

Genetic Endowments, Income Dynamics, and Wealth Accumulation Over the Lifecycle*

Daniel Barth[†]

Federal Reserve Board of Governors

Nicholas W. Papageorge[‡]

Johns Hopkins University, IZA and NBER

Kevin Thom[§]

University of Wisconsin - Milwaukee

Mateo Velásquez-Giraldo[¶]

Johns Hopkins University

ABSTRACT: We develop and estimate a life-cycle consumption savings model in which observed genetic variation is allowed to affect wealth accumulation through several distinct channels. We focus on genetic markers that predict educational attainment, aggregated into a predictive index called a polygenic score. Based on substantial descriptive evidence, we allow variation in these endowments to affect earnings, the disutility of labor, stock market participation costs, and idiosyncratic rates of return on risky investments. The model also incorporates endogenous retirement and a realistic social security system. Parameter estimates suggest that, in addition to earnings, genetic differences are significantly associated with risky asset returns, both of which contribute to wealth inequality. Counterfactual policy exercises indicate that two ways to lower costs of an aging population (extending the age of retirement or cutting social security benefits) have similar magnitudes and distributions of welfare costs even though the latter policy appears to reduce wealth differences between agents with different genetic endowments. This illustrates the importance of welfare calculations when evaluating how genes interact with policy, which is possible to do if we incorporate genetic data into structural models.

KEYWORDS: Wealth, Inequality, Genetic Endowments, Saving and Portfolio Choices, Retirement.

JEL CLASSIFICATION: D14, D15, D31, D63, G50, H31, H55, I38, J24, J26.

*First draft: July 1, 2022. Current draft: August 4, 2022. We are grateful for helpful comments from Christopher Carroll and Michale Keane along with seminar participants at the University of St. Andrews, the Board of Governors at the Federal Reserve, Wisconsin - Madison, Wisconsin - Milwaukee, Banco de la República de Colombia, Yale, SITE and NBER-Aging 2022 along with two conferences on genes and economic analysis: Genes, Social Mobility, and Inequalities across the Life-Course and Frontiers in Economic Analysis with Genetic Data. We also acknowledge excellent research assistance from Matthew Gonzalez. The usual caveats apply.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the Federal Reserve Board of Governors.

[†]dannybarth@gmail.com

[‡]papageorge@jhu.edu.

[§]kthom.work@gmail.com.

[¶]mvelasq2@jhu.edu.

1 Introduction

Over the past several decades, economists and other social scientists have paid considerable attention to rising income and wealth inequality in the developed world. A large literature documents numerous sources of heterogeneity affecting wealth accumulation over the life cycle, including differences in wages, financial literacy, entrepreneurial ability, health, and family structure. Genetic factors contribute substantially to this heterogeneity. Studies using a variety of data and methods consistently find a role for genetics in explaining cross-sectional variation in labor income and wealth (Cesarini et al., 2009; Cronqvist and Siegel, 2014, 2015; Barth, Papageorge, and Thom, 2020; Black et al., 2020; Fagereng, Mogstad, and Rønning, 2021). For example, one noteworthy twin study estimates that 39 percent of the variation in wealth can be attributed to genetic factors (Cronqvist and Siegel, 2015). While such studies quantify the overall contribution of genes, they reveal far less about the behavioral pathways or mechanisms through which genetic factors operate, limiting our understanding of how various policies may attenuate or exacerbate genetic influences and how genetic privilege may compound over time.

Newly available molecular genetic data allow researchers to measure important but previously unobserved sources of heterogeneity in life cycle outcomes. Most prominently, a series of landmark studies provide estimates of associations between individual genetic markers and educational attainment (Lee et al., 2018; Okbay et al., 2016; Rietveld et al., 2013). Recent work establishes these genetic markers as relevant for wealth accumulation; they not only predict educational attainment out of sample, but also predict wealth and financial behaviors even after controlling for education and family background (Barth, Papageorge, and Thom, 2020; Belsky et al., 2018). We incorporate this genetic variation into a structural model by linking it to economic primitives and argue that this facilitates the proper interpretation of these relationships and a better understanding of how policies shape genetic gradients. This approach permits counterfactual exercises that assess the extent to which policy changes concentrate costs or benefits on individuals with different traits, including genetic endowments. This is a natural extension of the existing literature that examines "gene-by-environment" or $G \times E$ interactions. Indeed, a utility-maximizing framework allows us to go beyond evaluations of $G \times E$ interactions in outcomes like consumption or wealth and to directly estimate these interactions in terms of welfare.

In this paper, we build and estimate a structural consumption-savings model that explicitly allows genetic variation to affect wealth accumulation through multiple channels. We measure genetic variation using a *polygenic score* developed to predict educational attainment (which we refer to as the EA score). Such scores are commonly used in statistical

analyses of complex traits to aggregate many individual genetic markers into a linear predictive index. The model incorporates a variety of features that have become standard in the consumption-savings literature: a realistic Social Security system, (transitory) income volatility, heterogeneous asset returns and participation costs, endogenous retirement, inheritances, bequests, and defined benefit retirement income. Motivated by substantial empirical evidence, in the model we allow variation in the EA score to influence the labor income process, the welfare costs of work, risky-asset returns, the financial cost of stock market participation, access to defined benefit pensions, and the distribution of inherited wealth. We estimate the model to match wealth, stockholding, and retirement patterns in the Health and Retirement Study.

The estimated model permits a decomposition of the EA score-wealth gradient, as well as ex-ante policy analyses that explore how less generous retirement benefits (achieved through either delayed or less generous benefits schedules) may affect genetic inequality in outcomes including wealth, consumption, retirement, and most importantly welfare. Our estimates suggest that differences in labor income and returns on risky financial assets together explain most of the gene-wealth gradient within education groups. Quantitatively, our estimates suggest that, compared to households in the lowest EA score decile, those in the top decile earn 3.9 percentage points higher returns from risky financial assets per year on average. By contrast, differences in inheritances and childhood SES together play a relatively small role in explaining the EA score-wealth gradients within education groups. The empirical results highlight the dynamic nature of genetic gradients in economic outcomes. Income and return differences compound cross-sectionally and over the life cycle. Near retirement, these interactions make agents at the lower end of the EA score distribution particularly vulnerable to changes that deteriorate the social safety net, such as cost-saving changes in the Social Security system.

We simulate two policy counterfactual reforms of the Social Security system. The fiscal challenges of an aging population have led to a number of proposals to reduce the cost of administering Social Security. (see, e.g., Diamond and Orszag, 2005; Feldstein, 2005; National Research Council, 2012). The size of the system, its redistributive nature, and the fact that it is the main source of income for a large fraction of elderly Americans (Dushi and Trenkamp, 2021) underline the importance of anticipating the effects of potential reforms and identifying the groups that could be most harmed. We analyze policies that exemplify two of the main types of reforms being considered: increases in the age of retirement and reductions in the level of benefits (Social Security Administration, 2022b). It is crucial to emphasize that our counterfactuals examine how policy changes would affect genetic gradients, but they do not involve targeting interventions to individuals with specific genotypes or otherwise

require individual-level genetic information. Targeting policies on the basis of polygenic scores is neither advisable nor sensible given that these scores do not predict *individual-level* outcomes with great precision. Rather, we argue that models enriched with population-level genetic heterogeneity can provide an effective diagnostic tool for assessing the distributional consequences of various reforms. It is important for policymakers to understand not only whether policies increase inequality, but also whether they do so on the basis of genetic endowments.

The first policy that we consider shifts the schedule of Social Security benefits five years forward, moving the earliest age at which individuals can claim benefits from 62 to 67. Our second counterfactual preserves the current benefit schedule’s timing but reduces all benefits by a fixed percentage. We set the reduction close to 28%, which raises the same revenue as the policy in our first counterfactual. In response to the benefit shift, some individuals work longer to counteract the loss of benefits, but median wealth declines starting at age 62. This decline is most pronounced among agents with low EA scores, which intensifies the relationship between the EA score and wealth as agents age.

Alternatively, benefit reductions increase the incentive to save for retirement, increasing wealth for agents until about age 72. This increase in savings means that financial sophistication plays a heightened role, with high EA-score individuals growing their additional savings more efficiently given their higher returns. Nonetheless, wealth increases in the years surrounding retirement are larger in relative terms for agents with lower EA scores. However, the increase in wealth comes at the cost of lower consumption, which disproportionately affects the utility of low-EA score households who earn lower average returns on savings. The consequence is that *welfare* losses are concentrated among individuals with lower EA scores. While the magnitude and distribution of welfare costs are similar across policies, the differences in the reduced-form association between genetic endowments and wealth would (incorrectly) suggest that genetic inequality has been reduced under benefit reductions relative to the status quo. Instead, genetic inequality along with dimension of well-being has worsened. These results highlight the usefulness of economic models for interpreting gene-by-environment interactions, and demonstrates how the assessment of welfare may be challenging without a rich model of economic behavior.

As our counterfactual results illustrates, studying genetic associations through the lens of a structural model offers two methodological contributions that highlight the synergy between economic modeling and behavioral genetics. First, such a framework is essential for the *ex-ante* analysis of gene-by-environment interactions. The literature connecting genetics and economic outcomes emphasizes understanding how environmental factors moderate or amplify the influence of genetic endowments. Since policy can modify the economic envi-

ronment, the study of $G \times E$ interactions can shed light on the scope for policy choices to shape social inequality arising from the random outcomes of the “genetic lottery.” (Harden, 2021). For example, existing studies provide evidence that educational reforms moderate genetic associations with completed schooling (Barcellos, Carvalho, and Turley, 2018), and that labor-market conditions and the availability of pensions may alter the EA score’s association with income and wealth (Barth, Papageorge, and Thom, 2020; Papageorge and Thom, 2020). However, such research has rarely, if ever, integrated genetics into optimizing models of economic behavior. As such, these retrospective $G \times E$ results are insufficient for predicting the effects of newly implemented or proposed policies. Understanding how policies moderate or exacerbate genetic gradients requires knowing more about the economic primitives affected by these parameters (e.g., earnings potential, financial sophistication, and family transfers), and the ways in which individuals of different genotypes will endogenously respond to changes in the environment.

Second, a utility-maximizing framework is critical for understanding the welfare implications of genetic gradients and $G \times E$ interactions. The presence of $G \times E$ interactions in an outcome like consumption or wealth is often interpreted in terms of welfare, but this may be misleading without an underlying behavioral model. For example, it may be tempting to conclude that policies that shrink gene-wealth gradients would also reduce genetic inequality in welfare. However, our counterfactual analyses offer an example of a policy (a Social Security benefit reduction) that slightly weakens the association between genes and wealth, but concentrates welfare losses on those with polygenic scores associated with disadvantage; that is, while the gene-wealth gradient is mildly flattened, the *gene-welfare* gradient is substantially *steepened*. Our analysis demonstrates how utility-maximizing models can be used to perform $G \times E$ analyses of expected lifetime welfare under different policies.

Our analysis connects to three main literatures. First, our paper is related to the literature exploring the degree to which the heritability of socioeconomic outcomes and behaviors is linked to biological pathways and is modulated by particular environments. Early evidence used twin studies to demonstrate that genetic variations have explanatory power for outcomes like earnings, risk-taking and giving, investment decisions and biases, and saving decisions (Taubman, 1976; Cesarini et al., 2009, 2010; Cronqvist and Siegel, 2014, 2015). More recently, advances in behavioral genetics have produced indices of individual-level genetic variations with out-of-sample explanatory power for complex socioeconomic outcomes (Rietveld et al., 2013; Okbay et al., 2016; Lee et al., 2018; Karlsson Linnér et al., 2019; Hill et al., 2019). This paper builds on previous results that have shown that a polygenic score for educational attainment predicts earnings, wealth, and portfolio decisions even after flexibly controlling for educational attainment (Barth, Papageorge, and Thom, 2020; Papageorge

and Thom, 2020). As is the case with other genetic associations (Barcellos, Carvalho, and Turley, 2018; Belsky et al., 2018; Muslimova et al., 2020; Ronda et al., 2022), past studies have found that the relationship between the EA score, income, and wealth is modulated by environmental factors such as labor-market conditions and pension arrangements.

The second literature uses life-cycle models to study the fiscal and welfare consequences of policies related to retirement. Wealth at retirement varies substantially and life-cycle models can provide benchmarks and tests to assess the adequacy of households' savings (Hubbard, Skinner, and Zeldes, 1995; Bernheim, Skinner, and Weinberg, 2001; Scholz, Seshadri, and Khitatrakun, 2006). Heterogeneity across dimensions like earning potential, life expectancy, and household structure is a crucial characteristic of models used to evaluate reforms to the Social Security system (Conesa and Krueger, 1999; Nishiyama and Smetters, 2007; Fuster, İmrohoroğlu, and İmrohoroğlu, 2007; Hairault, Langot, and Sopraseduth, 2008; Imrohoroğlu and Kitao, 2012). Indeed, the degree to which a household benefits or loses from a particular reform depends on its characteristics (Fuster, İmrohoroğlu, and İmrohoroğlu, 2003; Kaygusuz, 2015; Bagchi, 2019). Our model has a rich representation of heterogeneity along observable and unobservable dimensions (more than five thousand ex-ante agent types). We add to the literature by incorporating a novel measure of genetic endowments that affects financial sophistication, earnings, and the cost of work. Having a common driver of these characteristics generates compounding welfare effects from policy changes: the people who are most reliant on Social Security payments are those least well-positioned to offset changes using their private savings or extending their working lives. For each of the alternative policies that we analyze, we examine the distribution of expected welfare losses across the population and identify those who are most harmed.

Lastly, our paper relates to the literature examining households' portfolio choices. The workhorse models in this literature prescribe that every household should allocate a substantial fraction of their wealth to stocks (Merton, 1969; Samuelson, 1969; Cocco, Gomes, and Maenhout, 2005). Given the generalized low rates of stockholding that have been documented in various countries, later studies have incorporated financial costs to stockholding, heterogeneity in preferences, and tail-risks as ways to reconcile model predictions with household decisions (see, e.g., Vissing-Jorgensen, 2002; Gomes and Michaelides, 2005; Fagereng, Gottlieb, and Guiso, 2017; Catherine, 2021). Other studies have documented heterogeneity in the returns to wealth and financial sophistication, showing that sophistication covaries with characteristics like education (Calvet, Campbell, and Sodini, 2007; Fagereng et al., 2020). We contribute to the literature by presenting evidence that the EA score can be related to financial sophistication, which can lead to compounding effects of genetic endowments over the life cycle.

The remainder of this paper is organized as follows. Section 2 gives an overview of the genetic index that we use in this paper and its use in social science research. Section 3 describes our data and presents motivating summary statistics. Section 4 presents our model. Section 5 describes our strategy for estimating our model. Section 6 presents parameter estimates, the fit of the empirical patterns that we target, and the implications of the estimates. Section 7 uses the model to assess how socioeconomic outcomes, lifetime welfare, and their relationship with the EA score would change under two cost-saving changes to the Social Security system. Section 8 concludes.

2 Genome-Wide Association Studies and the EA score

2.1 Polygenic Scores

In this section, we describe how the EA polygenic score is constructed, what it measures, and some earlier results on its association with socioeconomic outcomes. More detailed accounts are found in past studies from economics. We also refer the reader to Beauchamp et al. (2011), Benjamin et al. (2012), and Visscher et al. (2017) for excellent reviews, upon which we base much of the information presented here.

The human genome consists of approximately three billion pairs of nucleotide molecules, called base pairs. Four molecules, adenine (A), cytosine (C), guanine (G), and thymine (T), combine to form one of two pairs, either an AT pair or a GC pair. For example, at a particular location in the genome, an individual may have two copies of the AT pair (AT,AT; one from each parent) or an (AT,CG) pair. Sequences of base pairs are called genes and govern bodily function through the synthesis of proteins. At over 99% of locations in the human genome people have the same base pairs. Locations where there is variation across humans are called *single nucleotide polymorphisms* (henceforth *SNPs*). A “reference allele” is the nucleotide pair that is most common at a given SNP, and an individual can thus have zero, one, or two of the reference allele. For example, at a particular SNP, if the most common base pair is AT, a person can have (AT, AT), (AT, GC), (GC, AT) or (GC, GC) and would have zero, one, one, or two copies of the reference allele, respectively.

Genome-wide association studies (GWAS) are empirical exercises that relate SNP-level data to behavior and outcomes, such as height, body mass, certain diseases, economic choices, or socioeconomic outcomes. The process consists of linearly regressing the outcome of interest onto count variables for the number of reference alleles each individual has at each SNP. Regressions are univariate and are run SNP-by-SNP. The coefficients estimated in a GWAS

are used to construct *polygenic scores*, which are linear indices based on the estimated regression coefficients associated with each SNP. This amounts to a weighted average index score, where the weights are the coefficients linking the SNPs to outcomes.¹ Loosely speaking, a higher polygenic score means an individual possesses more of the base pairs that are correlated with the outcome. Polygenic scores can be computed out of sample, with certain conditions, using the genetic data available in a given data set and the GWAS coefficients estimated from the GWAS discovery.

The polygenic score that we use in this paper is based on results from Lee et al. (2018), who conducted a GWAS for educational attainment that featured a discovery sample of over 1.1 million people. The score constructed by the authors explains 10.6% of the variation in the years of education of HRS respondents of European ancestry. This is a notable achievement because individuals from the HRS were not used in the GWAS discovery sample. We will refer to this polygenic score for educational attainment as the *EA3 polygenic score*, since it is the score based on the third generate GWAS estimation, or simply the *EA score*. Even though the EA score was constructed to predict years of education, Lee et al. (2018) showed that it also has out-of-sample predictive power for related outcomes, such as GPA and cognition. More recent studies have demonstrated that the EA score is robustly associated with more complex socioeconomic outcomes, such as labor income and wealth at retirement, even after flexibly controlling for educational attainment.

2.2 Limitations

There are various limitations and concerns to using genetic data in social scientific analysis, many of which have been discussed in earlier literature. Studies that established some of the basic empirical relationships on which this paper builds and that guide our modeling assumptions have been careful in addressing these concerns and others. First, it is wrong to conclude that the social outcomes that we study are purely biologically determined. Educational attainment, income, and wealth are all influenced by environmental factors (e.g., childhood economic advantages, parental investments, macroeconomic trends, life events, and social policies) along with genetic endowments and the interaction between the two. The EA score summarizes genetic variations that, on average and given current environments, predict educational attainment and other socioeconomic outcomes. Furthermore, despite being able to explain a moderate fraction of the aggregate variation in these outcomes, the EA score remains a poor individual-level predictor (Harden and Koellinger, 2020). For these reasons,

¹The regression procedure often makes adjustments to address issues like population stratification, multiple hypothesis testing, and correlation between SNPs. See Benjamin et al. (2012) for details.

we refrain from and caution against interpreting the EA score as “ability.”²

Second, with gene-outcome relationships causality is difficult to ascertain. Unadjusted relationships likely capture factors subsumed into an error term that are related to genetic endowments, such as family environments. We note, however, that many relationships between genes and outcomes, including the association between the EA score and education, hold in analyses with family fixed effects, though coefficients are smaller. Incorporating standard childhood SES variables into the analysis generates relationships that are close to those that rely on within-family variation in the EA score. Thus, many studies, including this one, incorporate such variables (e.g., father’s occupation and mother’s education). In doing so, we argue that we substantially mitigate the concern that any coefficient estimates using the EA score reflect environmental factors correlated with genetic endowments, such as household resources during childhood.

Third, a general limitation of polygenic scores is that they lose predictive power when applied to ethnic groups different from those represented in the sample of their GWAS (see Martin et al., 2017). In the case of the EA score, the discovery sample used by Lee et al. (2018) was comprised of individuals of European ancestry. This limits our ability to extrapolate the mechanisms that we study to people of different ethnic groups.³

Fourth, since the EA score is an out-of-sample predictor formed using a series of estimated linear models, it could be subject to measurement error and misspecification with respect to what would be a theoretical optimal genetic predictor of educational attainment. This would affect the empirical moments that we use to calibrate our models and our structural estimates as a result. Nevertheless, we would expect these issues to attenuate the relationships that we observe and model, making our estimates act as conservative bounds.

2.3 Social Scientific Research Using the EA Score

A common finding in studies that have associated the EA score with socioeconomic outcomes is that there are important interactions between the genetic endowments that the EA score summarizes and the environment, sometimes called gene-by-environment or $G \times E$ interactions. For instance, Papageorge and Thom (2020) and Ronda et al. (2022) find that the returns to the EA score in terms of educational attainment are higher for people growing

²Additionally, it has been found that various socioeconomic choices and outcomes like wages and occupational choices are best explained as functions of multiple dimensions of skills—e.g., cognitive, non-cognitive, motor—instead of a single measure of “ability” (see Willis and Rosen, 1979; Heckman, 1995; Heckman and Rubinstein, 2001; Heckman, Stixrud, and Urzua, 2006).

