC}$	1.0100	1.0429
95% Simulation Interval	(1.0064,1.0136)	(1.0278,1.0579)

Note: Table presents the results of the Monte Carlo exercise evaluating our model's ability to recover idiosyncratic permanent income. We report the mean and 95% simulation interval across Monte Carlo replications. Our model is able to closely identify the level of permanent income as our mean $\hat{\beta}^{MC}$ value is quite close to 1. Our filtering exercise slightly understates the shocks to permanent income, reflecting compounding of the slight deviations in levels.

Recovering individual level estimates of permanent risk. Finally, we perform an exercise to validate our model's ability to recover individual-level estimates of permanent income from the simulated data. We take our estimate of permanent income $\hat{z}_{i,t}$ for (simulated) individual i in period t, and compare to their actual simulated permanent income $z_{i,t}$. We summarize the accuracy of our method using an OLS equation:

$$z_{i,t} = \alpha + \beta^{MC} \hat{z}_{i,t} + \varepsilon_{i,t} \tag{13}$$

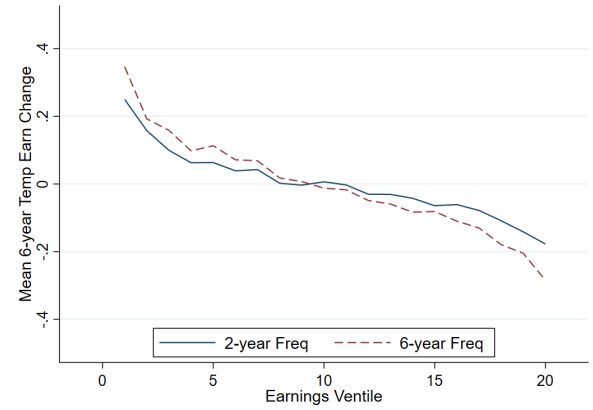
If $\beta^{MC} \approx 1$ then our model accurately recovers permanent income. The mean β^{MC} value and 95% simulation intervals from the exercise are given in Table 9. The mean value, 1.012, is very close to 1 which indicates that our model accurately recovers the individual-level shocks. We also report the coefficients of regressing changes in permanent income $z_{i,t} - z_{i,t-1}$ on changes in our point estimate of permanent income $\hat{z}_{i,t} - \hat{z}_{i,t-1}$. This coefficient, 1.049, is larger but this reflects a compounding of the slight difference in levels.

B Time aggregation

In this appendix we discuss the role of time aggregation. Our data is only ever observed in 5-year intervals. We deal with time aggregation in two ways. First, we note that it is possible to translate an income process with a 1-year time step into a 5-year time step by exploiting the random walk nature of our income process. Second, we show how the output of our filter for the implied shock to permanent income is very similar across different frequency of observations.

Figure 13: Persistent Income Shocks by Frequency of Observation

(a) Mean of Persistent Shocks



(b) Variance of Persistent Shocks



Note: The figure presents results for estimating our income process on the PSID where we observe the data every 2-years (blue, solid line) and every 6-years (red, dashed line). Panel (a) shows the mean of the shock to permanent earnings, and panel (b) shows the variance of shocks to permanent earnings

Time aggregation: theory. For ease of exposition, we omit drift terms and we assume that $\omega_{i,t}$ is iid. We first iterate the income process forward (ignoring non-filers)

$$z_{i,t} = Fz_{i,t-1} + \omega_{i,t} \implies z_{i,t} = z_{i,t-5} + \omega_{i,t} + \omega_{i,t-1} + \omega_{i,t-2} + \omega_{i,t-3} + \omega_{i,t-4}$$

We are therefore estimating the following income process on our data,

$$z_{i,t} = z_{i,t-5} + \widehat{\omega}_{i,t} \quad \text{where} \quad \widehat{\omega}_{i,t} = \omega_{i,t} + \omega_{i,t-1} + \omega_{i,t-2} + \omega_{i,t-3} + \omega_{i,t-4}$$

This implies that annualized variances are obtained by dividing our point estimates by the length of the interval, e.g. 5 in our case.

Time aggregation in practice. In this appendix, we use the PSID sample to examine how the frequency of observations impacts our estimation of the risk that individuals face.[65] In particular, to evaluate the impact of filtering at a 5-year horizon, we estimate our income process on the PSID for two horizons: 2-years and 6-years. We then compare the implied mean and variance of 6-year changes in permanent earnings (using our individual level estimates of permanent income) conditional on ventiles of recent earnings. We plot the conditional mean and variance of these changes in Figure 13 below.

[65] See Appendix C.5 for more details on how we estimate our income process on a PSID sample.

As Figure 13 verifies above, the conditional means (panel (a)) and variances (panel (b)) between filtering frequency are qualitatively very similar similar. Aside from a modest vertical shift, the shape of the variance profile is almost identical when filtered at different frequencies.[66]

C Additional Empirical Results

In this appendix, we present a series of additional empirical results.

C.1 Additional model estimates

In this section, we report additional estimates from the baseline model introduced in Section 3.1

Mean shocks to permanent income. In Figure 14 we present the mean shock to permanent income as a function of lagged income rank for employed individuals (Panel (a)) as well as non-employed/non-filers (Panel (b)). Among employed individuals, we observe substantial mean reversion: households at low lagged income ranks experience large positive mean shocks to permanent income, while higher-ranked households exhibit much smaller mean shocks. We additionally find some evidence of an increase in the mean shock to permanent income for households between the 80th and 99th percentiles from the 1970s to the 2010s. Among non-employed/non-filers, we find that the magnitude of permanent income declines increases with lagged income rank, implying that higher-income households experience larger permanent income losses during periods of non-employment or non-filing. We do not see much of a change in this pattern between the 1970s and the 2010s.

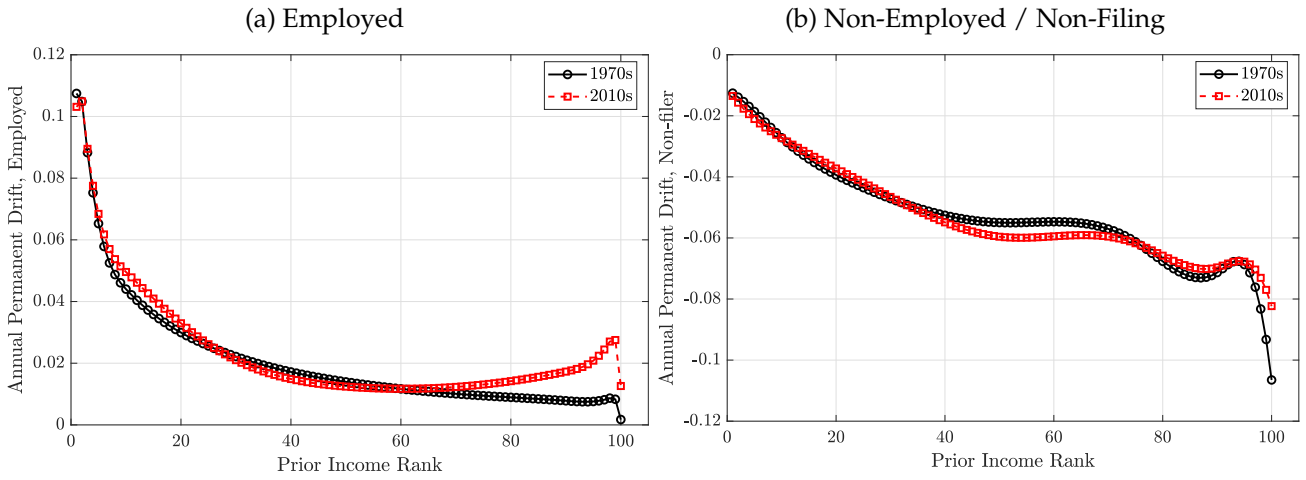
Initial draw of permanent income. Our method also produces an estimate of the initial dispersion in permanent income. In Figure 15, we plot the variance of the initial draw of permanent income by decade. The figure shows that there has been a steady increase in initial permanent income inequality over time.

