Genetic Ability, Wealth, and Financial Decision-Making*

Daniel Barth† Nicholas W. Papageorge‡ Kevin Thom§

December 20, 2016

Preliminary and incomplete. Please do not cite or circulate without express permission from the corresponding author.

Abstract: Recent advances in behavioral genetics have enabled the discovery of genetic scores linked to a variety of economic outcomes, including education. We build on this progress to demonstrate that the same genetic variants that predict educational attainment independently predict financial wealth in the Health and Retirement Study (HRS). This relationship is partly explained by higher earnings, but a substantial portion of this association cannot be explained mechanically by income flows or bequests. This leads us to explore the role of beliefs, financial literacy and portfolio decisions in explaining this genetic gradient in wealth. We show that individuals with lower genetic scores are more prone to reporting "extreme beliefs" (e.g., reporting that there is a 100% chance of a stock market decline in the near future) and they invest their savings accordingly (e.g., avoiding the stock market). Our findings suggest that genetic factors that promote human capital accumulation contribute to wealth disparities not only through education and higher earnings, but also through their impact on the ability to process information and make good financial decisions. The association between genetic ability and wealth is substantially lower among households receiving a defined benefit pension. Policies that transfer greater responsibility to individuals to manage their wealth might therefore exacerbate the consequences of labor market inequality.

*We thank Aysu Okbay for constructing and sharing the polygenic score used for HRS respondents. We are also grateful for helpful comments from Jess Benhabib, Dan Benjamin, David Cesarini, Steven Durlauf and Jon Faust along with seminar participants at the 2016 HCEO Conference on Genetics and Social Science. The usual caveats apply.

†Center for Economic and Social Research, University of Southern California. Email: dannybarth@gmail.com.

‡The Johns Hopkins University. Email: papageorge@jhu.edu.

§New York University. Email: kevin.thom@nyu.edu.

1 Introduction

Wealth inequality in the United States and many other countries is large and growing (Saez and Zucman, 2014; Jones, 2015). An extensive literature attempts to explain wealth inequality through variation in the earnings process, entrepreneurial talent, bequests, risk aversion and time-discounting. While these factors can generate a skewed wealth distribution, they fail to reproduce other key empirical features including the thickness of the tail.[1] A new wave of theoretical work argues that cross-sectional heterogeneity in the returns to wealth are required to match the basic features of the wealth distribution (Benhabib, Bisin, and Zhu, 2011; Benhabib and Bisin, 2016). This argument is supported by a growing empirical literature finding substantial persistent heterogeneity in such returns, with wealthier individuals enjoying greater returns (Fagereng et al., 2016; Bach, Thiemann, and Zucco, 2015). Since these persistent factors are treated as unobserved, little is known about their nature or the behavioral mechanisms through which they operate. Yet it may be highly policy relevant to understand, for example, whether this heterogeneity works through a preference channel (e.g. risk aversion), or through a skill (e.g. financial decision-making).

In this paper, we identify a biological source of wealth inequality: observed genetic ability related to human capital accumulation. In particular, we show that the same observed genetic markers that predict educational attainment also predict financial wealth in the Health and Retirement Study (HRS). Surprisingly, we demonstrate that this relationship is not merely driven by earnings or other income flows. Rather, the estimated association between genes and financial wealth is economically large and statistically significant after controlling for education, earnings, financial bequests and a host of other controls. Moreover, being able to observe a set of genetic variants implicated in wealth accumulation allows us to directly assess the mechanisms through which they operate and how they interact with policy-relevant environments. As discussed below, we provide novel evidence that the genetic endowments we study are linked to financial decision-making, stock market participation, financial literacy and more fundamentally, probabilistic thinking. These results can inform our interpretation of persistent differences in returns to wealth since they suggest that genetic endowments that predict education and, therefore, higher earnings also predict better financial decisions and wealth. Moreover, by implicating financial decision-making in wealth inequality, these findings also suggest that relatively straightforward policy tools, such as stronger public pension schemes, may help to reduce inequality and poverty among the

[1] Some models can match the thickness of the tail, but only under implausible assumptions about the level of heterogeneity in the earnings process. For example, Kindermann and Krueger (2014) require the top 0.25% of earners to earn 400-600 times more than the median earner. Empirically, this number is closer to 33.

elderly. This is especially relevant given the dramatic shift away from defined benefit retirement plans towards options that give individuals greater financial autonomy (Poterba and Wise, 1998).