³Related to this issue, it is common practice to address concerns that genetic relationships reflect population stratification by adjusting for the first 10 principal components of the full matrix of genetic data.

up in households with higher socioeconomic status. $G \times E$ interactions are an important part of research using genetics since environments are policy-variant. In some research, the aim is to directly consider how policies affect people with different genetic endowments. Schmitz and Conley (2017) show that Vietnam-era military conscription had an adverse effect on post-service schooling for veterans with below-average EA scores, but find no effect for those with above-average EA scores. Barth, Papageorge, and Thom (2020) also show that the relationship between the EA score and wealth is attenuated for people who have a defined-benefit pension plan. The interactions are not limited to the EA score: Barcellos, Carvalho, and Turley (2018) show that a policy that increased secondary education improved middle-age health outcomes and had a larger effect for people with a high genetic predisposition to obesity (measured using a polygenic score for body-mass index), thus reducing gaps in health outcomes associated with genetic endowments.

Disciplines other than economics, such as sociology, psychology, and behavioral genetics have long explored pathways through which traits and outcomes that economists would relate to human capital are produced and transmitted between generations. Studies in these areas recognize the importance of interactions between genetic endowments and environments for intellectual development, educational attainment, and socioeconomic outcomes (Scarr-Salapatek, 1971; Guo and Stearns, 2002; Turkheimer et al., 2003; Nisbett et al., 2012; Belsky et al., 2016, 2018). These interactions have been studied in different settings, at different stages of life and using different methods and, as shown by Tucker-Drob and Bates (2016), the level of complementarity between genetic endowments and childhood socioeconomic status in the production of academic achievement varies between countries, suggesting that the interaction can change as a function of societal conditions and policies. Therefore, examining these interactions can enhance our understanding of whether social policies reduce or entrench existing inequalities (Harden and Koellinger, 2020).

A natural next step in the use of genetic data in the social sciences is to incorporate the heterogeneity and interactions identified from studying polygenic scores into models of human capital accumulation, life-cycle decisions, and social policies. As Harden and Koellinger (2020) put it, social scientists can use polygenic scores as a type of “tracer dye” to reveal how predispositions, environments, policies, and life events interact in the production of adulthood outcomes, which is consistent with looking more carefully at dynamic life-cycle behavior. Moreover, while earlier work shows relationships between genetic endowments and socioeconomic differences, it cannot tell us how these relationships would change in counterfactual settings, for instance different social policies, labor-market conditions, tax systems, or incentives to education. Interactions between genetic data and previous policy changes allow the researcher to conduct ex-post policy evaluations, e.g., to understand how

education policy affected people with different EA scores (Barcellos, Carvalho, and Turley, 2018). In contrast, ex-ante policy analysis requires an explicit model that can be used to simulate counterfactual policies or environments (Wolpin, 2013). Genetic data combined with models of economic behavior permit examination of counterfactual gene-by-environment interactions.

To our knowledge, the only other paper that develops and estimates a structural model to examine genetic endowments related to education attainment is Houmark, Ronda, and Rosholm (2020).⁴ The authors incorporate genetic data into a model of skill formation and parental investments. The structure in their model means they can address questions about counterfactual policies, e.g., they can simulate the relationship between the EA score and children’s skills in a counterfactual environment where parental investments are held fixed across the population. In our case, we develop a model of lifecycle behavior where choices are an explicit result of dynamic utility maximization under uncertainty. Choices are optimal reactions to endowments (including genetic endowments) and the environment and endogenously adjust to changes in policies. The model also allows us to measure welfare so that we can understand how counterfactual policies not only affect optimal choices and outcomes, but also well-being for people with different genetic endowments. This, as we show, is important since policies that reduce inequality on some outcomes for people with different EA scores can increase inequality in terms of lifetime utility.

3 Data

3.1 The Health and Retirement Study and Our Analytical Sample

Our empirical analysis uses data from the Health and Retirement Study (henceforth HRS). The HRS surveys a representative sample of more than 20,000 Americans over the age of 50 and their spouses. The longitudinal design of the survey features biennial waves starting from 1992 and continuing until 2020, which provide information on respondents’ labor supply, income, wealth, financial decisions, retirement, mortality, and inheritances. The data set can also be linked to Social Security Administration records to provide life-cycle income and also includes retrospective questions about childhood socioeconomic status.⁵ Crucially, the HRS also contains genetic information via the collection of genetic samples from over 18,000 respondents starting in 2006. This has made it possible to construct various polygenic scores

⁴Biroli (2015) uses a structural model to examine genetic markers related to obesity.

⁵The HRS also collects information on a host of factors that are omitted from our analysis, including variables on health and family structure.

from the GWAS literature for the HRS respondents and to study associations between these scores and other HRS variables.

Data availability and issues surrounding the interpretation of polygenic scores detailed in the previous section place several restrictions on our sample. We begin with the sample used in Barth, Papageorge, and Thom (2020) which only includes households with members of European ancestry. As explained in (Martin et al., 2017), a set of technical issues prevent credible comparisons of polygenic scores across ethnic groups and associations between polygenic scores and outcomes tend to be weaker and noisier for individuals of non-European descent so incorporation of non-European households into our analysis would be misguided and could generate misleading conclusions about cross-ethnic group differences. The sample is also limited to households with non-missing data on wealth, stock market participation, and Social Security Administration earnings records. These requirements permit a maximum sample of 2,590 households (5,701 household-year observations) from the overall HRS sample of over 20,000 households and over 160,000 household-year observations.

We make further sample restrictions aligned to the structural model. In particular, the model is a unitary household model that abstracts from joint labor supply decisions and marriage dynamics. We focus on households that i) enter the HRS panel as married or partnered two-person male-female households; ii) that remain intact (except due to the death the female head of household); iii) are not observed earning income from jobs not covered by the Social Security Administration’s data, iv) receive at least 70 percent of their SSA earnings income from the male partner, and v) have non-missing genetic data for the male partner, which we take to be the genetic endowment of the household. This set of restrictions generates a sample of 870 households with wealth data observed for a total of 2,318 household-year observations.⁶

Table 1 presents basic descriptive statistics for the analytic sample. Approximately 30 percent of the sample has a college degree, making the sample more highly educated than the overall HRS population. Retirement rates increase from approximately 17 percent for ages 50-62 to about 59 percent for ages 63-67. By age 73, 86 percent of the sample has retired. Table 1 also presents basic descriptive statistics on the log of total prime-age SSA earnings. To construct this variable, we sum over male earnings over the entire age range 30-60. The median of total income over this age range is \$1.9 million, which would constitute an average of approximately \$59,000 per year in 2010 dollars. Table 1 also provides detailed descriptive statistics for wealth for household-year observations in which the male household member was aged 60-70. The median wealth in the sample is approximately \$372,000.

⁶Samples used for the inheritance and income processes are described in Appendix A.

3.2 Genetic Endowments and Family Environments

A natural concern is that genetic endowments are endogenous to family environments. The idea is that parents who provide their children with genetic material also provide them with family environments, including resources that could promote educational attainment and labor market outcomes. If so, estimates of coefficients relating the EA score to outcomes, including education, are likely to be upwardly biased. A host of studies have employed different methods to address this concern. For example, Trejo and Domingue (2018) and Belsky et al. (2018) rely on within-sibling variation in the EA score. This amounts to adjusting for a family fixed effect and nearly all relationships hold despite some differences. Another method is to control for a rich set of variables that describe childhood environments. In fact, Ronda et al. (2022) show that controlling for measures of family background can reduce the bias substantially.

Since the HRS does not permit within-family analysis, our approach is to incorporate information on childhood SES to essentially control for key dimensions of childhood socioeconomic status. Here, we discuss how we summarize this information into a single variable. In section 4, in which we specify the structural model, we discuss how this variable enters the model through our specification of unobserved heterogeneity that we allow to relate to the EA score.

To construct a summary SES measures, for the male members of each household in our sample, we estimate a cross-sectional regression of the following form:

$$\text{Educ}_i = b_0 + b_1 \text{EA}_i + b_2 X_i + e_i \quad (1)$$

Here X_i contains a vector of background variables that includes: mother’s years of schooling, father’s years of schooling, dummy variables for different subjective assessments of family SES growing up, dummy variables for categories of father’s occupation growing up, and dummy variables for subjective assessments of health in childhood. Table 1 presents descriptive statistics on the components of X_i . After estimating the above equation, we construct an index of childhood SES as: $\text{SES Score}_i = \widehat{b}_2 X_i$. We normalize this measure so that it has a mean of zero and a unit standard deviation.

Table 2 presents some basic regressions that highlight the critical relationships between the EA Score, wealth, and stock market participation that motivate our structural model. Panel A presents regressions of log household wealth on explanatory variables including the EA Score for household-year observations where the male household member is aged 60-70. All columns include the first 10 principal components of the SNP-level genetic data. Column

(1) includes the EA Score, the SES Score, and a dummy variable for college education as controls. The coefficients on all three variables are substantial and statistically significant. The results suggest that a one standard deviation higher EA Score is associated with 27 percent higher wealth. Column (2) adds the log of total prime-age income, and Column (3) adds a dummy variable for holding any stocks in a given household-year. Controlling for income reduces the coefficient on the EA Score to 0.20, while additionally controlling for stocks reduces the coefficient to 0.15. However, this coefficient remains significant, suggesting that in this sample, even controlling for college, childhood SES, life-time earnings, and stocks, a one standard-deviation higher EA score is associated with approximately 15 percent higher household wealth. Panel B of Table 2 presents regressions of a dummy variable for holding any stocks on the principal components, the EA score, college, and childhood SES. The estimates in Column (1) suggest that a one standard deviation higher EA score is associated with 6.8 percentage point increase in the likelihood of owning stocks. This coefficient is attenuated (0.049 v.s. 0.68), but remains highly statistically significant after controlling for lifetime earnings in Column (2).

4 Model

The model comprises of heterogeneous agents that live from age 21 to a maximum age of 90. Each year, an agent decides how much wealth to consume, how to allocate savings between a risky and a risk-free asset, and—once old enough—whether to retire or continue working for another year. Below we describe each of the components of the agents’ dynamic problem.

4.1 Ex-Ante Heterogeneity

Agents are indexed by subscript i . Agents enter the model with four dimensions of observable heterogeneity: birth year (BY_i), an indicator for college completion ($Coll_i$), the EA3 score (EA_i), and an indicator equal to one if the agent participates in a defined benefit pension plan (DB_i). We allow these characteristics to influence agents’ income, utility cost of work, expected returns on the risky asset (stocks), and the cost of stock market participation.

We also allow unobserved heterogeneity to affect agents’ earnings, stock market participation costs, and expected stock returns. We model these three dimensions of heterogeneity as a vector of individual-specific “fixed effects” $\vec{\zeta}_i = [\zeta_i^w, \zeta_i^F, \zeta_i^R]'$. These parameters account for heterogeneity along dimensions that we do not model directly. One particularly important source of heterogeneity is childhood socioeconomic status and rearing environment. If the EA score in part represents genetic endowments that are conducive to achieving higher

Table 1: SUMMARY STATISTICS

Variable	Mean	SD	N
Birth Year	1939.5	5.81	870
College	0.30	0.46	870
Retired			
50-62	0.17	0.37	3,350
63-67	0.59	0.49	1,786
68-72	0.78	0.42	1,393
73+	0.86	0.35	503
Total Prime-Age Income (\times \$1,000)			
Mean	1,984.48	.	870
Std. Dev.	802.85	.	870
25th Percentile	1,451.90	.	870
50th Percentile	1,928.22	.	870
75th Percentile	2,456.86	.	870
Household Wealth (\times \$1,000)			
Mean	716.34	.	2,318
Std. Dev.	926.29	.	2,318
10th Percentile	52.12	.	2,318
25th Percentile	150.90	.	2,318
50th Percentile	372.03	.	2,318
75th Percentile	872.71	.	2,318
90th Percentile	2,871.72	.	2,318
Any Stocks	0.68	0.47	2,318
Components of Childhood Socioeconomic Status (SES)			
Mother's Educ.	10.52	2.97	766
Father's Educ.	10.20	3.58	766
Family SES: Well Off	0.07	.	766
Family SES: Average	0.67	.	766
Family SES: Poor	0.24	.	766
Family SES: Varied	0.01	.	766
Father Job: Manger / Prof.	0.18	.	766
Father Job: Sales	0.07	.	766
Father Job: Clerical	0.03	.	766
Father Job: Service	0.04	.	766
Father Job: Manual / Operators	0.64	.	766
Father Job: Armed Forces	<0.01	.	766
Father Job: Don't Know	0.03	.	766
Father Job: Missing	<0.01	.	766
Child Health: Excellent	0.55	.	766
Child Health: Very Good	0.27	.	766
Child Health: Good	0.13	.	766
Child Health: Fair	0.03	.	766
Child Health: Poor	0.01	.	766

Notes: Summary statistics for demographics, income, wealth, and SES variables in our main analytical sample

levels of income and financial sophistication, then we would expect the average-, low-, and high-EA individuals to grow up in different environments. This indirect channel, known as

Table 2: Summary Statistics: Mechanisms

Panel A: log Wealth			
	[1]	[2]	[3]
EA Score	0.276 (0.054)	0.203 (0.050)	0.152 (0.045)
SES Score	0.294 (0.054)	0.222 (0.058)	0.169 (0.052)
College	0.527 (0.109)	0.331 (0.105)	0.260 (0.095)
log Prime Inc.		0.969 (0.195)	0.699 (0.180)
Any Stocks			1.058 (0.093)
N	2259	2259	2259

Panel B: Any Stocks		
	[1]	[2]
EA Score	0.068 (0.015)	0.049 (0.014)
SES Score	0.069 (0.016)	0.050 (0.015)
College	0.119 (0.032)	0.067 (0.030)
log Prime Inc.		0.255 (0.030)
N	2259	2259

Notes: This table reports results from models predicting log Wealth (Panel A), and a dummy variable for any stocks (Panel B). All regressions include the first 10 principal components of the genetic data as controls.

genetic nurture, could lead us to overestimate the influence of the EA score (Kong et al., 2018; Young et al., 2018; Ronda et al., 2022). To address this issue, we model unobserved heterogeneity as a combination of a fully random component (\vec{Z}_i) and a component that is correlated with the EA score through childhood socioeconomic status (SES_{*i*}),

$$\underbrace{\begin{bmatrix} \zeta_i^w \\ \zeta_i^F \\ \zeta_i^R \end{bmatrix}}_{\vec{\zeta}_i} = \underbrace{\begin{bmatrix} z_w \\ z_F \\ z_R \end{bmatrix}}_{\vec{z}} \times \text{SES}_i + \underbrace{\begin{bmatrix} Z_i^w \\ Z_i^F \\ Z_i^R \end{bmatrix}}_{\vec{Z}}, \quad \begin{bmatrix} Z_i^w \\ Z_i^F \\ Z_i^R \end{bmatrix} \sim \mathcal{N}(\vec{0}, \Sigma_Z) \quad (2)$$

where $\Sigma_{\mathcal{Z}} = \text{diag}[\sigma^2(\mathcal{Z}^w), \sigma^2(\mathcal{Z}^F), \sigma^2(\mathcal{Z}^R)]$ and

$$\text{SES}_i = \phi EA_i + \varepsilon_i, \quad \varepsilon_i \sim \mathcal{N}(0, \sigma_\varepsilon^2). \quad (3)$$

In Equations (2) and (3), \vec{z} , $\Sigma_{\mathcal{Z}}$, ϕ and σ_ε^2 are parameters to be estimated.

4.2 Utility

A surviving agent i derives period t utility from consumption and leisure through the utility function

$$u_{i,t}(C_t, \ell_t) = \frac{[C_t^\gamma (1 - 0.34 \times \ell_t)^{1-\gamma}]^{1-\omega}}{1-\omega} - d_{i,t} \times \ell_t,$$

where C_t is consumption and ℓ_t is a binary variable that indicates whether the agent is working ($\ell_t = 1$) or not ($\ell_t = 0$). This assumes that labor is indivisible and that it consumes 34% of an agent's endowment of time (8 of 24 hours per day), which we normalize to 1. A Cobb-Douglas function aggregates consumption and leisure with a weight $\gamma \in [0, 1]$ on consumption. The aggregate passes through a constant relative-risk-aversion function with coefficient of relative risk aversion ω . The term $d_{i,t} \times \ell_t$ is an additive utility cost of work. We allow this cost to vary with agents' education, EA score, and age:

$$d_{i,t} = d_0 + d_{\text{Coll}} \times \text{Coll}_i + d_{\text{EA}} \times \text{EA}_i + d_{\text{Age}} \times \max\{\text{Age}_{i,t} - 50, 0\}. \quad (4)$$

The heterogeneity of this cost is intended to capture potential differences in the physical intensity of jobs that may covary with certain traits.

Finally, to accommodate the fact that high-income households have higher saving rates and are slow to run down their wealth at old ages (Dynan, Skinner, and Zeldes, 2004), we include a "joy of giving" bequest motive (Carroll, 2002) that produces utility from end-of-life wealth. We use the functional form of De Nardi, French, and Jones (2010), in which a person who dies with savings $S_{i,t}$ receives utility

$$\varphi(S_{i,t}) = \theta \frac{(S_{i,t} + \kappa)^{1-\omega}}{1-\omega}, \quad (5)$$

where ω is the same coefficient of relative-risk-aversion as above, and θ and κ are parameters that we estimate and which govern the intensity of the motive and the degree to which bequests are a luxury good.

4.3 Labor Income

We model pre-tax labor income $\tilde{W}_{i,t}$ as the log-sum of a deterministic component that depends on individual characteristics and aggregate trends, the fixed unobserved heterogeneity draw \mathcal{Z}_i^w , and an agent- and time-specific shock $\epsilon_{i,t}^w$,

$$\ln \tilde{W}_{i,t} = f(\text{Age}_{i,t}, \text{EA}_i, \text{Coll}_i, \text{SES}_i, \text{DB}_i, \text{Year}_t, \text{Unemp}_t) + \mathcal{Z}_i^w + \epsilon_{i,t}^w, \quad (6)$$

where Unemp_t is the aggregate unemployment rate in period t . The wage shock is independent across both time and agents and is normally distributed, $\epsilon_{i,t}^w \sim \mathcal{N}(0, \sigma^2(\epsilon^w))$. We present our full specification and estimates for pre-tax labor income in Appendix A.

We use two types of labor income taxes: a constant-rate tax (as in Chai et al., 2011), and an additional proportional tax that applies up to a year-dependent maximum $\bar{T}(\text{Year})$ and represents Social Security taxes.⁷ We represent this tax scheme with a function $\tau_t(\cdot)$, that computes post-tax income as:

$$\tau_t(\tilde{W}) = \tilde{W}(1 - \tau^W) - \tau^{\text{FICA}} \times \min\{\tilde{W}, \bar{T}(\text{Year}_t)\}, \quad (7)$$

where τ^W and τ^{FICA} are the respective tax rates.

4.4 Retirement and Social Security

Once agents reach a minimum age they can decide to retire. This decision is irreversible and happens after receiving labor income, so that it takes effect in the following year. We impose the restriction that the first year of retirement has to occur in an age interval $[\text{Age}_0^R, \text{Age}_f^R]$, which we set to $[62, 80]$. Agents with defined benefit pension plans start receiving their payments in the first year in which they do not work. We assume that agents start claiming Social Security benefits the moment they stop working, or at the minimum claiming age ($\text{Age}_{\min}^{\text{SS}}$) if they stop working before that.⁸

Consistent with the payout policy in the U.S., SS benefits are a concave function of agents' average income over their highest-earning years and increase with each additional year of work up to a maximum. Our methodology for computing benefits follows that of the Social Security Administration; we provide a detailed description in Appendix C. Our main simplification is that we use expected (rather than realized) earnings when computing

⁷We use the Social Security Administration's historical maximum taxable incomes as $\bar{T}(\text{Year}_t)$.

⁸In our baseline scenario, the minimum claiming age and minimum retirement age are both 62, and therefore agents will always start claiming Social Security benefits the same year they stop working. This changes in our counterfactual policy experiments one of which shifts the minimum claiming age.

SS benefits. That is, in each period t in which the agent could choose to retire, instead of the agent looking backwards at her *realized* earnings (the top 35 years of which determine SS benefits), the agent anticipates receiving SS benefits that are determined by her *expected* top-35 yearly earnings as of period t , based on the income function described in Equation (6). This allows us to avoid including the additional state variable of realized cumulative earnings. Additionally, because the model has no permanent income shocks, transitory shocks over the life-cycle should largely cancel out, and expected earnings should be a good representative of realized earnings. We use $SSB_i(n)$ to denote the yearly benefits that person i would receive if she retired in the n -th possible year.