C.1.1 Additional results: estimates by self-employment status

[X] TO BE ADDED !! [X]

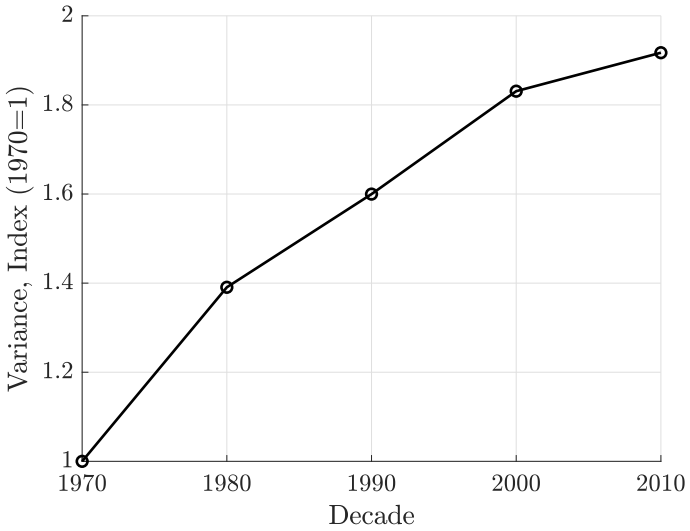
[66] Note the difference in the levels of variances reported in Figure 13 and 23 is because of the use of ventiles in Figure 13.

Figure 14: Mean of Shocks to Permanent Income



Note: Figure presents the results of estimating the income process in Section 2.1. Panel (a) plots the full profile of mean shocks to permanent income among the employed as a function of lagged income rank for the 1970s and 2010s. Panel (b) plots the full profile of mean shocks to permanent income among the non-employed/non-filing as a function of lagged income rank for the 1970s and 2010s. In both figures we present an annualized mean and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

Figure 15: Variance of initial draw of permanent income over time



Note: Figure presents the results of estimating the income process in Section 2.1, and show the variance of the initial draw of permanent income by decade. The y-axis is set such that the variance in 1970 is equal to one.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

C.2 Alternative income definitions

In this appendix, we show that our results are robust to a series of alternative income definitions.

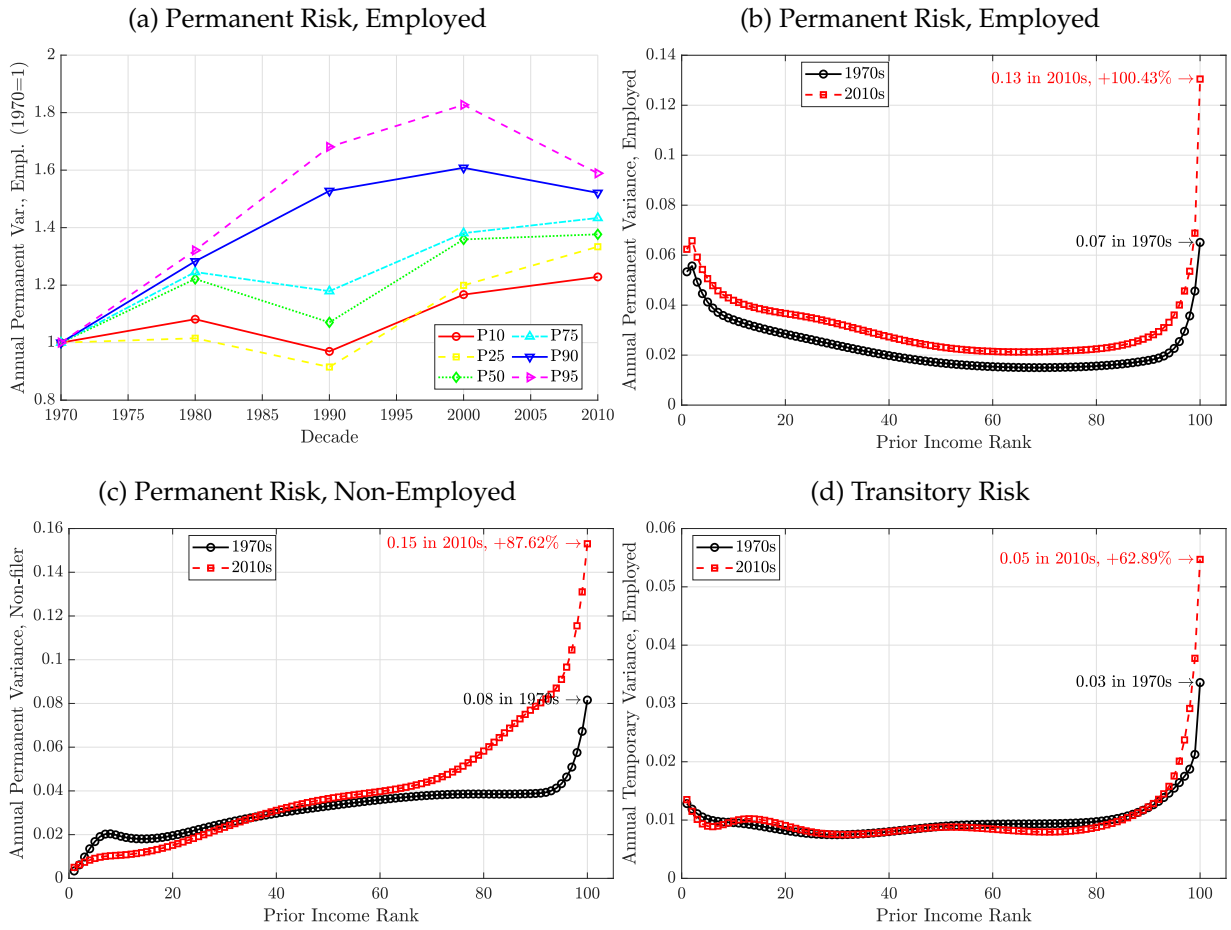
C.2.1 Adjusted Gross Income (AGI)

We first show that our results are robust to using adjusted gross income (AGI) as our measure of income in Figure 16. Panel (a) shows the evolution of permanent income risk among the employed using AGI for each decade from 1970 to 2010 for selected percentiles of the income distribution. The figure shows that households in the 95th and 90th percentiles of the income distribution have experienced the largest increases in income risk relative to households lower in the income distribution. Panel (b) shows that at the very top of the income distribution, permanent income risk has effectively doubled since the 1970s. Finally, panels (c) and (d) show that permanent income risk among the non-employed, as well as transitory income risk, have increased for high-income households since the 1970s.