To measure genetic endowments, we follow recent advances in behavioral genetics that have led to the identification of specific genetic markers that predict educational attainment. Results from the state of the art in this literature (Okbay et al., 2016) allow for the construction of a *polygenic score*—an index of these genetic markers—that robustly predicts education in the HRS (Papageorge and Thom, 2016). To explain the advantage of using an observed measure of genetic ability, it is useful to contrast our work with earlier studies linking genetic factors to economic behaviors the affect wealth accumulation. Results from twins studies have shown that genetic factors play a non-trivial role in explaining savings behavior and portfolio decisions (Cronqvist and Siegel, 2015; Cesarini et al., 2010).[2] However, it is generally difficult to learn about how specific genetic factors operate using twins data. For one, twins studies do not identify which particular markers matter for wealth, nor do they allow one to measure these genetic endowments at the individual-level. This means that analysis of their function is limited to variance decompositions related to mechanisms that are measured in existing twins data sets. Moreover, while testing hypotheses about specific mechanisms is conceptually possible using information on twins, in practice it requires large amounts of data to permit stratification by each potential mediating factor.[3] The development of observable polygenic scores helps to overcome these difficulties.

We present three main sets of findings. Our first results establish a robust relationship between the polygenic score and financial wealth. A one-standard-deviation increase in the score is associated with a 23 percent increase in household wealth (approximately $117,000). Education and household income account for a little more than half of this association. The relationship is only slightly moderated when we control for financial bequests, which could capture how individuals with higher polygenic scores have parents who are more highly educated and therefore wealthier. Despite these controls, the polygenic score continues to exhibit a strong and economically large relationship with financial wealth. This leads us to explore how these genetic factors relate to investment decisions. A classic example of such a decision is stock market participation (Vissing-Jørgensen, 2002). We demonstrate that the polygenic score predicts stock market participation even after we have controlled

[2]For example, using the Swedish Twin Registry, Cesarini et al. (2010) demonstrate that about 25% of individual variation in portfolio risk is attributable to genetic variation while Cronqvist and Siegel (2015) show that 35% of variation in propensity to save has a genetic basis.

[3]Variance decomposition exercises such as twins studies treat genes as unobserved factors. Learning about interactions between observed and unobserved factors is generally difficult and rests on modeling assumptions.

for education, wealth, and wages, and that stock ownership significantly moderates the relationship between the polygenic score and wealth.[4] Since exposure to risk is one factor that affects the returns to wealth, this provides some evidence that the genetic endowments captured by this score may serve as a genetic micro-foundation for the type of heterogeneity described in Benhabib (et al 2011).

Having shown evidence that at least part of the genetic gradient in financial wealth can be explained by financial decision-making, our second set of results exploits rich data on subjective macroeconomic expectations to provide evidence on one mechanism at play. We show that the polygenic score is associated with expectations that are objectively more correct (Dominitz and Manski, 2007).[5] Furthermore, a lower score is associated with beliefs about the economy that are heaped on probabilities of 0% or 100% (a phenomenon we refer to as "extreme beliefs"), which suggests difficulty with probabilistic thinking. Importantly, these extreme beliefs are not purely the result of respondent confusion or measurement error; households that report extreme beliefs make choices that are consistent with their beliefs. For example, households that report a 100% probability of a stock market increase over the following year are 20 percentage points more likely to invest in stocks than those that report a 0% probability.[6] We do not claim that this is the sole mechanism linking genetic ability, decision-making, and wealth. However, the results on beliefs offer a clear channel driving some individuals to make consistently better financial decisions and to access higher returns to wealth.

Our third set of findings considers policy implications of our first two sets of findings. If genetic ability affects the quality of financial decision-making, policies that grant individuals greater discretion in how they invest and manage financial decisions could exacerbate ability-based inequality. Comparing wealth outcomes of individuals with and without defined benefit pensions offers some evidence on this. In particular, we show that the gene-wealth gradient is much stronger for people who do not have such pensions as a source of retirement income. Our interpretation is that lower genetic scores are more problematic when

[4]We also control for a measure of risk aversion, which previous research suggests has a genetic basis (Cesarini et al., 2009). Unfortunately, we can only do so for a subsample of individuals for whom this information is available. Doing so does not change parameter estimates, but we believe we may be underpowered to detect differences.

[5]Hurd (2009) provides a review of subjective probabilities reported in household surveys such as the HRS. A number of researchers have studied used the HRS to study cognition, probabilistic thinking and investment decisions (Lillard and Willis, 2001; Kézdi and Willis, 2009, 2003). Another set of related studies focuses on cognitive decline and retirement decisions (Rohwedder and Willis, 2010; Kézdi and Willis, 2013; Delavande et al., 2006; Delavande, Rohwedder, and Willis, 2008).