In addition to social security benefits, retirees who have a defined benefit pension plan ($DB_i = 1$) receive pension income. We model the annual amount of DB pension payments (DBf_i) as a log-linear function of the time-invariant components of income, EA_i and $Coll_i$, which we estimate from the data. This is similar to our treatment of SS income: to avoid the introduction of an additional state variable, we tie DB retirement benefits to the predictable component of earnings rather than realized earnings. We present the full specification and estimates in Appendix A.

Both Social Security benefits and defined benefit pension flows are taxed at a constant rate τ^s .

4.5 Inheritances

Every period, agents receive inheritances ($Inher_{i,t}$) that follow the agent- and age-specific stochastic process

$$Inher_{i,t} = \begin{cases} C.Inher_i(Age_{i,t}), & \text{With probability } P_i^I(Age_{i,t}) \\ 0, & \text{With probability } 1 - P_i^I(Age_{i,t}). \end{cases} \quad (8)$$

Both the probability of receiving an inheritance P_i^I and the value of inheritances conditional on reception $C.Inher_i$ depend on $Coll_i$ and EA_i . We estimate both functions using our HRS sample; their specification and parameter estimates can be found in Appendix A.

4.6 Financial Assets

At the end of each period, agents decide how to allocate their savings between two financial assets: a risk-free bond with return factor R , and a risky asset representing the stock market with an agent- and time-specific return factor $\tilde{R}_{i,t}$. Short-selling of either asset is not

permitted.⁹

For agent i , the risky return factor follows the process

$$\ln \tilde{R}_{i,t} = \ln R_t^{\text{SP500}} - \mu^{\text{SP500}} \times g(r_0 + r_{\text{Coll}} \times \text{Coll}_i + r_{\text{EA}} \times \text{EA}_i + \zeta_i^R) \quad (9)$$

where $\ln R_t^{\text{SP500}}$ and μ^{SP500} are the log-return of the S&P 500 stock market index in year t and its mean log-return, respectively. The factor $g(\cdot)$ captures agents' degree of inefficiency when investing in risky assets. It takes the form of a logistic function

$$g(x) = \frac{e^x}{1 + e^x},$$

which ranges from 0 to 1. An efficient agent ($g \approx 0$) will replicate the market's returns. An inefficient agent ($g \approx 1$) will have expected log-returns close to 0. This parsimonious specification is similar to that of Lusardi, Michaud, and Mitchell (2017). Note that all agents face the same risk on the risky asset return, only the mean return is agent-specific.

We conceive of heterogeneity in risky asset returns as proxying for sound investment decisions. Paying higher fees on mutual fund investments or excessive trading (and resulting taxes), for instance, would degrade the average return earned in the market. Inopportune market-timing strategies, such as buying during periods of high price-earnings ratios and selling during periods of low price-earnings ratios, would also be detrimental to returns. Further, because risk is not agent-specific, lower expected returns also mean lower Sharpe ratios. From this perspective, a (roughly) equivalent interpretation of lower expected returns would be higher risk for a given level of expected return.¹⁰ Greater risk exposure conditional on expected returns would be consistent with poor diversification. Agents assume stock returns follow a normal distribution given by $\ln R_t^{\text{SP500}} \sim \mathcal{N}(\mu^{\text{SP500}}, \sigma^{\text{SP500}})$.

To own the risky asset, agents must pay a per-period monetary participation cost that represents the administrative and opportunity costs of managing investments (Vissing-Jorgensen, 2002). The cost F_i depends on ex-ante demographic characteristics

$$\ln F_i = f_0 + f_{\text{Coll}} \times \text{Coll}_i + f_{\text{EA}} \times \text{EA}_i + \zeta_i^F. \quad (10)$$

Capital gains are taxed at a constant rate τ^c .

⁹This means borrowing is prohibited as borrowing constitutes a short position in the risk-free bond.

¹⁰In the simple case of static mean-variance preferences, the optimal investment in the risk portfolio is given by $\omega_p^* = \frac{E[R_p] - r_f}{\sigma_p^2 \gamma}$, where γ is risk aversion. It is clear in this setting that expected return and variance are substitutes; decreasing $E[R_p]$ is equivalent to increasing σ_p^2 .

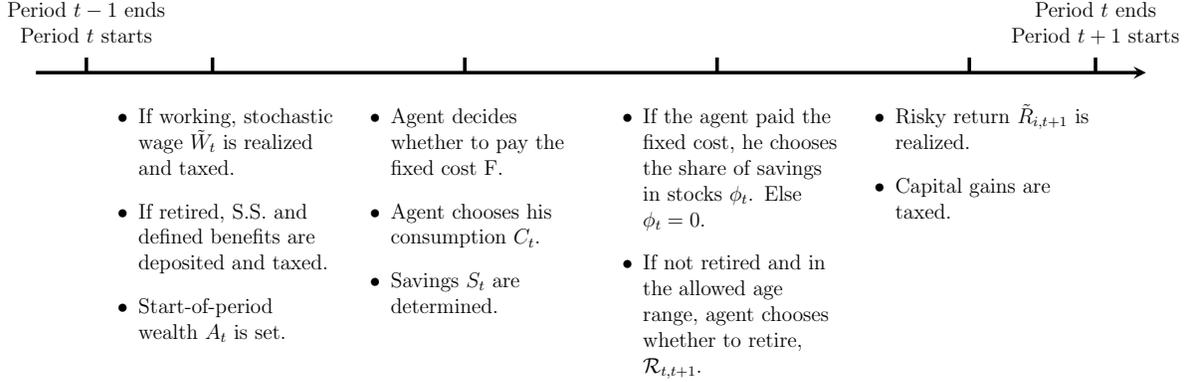


Figure 1: Timing of decisions and shock realizations.

4.7 Recursive Representation and Timing Summary

Agents maximize their expected discounted lifetime utility using a discount factor β and taking into account their probability of survival δ_t . An agent's possible choice variables at a given time are his consumption $C_{i,t}$, the fraction of his savings allocated to the risky asset $\phi_{i,t}$, and his retirement status next period $\mathcal{R}_{i,t+1}$. His state vector consists of his beginning-of-period wealth $A_{i,t}$ and his retirement status $\mathcal{R}_{i,t}$. Retirement status is a discrete variable that follows the convention

$$\mathcal{R}_{i,t} = \begin{cases} 0 & \text{If the agent has not retired yet} \\ n & \text{If the agent retired in the } n\text{-th possible period} \end{cases} \quad (11)$$

for $n = 1, 2, \dots, \text{Age}_f^R - \text{Age}_0^R + 1$. Figure 1 summarizes the timing of decisions and shocks in our model.

The choices and constraints that an agent faces depend on his age and retirement status.

The value function of an agent with the option to retire is

$$V_{i,t}(A_{i,t}, \mathcal{R}_{i,t} = 0) = \max_{C_{i,t}, \phi_{i,t}, \mathcal{R}_{i,t+1}} u(C_t, \ell = 1) + \beta \delta_t \mathbb{E}_t [V_{i,t+1}(A_{t+1}, \mathcal{R}_{i,t+1})] + \delta_t \varphi(S_{i,t})$$

Subject to:

$$0 \leq C_{i,t}, \quad 0 \leq S_{i,t}, \quad 0 \leq \phi_{i,t} \leq 1$$

$$\mathcal{R}_{i,t+1} \in \{0, (\text{Age}_{i,t} + 1 - \text{Age}_0^R) + 1\}$$

$$S_{i,t} = A_{i,t} - C_{i,t} - F_i \times \mathbf{1}[\phi_{i,t} > 0],$$

$$A_{i,t+1} = \left\{ (1 - \tau^c) \left[\phi_{i,t} \tilde{R}_{i,t+1} + (1 - \phi_{i,t}) R \right] + \tau^c \right\} \times S_{i,t} +$$

$$\mathbf{1}[\mathcal{R}_{i,t+1} = 0] \times \tau_t (\tilde{W}_{i,t+1}) +$$

$$\mathbf{1}[\mathcal{R}_{i,t+1} > 0] \times (1 - \tau^s) \times \text{DBf}_i +$$

$$\mathbf{1}[\mathcal{R}_{i,t+1} > 0, \text{Age}_{i,t+1} \geq \text{Age}_{\min}^{SS}] \times (1 - \tau^s) \times \text{SSB}_i(\mathcal{R}_{i,t+1}) +$$

$$\text{Inher}_{i,t+1},$$

(12)

where $S_{i,t}$ denotes the agent's savings and $\delta_t \equiv 1 - \delta_t$ is the probability of death. The recursive representation for the cases where the agent is not yet able to retire or is retired can be found in Appendix D.

5 Estimation

The model described in Section 4 contains parameters that are estimated as well as parameters whose values we take from the existing literature or other sources (such as the U.S. Social Security tax and benefits schedules). The values of non-estimated parameters are shown in Table 3.¹¹

We estimate the model in two steps. In the first step, we directly estimate the specifications of wages (Equation 6), inheritances (Equation 8), and defined benefit pension flows on the HRS sample data. This pins down model parameters that we estimate outside of the method of moments machinery. The details and results of the first step can be found in Appendix A.

In the second step, we use the method of simulated moments (MSM) to estimate the

¹¹The data that we use for our calibration of the distribution of stock-market returns comes from the accompanying data file to Chapter 26 of Shiller (1990), which is available on the author's website. The file, as of June 2020, is located at <http://www.econ.yale.edu/shiller/data/chapt26.xlsx>

Table 3: Non-estimated parameters.

Parameter	Symbol	Value	Source
Coef. of relative risk aversion	ω	1.6	Lusardi, Michaud, and Mitchell (2017)
Consumption’s share of utility	γ	0.39	Imrohoroğlu and Kitao (2012)
Yearly risk-free return factor	R	1.015	
Yearly utility discount factor	β	0.96	
Yearly survival probabilities	$\{\delta_t\}_{t=1}^T$		SSA Life Tables. Based on males born in 1940.
Mean of real S&P500 log-returns	μ^{SP500}	0.0645	Historical, 1936-2018
Std. Dev. of real S&P500 log-returns	σ^{SP500}	0.1649	Historical, 1936-2018
Minimum and maximum retirement years	$\{\text{Age}_0^R, \text{Age}_f^R\}$	$\{62, 80\}$	
Minimum and “full” S.S. benefit claiming years	$\{\text{Age}_{\min}^{SS}, \text{Age}_{\text{Full}}^{SS}\}$	$\{62, 67\}$	
Income tax rate (soc. sec)	τ^{FICA}	0.06	
Income tax rate	τ^W	0.24	
Capital gains tax rate	τ^c	0.20	
S.S. benefits tax rate	τ^s	0.15	

parameters that govern the financial costs to stock market participation (Equation 10), inefficiencies in risky investments (Equation 9), the disutility of work (Equation 4), the dispersion of unobserved heterogeneity (Equation 2), the bequest motive (Equation 5), and the influence of childhood socioeconomic status on the unobserved heterogeneity draws for costs and returns (Equation 2).¹² The full set of parameters that we estimate internally is

$$\Theta = \{f_0, f_{\text{Coll}}, f_{\text{EA}}, \sigma(\mathcal{Z}^F), r_0, r_{\text{Coll}}, r_{\text{EA}}, \sigma(\mathcal{Z}^r), d_0, d_{\text{Coll}}, d_{\text{EA}}, d_{\text{Age}}, \theta, \kappa, z_F, z_R\}. \quad (13)$$

Specifically, define \mathcal{M} as the vector empirical sample moments to be matched in the estimation. The MSM methodology begins by generating counterparts to \mathcal{M} from a simulated sample, generated conditional on a candidate set of parameters $\tilde{\theta}$. To generate the synthetic sample, we match the sample data’s joint distributions of birth years, college completion, EA3 polygenic scores, and defined benefit pension plan participation. To do so, we discretize the EA score by its deciles and group individuals into five-year birth cohorts from 1915 to 1960. This gives us 400 possible combinations of observable characteristics: ten EA score

¹²Since income is observable, we estimate the parameters pertaining to unobserved heterogeneity of income ($z_W, \sigma(\mathcal{Z}^W)$) directly. See Appendix A.

deciles, times ten birth cohorts, times two educational attainment levels (a completed college degree or not), times two possible pension possibilities (participating in a defined benefit plan or not). Only 190 out of the 400 combinations are populated in the sample data. Denote by N_q the number of observations in each of the $q = 1, 2, \dots, 190$ populated combinations. We expand the synthetic sample by multiplying N_q by 10 for each q . Note that this expands the synthetic sample by a factor of 10 but preserves the relative distribution of characteristics.

We then introduce unobserved heterogeneity. Given a candidate set of parameters and an agent's EA score, Equations (2) and (3) imply:

$$\vec{\zeta}_i | EA_i \sim \mathcal{N}(z\phi EA_i, \sigma_{\text{SES}}^2 z z' + \Sigma_Z). \quad (14)$$

We discretize the distribution in Equation (14) using 27 equiprobable points for $\vec{\zeta}_i$ (three equiprobable values for each of the three sources of unobserved heterogeneity, for a total of 27 possible triplets). Lastly, we simulate 27 agents, one for each of the 27 possible unobserved heterogeneity draws, for each of the $N_q \times 10$ agents that populate the 190 populated bins.¹³ That is, each of the 190 non-empty combinations of observable characteristics in the HRS sample will contain 27 different simulated agents, each with different unobserved-heterogeneity draws, replicated $N_q \times 10$ times. Expanding the 190 groups of observable characteristics with the 27 EA-specific draws of unobserved heterogeneity leaves us with 5,130 ex-ante types of agents. These types constitute the full structure of ex-ante heterogeneity.

We then solve the life-cycle problem by backward induction for the 5,130 types, which delivers policy functions for choice variables conditional on states. Appendix E presents details on how we solve the model. We then simulate the lives of our entire population of agents. We generate draws of lifetime labor income following Equation (6) and mortality draws from the survival rates $\{\delta\}_{t=21}^{90}$. For stock-market returns R_t^{SP500} , we take the realized annual return of the S&P 500 index for the relevant year, which is determined by each agent's birth-year and age. This produces a random sample of synthetic life cycles, and we calculate counterparts of our targeted moments on the simulated data. Note that this entire procedure is conditional on *a given set of parameters* $\tilde{\theta}$. Denote by $\widehat{\mathcal{M}}(\tilde{\theta})$ the simulated moments conditional on $\tilde{\theta}$.

We define the loss function for a set of parameters $\tilde{\Theta}$ given K empirical moments as:

$$L(\tilde{\theta}) = \sum_{k=1}^K \left(\frac{\widehat{\mathcal{M}}(\tilde{\theta})_k - \mathcal{M}_k}{\mathcal{M}_k} \right)^2. \quad (15)$$

¹³We set the socioeconomic status of simulated agents to $E[\text{SES}_i | EA_i, \vec{\zeta}_i]$. See Appendix B for details.

We estimate the model by minimizing the loss function¹⁴:

$$\hat{\theta} = \arg \min_{\theta} L(\theta).$$

5.1 Empirical Moments to Match

We choose the following empirical moments, which combine moments of relevant distributions (such as the distribution of wealth) and parameters from regressions of relevant outcomes (such as stock market participation) on covariates, in order to identify the estimated model parameters:

1. *Wealth distribution*: The mean and percentiles 10, 25, 50, 75, and 90 percentiles of the wealth distribution for people in the age range 60 – 70.
2. *Wealth regression coefficients*: The coefficients from a regression of log-wealth on the EA score EA_i , our index of childhood socioeconomic status SES_i , a binary indicator of college completion $Coll_i$, the log “earned income” $\ln \text{Earned Income}_i$ which we define as the sum of labor income earned between the ages of 30 and 60, a binary indicator of stock ownership $Stocks_{i,t}$, and age. We use only people aged 60 to 70 that report strictly positive wealth. We target the coefficients on $Coll_i$, EA_i , SES_i , and $Stocks_{i,t}$.
3. *Stock ownership rate*: The fraction of individual-year observations in our sample that report owning any stocks ($Stocks_{i,t} = 1$). We consider people aged 60 to 70.
4. *Stock ownership regression*: The coefficients from a regression of the stock ownership indicator $Stocks_{i,t}$ on the EA score EA_i , our index of childhood socioeconomic status SES_i , a binary indicator of college completion $Coll_i$, the log of earned income $\ln(\text{Earned Income}_i)$, and age. We include only people aged 60 to 70 that report strictly positive wealth and do not include the coefficient on age as a moment to match.
5. *Retirement rates*: The fraction of people in each of the following age brackets who are retired: [50, 62], [63, 67], [68, 72] and ≥ 73 .
6. *Retirement regression*: The coefficients from a regression of an indicator for whether an individual is retired $Retired_{i,t}$ on the EA score EA_i , our index of childhood socioeconomic status SES_i , a binary indicator of college completion $Coll_i$, and age. We consider

¹⁴We solve the minimization problem using the Nelder-Mead algorithm with multiple starting points. We start by evaluating the objective function in 20,000 initial points and then use the Nelder-Mead algorithm starting at the 25 points with the smallest loss functions. We take the result with the lowest associated loss as our estimate.

people aged 61 to 75 and only use the coefficients on Coll_i and EA_i as moments in the estimation.

5.2 Identification

In this section we briefly discuss how the variation in the empirical moments identifies the set of estimated parameters Θ . Identification is based on the joint distributions of moments, and generally individual moments do not alone identify a specific parameter. We do, however, select moments intentionally based on how their variability informs the model parameters.

First, consider the parameters associated with stock market participation costs ($f_0, f_{\text{Coll}}, f_{\text{EA}}, \sigma(\mathcal{Z}^F)$) and stock returns ($r_0, r_{\text{Coll}}, r_{\text{EA}}, \sigma(\mathcal{Z}^r)$). Intuitively, the joint distribution of income, wealth, and stock market participation identifies the cost of stock market participation and returns on risky assets. Note first that the income process is estimated directly, outside the MSM procedure. Then, holding risky asset returns fixed, the association between stock market participation and income identifies the distribution of participation costs, since higher average costs will predict lower participation at lower levels of income, and a larger variance in unobserved heterogeneity will imply a flatter relationship between income and stock market participation. Similarly, holding income and participation costs fixed, the distribution of wealth determines the returns on risky assets, since in the model the growth rate of savings is determined entirely from the return on assets.

The coefficient on SES in the wealth regression identifies z_R , the part of unobserved heterogeneity in stock returns that is determined by SES . Note that the coefficient estimate is not a direct estimate of z_R ; z_R is the unobserved heterogeneity in *returns* that are linked to SES , not wealth. However, the parameter z_R will affect the association between SES and wealth, and therefore the empirical association between SES and wealth implicitly identifies the association between returns and SES . Similarly, the coefficient on SES in the stock market participation regression identifies z_F , the parameter that determines the effect of SES on unobserved stock market participation cost heterogeneity.

In addition to the stock market parameters, the wealth distribution also identifies the two parameters governing the bequest motive, θ and κ . Without a bequest motive, households would spend down their wealth aggressively as they aged and the wealth distribution would be less skewed at the top. Because there are more wealth distribution moments than are needed to identify the coefficients related to stock market returns, the additional information contained in the distribution of wealth identifies the parameters governing households bequest incentives.¹⁵

¹⁵We have experimented with alternative models that omit the bequest motive. The key difference is that

Next, consider coefficients on the disutility of labor ($d_0, d_{\text{Coll}}, d_{\text{EA}}, d_{\text{Age}}$). Conditional on income and wealth, the retirement decision is influenced by how (un)enjoyable is working. The rates of retirement at each of the four age ranges therefore identifies the distribution of labor utility costs.

Finally, the distributions of stock market participation costs, stock market returns, and disutility of labor are each allowed to be influenced by covariates, in particular the EA score and education. The coefficients from regressions of wealth, stock market participation, and retirement on the EA score and education constitute moments that identify how these distributions shift with these controls.

6 Results

6.1 Model Fit, Parameter Estimates, and Their Implications

Table 4 presents parameter estimates and their standard errors. To find standard errors, we calculate the targeted moments for 50 different bootstrapped sub-samples of our analytical sample data. Within each subsample, we select the parameter vector from a pre-specified grid of 20,000 parameter vectors that minimizes the loss function (Equation 15) between the selected vector the vector of targeted moments. This gives 50 sets of parameter vectors, one for each subsample, each of which minimizes the loss function within that particular subsample. We compute standard errors as the standard deviation of each parameter over the 50 resulting parameter vectors.¹⁶

Table 5 shows that the estimated model is able to match most of the empirical patterns remarkably well. We closely match the distribution of wealth up to the 75th percentile, but under-predict wealth a bit in the upper tail of the distribution. We also match relationships between the EA polygenic score, college attendance, wealth, and stock ownership as captured by the auxiliary regressions. The overall stock ownership rate is also close to the sample rate. For retirement, while the model fits the auxiliary regression that relates retirement to the EA score and college attendance, the age-binned retirement rates show slight discrepancies arising mainly from over-estimating the fraction of agents that retire between ages 63 and 67.