C.2.2 Wage and Salary Income (WSI)

Next, we examine the evolution of income risk across the income distribution and over time using wage and salary income (WSI) as our measure of income. Figure 17 presents the results of estimating the income process described in Section 2.1 with income measured using WSI. Panel (a) shows the evolution of permanent income risk across the income distribution since the 1970s for selected percentiles. The figure shows that high-income households have experienced the largest increase in permanent income risk while employed. Panel (b) shows that permanent income risk has increased across the income distribution between the 1970s and the 2010s, with the largest increase occurring among high earners. In particular, for individuals at the very top of the income distribution, permanent income risk has increased by nearly 60%. Panel (c) shows that permanent income risk has increased during spells of non-employment for high earners and has effectively doubled at the very top of the income distribution. Panel (d) plots the profile of transitory income risk in the 1970s and the 2010s and shows that, when income is measured using WSI, transitory risk is very similar across the two time periods except in the right tail of the distribution. Taken together, these results show that when income is measured using WSI, both permanent and transitory income risk have increased for top earners.

Figure 16: Evolution of Income Risk Using Adjusted Gross Income (AGI)

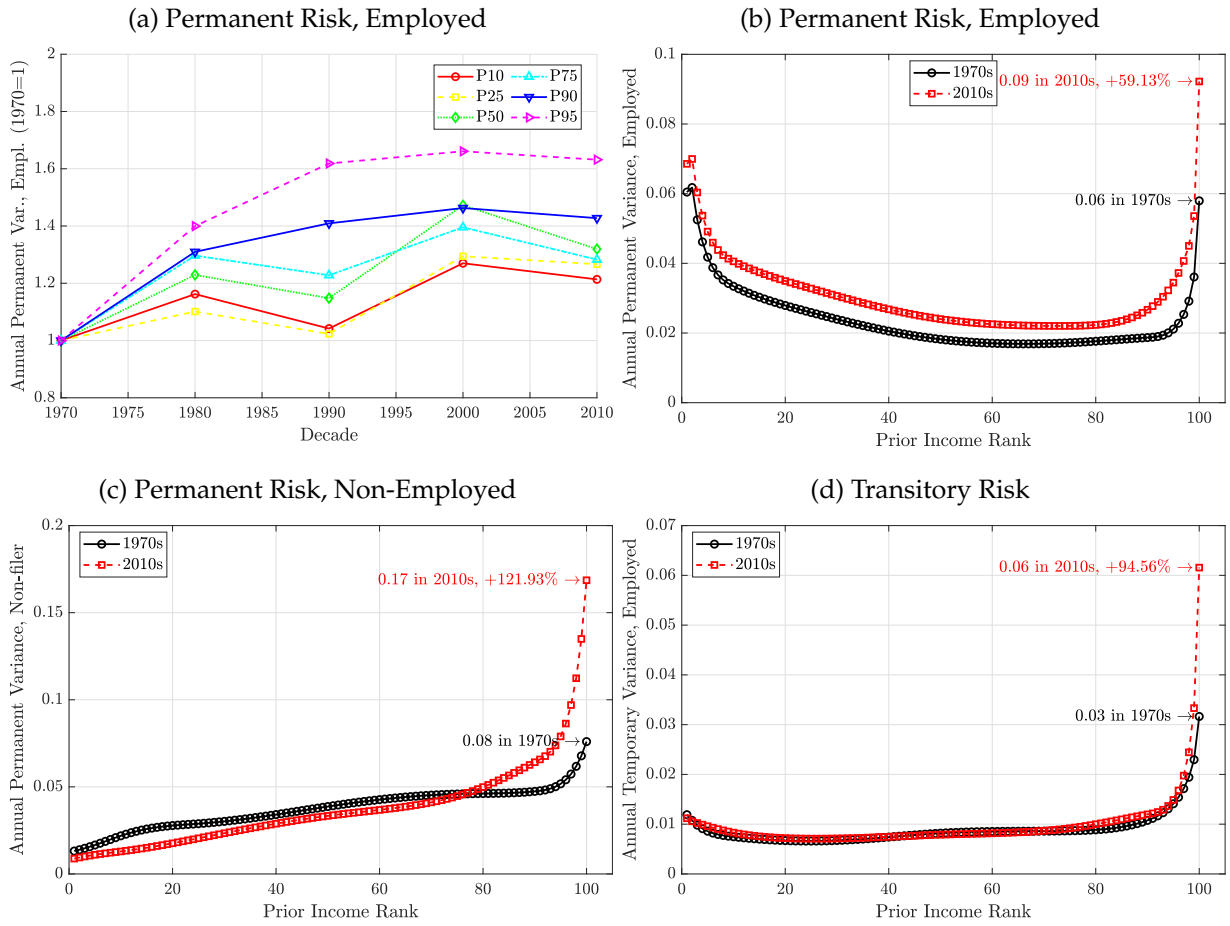


Note: Figure presents the results of estimating the income process in Section 2.1 where our measure of income is adjusted gross income (AGI). Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income, while panel (c) plots the full profile of shocks to permanent income among the non-employed. Panel (d) plots the full profile of shocks to transitory income. In panels(b)-(d) figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers. Note in panels (b)-(d) we annualized variances.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

C.2.3 Removing Capital Gains

In this appendix, we examine whether our estimates of income risk are sensitive to the inclusion of capital gains in our income measure. Because capital gains can only be separately identified in our data beginning in 1994, we restrict the estimation sample to 1994 onward for this exercise. We then re-estimate our income process allowing the variance of income shocks to vary by

Figure 17: Evolution of Income Risk Using Wage and Salary Income (WSI)