According to estimates in Table 4, both the EA score and education have negligible effects on the cost of stock market participation, F_i . This is confirmed by Figure 2, which plots the

we fail to match upper-tail wealth, but key qualitative results are unchanged, i.e., our specification of the bequest motive does not drive our main results.

¹⁶The 20,000 pre-specified parameter vectors are the initial grid that we use in our main estimation routine.

Table 4: Internally-estimated parameters.

Participation cost			
$\ln F_i = f_0 + f_{\text{Coll}} \times \text{Coll}_i + f_{\text{EA}} \times \text{EA}_i + \zeta_i^F$			
f_0	f_{Coll}	f_{EA}	
-0.9867	0.0311	0.0066	
(0.2092)	(0.0369)	(0.0143)	
Risky asset returns			
$\ln \tilde{R}_{i,t} = \ln R_t^{\text{SP500}} - \mu^{\text{SP500}} \times g(r_0 + r_{\text{Coll}} \times \text{Coll}_i + r_{\text{EA}} \times \text{EA}_i + \zeta_i^R)$			
r_0	r_{Coll}	r_{EA}	
-0.0366	-1.1055	-0.6610	
(0.0608)	(0.2568)	(0.1326)	
Disutility from work			
$d_{i,t} = d_0 + d_{\text{Coll}} \times \text{Coll}_i + d_{\text{EA}} \times \text{EA}_i + d_{\text{Age}} \times \max\{\text{Age}_{i,t} - 50, 0\}$			
d_0	d_{Coll}	d_{EA}	d_{Age}
0.3961	-0.0052	-0.0033	-0.0241
(0.0816)	(0.0015)	(0.0009)	(0.0062)
Unobserved heterogeneity			
$\vec{\zeta}_i = \vec{z} \times \text{SES}_i + \vec{Z}_i$			
$\ln \sigma(\mathcal{Z}^F)$	$\ln \sigma(\mathcal{Z}^r)$	z_F	z_R
1.1838	-4.0326	-0.0434	-0.7250
(0.5698)	(1.3845)	(0.0475)	(0.1451)
Bequest motive			
$\varphi(S_{i,t}) = \theta(S_{i,t} + \kappa)^{1-\omega} / (1 - \omega)$			
$\ln \kappa$	$\ln \theta$		
7.0638	6.9423		
(0.2474)	(0.5353)		

This table presents parameter estimates from the method of simulated moments. See the main text for details about the model and targeted moments. Standard errors are reported in parentheses and calculated using a bootstrap approximation that is also discussed in the main text.

Table 5: Targeted moments in the HRS and in the estimated model.

Wealth distribution						
<i>Thousands of dollars</i>						
	P_{10}	P_{25}	P_{50}	P_{75}	P_{90}	Mean
Data	52	151	372	873	1819	716
Model	52	166	427	822	1379	605

Wealth regression				
Coefficients from: $\ln \text{Wealth}_{i,t} = X_{i,t}\beta + \varepsilon_{i,t}$				
	EA	SES	Coll	Stocks
Data	0.15	0.17	0.26	1.06
Model	0.17	0.18	0.26	1.12

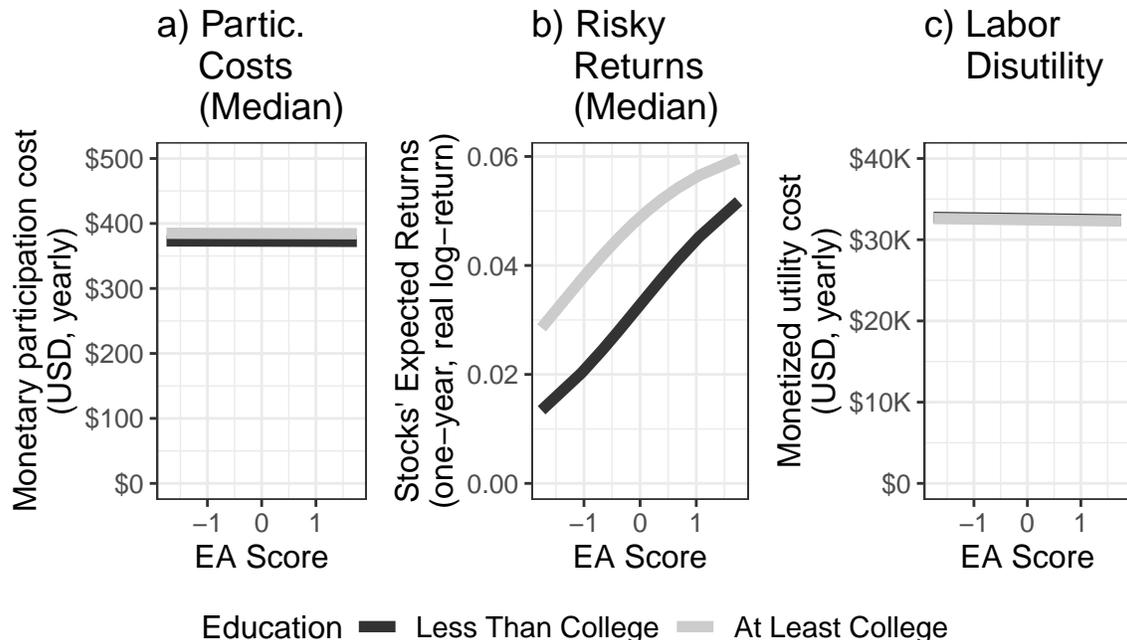
Stock-ownership regression				
Coefficients from: $\text{Stocks}_{i,t} = X_{i,t}\beta + \varepsilon_{i,t}$				
	EA	SES	Coll	ln E. Inc.
Data	0.049	0.050	0.067	0.255
Model	0.046	0.052	0.069	0.220

Stock ownership rate	
<i>Ages 60 to 70</i>	
	Rate
Data	0.68
Model	0.72

Retirement regression		
Coefficients from: $\text{Retired}_{i,t} = X_{i,t}\beta + \varepsilon_{i,t}$		
	EA	Coll
Data	-0.0059	-0.0564
Model	-0.0059	-0.0575

Retirement rates				
Fraction of individuals retired by age bracket				
	[50, 62]	[63, 67]	[68, 72]	≥ 73
Data	0.17	0.59	0.78	0.86
Model	0.11	0.69	0.91	1.00

This table reports the moments that we target in estimation calculated both using our HRS analytical sample and using simulations from the model with the estimated parameter values. See the main text for variable definitions, sample descriptions, and complete specifications of the controls included in $X_{i,t}$ for each regression.



The disutility of labor is taken at age 50 and monetized (see the main text for details). We find the participation costs, expected log returns, and disutility from work at age 50 for each of our 5, 130 agent types and weight them by the number of agents of each type in our simulated population. The figure presents medians at different levels of the EA score and educational attainment. Labor disutility depends only on agent's age, EA score, and educational attainment so there is no remaining heterogeneity after conditioning on these characteristics.

Figure 2: Participation costs, risky asset returns and disutility of work, and their relationship with the EA score and education.

median participation costs, expected risky returns, and disutility of labor against the EA score for both levels of education. Panel a) shows no meaningful variation of the median annual participation cost with either the EA score or education. We estimate this cost to range between \$372 and \$385, which is similar to estimates in recent studies in household finance (see e.g., Fagereng, Gottlieb, and Guiso, 2017; Catherine, 2021).

Risky asset returns, however, appear to be significantly affected by both the EA score and education. Table 4 and Panel b) of Figure 2 show that between the lowest and highest EA deciles, the median expected log-return on stocks increases from 1.3% to 5.2% for those without a college degree, and from 2.8% to 6% for those with a college degree. The direction of these estimated relationships is consistent with past studies (e.g., Calvet, Campbell, and Sodini, 2007; Fagereng et al., 2020) that have found that returns to wealth are heterogeneous and correlated with wealth, income and education, each of which covaries with the EA score. As discussed in Section 4, our specification of risky returns implies that agents receive different compensation for taking risk. An alternative to assessing simple returns is to examine risk-adjusted returns by way of the Sharpe ratio. In particular, we calculate the

Table 6: Implied Sharpe ratios and relative Sharpe ratio losses.

Variable	Mean	Std. Dev.	Percentiles					
			25th	50th	75th	90th	95th	99th
<i>All simulated agents</i>								
Sharpe Ratio	0.207	0.088	0.134	0.211	0.280	0.324	0.341	0.356
RSRL	0.440	0.239	0.238	0.426	0.635	0.792	0.838	0.908
<i>Stockholders at age 65</i>								
Sharpe Ratio	0.220	0.086	0.143	0.231	0.294	0.332	0.342	0.356
RSRL	0.403	0.234	0.200	0.372	0.611	0.746	0.801	0.907

The table reports summary statistics of the Sharpe ratio of the risky asset that is available to our simulated agents. The relative Sharpe ratio loss (RSRL) measures how far an agent’s risky investments are from the maximum Sharpe ratio available, which in our case is that of the S&P500.

relative Sharpe ratio loss (RSRL) as in Calvet, Campbell, and Sodini (2007):

$$\text{Sharpe}_i = \frac{\mathbb{E}[\tilde{R}_{i,t}] - R}{\sqrt{\mathbb{V}(\tilde{R}_{i,t})}}, \quad \text{RSRL}_i = 1 - \frac{\text{Sharpe}_i}{(\mathbb{E}[R_t^{\text{SP500}}] - R) / \sqrt{\mathbb{V}(R_t^{\text{SP500}})}}.$$

Table 6 presents summary statistics of the distribution of Sharpe ratios and RSRLs for the full population of simulated agents and for the group of agents who own stocks at age 65. The table shows that agents in our model receive very different compensations for their financial risk-taking, and that most of them are far from the performance of the benchmark return R_t^{SP500} . The median agent’s available risky investments deliver compensation per unit of risk that is 42.6% lower than that of the benchmark S&P500, and one tenth of agents earn compensations almost 80% lower. We cannot directly construct observed Sharpe ratios and RSRLs for the HRS sample since doing so would require more detailed information on investments. However the ranges we calculate using the estimated model are similar to those estimated by Calvet, Campbell, and Sodini (2007) using Swedish administrative data.¹⁷

We note also that the Sharpe ratios of stock market participants and non-participants are similar, although participants do fair somewhat better, especially at the lower end of the RSRL distribution. This is because estimated participation costs end up having only a small effect on participation — the relatively mild median participation cost means that expected returns won’t need to be excessively high to rationalize stock market entry for median-cost households. For high cost households, the cost is so severe that evening earning

¹⁷Calvet, Campbell, and Sodini (2007) estimate RSRLs whose 25th and 99th percentiles are $\{0.29, 0.89\}$, $\{0.07, 0.85\}$, or $\{-0.16, 0.82\}$ depending on the benchmark used for comparison.

the efficient market return is not enough to incentive entry. The estimation then suggests that the roughly 30% of non-participants are dissuaded from participation due to excessive costs, but their returns were they to participate would not be noticeably different from the population.

As reported in previous studies, our income estimates imply a positive and significant association between the EA polygenic score and labor income even after controlling for educational attainment. Figure 3 depicts the estimated age-earnings profiles of median income for agents with different EA scores and levels of education. College graduates exhibit higher incomes and a different life-cycle earnings profile (peaking later in life) compared to agents without a college degree. For both levels of education, though, incomes are substantially higher for agents with higher EA scores. For those without a college degree, earnings peak around the age of 45, when the median earnings of those in the 4th and 7th decile of the EA score are \$40,651 and \$42,975 respectively. For those with a college degree, earnings peak later, around the age of 50, when the median earnings of those in the 4th and 7th decile of the EA score are \$48,989 and \$51,367, respectively. Because the income model has a large number of controls, our estimates of the income process are reported in Appendix A.

The estimated additive component of disutility from labor varies little with education, the EA score, and age. We monetize the disutility to convey its magnitude. For a utility cost $d_{i,t} > 0$, we find the monetary value $m_{i,t}$ that an agent with no dynamic considerations and a baseline consumption of \$40,000 would be willing to pay to avoid the cost $d_{i,t}$.¹⁸ Formally, $m_{i,t}$ solves

$$u_{i,t}(C_t = \$40k, \ell_t = 1) = u_{i,t}(C_t = \$40k - \$m_{i,t}, \ell_t = 1) + d_{i,t}.$$

Panel c) of Figure 2 displays monetized costs of work $m_{i,t}$ at age 50 against the EA score for both education groups. The figure confirms that there is no meaningful heterogeneity in our estimated disutility of work.

The disutility of work parameter would most clearly affect retirement. Holding wealth and expected retirement income constant, a greater disutility of work would imply earlier retirement rates. The null results of disutility of work suggest the model can fit the observed retirement patterns based on variation in wealth, income, and expected retirement income, without needing to assign additional power to unenjoyable labor provision. This is further evidenced by the strong fit of the retirement regression moments.

¹⁸Our specification of the utility cost (Equation 4) is age-dependent. We base our calculations on its values at age 50.

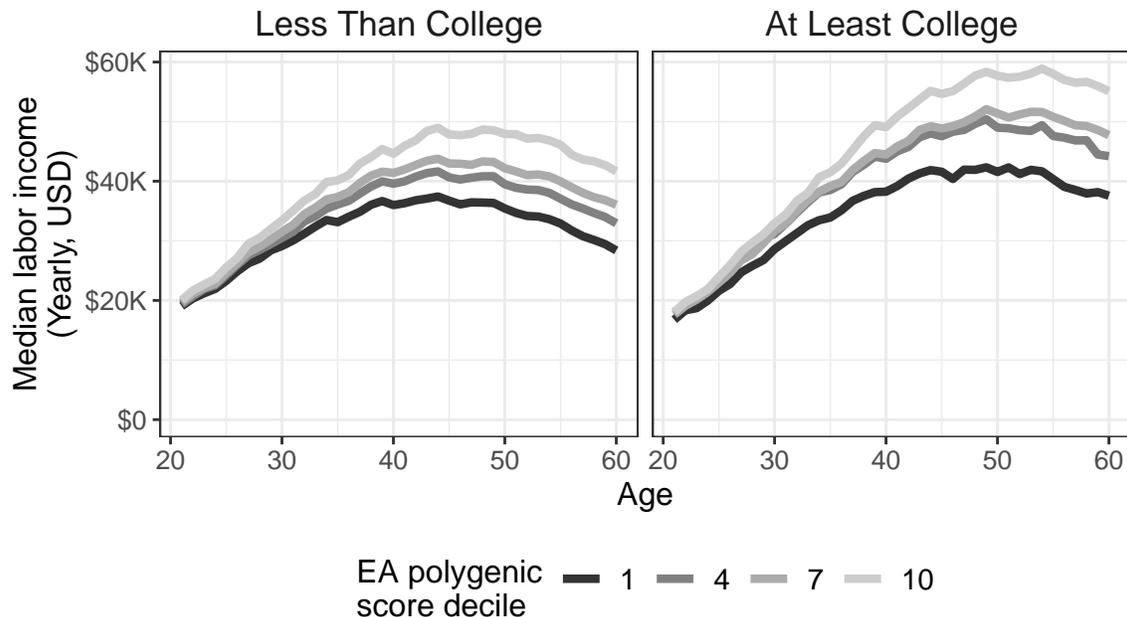


Figure 3: Median life-cycle income for different EA polygenic score deciles.

6.2 Life-Cycle Choices and Outcomes

The estimates suggest no meaningful relationship between the EA polygenic score and either stock market participation costs or disutility from work, but does suggest sizable associations with labor incomes and risky-asset returns, even at comparable levels of education. This section examines how these differences influence other choices and outcomes, such as consumption, retirement, stock-market participation, and wealth across the life cycle. Figure 4 presents the simulated age-profiles of these choices and outcomes for agents in different deciles of the EA polygenic score.

The greater average incomes and more efficient investments of agents with higher EA polygenic scores afford them higher consumption throughout their lives. Differences in median consumption become noticeable around age 35. By age 40, the difference between the median consumption of agents in the 4th and 7th deciles of the EA score reaches \$970 per year; this gap grows to \$3,495 by age 60, and \$4,427 by age 80.

Despite higher consumption, agents with higher EA scores accumulate more wealth. Figure 4 shows that an EA score-wealth relationship emerges early in life and widens as agents age. The median wealth balances of agents in the 4th and 7th decile of the EA polygenic score are \$45,175 and \$56,928 at age 40, \$297,997 and \$410,610 at age 60, and \$344,424 and \$523,125 at age 80. The growth in wealth differences result from the combination of

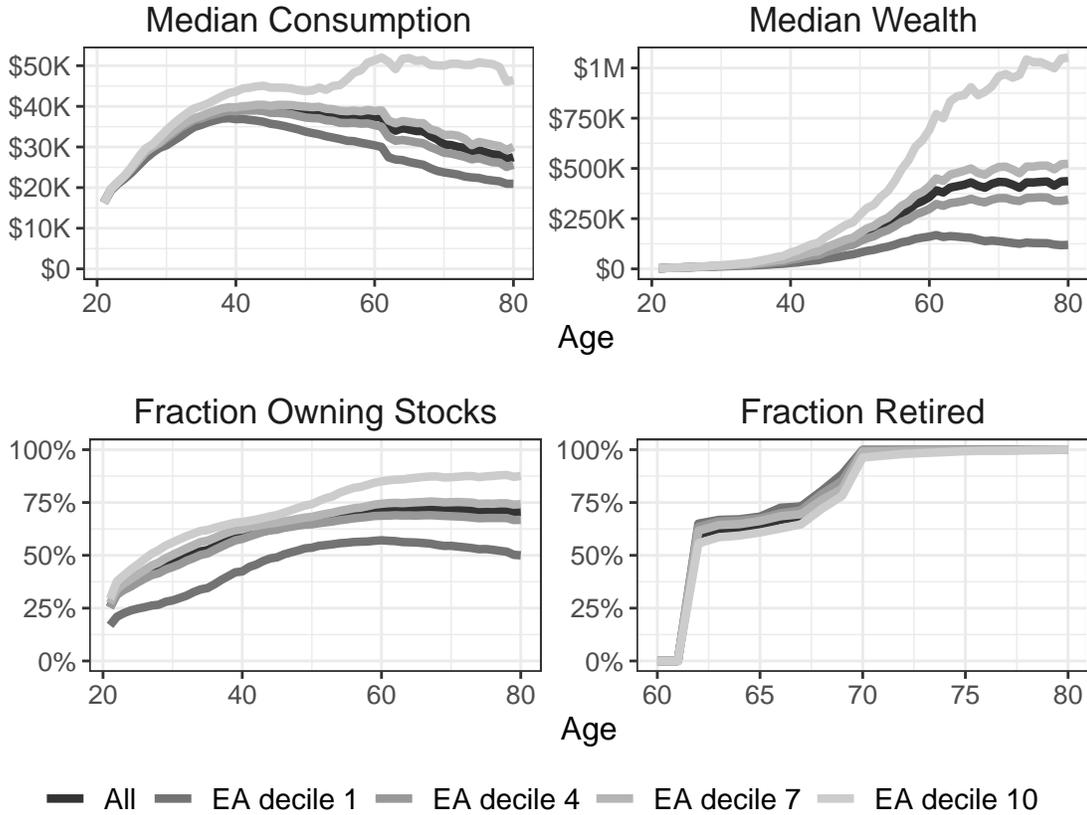
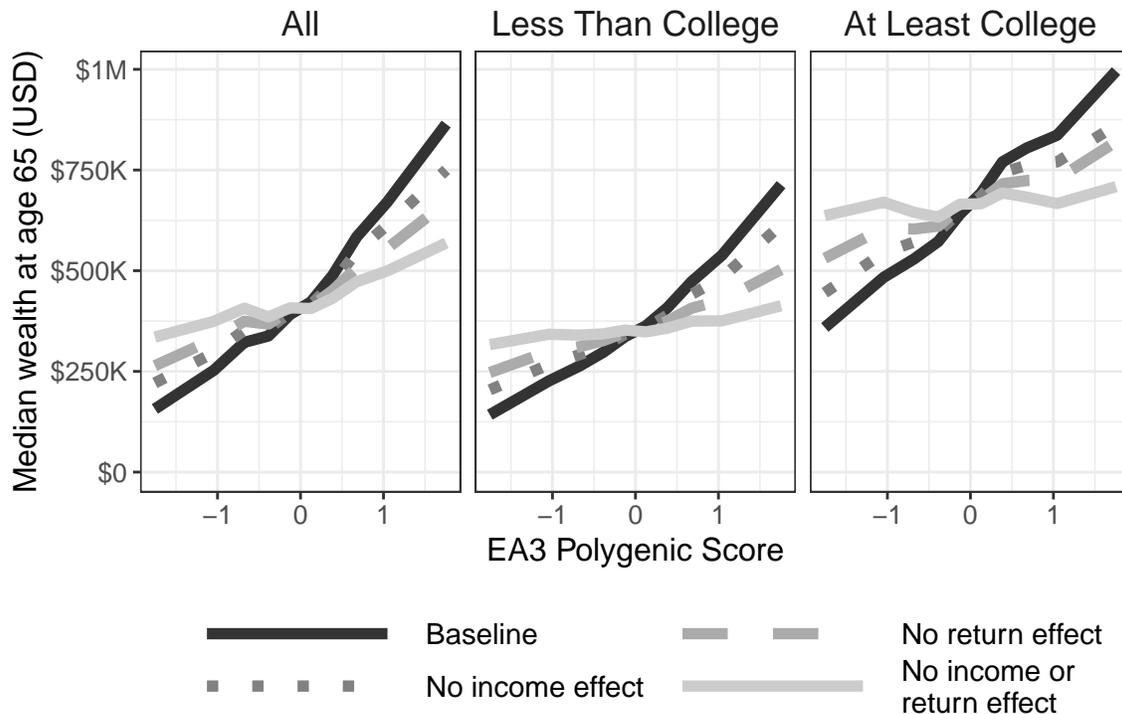


Figure 4: The choices and economic outcomes of agents with different EA scores.

higher labor earnings and higher returns on savings, which compound to magnify differences at later ages. As Figure 4 shows, stock market participation is monotonically increasing in the EA score and peaks at age 67 for those with EA scores the 4th and 7th deciles, reaching 69.0% and 75.5%, respectively.

Despite the lower levels of wealth and similar levels of labor disutility, differences in the retirement decisions of agents with different EA scores appear quite small. Agents with higher EA scores retire at later ages on average. Our model matches this fact: 63.5% of agents in the 4th decile of the EA score retire at the minimum age of 62 and 71.3% have retired by age 67; for the 7th decile of the EA score these numbers are 61.6% and 69.5% respectively. There are a few possibilities that could explain this finding. First, lower EA score individuals will receive retirement income that represents a higher fraction of their lifetime earnings due to the progressivity of the Social Security system. Second, the labor earnings lost in retirement are lower for low EA score agents because they earn less income in the labor market on average. These two effects seemingly dominate the lower levels of wealth and consumption in retirement also experienced by low EA score agents.



“Baseline” corresponds to conditional medians calculated using the simulations of our estimated model. “No income effect” corresponds to simulations from a version of the model in which we set the EA score’s coefficients in our income specification to 0, leaving the other components of the model unchanged. “No return effects” sets the EA score’s coefficient on expected risky-asset log-returns to 0. “No income or return effect” sets both sets of coefficients to 0.

Figure 5: The EA score-wealth gradient and its sources.

The foregoing discussion highlights how the estimated effects of the EA score on labor income and financial proficiency compound to generate differences in wealth, stock market participation, and retirement. Among these outcomes, wealth has received the greatest attention as a vast literature in economics has worked to understand key sources of wealth disparities. Barth, Papageorge, and Thom (2020) demonstrate that there is a robust association between the EA score and wealth at retirement. The model developed here allows us to disentangle multiple potential channels of wealth accumulation and to evaluate their importance in driving this relationship. A first driver of the relationship between the EA score and wealth is education. The EA score was built with the purpose of predicting educational attainment and therefore agents with higher EA scores will, on average, have more years of education that translate to greater earnings and wealth. After controlling for education by grouping individuals with a college degree and no college degree, the difference in median wealth at age 65 between the 1st and 10th EA score deciles falls from an unconditional value of \$707,423 to \$571,332 for those without a college degree and \$636,537 for those with a college degree.

In Figure 5, we further decompose the effects on wealth while continuing to control for education. To do so, we examine how the relationship between the EA score and wealth would change under different environments. First, we solve and simulate our model under an environment in which the income specification has all coefficients that multiply the EA score to 0, leaving the rest unchanged. The resulting wealth gradients are labeled “No income effect.” The differences in median wealth between agents in the highest and lowest EA score deciles fall to \$409,547 for those without a college degree and \$416,348 for college graduates. Restoring the income process and eliminating the direct effect of the EA score on returns by setting $r_{EA} = 0$ in Equation 9 results in the wealth gradient labeled “No return effect.” The differences in median wealth now fall to \$255,572 for those without a college degree and \$289,576 for college graduates, a significantly greater reduction than the results from removing the direct income effect. Finally, we remove the direct effects of the EA score on both income and returns. This reduces the gaps in median wealth between the top and bottom EA deciles to \$96,612 for those without a college degree and \$72,090 for college graduates, just 16.9% and 11.3% of the baseline difference. The small remaining difference is due to the combination of other direct effects, such as the EA score effect on the disutility of work and participation costs, as well as indirect effects, including the correlation of the EA score and childhood socioeconomic status, the influence of the EA score on the inheritance process, and covariation between the EA score and other characteristics such as birth year and defined benefit pension arrangements.

7 Policy Experiments

The past few decades have seen multiple proposals to reform the Social Security system in order to address the fiscal challenges presented by an aging population. These challenges are a topic of current political and economic debate, with multiple proposals presented every year.¹⁹ To reduce the cost of the system, a policy would have to collect more resources from young workers, reduce the benefits distributed to retirees, or some combination of the two. Such changes can affect welfare on multiple dimensions. A more complete understanding of the ways in which compounding economic disadvantages interact with different policy regimes can aid policymakers in finding changes that reduce costs but avoid concentrating welfare losses among people who are least able to bear them.

To explore how the generosity of the retirement social safety net could interact over the life cycle with various dimensions of individual-level heterogeneity, including genetic

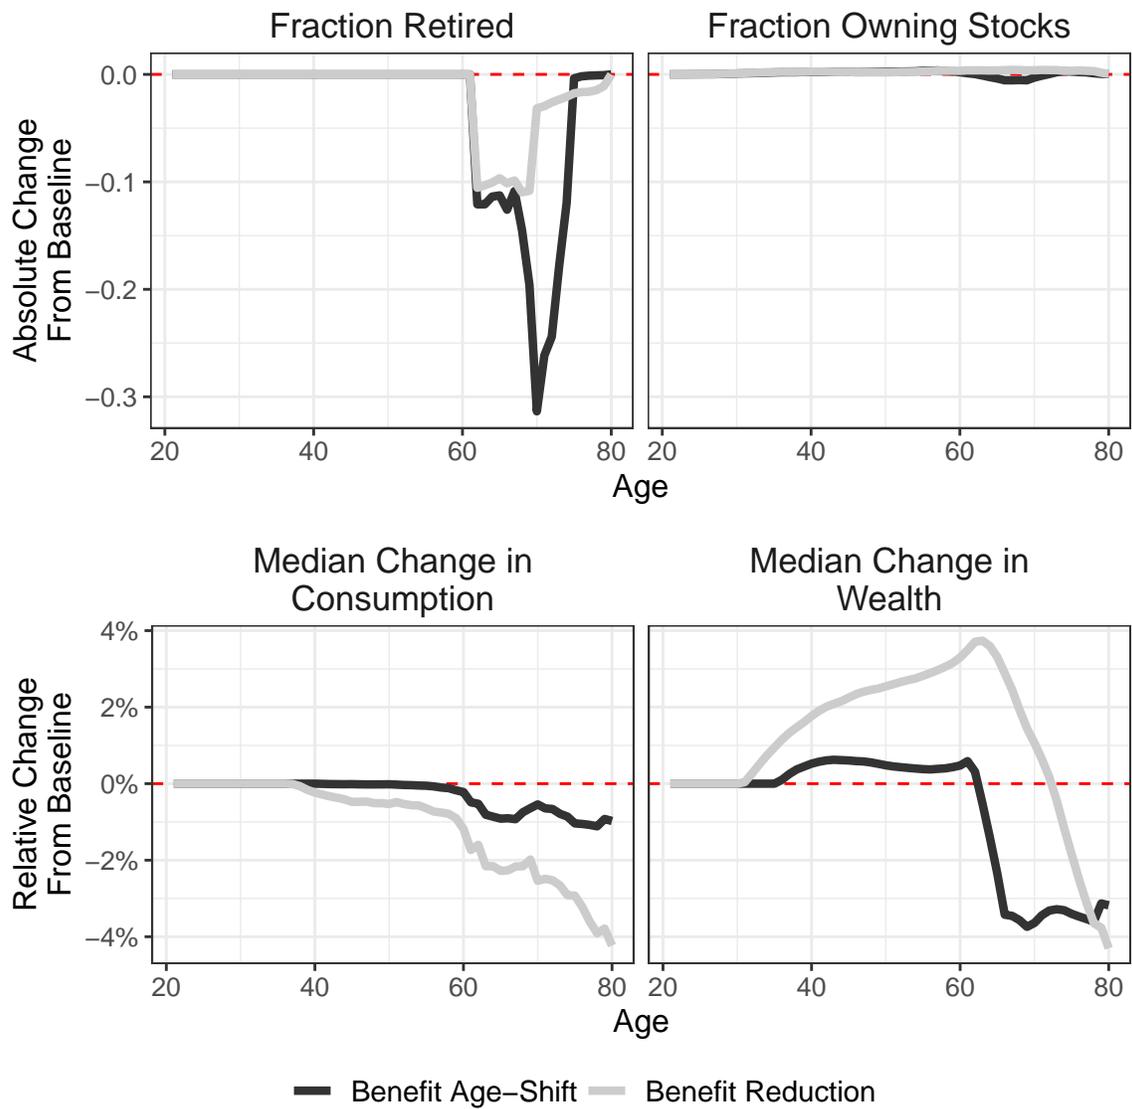
¹⁹The Office of the Chief Actuary of the Social Security Administration maintains a repository of fiscal analyses of the various proposals presented by congress (Social Security Administration, 2022a).

endowments, this section compares two cost-saving policies. The first policy, which we call the “benefit shift” policy, amounts to a five-year forward shift in the Social Security benefit schedule. The minimum benefit-claiming age increases from 62 to 67, and the benefit calculation formulas move forward by five years. If agent i retiring at age a currently receives yearly benefits equal to $SSB_i(a)$, under the benefit shift policy he receives $SSB_i(a - 5)$. This policy mimics many proposals suggesting increases in the age of benefit claiming that would encourage people to work longer. Net costs decrease because people will either work for longer (and reduce the number of years in which they claim benefits), or would retire before the full benefit claiming age and receive lower benefits than they would in the baseline scenario. Importantly, we assume agents are aware of this policy starting at age 21, when the model begins, i.e., we do not model policy shifts later in the life cycle or uncertainty about policy regimes.

The second policy, which we call the “benefit reduction policy” represents a broad class of provisions that would alter the way in which benefits are calculated, lowering the annual payments to some or all future retirees. We analyze a simple policy of reducing all benefits by a constant fraction ψ . If agent i retiring at age a receives benefits $SSB_i(a)$ in the baseline scenario, this policy would make their benefits $(1 - \psi)SSB_i(a)$. In order to make the two policy changes comparable in fiscal terms, we find the reduction rate ψ that would produce the same increase in revenue as the previous policy change. We present details about the revenue measure in Appendix F. We find that the reduction that makes the two policies fiscally equivalent is $\psi = 0.2789$. As with the benefit shift policy, we assume agents are fully aware of the benefit reduction policy when the model begins.

7.1 The Impact of Policy on Choices and Outcomes

To begin, we examine how the counterfactual benefit shift and benefit reduction policies affect median behavior and outcomes. While both policies effectively reduce benefits, they differ in how this is accomplished and thus in the margins along which they incentivize people to optimally adjust. Figure 6 depicts how these policies affect median consumption, wealth, retirement, and stockholding across their life cycle. The figure shows deviations from the baseline scenario. Under the benefit shift policy, retirement occurs substantially later. For example, the fraction of agents retired at age 70 drops by over 30 percentage points. The fraction owning stocks remains largely unchanged, suggesting the policy does not affect stock market participation. Even though some agents work longer, which implies higher labor income at later ages, the drop in benefits leads to a net decline in consumption for the median agent of roughly one percentage point starting at age 62. Increases in wealth



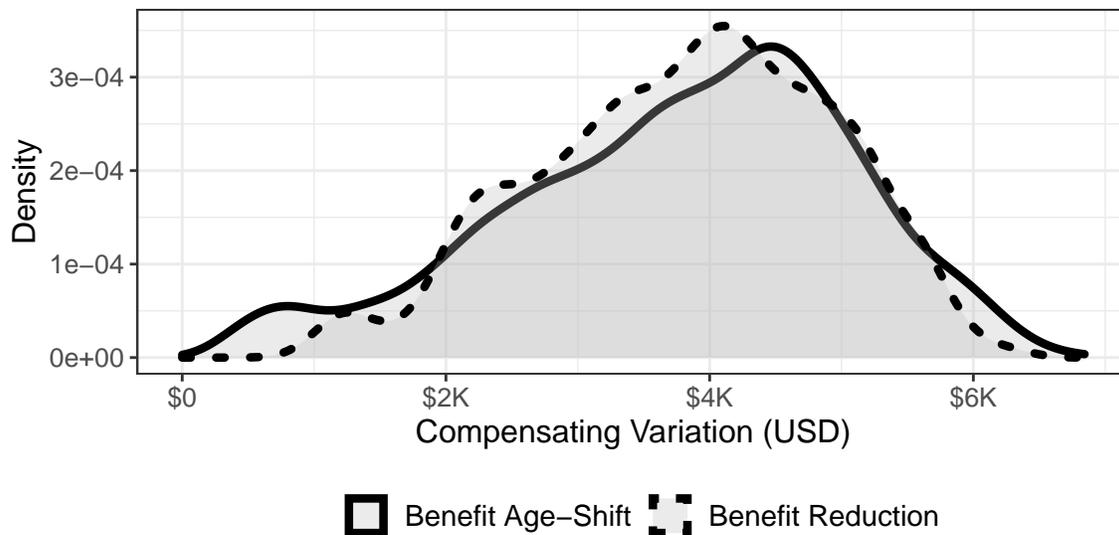
We simulate identical populations of agents that experience identical sequences of shocks under the status quo and the two alternative social security changes. We find the relative consumption and wealth changes with respect to the baseline scenario agent-by-agent and report the medians of those changes at every age. We also find the fraction of agents who are retired and who are holding stocks at every age and report the differences with respect to the baseline scenario.

Figure 6: Life-cycle adjustments to Social Security changes.

start around age 35 as agents save more in anticipation of the delay in benefits, followed by a median decline of nearly four percentage points starting at age 62 driven by agents who retire but do not yet receive benefits.

In contrast to the benefit shift policy, the benefit reduction policy leads to less postponement of retirement. It does not cause a change in stock market participation, again suggesting that portfolio allocations remain largely unchanged. However, it does induce a larger reduction in consumption than under the benefit shift policy as agents increase their savings in anticipation of reduced benefits. As we will examine in greater detail below, these differences arise because of variation in returns to savings and in access to non-labor income through defined benefit pensions, which some agents can access if they choose not to postpone retirement. This larger shift in consumption means that, under the benefit reduction, wealth increases more dramatically than under then benefit shift policy. This increase begins at age 30 and peaks at nearly 4% at age 62. After that, the benefit reduction leads to a rapid decline in wealth so that by age 80 the median agent has 4% less wealth than under the baseline.

7.2 Welfare Implications



The compensating variations for both policies are computed at age 21 and from a starting wealth of \$20,000. We find the compensating variation for each of our 5,130 agent types and weight them by the number of agents of each type in our simulated population. The figure presents kernel density estimates of the distribution of these compensating variations.

Figure 7: Compensating variations for both policies across the population.

These results suggest that the two policies have different effects on the median agent.

While working longer or consuming less are both costly, it is not clear which policy entails the largest welfare losses. This question is central because the policies generate the same savings to the Social Security system. To examine welfare, we calculate compensating variations for agents who are 21 and have a starting wealth of \$20,000, but who vary on other observable characteristics to match our empirical sample. Concretely, if $V_{i,21}(\cdot)$ is the baseline value function of agent i at age 21 and $V_{i,21}^x(\cdot)$ is his value function under the alternative policy scenario x , the compensating variation from an initial wealth of W , $CV_i^x(W)$, solves:

$$V_{i,21}(W) = V_{i,21}^x(W + CV_i^x(W)).$$

$CV_i^x(W)$ is the monetary transfer that would be required to restore person i 's expected lifetime welfare to what it was before the policy change. Therefore, $CV_i^x(W)$ will be higher for policies that bring greater reductions to i 's lifetime welfare. We plot the distributions of welfare costs for each policy and across agents in Figure 7. The figure reveals several points. First, there is substantial heterogeneity in welfare costs ranging from roughly zero to over \$6,500. Second, under a social welfare function that weighted agents equally, it is not clear which policy is more favorable since the cost distributions largely overlap. The mean for the benefit shift and benefit reduction policies are \$3,780 and \$3,820, respectively. Other empirical moments are also similar. For example, standard deviations are \$1,300 and \$1,100 respectively, which means it is not the case that similar means obscure fat tails for one policy versus the other.

To further explore the welfare implications we turn to an examination of heterogeneity. For each policy, we examine how different magnitudes of the cost of the policy (measured as the compensating variation described above) relate to different variables and model estimates. Results are in Table 7 where each row is a quartile of the policy-cost distribution. Among agents who incur low welfare losses (the first quartile of the distribution of CV) from the benefit shift policy, 78% have a college degree. In contrast, among agents who suffer the highest losses (fourth quartile) only 1.4% have a college degree. The corresponding percentages for the benefit reduction policy are 80% and 0.1%. The basic pattern that agents at a disadvantage incur higher costs of policies holds for different measures, including the EA score and expected returns. For example, agents with lower expected returns tend to face the highest costs of both policies as both groups will find it difficult to recoup the costs of a less generous Social Security system. Likewise, agents with lower EA scores face higher costs for a host of reasons.

The distribution of outcomes for the most affected groups is not always intuitive. For instance, the relationship between welfare costs and defined benefit pensions is not entirely

obvious. Under the benefit reduction policy, the relationship is not even monotonic, reflecting that agents with this type of pension scheme have a guaranteed source of income, but also tend to earn less. Under the benefit shift policy, one might expect lower welfare losses among those with a defined benefit pension, but the opposite is the case. More affected individuals also have larger stock market participation costs, as shown in quartile 4 of the benefit reduction policy. Because the distribution of participation costs implied by our estimates and heterogeneity structure is tri-modal, with the modes close to \$10, \$375, and \$13,000, with little variation in between, and because medians are plotted within each quartile in Table 7, the difference in median costs is stark. However, in unreported analysis, we also find that the fraction of individuals with the highest participation costs are much lower for lower quartiles, suggesting costs are rising on average as CV grows.

Table 7: Characteristics of agents with different welfare losses.

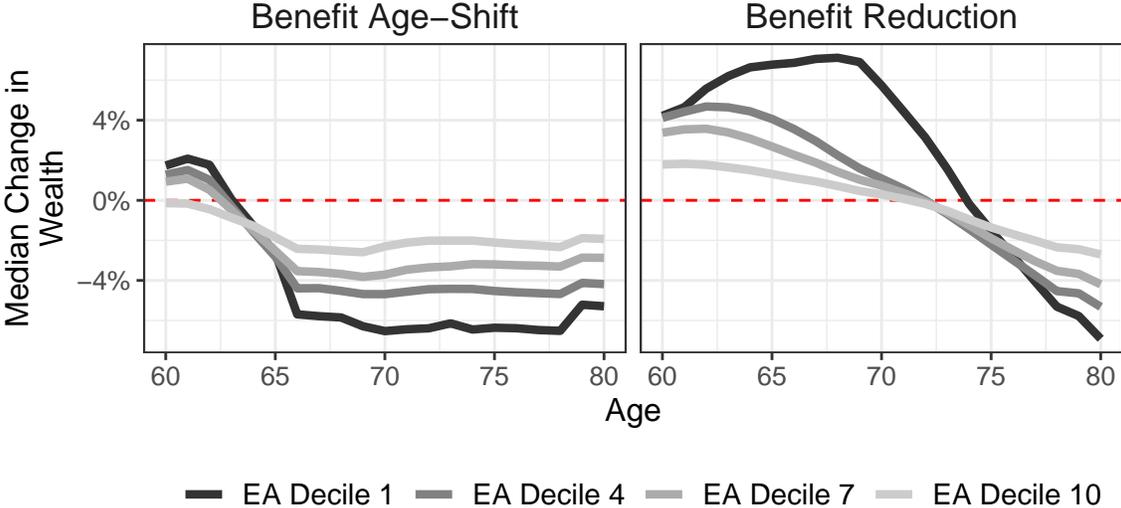
Compensating Variation	Fraction			Median		
	College	D.B. Pen.	EA Score	Part. Cost (F)	Exp. Returns $E[\ln \tilde{R}]$	Lab. Disut. (d)
<i>A. Benefit Age-Shift</i>						
Quart. 1	0.784	0.178	0.677	0.384	0.054	0.390
Quart. 2	0.262	0.369	0.126	0.373	0.042	0.394
Quart. 3	0.053	0.369	-0.126	0.373	0.031	0.397
Quart. 4	0.014	0.498	-1.045	0.373	0.017	0.400
<i>B. Benefit Reduction</i>						
Quart. 1	0.800	0.361	0.677	0.374	0.054	0.390
Quart. 2	0.256	0.424	0.126	0.373	0.041	0.394
Quart. 3	0.056	0.348	-0.126	0.373	0.029	0.397
Quart. 4	0.001	0.281	-1.045	13.123	0.019	0.400

The compensating variations for both policies are computed at age 21 and from a starting wealth of \$20,000. We find the compensating variation for each of our 5,130 agent types and weight them by the number of agents of each type in our simulated population. For each policy, we divide agents into quartiles according to their welfare losses as measured by their compensating variation. The table reports summary statistics of the characteristics of agents in each of the group. The disutility of labor varies with age; we report the median of its value at age 50.