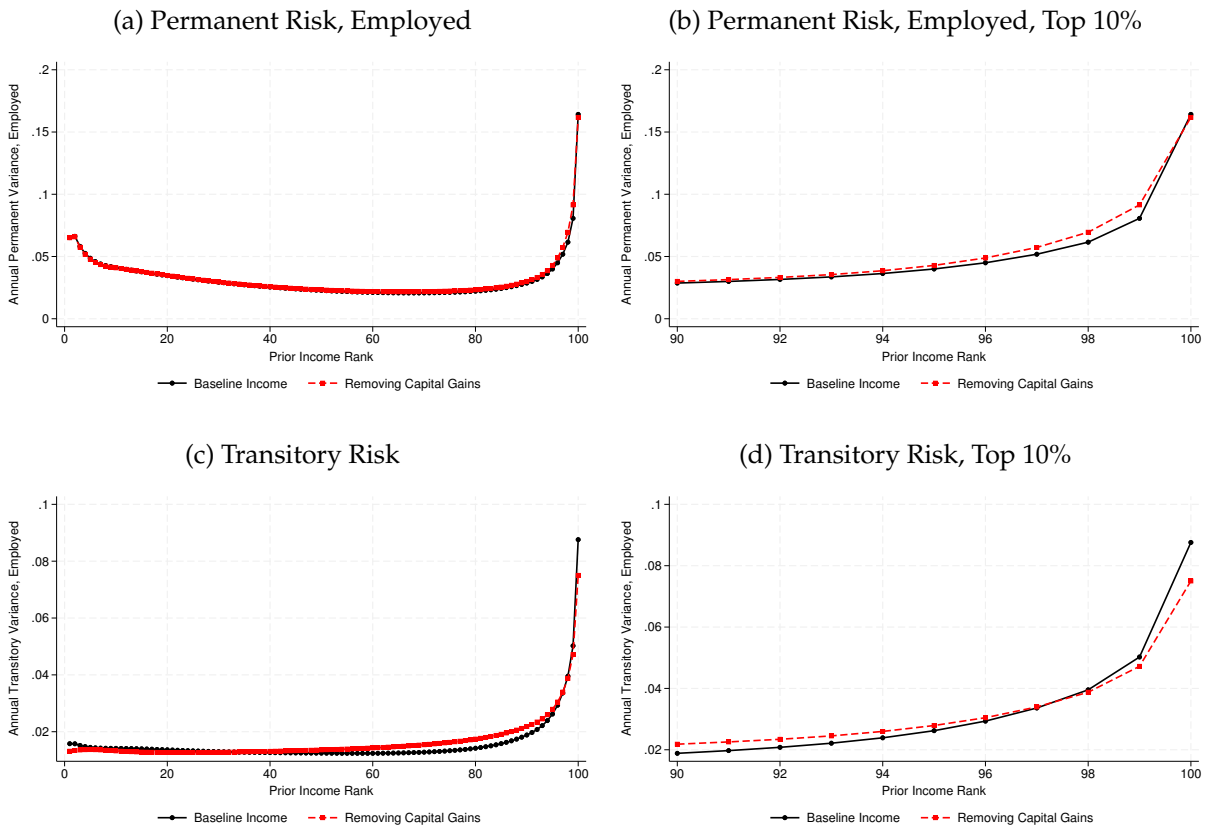


Note: Figure presents the results of estimating the income process in Section 2.1 where our measure of income is wage and salary income (WSI). Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income, while panel (c) plots the full profile of shocks to permanent income among the non-employed. Panel (d) plots the full profile of shocks to transitory income. In panels(b)-(d) figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers. Note in panels (b)-(d) we present an annualized variance.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

lagged income rank and age, but not by decade.

Figure 18 compares estimates obtained using our baseline income measure (black solid line with circle markers) to estimates obtained when capital gains are removed from income (red dashed line with square markers). Panels (a) and (b) show that the estimated profile of permanent income risk by lagged income rank is nearly identical whether or not capital gains are

Figure 18: Evolution of Income Risk Removing Capital Gains



Note: Figure compares estimates of the income process described in Section 2.1 using two income measures: baseline income (black solid line with circle markers) and baseline income excluding capital gains (red dashed line with square markers). Because capital gains can only be separately identified beginning in 1994, the estimation sample for this exercise is restricted to 1994 onward. In this specification, the variance of income shocks is allowed to vary flexibly with age and lagged income rank but is restricted to be constant across decades. Panels (a) and (b) plot the estimated variance of permanent income shocks among employed individuals as a function of lagged income rank. Panels (c) and (d) plot the corresponding estimates for transitory income shocks. Panels (b) and (d) zoom in on the top 10 percent of the lagged income distribution. All variances are reported at an annual frequency.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

included in the income measure. Even at the very top of the income distribution, excluding capital gains has little effect on the estimated variance of permanent shocks. Panels (c) and (d) present the corresponding estimates for transitory risk. The overall shape of the transitory risk profile is again very similar across specifications. However, transitory risk at the very top of the income distribution is modestly lower when capital gains are excluded. This pattern is consistent with capital gains primarily contributing additional transitory volatility among top income households rather than affecting estimates of persistent income risk.

C.3 Random sample of 1040s

In this appendix, we examine the evolution of income risk across the income distribution and over time using a 2% random sample of individuals from the 1040s data and our baseline measure of income.[67] Figure 19 presents the results of estimating the income process described in Section 2.1 on this 2% random sample. Panel (a) plots the evolution of permanent income risk among the employed for selected percentiles of the lagged income distribution. The figure shows that permanent income risk has increased across the income distribution between the 1970s and the 2010s, with the largest increases occurring among individuals in the upper tail of the income distribution. Panel (b) shows the full profile of permanent income risk among the employed by lagged income rank in the 1970s and the 2010s. The figure shows that the increase in permanent income risk is pronounced among high-income individuals. Panel (c) shows that permanent income risk among the non-employed has increased for individuals above roughly the 60th percentile of the income distribution, with especially large increases at the top of the income distribution. Finally, panel (d) shows that transitory income risk is very similar between the 1970s and the 2010s for most of the income distribution, but increases noticeably in the right tail, particularly above the 95th percentile. Taken together, these results show that our baseline findings are robust when estimating the income process using a random sample of 1040s.

C.4 Additional Results: Composition.

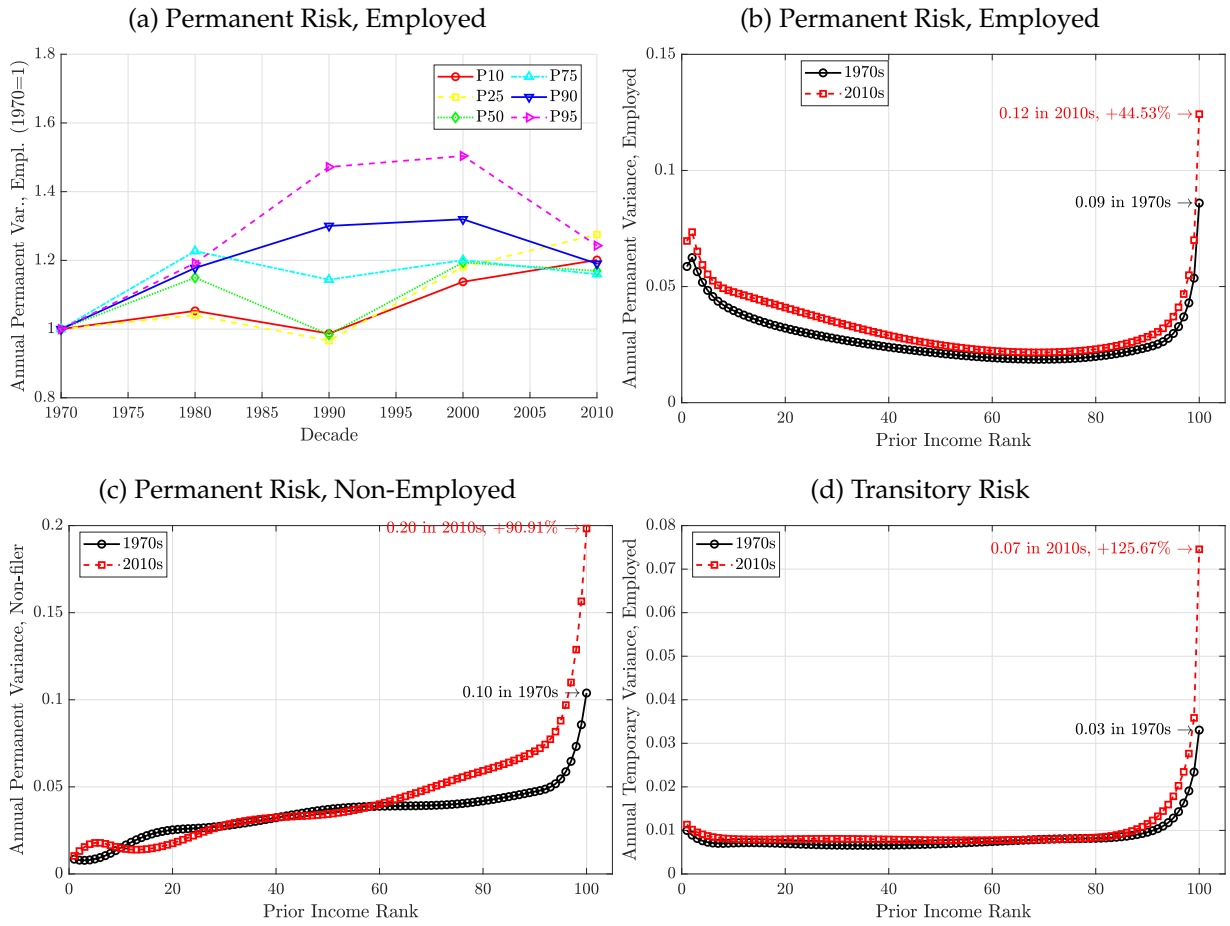
In this appendix, we summarize a set of additional results to highlight that our results are robust to different methods of controlling for changes in composition of workers within a household (Appendix C.4.1), occupation (Appendix C.4.2) and age (Appendix C.4.3).