7.3 Gene by Outcome vs. Gene by Welfare Associations

The prior section showed evidence of heterogeneity in the welfare consequences of both policies under consideration, including differences by EA score and by expected returns, which the EA score influences. In this section, we examine the behavioral and welfare implications of both policies, focusing on differences along the EA score distribution.

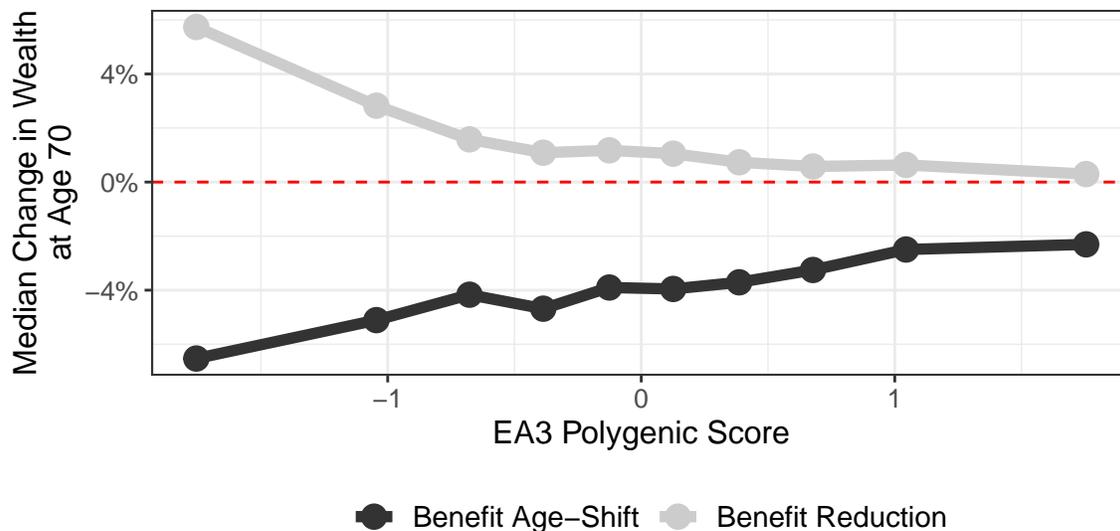
To begin, for each policy, we consider median wealth changes over the life cycle for the first, fourth, seventh and tenth deciles of the EA score distribution, which we plot in Figure 8. The benefit shift leads to declines in wealth that are particularly strong among agents with low EA scores. In contrast, the benefit reduction policy leads to striking increases in wealth, especially among those with the lowest EA scores. In other words, the benefit reduction policy decreases wealth inequality along the EA score distribution for much of the life cycle, while the benefit shift policy increases wealth inequality over much of the life cycle. If we consider policy-induced wealth changes at a fixed age, for instance age 70, the benefit reduction policy increases wealth across the board and reduces wealth inequality while the benefit shift policy decreases wealth across the board, but increases wealth inequality. In other words, the benefit reduction policy weakens the relationship between genetic endowments and wealth at age 70 while the benefit-shift policy steepens it. This illustrates that genetic relationships with outcomes are not immutable “laws and nature,” but are functions of policy environments and optimal decisions. Moreover, the flattening of the gene-wealth gradient might be an attractive feature if the reduction in inequality is considered a desirable outcome.



We simulate identical populations of agents that experience identical sequences of shocks under the status quo and the two alternative Social Security changes. We find the relative wealth changes with respect to the baseline scenario agent-by-agent and report the medians of those changes at every age for different levels of the EA score.

Figure 8: Wealth response to policy changes by EA3 polygenic score.

Figure 10, however, demonstrates that this interpretation is perilous. Figure 10 plots the distribution of compensating variations (calculated using the same assumptions as in Figure 7) for agents with different EA scores. From the perspective of *welfare*, both the



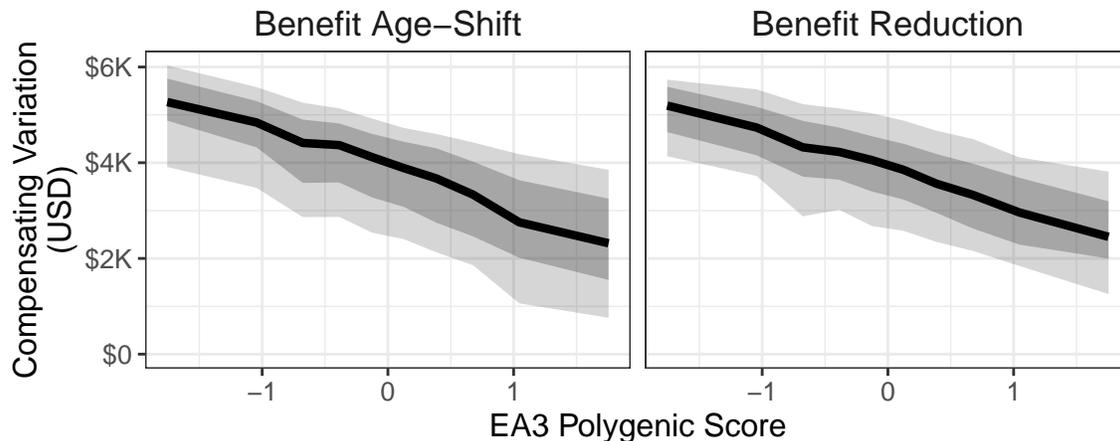
We simulate identical populations of agents that experience identical sequences of shocks under the status quo and the two alternative Social Security changes. We find the relative wealth changes at age 70 with respect to the baseline scenario agent-by-agent and report the medians of those changes for different levels of the EA score.

Figure 9: Social Security changes and the EA-wealth gradient.

benefit reduction and benefit shift policies are similarly costly, and harm low EA score agents disproportionately; this is despite the flattening of the gene-wealth gradient under the benefit reduction policy. Put differently, the reduced-form relationships between genes and economic outcomes, such as wealth, may be a poor proxy for the underlying association with well-being. This insight also suggests that assessing the effects of various policies on gene-outcome associations may obscure the effects on inequality.

What motivates the disconnect between the flattening gene-wealth gradient and steepening gene-welfare gradient under the benefit reduction policy? The association between genes and other endogenous choices, such as retirement age and consumption, can provide clues. Figure 11 examines changes to consumption by EA decile. We plot changes starting at age 60. The benefit shift policy leads to relatively small drops (about 1%) in consumption, although lower EA score agents generally reduce consumption by more. In contrast, the benefit reduction policy induces a sizeable drop in consumption which is much larger for lower EA scores. The bottom decile of the EA score distribution exhibits a 2% drop in consumption at age 60 that rises in magnitude to a 5% drop by age 80. The top EA decile shows a much more moderate decline in consumption.

This suggests one important source of the EA score-welfare penalty in the benefit reduction scenario. Low EA score households are forced to accumulate resources at a much greater rate to replace the lost Social Security benefits. This serves to moderately flatten



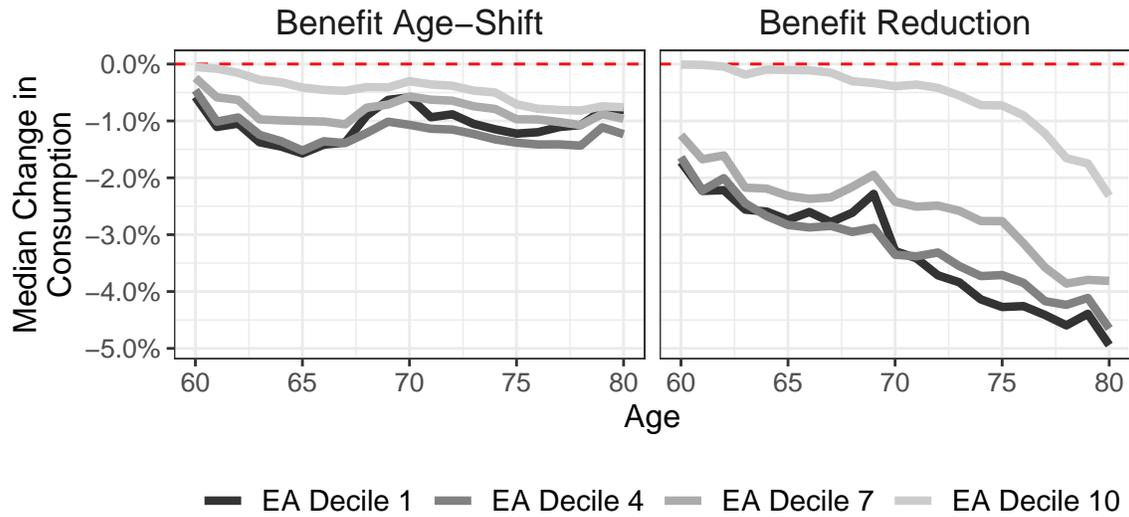
The compensating variations for both policies are computed at age 21 and from a starting wealth of \$20,000. We find the compensating variation for each of our 5,130 agent types and weight them by the number of agents of each type in our simulated population. The figure depicts percentiles of the distribution of these compensating variations at different levels of the EA score. The solid line corresponds median. Inner shaded areas cover observations between the 25th and 75th percentiles. Outer shaded areas cover observations between the 10th and 90th percentile.

Figure 10: Distribution of compensating variations at different levels of the EA score.

the gene-wealth gradient. But, low EA score individuals also earn lower rates of return on accumulated savings, which implies the reduction in consumption during the working life must be more severe for low EA scores. Thus, while low EA score agents accumulate greater wealth under this policy, relative to high EA score agents, this accumulation is supported by a disproportionately large decrease in consumption, and in turn, welfare.

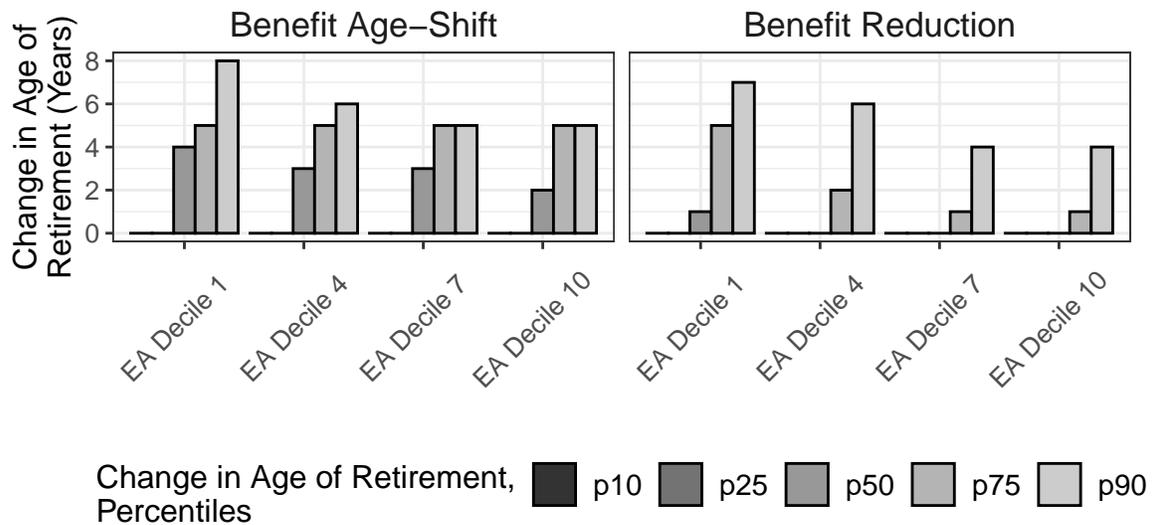
Figure 12 shows an additional source of welfare disparity that results from extended working lives. In both policies, there is a larger positive change in labor supply for agents with lower EA scores, suggesting their need to recoup lost benefits through a longer working life. This is in addition to the dramatic decline in consumption under the benefit reduction policy. However, the benefit age shift induces stronger responses in employment. The median increase in retirement age for agents in the lowest EA decile is four years in the benefit shift policy and only one year in the benefit reduction policy. The higher EA score deciles show significant extended employment lives under the benefit shift policy as well, although the effect is much more muted in the benefit reduction setting. Because working has a utility cost, extended working lives further contribute to declines in welfare.

The key takeaway from this exercise is that observed gene-outcome associations are the consequence of a multitude of factors that may interact in complex ways. A model of economic behavior is needed to make sense of these associations, since endogenous responses to changes in the environment may have nuanced effects on observed outcomes. Weakening associations between genes and outcomes may hide and strengthening association between



We simulate identical populations of agents that experience identical sequences of shocks under the status quo and the two alternative Social Security changes. We find the relative consumption changes with respect to the baseline scenario agent-by-agent and report the medians of those changes at every age for different levels of the EA score.

Figure 11: Consumption response to policy changes by EA3 polygenic score.



We simulate identical populations of agents that experience identical sequences of shocks under the status quo and the two alternative Social Security changes. We find the age of retirement of every agent in every scenario and find the difference with respect to the baseline scenario for both alternative policies. This table reports percentiles of these individual-level changes in the age of retirement. We consider only agents who did not die before retiring in any of the policy scenarios. The figure only considers agents without defined benefit pension plans.

Figure 12: Retirement responses of agents without D.B. pension plans.

genes and welfare. The model we develop here allows us to assess the effect of policy changes on the *gene-welfare* gradient, which should be of primary interest to those interested in understanding how policies may affect individuals with different genetic endowments differently.

8 Conclusion

We build and estimate a structural consumption-savings model that explicitly incorporates genetic variation that is summarized as a polygenic score for educational attainment, which we refer to as the EA score. The model includes realistic Social Security and tax schemes, and heterogeneity in labor market outcomes, stock market participation and returns, inheritances, and defined benefit pensions. Based on descriptive evidence established in earlier literature we allow the EA score to affect endogenous outcomes, such as wealth and retirement through a multitude of channels. In particular, the EA score affects labor income, the welfare costs of work, stock market returns, stock market participation costs, access to defined benefit pensions, and the distribution of inheritances. We use the model to conduct cost-saving counterfactuals that reduce the generosity of Social Security benefits in order to assess the ways in which policy changes may have different effects on people with different genetic endowments.

Our model estimates suggest that, compared to households in the lowest EA score decile, those in the top decile earn significantly higher expected returns on stocks—3.9 percentage points higher on average per year—and earn lower incomes in the labor market. We find little association between the EA score and the costs of stock market participation or the disutility of work. These differences compound to make people at the lower end of the EA score distribution particularly vulnerable to changes that deteriorate the social safety net, such as cost-saving changes in the Social Security system.

Indeed, in the counterfactual policy environments we study, we find that both policies have similar welfare costs in general and that both have higher costs for people with less advantageous genetic endowments. Thus, both policies disproportionately harm people who are already at a disadvantage. However, we also find that one of the policies weakens the relationship between genetic endowments and wealth, which may lead a policy-maker to favor it. This illustrates the importance of conducting welfare analyses when evaluating policies.

More broadly, our findings highlight that gene-by-environment interactions are a function of the policy and economic landscape, and are not immutable properties of nature.

Additionally, the interactions between genes, outcomes, and policies are complex and multifaceted, and a careful assessment of consequences is needed to determine the true costs and distributional consequences of alternative policy regimes.

References

- Bagchi, Shantanu. 2019. “Differential Mortality and the Progressivity of Social Security.” *Journal of Public Economics* 177:104044.
- Barcellos, Silvia H., Leandro S. Carvalho, and Patrick Turley. 2018. “Education Can Reduce Health Differences Related to Genetic Risk of Obesity.” *Proceedings of the National Academy of Sciences* 115 (42):E9765–E9772.
- Barth, Daniel, Nicholas W. Papageorge, and Kevin Thom. 2020. “Genetic Endowments and Wealth Inequality.” *Journal of Political Economy* 128 (4):1474–1522.
- Beauchamp, Jonathan P., David Cesarini, Magnus Johannesson, Matthijs J. H. M. van der Loos, Philipp D. Koellinger, Patrick J. F. Groenen, James H. Fowler, J. Niels Rosenquist, A. Roy Thurik, and Nicholas A. Christakis. 2011. “Molecular Genetics and Economics.” *Journal of Economic Perspectives* 25 (4):57–82.
- Belsky, Daniel W., Benjamin W. Domingue, Robbee Wedow, Louise Arseneault, Jason D. Boardman, Avshalom Caspi, Dalton Conley, Jason M. Fletcher, Jeremy Freese, Pamela Herd, Terrie E. Moffitt, Richie Poulton, Kamil Sicinski, Jasmin Wertz, and Kathleen Mulvan Harris. 2018. “Genetic Analysis of Social-Class Mobility in Five Longitudinal Studies.” *Proceedings of the National Academy of Sciences* 115 (31):E7275–E7284.
- Belsky, Daniel W., Terrie E. Moffitt, David L. Corcoran, Benjamin Domingue, HonaLee Harrington, Sean Hogan, Renate Houts, Sandhya Ramrakha, Karen Sugden, Benjamin S. Williams, Richie Poulton, and Avshalom Caspi. 2016. “The Genetics of Success: How Single-Nucleotide Polymorphisms Associated With Educational Attainment Relate to Life-Course Development.” *Psychological Science* 27 (7):957–972.
- Benjamin, Daniel J., David Cesarini, Christopher F. Chabris, Edward L. Glaeser, David I. Laibson, Vilmundur and Guðnason, Tamara B. Harris, Lenore J. Launer, Shaun Purcell, Albert Vernon Smith, Magnus and Johannesson, Patrik K.E. Magnusson, Jonathan P. and Beauchamp, Nicholas A. Christakis, Craig S. and Atwood, Benjamin Hebert, Jeremy

- Freese, Robert M. Hauser, Taissa S. Hauser, Alexander and Grankvist, Christina M. Hultman, and Paul Lichtenstein. 2012. “The Promises and Pitfalls of Genoeconomics.” *Annual Review of Economics* 4 (1):627–662.
- Bernheim, B. Douglas, Jonathan Skinner, and Steven Weinberg. 2001. “What Accounts for the Variation in Retirement Wealth among U.S. Households?” *American Economic Review* 91 (4):832–857.
- Biroli, Pietro. 2015. “Genetic and Economic Interaction in Health Formation: The Case of Obesity.” *Working Paper* :70.
- Black, Sandra E, Paul J Devereux, Petter Lundborg, and Kaveh Majlesi. 2020. “Poor Little Rich Kids? The Role of Nature versus Nurture in Wealth and Other Economic Outcomes and Behaviours.” *The Review of Economic Studies* 87 (4):1683–1725.
- Calvet, Laurent E., John Y. Campbell, and Paolo Sodini. 2007. “Down or Out: Assessing the Welfare Costs of Household Investment Mistakes.” *Journal of Political Economy* 115 (5):707–747.
- Carroll, Christopher D. 2002. “Why Do the Rich Save So Much?” In *Does Atlas Shrug? The Economic Consequences of Taxing the Rich*, edited by Joel B. Slemrod. Harvard University Press.
- Catherine, Sylvain. 2021. “Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle.” *The Review of Financial Studies* :hhab136.
- Cesarini, David, Christopher T. Dawes, Magnus Johannesson, Paul Lichtenstein, and Björn Wallace. 2009. “Genetic Variation in Preferences for Giving and Risk Taking*.” *The Quarterly Journal of Economics* 124 (2):809–842.
- Cesarini, David, Magnus Johannesson, Paul Lichtenstein, Örjan Sandewall, and Björn Wallace. 2010. “Genetic Variation in Financial Decision-Making.” *The Journal of Finance* 65 (5):1725–1754.
- Chai, Jingjing, Wolfram Horneff, Raimond Maurer, and Olivia S. Mitchell. 2011. “Optimal Portfolio Choice over the Life Cycle with Flexible Work, Endogenous Retirement, and Lifetime Payouts*.” *Review of Finance* 15 (4):875–907. URL <https://doi.org/10.1093/rof/rfr016>.
- Cocco, João F., Francisco Gomes, and Pascal J. Maenhout. 2005. “Consumption and Portfolio Choice over the Life Cycle.” *The Review of Financial Studies* 18 (2):491–533.

- Conesa, Juan C. and Dirk Krueger. 1999. "Social Security Reform with Heterogeneous Agents." *Review of Economic Dynamics* 2 (4):757–795.
- Cronqvist, Henrik and Stephan Siegel. 2014. "The Genetics of Investment Biases." *Journal of Financial Economics* 113 (2):215–234.
- . 2015. "The Origins of Savings Behavior." *Journal of Political Economy* 123 (1):123–169.
- De Nardi, Mariacristina, Eric French, and John B. Jones. 2010. "Why Do the Elderly Save? The Role of Medical Expenses." *Journal of Political Economy* 118 (1):39–75.
- Diamond, Peter A. and Peter R. Orszag. 2005. "Saving Social Security." *Journal of Economic Perspectives* 19 (2):11–32.
- Dushi, Irena and Brad Trenkamp. 2021. "Improving the Measurement of Retirement Income of the Aged Population." SSRN Scholarly Paper 3765468, Social Science Research Network, Rochester, NY.
- Dynan, Karen E., Jonathan Skinner, and Stephen P. Zeldes. 2004. "Do the Rich Save More?" *Journal of Political Economy* 112 (2):397–444.
- Fagereng, Andreas, Charles Gottlieb, and Luigi Guiso. 2017. "Asset Market Participation and Portfolio Choice over the Life-Cycle." *The Journal of Finance* 72 (2):705–750.
- Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2020. "Heterogeneity and Persistence in Returns to Wealth." *Econometrica* 88 (1):115–170.
- Fagereng, Andreas, Magne Mogstad, and Marte Rønning. 2021. "Why Do Wealthy Parents Have Wealthy Children?" *Journal of Political Economy* 129 (3):703–756.
- Feldstein, Martin. 2005. "Structural Reform of Social Security." *Journal of Economic Perspectives* 19 (2):33–55.
- Fuster, Luisa, Ayşe İmrohoroğlu, and Selahattin İmrohoroğlu. 2003. "A Welfare Analysis of Social Security in a Dynastic Framework*." *International Economic Review* 44 (4):1247–1274.
- . 2007. "Elimination of Social Security in a Dynastic Framework." *The Review of Economic Studies* 74 (1):113–145.
- Gomes, Francisco and Alexander Michaelides. 2005. "Optimal Life-Cycle Asset Allocation: Understanding the Empirical Evidence." *The Journal of Finance* 60 (2):869–904.

- Guo, Guang and Elizabeth Stearns. 2002. “The Social Influences on the Realization of Genetic Potential for Intellectual Development*.” *Social Forces* 80 (3):881–910.
- Hairault, Jean-Olivier, François Langot, and Thepthida Sopraseuth. 2008. “Quantifying the Laffer Curve on the Continued Activity Tax in a Dynastic Framework*.” *International Economic Review* 49 (3):755–797.
- Harden, K. Paige and Philipp D. Koellinger. 2020. “Using Genetics for Social Science.” *Nature Human Behaviour* 4 (6):567–576.
- Harden, Kathryn Paige. 2021. *The Genetic Lottery: Why DNA Matters for Social Equality*. Princeton University Press, 1st ed.
- Heckman, James J. 1995. “Lessons from the Bell Curve.” *Journal of Political Economy* 103 (5):1091–1120.
- Heckman, James J. and Yona Rubinstein. 2001. “The Importance of Noncognitive Skills: Lessons from the GED Testing Program.” *American Economic Review* 91 (2):145–149.
- Heckman, James J., Jora Stixrud, and Sergio Urzua. 2006. “The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior.” *Journal of Labor Economics* 24 (3):411–482.
- Hill, W. David, Neil M. Davies, Stuart J. Ritchie, Nathan G. Skene, Julien Bryois, Steven Bell, Emanuele Di Angelantonio, David J. Roberts, Shen Xueyi, Gail Davies, David C. M. Liewald, David J. Porteous, Caroline Hayward, Adam S. Butterworth, Andrew M. McIntosh, Catharine R. Gale, and Ian J. Deary. 2019. “Genome-Wide Analysis Identifies Molecular Systems and 149 Genetic Loci Associated with Income.” *Nature Communications* 10 (1):5741.
- Houmark, Mikkel, Victor Ronda, and Michael Rosholm. 2020. “The Nurture of Nature and the Nature of Nurture: How Genes and Investments Interact in the Formation of Skills.” SSRN Scholarly Paper ID 3708642, Social Science Research Network, Rochester, NY.
- Hubbard, R. Glenn, Jonathan Skinner, and Stephen P. Zeldes. 1995. “Precautionary Saving and Social Insurance.” *Journal of Political Economy* 103 (2):360–399.
- Imrohoroglu, Selahattin and Sagiri Kitao. 2012. “Social Security Reforms: Benefit Claiming, Labor Force Participation, and Long-Run Sustainability.” *American Economic Journal: Macroeconomics* 4 (3):96–127.

- Karlsson Linnér, Richard, Pietro Biroli, Edward Kong, S. Fleur W. Meddens, Robbee Wedow et al. 2019. “Genome-Wide Association Analyses of Risk Tolerance and Risky Behaviors in over 1 Million Individuals Identify Hundreds of Loci and Shared Genetic Influences.” *Nature Genetics* 51 (2):245–257.
- Kaygusuz, Remzi. 2015. “Social Security and Two-Earner Households.” *Journal of Economic Dynamics and Control* 59:163–178.
- Kong, Augustine, Gudmar Thorleifsson, Michael L. Frigge, Bjarni J. Vilhjalmsón, Alexander I. Young, Thorgeir E. Thorgeirsson, Stefania Benonisdóttir, Asmundur Oddsson, Bjarni V. Halldorsson, Gisli Masson, Daniel F. Gudbjartsson, Agnar Helgason, Gyda Bjornsdóttir, Unnur Thorsteinsdóttir, and Kari Stefansson. 2018. “The Nature of Nurture: Effects of Parental Genotypes.” *Science* 359 (6374):424–428.
- Lee, James J., Robbee Wedow, Aysu Okbay, Edward Kong, Omeed Maghizian et al. 2018. “Gene Discovery and Polygenic Prediction from a Genome-Wide Association Study of Educational Attainment in 1.1 Million Individuals.” *Nature Genetics* 50 (8):1112–1121.
- Lusardi, Annamaria, Pierre-Carl Michaud, and Olivia S. Mitchell. 2017. “Optimal Financial Knowledge and Wealth Inequality.” *Journal of Political Economy* 125 (2):431–477.
- Martin, Alicia R., Christopher R. Gignoux, Raymond K. Walters, Genevieve L. Wojcik, Benjamin M. Neale, Simon Gravel, Mark J. Daly, Carlos D. Bustamante, and Eimear E. Kenny. 2017. “Human Demographic History Impacts Genetic Risk Prediction across Diverse Populations.” *The American Journal of Human Genetics* 100 (4):635–649.
- Merton, Robert C. 1969. “Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case.” *The Review of Economics and Statistics* 51 (3):247–257.
- Muslimova, Dilnoza, Hans van Kippersluis, Cornelius A. Rietveld, Stephanie von Hinke, and Fleur Meddens. 2020. “Dynamic Complementarity in Skill Production: Evidence From Genetic Endowments and Birth Order.” SSRN Scholarly Paper 3748468, Social Science Research Network, Rochester, NY.
- National Research Council. 2012. *Aging and the Macroeconomy: Long-Term Implications of an Older Population*.
- Nisbett, Richard E., Joshua Aronson, Clancy Blair, William Dickens, James Flynn, Diane F. Halpern, and Eric Turkheimer. 2012. “Intelligence: New Findings and Theoretical Developments.” *American Psychologist* 67 (2):130–159.

- Nishiyama, Shinichi and Kent Smetters. 2007. “Does Social Security Privatization Produce Efficiency Gains?*” *The Quarterly Journal of Economics* 122 (4):1677–1719.
- Okbay, Aysu, Jonathan P. Beauchamp, Mark Alan Fontana, James J. Lee, Tune H. Pers et al. 2016. “Genome-Wide Association Study Identifies 74 Loci Associated with Educational Attainment.” *Nature* 533 (7604):539–542.
- Papageorge, Nicholas W and Kevin Thom. 2020. “Genes, Education, and Labor Market Outcomes: Evidence from the Health and Retirement Study.” *Journal of the European Economic Association* 18 (3):1351–1399.
- Rietveld, Cornelius A., Sarah E. Medland, Jaime Derringer, Jian Yang, Tõnu Esko et al. 2013. “GWAS of 126,559 Individuals Identifies Genetic Variants Associated with Educational Attainment.” *Science* 340 (6139):1467–1471.
- Ronda, Victor, Esben Agerbo, Dorte Bleses, Preben Bo Mortensen, Anders Børghlum, Ole Mors, Michael Rosholm, David M. Hougaard, Merete Nordentoft, and Thomas Werge. 2022. “Family Disadvantage, Gender, and the Returns to Genetic Human Capital*.” *The Scandinavian Journal of Economics* 124 (2):550–578.
- Samuelson, Paul A. 1969. “Lifetime Portfolio Selection By Dynamic Stochastic Programming.” *The Review of Economics and Statistics* 51 (3):239–246.
- Scarr-Salapatek, Sandra. 1971. “Race, Social Class, and IQ.” *Science* 174 (4016):1285–1295.
- Schmitz, Lauren L. and Dalton Conley. 2017. “The Effect of Vietnam-era Conscription and Genetic Potential for Educational Attainment on Schooling Outcomes.” *Economics of Education Review* 61:85–97.
- Scholz, John Karl, Ananth Seshadri, and Surachai Khitatrakun. 2006. “Are Americans Saving “Optimally” for Retirement?” *Journal of Political Economy* 114 (4):607–643.
- Shiller, Robert J. 1990. *Market Volatility*. Cambridge, MA, USA: MIT Press.
- Social Security Administration. 2019. “Annual Statistical Supplement to the Social Security Bulletin, 2018.”
- . 2022a. “Proposals to Change Social Security.” <https://www.ssa.gov/oact/solvency/index.html>.
- . 2022b. “Summary of Provisions That Would Change the Social Security Program.” Tech. rep., Social Security Administration.

- Taubman, Paul. 1976. “The Determinants of Earnings: Genetics, Family, and Other Environments: A Study of White Male Twins.” *The American Economic Review* 66 (5):858–870.
- Trejo, Sam and Benjamin W. Domingue. 2018. “Genetic Nature or Genetic Nurture? Introducing Social Genetic Parameters to Quantify Bias in Polygenic Score Analyses.” *Biodemography and Social Biology* 64 (3-4):187–215.
- Tucker-Drob, Elliot M. and Timothy C. Bates. 2016. “Large Cross-National Differences in Gene \times Socioeconomic Status Interaction on Intelligence.” *Psychological Science* 27 (2):138–149.
- Turkheimer, Eric, Andreana Haley, Mary Waldron, Brian D’Onofrio, and Irving I. Gottesman. 2003. “Socioeconomic Status Modifies Heritability of IQ in Young Children.” *Psychological Science* 14 (6):623–628.
- Visscher, Peter M., Naomi R. Wray, Qian Zhang, Pamela Sklar, Mark I. McCarthy, Matthew A. Brown, and Jian Yang. 2017. “10 Years of GWAS Discovery: Biology, Function, and Translation.” *The American Journal of Human Genetics* 101 (1):5–22.
- Vissing-Jorgensen, Annette. 2002. “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures.” Working Paper 8884, National Bureau of Economic Research.
- Willis, Robert J. and Sherwin Rosen. 1979. “Education and Self-Selection.” *Journal of Political Economy* 87 (5, Part 2):S7–S36.
- Wolpin, Kenneth I. 2013. *The Limits of Inference without Theory*. The MIT Press.
- Young, Alexander I., Michael L. Frigge, Daniel F. Gudbjartsson, Gudmar Thorleifsson, Gyda Bjornsdottir, Patrick Sulem, Gisli Masson, Unnur Thorsteinsdottir, Kari Stefansson, and Augustine Kong. 2018. “Relatedness Disequilibrium Regression Estimates Heritability without Environmental Bias.” *Nature Genetics* 50 (9):1304–1310.

Appendix A Income Process Specification and Estimates

Our model of pre-tax labor income is

$$\ln \tilde{W}_{i,t} = f(\text{Age}_{i,t}, \text{EA}_i, \text{Coll}_i, \text{SES}_i, \text{DB}_i, \text{Year}_t, \text{Unemp}_t) + \mathcal{Z}_i^w + \epsilon_{i,t}^w.$$

We let $f(\cdot)$ be a linear function of its arguments, some of their interactions, and some simple transformations. We estimate this equation with fixed-effects regressions using the matched Social Security Administration earnings records that have been linked to the HRS.

The sample for our earnings regressions consists of all person-year observations for individuals in our main wealth sample with non-missing SSA earnings data. Our dependent variable is the natural logarithm of real earnings, taken from the SSA Summary earnings files. We adjust these amounts to replace top-coded values with estimates of the expected value of top-coded amounts as in Barth, Papageorge, and Thom (2020). We consider observations in the age range 22-65 or 22-retirement age for those with non-missing retirement data. We also restrict our sample to person-year observations with earnings of at least \$10,000 in real income (with 2010 as the base monetary year). The remaining sample consists of 31,883 person-year observations from 869 unique individuals.

Table A1 presents our estimates. SES.Miss is an indicator for individuals for whom the childhood socioeconomic status index is missing. DB.Ever is an indicator for whether the individual was ever observed receiving a defined benefit pension when retired. DB.Always is an indicator for whether the individual was always observed receiving a defined benefit pension when retired. DB.Miss is an indicator for individuals for which the information about defined benefit pensions is missing. In all specifications, we control for the 10 first principal components of genetic data as released in the SSGAC file for the EA polygenic score (Lee et al., 2018). We use specification “C” in our baseline results.

Appendix A.1 Inheritances

Our model for inheritances is

$$\text{Inher}_{i,t} = \begin{cases} \text{C.Inher}_i(\text{Age}_{i,t}), & \text{With probability } P_i^I(\text{Age}_{i,t}) \\ 0, & \text{With probability } 1 - P_i^I(\text{Age}_{i,t}). \end{cases}$$

Table A1: Income Process Estimates

	A	B	C
Const.	8.9706	8.5911	8.6266
Age	0.0854	0.1117	0.1092
Age ² /100	-0.0927	-0.1242	-0.1250
Coll	0.1182	-0.1649	-0.3027
EA	0.0515	0.0083	-0.0405
SES	0.0762	0.0156	0.0057
SES.Miss.	0.3673	0.0628	-0.0051
EA × SES.Miss.	-0.0230	-0.0228	-0.0199
Age × Coll			0.0099
Age × EA			0.0022
Age × SES			0.0016
Age × SES.Miss.			0.0085
Year	0.0189	0.0047	0.0075
max{0, Year - 1980}	-0.0217		
Year × Coll		0.0089	
Year × EA		0.0014	
Year × SES		0.0019	
Year × SES.Miss.		0.0097	
DB.Ever	0.1342	0.1311	0.1342
DB.Always	0.0407	0.0432	0.0411
DB.Miss.	0.0916	0.0768	0.0864
Unemp	-0.1051	-0.0642	-0.0648
Unemp ²	0.0049	0.0033	0.0033
Princ. Comp. Gen.	✓	✓	✓
Num. obs.	31883	31883	31883
$\sigma(\mathcal{Z}_i^w)$	0.3783	0.3498	0.3740
$\sigma(\epsilon_{i,t}^w)$	0.3906	0.3867	0.3867
Sigma OLS	0.4977	0.4949	0.4947

For all columns, the dependent variable is the natural logarithm of real earnings. See the main text for sample and variable definitions. All regressions include the first ten principal components of the full matrix of genetic information as controls.

Table A2: Inheritance Process Estimates

	P^I , LPM	P^I , Probit	$\ln(\text{C.Inher})$, OLS
Const.	-0.07 (0.01)	-6.03 (0.24)	11.70 (0.95)
Age	0.00 (0.00)	0.13 (0.01)	-0.04 (0.03)
Age ² /100	-0.00 (0.00)	-0.10 (0.01)	0.03 (0.03)
Coll	0.01 (0.00)	0.13 (0.03)	0.53 (0.10)
EA	0.00 (0.00)	0.02 (0.01)	0.15 (0.05)
Princ. Comp. Gen.	✓	✓	✓
Num. obs.	49000	49000	1056
Std. dev. Error			1.53

The first column presents coefficient estimates from a linear probability model in which the dependent variable is a binary indicator of whether an inheritance is received. The second column presents coefficient estimates from a probit model in which the dependent variable is a binary indicator of whether an inheritance is received. The third column presents coefficient estimates from a linear model in which the dependent variable is the natural logarithm of inheritances for those who receive them. All regressions include the first ten principal components of the full matrix of genetic information as controls. See the main text for sample definitions.

We estimate both the probability of receiving an inheritance P^I and the value of an inheritance conditional on its reception C.Inher as functions of a person’s age, education, and EA score. Table A2 presents our estimates. The first and second column model the probability of receiving an inheritance using a linear probability model and a probit model, respectively. The third column presents OLS estimates for the log-linear model we use for the value of inheritances conditional on reception. In all specifications, we control for the 10 first principal components of genetic data as released in the SSGAC file for the EA polygenic score (Lee et al., 2018).

The samples for each of the regressions are the following. First, using retrospective survey items about the timing of receiving an inheritance, we construct a panel of 49,000 household-year observations for the 870 households in our main wealth sample. Each household-year corresponds to a person-year for the male household member when he was between the ages of 22 and 80. We use this sample to estimate the models of the probability of receiving an inheritance, P^I . For the value of inheritances conditional on reception, we use the sub-sample of 1056 household-year observations for 575 households corresponding to periods when they received a non-zero inheritance.

Table A3: Defined benefit pension estimates

	A	B
Const.	-11.89	-14.10
BY	0.01	0.01
Coll	0.66	0.69
EA	0.11	
EA Quint. 2		0.11
EA Quint. 3		0.23
EA Quint. 4		0.39
EA Quint. 5		0.20

In both columns, the dependent variable is the natural logarithm of income from defined benefit pension plans. In the first specification, the EA score enters linearly; in the second, it enters discretized into quintile indicators. BY stands for birth year and Coll is a binary indicator of college completion.

Appendix A.2 Defined Benefit Flows

Table A3 presents our estimates of the defined-benefit pension-flow specifications. We use specification “B” in our structural model. Details about sample and variable definitions are available from the authors upon request.

Appendix B Socio-Economic Status and Unobserved Heterogeneity

Assume

$$\underbrace{\begin{bmatrix} \zeta_i^w \\ \zeta_i^F \\ \zeta_i^R \end{bmatrix}}_{\vec{\zeta}_i} = \underbrace{\begin{bmatrix} z_w \\ z_F \\ z_R \end{bmatrix}}_z \times \text{SES}_i + \underbrace{\begin{bmatrix} Z_i^w \\ Z_i^F \\ Z_i^R \end{bmatrix}}_{\vec{Z}}, \quad \begin{bmatrix} Z_i^w \\ Z_i^F \\ Z_i^R \end{bmatrix} \sim \mathcal{N}(\vec{0}, \Sigma_Z)$$

and

$$\text{SES}_i = \phi EA_i + \varepsilon_i^{\text{SES}}, \quad \varepsilon_i^{\text{SES}} \sim \mathcal{N}(0, \sigma_{\text{SES}}^2).$$

We can write the reduced form relationship $\vec{\zeta}_i = z\phi EA_i + z\varepsilon_i^{\text{SES}} + \vec{Z}_i$ and from it derive the fact

$$\zeta_i | EA_i \sim \mathcal{N}(z\phi EA_i, \sigma_{\text{SES}}^2 z z' + \Sigma_Z). \quad (16)$$

To generate the unobserved heterogeneity draws of an agent with EA score EA_i , we discretize the distribution in Equation 16 using 27 equiprobable points.