C.4.1 Alternative Residualization.

In this appendix, we show that our results are robust to altering our process for residualizing earnings to including fixed effects for the interactions between the number of adults filing and year. This specification allows the relationship between income and household filing status to vary flexibly over time and therefore helps control for changes in the gender and household composition of the labor market. Figure 20 presents the results. Panel (a) plots the evolution of permanent income risk among the employed by decade for selected percentiles of the lagged income distribution. Similar to our baseline results, permanent income risk increases

[67] For this exercise, we randomly sample individuals. In our data, we observe the scrambled Social Security number (PIK) of a spouse when present. In cases where both members of a two-adult household are drawn into the sample, we assign each individual a weight of $1/2$ in the estimation.

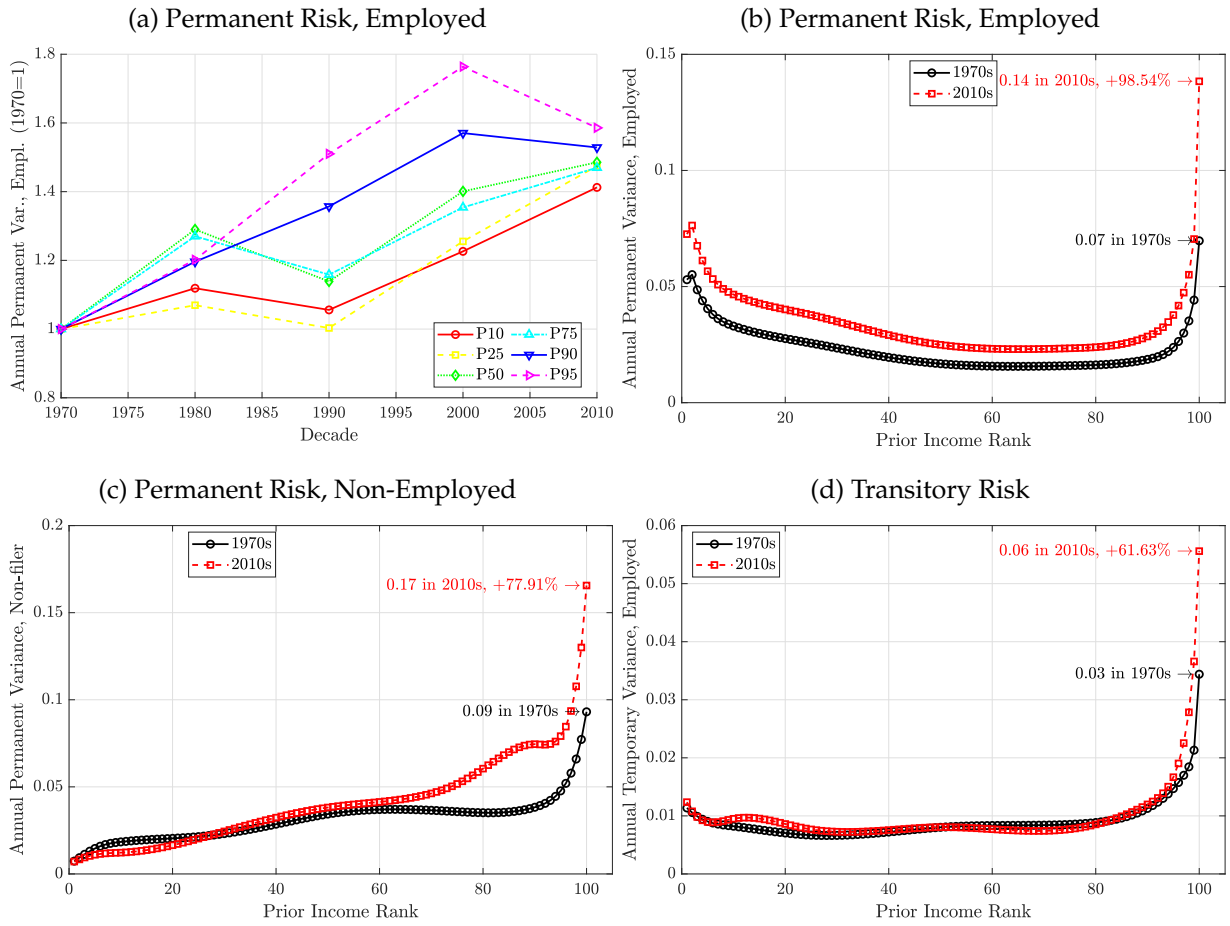
Figure 19: Persistent and Transitory Risk Using Random Sample of 1040s



Note: Figure presents the results of estimating the income process in Section 2.1 using our baseline measure of income on a 2% random sample of 1040s. Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income, while panel (c) plots the full profile of shocks to permanent income among the non-employed. Panel (d) plots the full profile of shocks to transitory income. In panels(b)-(d) figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers. Note in panels (b)-(d) we present an annualized variance.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

over time across the distribution, with the largest increases occurring for individuals in the upper tail. Panel (b) shows the full profile of shocks to permanent income while employed for the 1970s and 2010s, and shows that the increase in permanent income risk is especially large at the very top of the income distribution. Panel (c) shows that permanent income risk among the non-employed also increases substantially for individuals in the upper half of the income

Figure 20: Evolution of Income Risk, Alternative Residualization



Note: Figure presents the results of estimating the income process in Section 2.1 using our baseline measure of income, where our process for residualizing income includes a set of fixed effects interacting year with number of adults filing. Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income, while panel (c) plots the full profile of shocks to permanent income among the non-employed. Panel (d) plots the full profile of shocks to transitory income. In panels(b)-(d) figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers. Note in panels (b)-(d) we present an annualized variance.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

distribution. Finally, panel (d) plots the profile of transitory income risk in the 1970s and the 2010s. As in our baseline results, transitory income risk is similar across the two time periods for most of the income distribution but increases among households at the top of the income distribution.