Additionally, we can write

$$\begin{bmatrix} \text{SES}_i \\ \vec{\zeta}_i \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ z & I \end{bmatrix} \begin{bmatrix} \text{SES}_i \\ \vec{Z}_i \end{bmatrix}$$

and from this relationship we know that

$$\begin{bmatrix} \text{SES}_i \\ \vec{\zeta}_i \end{bmatrix} | EA_i \sim \mathcal{N} \left(\begin{bmatrix} \phi EA_i \\ z \phi EA_i \end{bmatrix}, \begin{bmatrix} \sigma_{\text{SES}}^2 & \sigma_{\text{SES}}^2 z' \\ \sigma_{\text{SES}}^2 z & \sigma_{\text{SES}}^2 z z' + \Sigma_Z \end{bmatrix} \right).$$

Finally, from the conditional distribution of a partitioned normal, we know

$$E[\text{SES}_i | EA_i, \vec{\zeta}_i] = \phi EA_i + (\sigma_{\text{SES}}^2 z') (\sigma_{\text{SES}}^2 z z' + \Sigma_Z)^{-1} (\vec{\zeta}_i - z \phi EA_i). \quad (17)$$

We use Equation 17 to assign a socioeconomic status score to each of our simulated agents given their EA_i and $\vec{\zeta}_i$ draws.

Appendix C Computing Social Security Benefits

This section describes the procedure we follow for calculating the Social Security benefit schedules in our model. We aim to replicate the real methodology used by the Social Security Administration (Social Security Administration, 2019, Section 2) with simplification that ease the computational costs of solving our model.

Appendix C.1 The Average Indexed Monthly Earnings AIME

The basis of Social Security benefits are a worker's average earnings over his 35 highest-earning years. Tracking this object precisely would require us to incorporate multiple additional state variables to our model. Instead, we make the simplification that agent i 's benefits are calculated based on the wage trajectory that someone with i 's characteristics would *expect*. Formally, following Equation 6, i 's realized pre-tax wages at time t would be

$$\tilde{W}_{i,t} = \exp \left\{ f(\text{Age}_{i,t}, EA_i, \text{Coll}_i, \text{SES}_i, \text{DB}_i, \text{Year}_t, \text{Unemp}_t) + \mathcal{Z}_i^w + \epsilon_{i,t}^w \right\}.$$

We define his expected pre-tax wages as

$$\tilde{W}_{i,t}^E = \exp \left\{ f(\text{Age}_{i,t}, EA_i, \text{Coll}_i, \text{SES}_i, \text{DB}_i, \text{Year}_t, \text{Unemp}) + \mathcal{Z}_i^w + \frac{\sigma^2(\epsilon^w)}{2} \right\},$$

where $\bar{\text{Unemp}}$ is the average historical unemployment rate. The computational convenience comes from the fact that $\tilde{W}_{i,t}^E$ is the same for every agent with the same ex-ante characteristics as i and therefore tracking it requires no additional state variables.

For some given retirement age $A \in [\text{Age}_0^R, \text{Age}_f^R]$, denote with $\tilde{W}_{i,A}^*$ the 35 greatest elements of the set

$$\left\{ \min \left[\tilde{W}_{i,t}^E, \bar{T}(\text{Year}_t) \right] \right\}_{t=t(21)}^{t(A)-1}$$

where $t(a)$ is an auxiliary function indicating the simulation period at which the agent reaches age a . The set represents the 35 highest-earning years of i up to age A capped at the time-varying maximum taxable earnings. For each retirement age, we find the average indexed monthly earnings as

$$\text{AIME}_i^A = \frac{1}{12} \times \frac{1}{35} \sum_{W \in \tilde{W}_{A,i}^*} W.$$

Appendix C.2 Primary Insurance Amount and Adjustments

The primary insurance amount (PIA) is the monthly benefit that a person retiring at the *full retirement age* of 67 would receive. It is a concave piece-wise linear function of the AIME; it achieves its progressiveness by having decreasing replacement rates for different income brackets. We use the bracket limits and replacement rates defined by the SSA, which the following equation summarizes

$$\begin{aligned} \text{PIA}(\text{AIME}_i^A) = & \min\{\text{AIME}_i^A, 895\} \times 0.90 + \\ & \max\{\min\{5397 - 859, \text{AIME}_i^A - 859\}, 0\} \times 0.32 + \\ & \max\{\text{AIME}_i^A - 5397, 0\} \times 0.15. \end{aligned} \quad (18)$$

The final step is to adjust the PIA depending on whether a person retires before or after the full retirement age. Benefits are reduced by 5/9 percentage points per month of early retirement up to 36 months, and by 5/12 percentage points for each additional month over 36. On the other hand, for agents retiring after the full retirement age, benefits increase by 16/24 percentage points for each month not receiving benefits up to age 70, when benefits stop increasing. Formally, for an age of retirement A , defining $H = A - 67$, we obtain yearly benefits as

$$\text{SSB}_i^A = \begin{cases} 12 \times \text{PIA}_i^A \times \left(1 - \frac{5}{9} \frac{\min\{12|H|, 36\} + \frac{5}{12} \max\{12|H| - 36, 0\}}{100} \right), & \text{If } H < 0 \\ 12 \times \text{PIA}_i^A \times \left(1 + \frac{16}{24} \frac{\min\{H, 3 \times 12\}}{100} \right), & \text{If } H \geq 0 \end{cases}. \quad (19)$$

Appendix D Recursive Representation of the Agent's Problem

This section presents all the dynamic optimization problems that our model's agents can face depending on their age and retirement status. All symbols and variables follow the notation introduced in Section 4. We drop the individual sub-indices i in this section.

- Between ages 21 and $\text{Age}_0^R - 2$, the agent has not retired and does not have the option to retire next period. Therefore, his problem is

$$V_t(A_t, \mathcal{R}_t = 0) = \max_{C_t, \phi_t} u(C_t, \ell = 1) + \beta \delta_t \mathbb{E}_t [V_{t+1}(A_{t+1}, \mathcal{R}_{t+1} = 0)] + \delta_t \varphi(S_t)$$

Subject to:

$$0 \leq C_t, \quad 0 \leq S_t, \quad 0 \leq \phi_t \leq 1$$

$$S_t = A_t - C_t - F_i \times \mathbf{1}[\phi_t > 0],$$

$$A_{t+1} = \left\{ (1 - \tau^c) \left[\phi_t \tilde{R}_{t+1} + (1 - \phi_t) R \right] + \tau^c \right\} \times S_t + \tau_t(\tilde{W}_{t+1}) + \text{Inher}_{t+1}.$$

- Non-retired agents whose age satisfies $\text{Age}_t \in [\text{Age}_0^R - 1, \text{Age}_f^R - 2]$ have the option of deciding to retire (with the choice becoming effective next period). Their problem is:

$$V_t(A_t, \mathcal{R}_t = 0) = \max_{C_t, \phi_t, \mathcal{R}_{t+1}} u(C_t, \ell = 1) + \beta \delta_t \mathbb{E}_t [V_{t+1}(A_{t+1}, \mathcal{R}_{t+1})] + \delta_t \varphi(S_t)$$

Subject to:

$$0 \leq C_t, \quad 0 \leq S_t, \quad 0 \leq \phi_t \leq 1$$

$$\mathcal{R}_{t+1} \in \{0, (\text{Age}_t + 1 - \text{Age}_0^R) + 1\}$$

$$S_t = A_t - C_t - F \times \mathbf{1}[\phi_t > 0],$$

$$A_{t+1} = \left\{ (1 - \tau^c) \left[\phi_t \tilde{R}_{t+1} + (1 - \phi_t) R \right] + \tau^c \right\} \times S_t +$$

$$\mathbf{1}[\mathcal{R}_{t+1} = 0] \times \tau_t(\tilde{W}_{t+1}) +$$

$$\mathbf{1}[\mathcal{R}_{t+1} > 0] \times (1 - \tau^s) \times \text{DBf} +$$

$$\mathbf{1}[\mathcal{R}_{t+1} > 0, \text{Age}_{t+1} \geq \text{Age}_{\min}^{SS}] \times (1 - \tau^s) \times \text{SSB}(\mathcal{R}_{t+1}) +$$

$$\text{Inher}_{t+1}.$$

- At $\text{Age}_t = \text{Age}_f^R - 1$ an agent who has not retired is forced to retire since he must be retired at age Age_f^R . Thus, his problem is:

$$V_t(A_t, \mathcal{R}_t = 0) = \max_{C_t, \phi_t} u(C_t, \ell = 1) + \beta \delta_t \mathbb{E}_t [V_{t+1}(A_{t+1}, \mathcal{R}_{t+1})] + \delta_t \varphi(S_t)$$

Subject to:

$$0 \leq C_t, \quad 0 \leq S_t, \quad 0 \leq \phi_t \leq 1$$

$$\mathcal{R}_{t+1} = \text{Age}_f^R - \text{Age}_0^R + 1$$

$$S_t = A_t - C_t - F \times \mathbf{1}[\phi_t > 0],$$

$$A_{t+1} = \left\{ (1 - \tau^c) \left[\phi_t \tilde{R}_{t+1} + (1 - \phi_t) R \right] + \tau^c \right\} \times S_t + \\ (1 - \tau^s) \times \text{DBf} + \\ \mathbf{1}[\text{Age}_{t+1} \geq \text{Age}_{\min}^{\text{SS}}] \times (1 - \tau^s) \times \text{SSB}(\mathcal{R}_{t+1}) + \\ \text{Inher}_{t+1}.$$

- For a retired agent who retired in the n th possible period ($\mathcal{R}_t = n > 0$), the recursive problem for a non-terminal period is:

$$V_t(A_t, \mathcal{R}_t = n) = \max_{C_t, \phi_t} u(C_t, \ell = 0) + \beta \delta_t \mathbb{E}_t [V_{t+1}(A_{t+1}, \mathcal{R}_{t+1} = n)] + \delta_t \varphi(S_t)$$

Subject to:

$$0 \leq C_t, \quad 0 \leq S_t, \quad 0 \leq \phi_t \leq 1$$

$$S_t = A_t - C_t - F \times \mathbf{1}[\phi_t > 0],$$

$$A_{t+1} = \left\{ (1 - \tau^c) \left[\phi_t \tilde{R}_{t+1} + (1 - \phi_t) R \right] + \tau^c \right\} \times S_t + \\ (1 - \tau^s) \times \text{DBf} + \\ \mathbf{1}[\text{Age}_{t+1} \geq \text{Age}_{\min}^{\text{SS}}] \times (1 - \tau^s) \times \text{SSB}(\mathcal{R}_{t+1}) + \\ \text{Inher}_{t+1}.$$

- In the terminal period, the agent simply allocates his assets between consumption and bequests. He won't be working, provided that the maximum age of retirement is lower

than the oldest possible age. His problem is:

$$V_T(A_T, \mathcal{R}_T) = \max_{C_T} u(C_T, \ell = 0) + \varphi(S_T)$$

Subject to:

$$0 \leq C_T, \quad 0 \leq S_T$$

$$S_T = A_T - C_T$$

Appendix E Solution of the Life-Cycle Model

We solve the model using value-function iteration for each of the 5,130 agent types. This section describes how we solve each agent's dynamic problem. We drop the i sub-indices for compactness.

Appendix E.1 Grids and Discretizations

Agents have two continuous choices: consumption C and the risky-asset portfolio share ϕ . We discretize C expressing it as a fraction of the assets available for consumption, starting at 0.005 until 1.000 in increments of 0.005. We discretize the risky asset portfolio share using 11 points, from 0.0 to 1.0 in increments of 0.1. Finally, we construct a grid that we use to represent discretizations of various versions of wealth (e.g., start-of period, savings) and generically refer to it as the wealth grid. The wealth grid has 51 points spanning $[10^2, 10^7]$ USD and is denser at lower values.

We also discretize the two normal random variables over which agents form expectations, the income shock ϵ^w and the market log-return $\ln R^{\text{SP500}}$. We use equiprobable grids with 9 points for each.

Appendix E.2 Backward Induction

Appendix E.2.1 Terminal Period

An agent that reaches the terminal age of 90 faces no portfolio decision, he simply allocates his wealth between consumption and bequests. His optimal solution will be to consume all of his available wealth. Since 90 is also larger than the maximum age of retirement, we know

the agent will be retired. Therefore, the agent's value function is

$$V_{90}(A_{90}, \mathcal{R}_{90}) = \max_{C_{90}} u(C_{90}, \ell = 0) + \varphi(S_{90})$$

Subject to:

$$0 \leq C_{90}, \quad 0 \leq S_{90}$$

$$S_{90} = A_{90} - C_{90}$$

We solve this problem for every A_{90} in our wealth grid with consumption taking the values of our discretized proportional grid. The solution is the same for every retirement status \mathcal{R}_{90} . We use the solutions to construct linear interpolators for $V_{90}(\cdot, \mathcal{R}_{90})$ and move on to non-terminal periods.

Appendix E.2.2 Non-Terminal Periods

Given value function $V_{t+1}(\cdot, \cdot)$, we must solve for $V_t(\cdot, \cdot)$.

We start by defining the function $\mathbf{E}max_{t+1}(\cdot, \cdot, \cdot)$ as

$$\begin{aligned} \mathbf{E}max_t(S_t, \phi_t, \mathcal{R}_{t+1}) &= E_t [V_{t+1}(A_{t+1}, \mathcal{R}_{t+1})] \\ &\text{where} \\ X_{t+1} &= \left\{ (1 - \tau^c) \left[\phi_t \tilde{R}_{t+1} + (1 - \phi_t) R \right] + \tau^c \right\} \times S_t \\ Y_{t+1} &= \text{Inher}_{t+1} + \mathbf{1}[\mathcal{R}_{t+1} = 0] \times \tau_t(\tilde{W}_{t+1}) + \\ &\quad \mathbf{1}[\mathcal{R}_{t+1} > 0] \times (1 - \tau^s) \times \text{DBf} + \\ &\quad \mathbf{1}[\mathcal{R}_{t+1} > 0, \text{Age}_{t+1} \geq \text{Age}_{\min}^{SS}] \times (1 - \tau^s) \times \text{SSB}(\mathcal{R}_{t+1}) + \\ A_{t+1} &= X_{t+1} + Y_{t+1} \end{aligned} \tag{20}$$

$\mathbf{E}max$ represents the continuation value from each of the possible states in which the agent can end period t . In Equation 20, we have separated A_{t+1} into wealth coming from past savings X_{t+1} and current income Y_{t+1} . These two components are independent from the point of view of period t and we take advantage of this fact to approximate $\mathbf{E}max$.

We first define the intermediate function $\mathcal{Q}_{t+1}(\cdot, \cdot)$ as

$$\mathcal{Q}_t(X_{t+1}, \mathcal{R}_{t+1}) = E[V_{t+1}(X_{t+1} + Y_{t+1}, \mathcal{R}_{t+1}) | X_t],$$

where the expectation is being taken over Inher_{t+1} and income shock ϵ_{t+1}^w . For every feasible

value of \mathcal{R}_{t+1} , we evaluate this function with X_{t+1} taking the values in our wealth grid and use the results to construct linear interpolators of $\mathcal{Q}_{t+1}(\cdot, \mathcal{R}_{t+1})$.

Then, for every feasible \mathcal{R}_{t+1} and every ϕ_t , we evaluate

$$\mathbf{E}\max_{t+1}(S_t, \phi_t, \mathcal{R}_{t+1}) = E[\mathcal{Q}_{t+1}(X_{t+1}, \mathcal{R}_{t+1})|S_t, \phi_t]$$

with S_t taking all the values in our wealth grid and the expectation being taken over market returns only. We use the results to construct linear interpolators of $\mathbf{E}\max_{t+1}(\cdot, \phi_t, \mathcal{R}_{t+1})$.

We now turn to finding the agent's optimal choices. It is useful to note at this point that we can use Equation 20 to re-express the agent's problem as

$$V_t(A_t, \mathcal{R}_t) = \max_{C_t, \phi_t, \mathcal{R}_{t+1}} u(C_t, \ell) + \beta\delta_t \mathbf{E}\max_{t+1}(S_t, \phi_t, \mathcal{R}_{t+1}) + \delta\varphi(S_t)$$

Subject to:

$$\begin{aligned} 0 \leq C_t, \quad 0 \leq S_t, \quad 0 \leq \phi_t \leq 1 \\ S_t = A_t - C_t - F \times \mathbf{1}[\phi_t > 0], \end{aligned}$$

with the feasible options for \mathcal{R}_{t+1} depending on the agent's age and current retirement status.

We can see the choice of $\{C_t, \phi_t, \mathcal{R}_{t+1}\}$ as happening in two steps: the agent first commits himself to a ϕ_t and pays the cost F if necessary and then, conditioning on that choice, picks $\{C_t, \mathcal{R}_{t+1}\}$. The problem and value function of the agent who takes ϕ_t and his net-of-fixed-cost wealth \tilde{A}_t as given is

$$\tilde{V}_t(\tilde{A}_t, \phi_t, \mathcal{R}_t) = \max_{C_t, \mathcal{R}_{t+1}} u(C_t, \ell) + \beta\delta_t \mathbf{E}\max_{t+1}(S_t, \phi_t, \mathcal{R}_{t+1}) + \delta\varphi(S_t)$$

Subject to:

$$\begin{aligned} 0 \leq C_t, \quad 0 \leq S_t \\ S_t = \tilde{A}_t - C_t. \end{aligned}$$

We solve this problem for all combinations of \tilde{A}_t in our wealth grid, ϕ_t on its grid, and the feasible \mathcal{R}_t . The values we consider for C_t are our consumption (proportional) grid times \tilde{A}_t . We use the results to construct linear interpolators for $\tilde{V}_t(\cdot, \phi_t, \mathcal{R}_t)$ for every combination of $\{\phi_t, \mathcal{R}_t\}$.

Finally, we can express the value function as

$$V_t(A_t, \mathcal{R}_t) = \max_{\phi_t} \tilde{V}_{t+1}(\tilde{A}_t, \phi_t, \mathcal{R}_{t+1})$$

Subject to:

$$\begin{aligned} 0 &\leq \tilde{A}_t, \quad 0 \leq \phi_t \leq 1 \\ \tilde{A}_t &= A_t - F \times \mathbf{1}[\phi_t > 0]. \end{aligned}$$

For every feasible \mathcal{R}_{t+1} , we solve this problem at every A_t in our wealth grid and use the results to construct linear interpolators for $V_t(\cdot, \mathcal{R}_t)$. With these interpolators, we move on to period $t - 1$ and repeat the process.

Appendix F Steady State Revenues From Changes to the S.S. System

Appendix F.1 The Steady-State Cost of Social Security

Parameterizing our counterfactuals policies in a way that makes them comparable requires us to calculate what would be the per-period per-person cost of a Social Security system in *steady state*. To calculate this cost, we start by grouping all the age-invariant characteristics of our model—EA score, college attendance, childhood socioeconomic status, birth year, pension arrangement, and unobserved heterogeneity draws—in a vector that we call \vec{H}_i for person i . Then, we use $\mathcal{C}(a, \vec{H})$ to denote the expected net revenue that the government will collect from a person with age-invariant characteristics \vec{H} when he is a -years old through the social security system. We consider the net revenue to be FICA taxes paid on income if working, or the negative of social security benefits if retired.

We think of a *steady state* as a point in which

- The distribution of age-invariant characteristics over the agents born each period (which we will denote $F_{\vec{H}}$) is constant.²⁰ The number of agents born each period is also constant and denoted with N_{ss} .
- All of the model's components (shock distributions, earning patterns) have stabilized so that $\mathcal{C}(a, \vec{H})$ is not time-varying.

²⁰For a characteristic like birth year, which will not stay constant, we interpret this definition as pertaining to the distribution of the individual characteristics associated with being born in a particular year that matter in our model.

If we denote p_a the probability that a person survives to age a , we know that in any steady state period there will be $p_a \times N_{ss}$ agents of age a alive. Since in our model survival probabilities are not related to characteristics \vec{H} , the distribution of characteristics across agents of any age at any steady-state time will be $F_{\vec{H}}$. Therefore, the total per-period expected net Social Security revenue will be

$$\sum_{a=21}^{90} \left(p_a N_{ss} \times \int \mathcal{C}(a, h) dF_{\vec{H}}(h) \right) \propto \sum_{a=21}^{90} p_a \times E_{\vec{H}}[\mathcal{C}(a, \vec{H})].$$

We approximate the right hand side of the previous equation using our simulated sample \mathcal{I}

$$\sum_{a=21}^{90} \frac{1}{|\mathcal{I}|} \sum_{i \in \mathcal{I}} c_{i,a}, \tag{21}$$

where $c_{i,a}$ is the net Social Security revenue collected from agent i at age a (zero if the agent is dead). This approximation relies on assuming that our sample can produce a good approximation of the steady-state average Social Security revenue at every age.