C.4.2 Ranking Income Within Occupation.

In this appendix, we show that our results are robust to ranking workers within their two-digit occupation. We use an individual's first reported occupation in the CPS-ASEC to classify an individual into an occupation, where we classify occupations using the classification system from Autor and Dorn (2013). We then rank individuals based on their residual log income within their two-digit occupation as well as by age and year. Figure 21 presents the results. Panel (a) plots the evolution of permanent income risk among the employed by decade for selected percentiles of the occupation-specific lagged income distribution. Similar to our baseline results, permanent income risk increases over time across the distribution, with the largest increases occurring for individuals in the upper tail. Panel (b) plots the full profile of permanent income risk among the employed as a function of lagged income rank within occupation in the 1970s and the 2010s. The figure shows that permanent income risk increases across most of the distribution, but the increase is particularly pronounced among individuals at the very top of the distribution, where permanent income risk rises by nearly 100 percent. Panel (c) shows that permanent income risk among the non-employed also increases substantially for individuals in the upper half of the income distribution when ranking is performed within occupation. Finally, panel (d) plots the profile of transitory income risk in the 1970s and the 2010s. Consistent with our baseline results, transitory income risk is similar across the two time periods for most of the income distribution but increases in the right tail. Taken together, these results show that the increase in income risk for high-income individuals is not driven by differences across occupations but instead reflects changes within occupations.

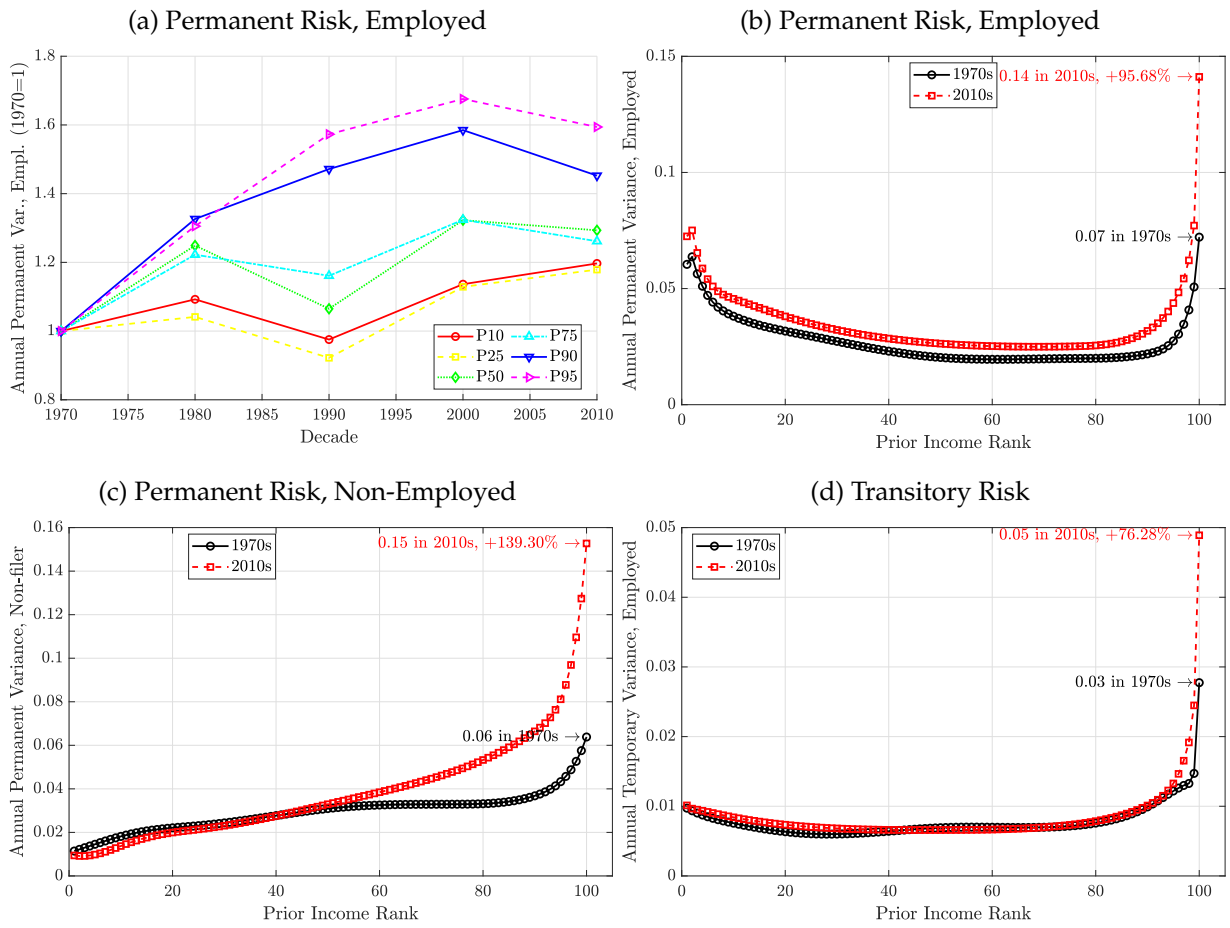
C.4.3 Profiles by Age.

Finally, in this appendix we examine whether there is heterogeneity in the evolution of permanent income risk across the income distribution and over time by the *age* of the individual. To do so, we adjust the income process presented in Section 2.1 so that the basis functions that are allowed to vary by decade and lagged income rank are also allowed to vary by whether the individual is younger than age 40 or age 40 and older. In this estimation, we remove the inclusion of an age quadratic from the specification. We present the resulting variance of permanent income for the employed in Figure 22.[68]

The left panel of Figure 22 shows the profile of permanent income risk for individuals under the age of 40 for the 1970s (black, solid line with circle markers) and the 2010s (red, dashed line with square markers) by lagged income rank. The figure shows that the U-shaped nature

[68]The results for the other parameters of the income process are available upon request.

Figure 21: Evolution of Income Risk, Ranking Within 2-Digit Occupation

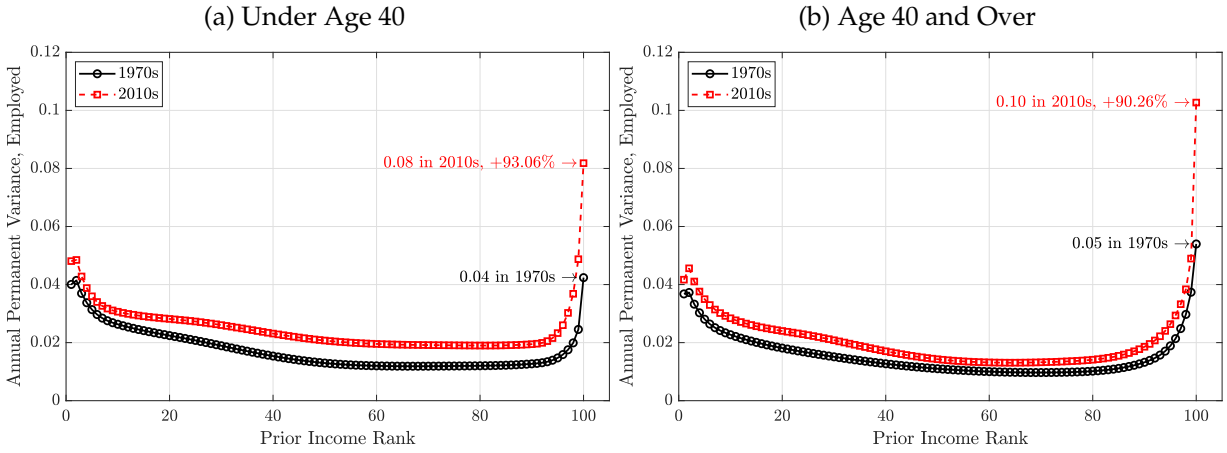


Note: Figure presents the results of estimating the income process in Section 2.1 using our baseline measure of income and ranking individuals within their 2-digit occupation. Panel (a) plots the time series of the variance of shocks to permanent income (y-axis) among the employed by decade from the 1970s to the 2010s (x-axis) for selected percentiles of the lagged income distribution. In this figure, we have normalized the values on the y-axis to be equal to 1 in 1970s. Panel (b) plots the full profile of shocks to permanent income among the employed as a function of lagged income, while panel (c) plots the full profile of shocks to permanent income among the non-employed. Panel (d) plots the full profile of shocks to transitory income. In panels(b)-(d) figures we present an annualized variance and the 1970s correspond to the black, solid line with circle markers, while the 2010s correspond to the red, dashed line with square markers. Note in panels (b)-(d) we present an annualized variance.

Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

of permanent income risk is present among younger individuals, and that among the young the largest increases in risk occur for higher earners. At the top of the distribution, permanent income risk has increased by approximately 93%. Panel (b) shows the results for permanent income risk among individuals age 40 and older. We obtain similar results for older individuals, both in terms of the U-shaped relationship between risk and income rank and in terms of

Figure 22: Permanent Income Risk Among Employed by Income Rank and Age



Note: Figure presents the results of estimating the income process in Section 2.1 using our baseline measure of income, where the profile of shocks as a function of lagged income rank is allowed to vary by age (under 40, or 40 and over) as well as decade. Panel (a) plots the full profile of shocks to permanent income among the employed as a function of lagged income for individuals under the age of 40, while panel (b) plots the full profile of shocks to permanent income among the employed for individuals over the age of 40. Note in both panels we present an annualized variance.
Source: 1973, 1979, 1981-1991, 1994, and 1996-2020 Current Population Survey Annual Social and Economic Supplement linked to 1040s in the years 1969, 1974, 1979, 1984, 1989, 1994, 1999, 2004, 2009, 2014, 2019.

the increase in risk being concentrated at the top of the income distribution. In particular, for individuals at the top of the income distribution permanent income risk has increased by 90%. We do, however, find that the level of income risk for high-income individuals is higher among older individuals. Taken together, these results indicate that the increase in permanent income risk among high-income individuals has occurred for both younger and older individuals.

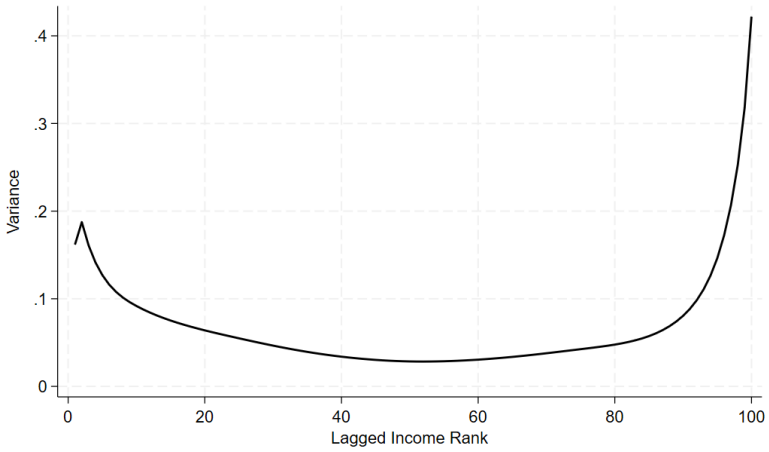
C.5 PSID Estimation and Results

In this appendix, we present the results of estimating our income process on data from the PSID. We show that we obtain a similar U-shaped pattern in permanent income risk across the income distribution using the publicly available data from the PSID.

PSID Sample. We use data from the PSID for the years 1971 to 2013.[69] To mimic the 5-year observation structure of our tax data, we impose that we observe earnings every 6-years in our PSID dataset. To maintain comparability with our baseline results, we impose similar sampling criteria on the PSID where possible. In particular, for our PSID sample we require: (1) individu-

[69]We use the PSID earnings histories that are made available from the CNEF for the years 1971 to 2013. We thank them for making this data available to us.

Figure 23: Permanent income risk in the PSID



als to be between the ages of 25 and 60 (inclusive), and (2) to have earnings over the minimum earnings cutoff twice and in 50% of years between when they first and last satisfy the cutoff. We additionally drop the SEO oversample in our estimation.

In Figure 23, we present the variance of shocks to permanent earnings as a function of lagged earnings ranks. The figure shows that in the PSID, we obtain a similar pattern of permanent income risk being U-shaped. Thus, the result that high and low income individuals face higher income risk relative to the middle of the income distribution is not driven by the nature of our administrative earnings data.

D Capitalizing interest and dividend income

In this appendix, we discuss how we capitalize interest and dividend income to arrive at a measure of wealth. We follow Smith, Zidar, and Zwick (2023) and impute household-level wealth from components of income, observable in 1040 tax forms, that originate from returns on components of wealth. Specifically, we capitalize taxable interest income and dividend income from the 1040 forms into imputed fixed-income and C-corp wealth. The strategy can be described as follows.

Let $w^c_{i,t}$ be household i's wealth in component c at date t. In the 1040 form, household i reports the income generated from this wealth component, $I^c_{i,t} = r^c_{i,t} w^c_{i,t}$, where $r^c_{i,t}$ is the net return of household i's wealth in component c at date t, which may vary across households due to differing portfolios within component c. If the net return $r^c_{i,t}$ for a household is known, then the household's wealth can be inferred as simply $w^c_{i,t} = \kappa^c_{i,t} I^c_{i,t}$ where $\kappa^c_{i,t} = \frac{1}{r^c_{i,t}}$ is the capitalization factor for household i's wealth in component c at date t.

Smith, Zidar, and Zwick (2023) compute capitalization factors for taxable interest income and dividend income by constructing portfolios and computing net returns on fixed-income and C-corp wealth, respectively. For imputing C-corp wealth, we do not observe capital gains so we use their capitalization factor restricted to load solely on dividend income.[70] This is homogenous across individuals but varies over time. To capitalize fixed income wealth, we begin with their 3-tier classical minimum distance (CMD) specification, which is heteregenous across individuals as sorted by interest income. We avoid sharp discontinuities and nonmonotinicities in imputed wealth implied by taxable interesst income at the cutoffs defining the 3 tiers by interpolating. Specifically, we assign the lowest-tier capitalization factor to the households below the 90th percentile of taxable fixed income wealth. We assign the 99th percentile to have the second-tier capitalization factor, and linearly interpolate between the first- and second-tier capitalization factor for households with fixed interest income percentiles between 90-99. We then assign the highest-tier capitalization factor to the 99.9th percentile, and linearly interpolate between the second-tier and highest-tier capitalization factors between the 99-99.9th percentile. We finally assign the highest-tier capitalization factor to those above the 99.9th percentile.

The result of this procedure is capitalized fixed-income and C-corp wealth for individuals in our sample, the sum of which we take to be our measure of household wealth. This can be best thought of as being the liquid components of income-generating wealth - we do not observe housing or pension wealth. Our wealth measure begins in 1974, when the taxable interest income and dividend income 1040 boxes begin to be observed in our data, and end in 2014 due to our capitalization factor series ending in 2016.

E Theoretical motivation

Before discussing our empirical results, it is instructive to discuss some simple theoretical intuition for the tight connection between the risk-free rate and the level of uninsurable idiosyncratic risk. Without loss of generality, we can decompose any agent's log consumption growth as follows

$$\Delta \log c_{i,t+1} = E_t \Delta \log c_{i,t+1} + \omega_{i,t}[\Delta \log C_{t+1} - E_t \Delta \log C_{t+1}] + \eta_{i,t+1} \equiv \mu_{i,t} + \omega_{i,t}\epsilon_{t+1} + \eta_{i,t+1},$$

where C_t is aggregate consumption, ω_t is the linear projection of person i's consumption on the innovation in aggregate consumption, and $\eta_{i,t+1}$ is a residual term which has mean zero and is uncorrelated with aggregate consumption. Next, suppose that an investor has CRRA

[70] We thank Eric Zwick for sending us this capitalization factor.

preferences with risk aversion parameter σ and discount factor β. Then, the following equation holds for any agent who saves in the risk free asset:

$$\frac{1}{R_{f,t}} = E_t\left[\beta\left(\frac{c_{i,t+1}}{c_{i,t}}\right)^{-\sigma}\right] = \beta E_t\Bigg[\exp\left(-\sigma[\mu_{i,t}+\omega_{i,t}\epsilon_{t+1}]\right)\underbrace{E_t[\exp(-\sigma\eta_{i,t+1})|\epsilon_{t+1}]}_{\equiv\exp(-\sigma\vartheta_{i,t+1})}\Bigg],$$

where the second equality uses our decomposition and the law of iterated expectations to define the "certainty equivalent" $\nu_{i,t+1}$ which converts the expected level of marginal utility associated with residual idiosyncratic risk $\eta_{i,t+1}$ (conditional on ϵ_{t+1}) into an equivalent movement in aggregate consumption. Rearranging, we can rewrite the risk-free rate for any agent as

$$\log R_{f,t} = \underbrace{-\log\beta}_{\substack{\text{time}\\ \text{discounting}}} + \underbrace{\sigma\mu_{it}}_{\substack{\text{expected}\\ \text{growth}}} + \underbrace{\sigma E_t(\vartheta_{i,t+1})}_{\substack{\text{expected CE over}\\ \text{idiosyncratic risk}}} - \underbrace{L_t[-\sigma(\omega_{i,t}\epsilon_{t+1}+\vartheta_{i,t+1})]}_{\substack{\text{uncertainty about aggregate and}\\ \text{state-dependent idiosyncratic risk}}},$$

where $L_t[X_{t+1}] \equiv \log E_t \exp(X_{t+1}) - E_t X_{t+1}$ denotes the entropy of the random variable X_{t+1}. If for instance idiosyncratic risk is close to state independent and ϵ_{t+1} is normally distributed, then $L_t[-\sigma(\omega_{i,t}\epsilon_{t+1}+\vartheta_{i,t+1})] = 1/2\sigma^2\omega_{i,t}^2 Var[\epsilon_{t+1}]$.

While the first and second terms, as well as the entropy of the aggregate shock ϵ_{t+1}, feature very prominently in the discussion of potential drivers of recent declines in interest rates, our focus instead relates to the third term. In a world in which the level of idiosyncratic consumption risk faced by savers – especially the wealthy who own a disproportionate share of financial assets – has increased, the Euler equation would suggest that we would anticipate a decline in real interest rates.[71] Since marginal utility is convex, mean preserving spreads in future consumption increase the desire to save today, a force which is offset by declines in equilibrium real interest rates.

In the remainder of the paper, we first document a sizable rise in the income risk faced by the very wealthy, a result which obtains for both broad (adjusted gross income) and narrower (wage and salary) sources of income. Specifically, we will decompose household income into independent permanent and transitory components, with shocks $\nu_{i,t}$ and $\omega_{i,t+1}$ and characterize the time-variation in their variances. To see the relation between these components and the risk-free rate, we can additionally layer in the Blundell, Pistaferri, and Preston (2008) approximation for consumption, we can write $\eta_{i,t+1} \approx \pi_{i,t}\nu_{i,t} + \gamma_{i,t}\pi_{i,t}\omega_{i,t} + \xi_{i,t+1}$, where $\pi_{i,t}$ is the ratio of the present discounted value of human capital to total wealth, $\gamma_{i,t}$ is an age-dependent factor

[71] This intuition is not specific to assuming CRRA preferences. A similar decomposition obtains with more general utility functions such as Epstein-Zin preferences, albeit with different coefficients and additional terms capturing continuation utility.

capturing the share of the year $t+1$ income to total human wealth, and $\xi_{i,t+1}$ is an orthogonal component capturing either measurement error or non income-related drivers of consumption. If we apply the Arrow-Pratt approximation (exact for normally-distributed shocks) to compute the certainty equivalent $\vartheta_{i,t+1}$, then

$$E_t[\vartheta_{i,t+1}] \approx -\frac{\sigma}{2}\left[\pi_{i,t}^2 Var_t[\nu_{i,t+1}] + \pi_{i,t}^2\gamma_{i,t}^2 Var_t[\omega_{i,t+1}] + Var_t[\xi_{i,t+1}]\right].$$

Given that estimates of $\gamma_{i,t}$ are typically very small, three key determinants of this component are 1) the level of risk aversion σ, 2) the ratio of human wealth to total wealth (i.e., the pass through of a permanent income shock to consumption), and 3) the variance of the permanent shock $Var_t[\nu_{i,t+1}]$. While $\pi_{i,t}$ is not directly observable, it is estimable with panel data on consumption or by specifying a model of consumption and income. Moreover, it is worth noting that for most working-age households, the sum of labor and entrepreneurial income provides the vast majority of total disposable income, even among the very wealthy. Hence, $\pi_{i,t}$ is unlikely to be small.

We then conduct a quantitative experiment which considers the potential impact of the rise in income risk we document for the equilibrium interest rate in a fairly standard model. In addition to an assumption on preferences which enables us to quantify the relevant terms in the Euler equation, we also take a stand on the savings technology available to agents in the model, which provides a mapping between observable income risk and the (not easily observable) *consumption* risk which appears in the Euler equation. While specific magnitudes of the predicted risk free rate decline will of course vary with different assumptions on preferences, the intuition discussed above related to the Euler equation is quite general. Indeed, one of the most robust properties of economies with incomplete markets is that they typically feature lower interest rates relative to comparable models which assume complete markets. Analogously, increasing persistent idiosyncratic risk will tend to increase demand for savings and lower real interest rates in these models, all else constant.