[6]Lillard and Willis (2001) also recognize that focal point beliefs are predictive of investment behavior. This highlights an important distinction: individuals who report extreme beliefs do not necessarily make irrational choices at random, but rather make potentially optimal choices based on incorrect beliefs.

individuals are tasked with making their own saving and investment decisions. Individuals receiving defined pensions typically have much less control over how retirement funds are invested and managed relative to those with defined contribution plans or other arrangements. Although we find no association between the genetic score and having a defined benefit plan, the genetic gradient in wealth is about half as large in the sample of defined benefit recipients. Such evidence highlights the possible costs of policy proposals, such as the privatization of Social Security, that would increase latitude in financial decision-making. Correspondingly, this suggests that some groups may benefit from institutions that make financial decisions on their behalf. Promoting better financial decisions represents an alternative to such paternalism. However, this is only viable if the genetic gradients we describe are easily modified by household resources and other environmental factors. To explore this possibility, we assess whether individuals with low genetic scores who are born into more advantageous environments are less likely to report "extreme beliefs". Although measures of childhood socioeconomic status (SES) have been shown to moderate the relationship between these genetic endowments and college completion in the HRS (Papageorge and Thom, 2016), we do not find that SES plays such a moderating role for wealth or the accuracy of subjective expectations. This provides a note of caution for policy debates on inequality. Some drivers of inequality might not be addressed through standard redistributive policies.

We demonstrate that genetic endowments related to education are associated with wealth and behaviors that increase the returns to wealth. This suggests that the persistent heterogeneity described in Benhabib, Bisin, and Zhu (2011) may at least partially operate through genetic endowments linked to human capital accumulation, financial decision-making, and expectations. The insights we gain by combining the polygenic score to HRS data enriches the body of knowledge that has emerged from twins studies. For example, we note that our analysis both complements and contrasts with the findings of Cronqvist and Siegel (2015), who using twins to study a genetic basis for savings behavior. Our results validate and buttress the point that genetic factors drive heterogeneity in wealth. However, Cronqvist and Siegel (2015) find no genetic correlation between savings and education, and suggest that genetic factors related to wealth may operate through self control because of genetic correlations between savings, smoking, and BMI. In contrast, the polygenic score that we study captures exactly those genetic variants that are associated with education and shows that they are related to wealth. This suggests that the polygenic score we study measures genetic endowments that are distinct from, and operate through different mechanisms, than those genes that matter for savings. Indeed, the combination of results here and in Cronqvist and Siegel (2015) support the existence of multiple clusters of genetic endowments linked to wealth—some operating through savings behavior, and others potentially operating through

financial decision-making and expectations formation. Moreover, while the genetic endowments that predict savings surely impact wealth, they do not necessarily induce persistent heterogeneity in the returns to wealth. That is, the genetic endowments reflected in the polygenic score we study here may be particularly relevant for the heterogeneity described in (Benhabib, Bisin, and Zhu, 2011).

The remainder of this paper is organized as follows. Section 2 provides some background on the genetic index used in this paper. Section 3 introduces the data used in this project, paying particular attention to the construction of wealth measures. Section 4 studies the relationship between the polygenic score and wealth. Section 5 discusses the polygenic score, financial decision-making, beliefs about the economy and probabilistic thinking. Section 6 discusses some policy implications of our results. Section 7 concludes.

2 Genetic Data and GWAS

We measure labor market ability using a polygenic score that predicts educational attainment. Since economists are only just beginning to use genetic data, we first provide a brief introduction to the genetic data we use, as well as its advantages and shortcomings. While we only present some main ideas here, Appendix A offers much more detail, and a host of references for the interested reader.

The first point concerns problems linking genetic data to economic behavior. An individual's genome consists of 23 pairs of chromosomes, one from each parent. Each chromosome can be divided into sections that are functionally related, called *genes*. Each gene is comprised millions of *nucleotide pairs* — these are the "rungs on the ladder" in illustrations of our DNA — and such pairs can take only one of two values.[7] Further, humans only differ from one another in a few million of these pairs (less than 1% of the total). The pairs in which humans may differ are called *single nucleotide polymorphisms*, or SNPs (pronounced "snips").

Once scientists could observe these SNPs (when the human genome was sequenced about 15 years ago) researchers began to link specific SNPs to physical characteristics (e.g., hair color), but also to behavior (e.g., smoking). However, with millions of SNPs, it is not clear how to identify which SNPs are relevant for a particular trait. An initial remedy was to use so-called "candidate genes", identified by theories about biological processes likely to be

[7]Each pair can be an adenine-thymine (AT) pair or a guanine-cytosine (GC) pair. While each rung will be one of these two molecules, the rungs might differ in terms of their relationship to the sides or rails of the ladder. That is, we might have an AT molecule or a TA molecule, where different ends of the rung are connected to different sides of the ladder.

important for the behavior or outcomes of interest. A multiple-hypothesis testing problem arose, however, given the sheer number of possible candidates, even within a particular gene. This was such a problem that many encouraging results turned out to be false positives (Hewitt, 2012; Benjamin et al., 2012).

These challenges led behavioral geneticists to a new approach, known as genome-wide association studies (or GWAS). In a GWAS, researchers embrace an atheoretical approach and test each SNP individually for a relationship with the outcome of interest. Stringent controls are applied to account for multiple hypothesis testing. Essentially, all SNPs are regressed one by one, along with a set of essential controls. The GWAS revolution has produced a number of robust, credible results, including the discovery of the most well known genetic variant associated with obesity and several markers associated with smoking (Bierut, 2010; Thorgeirsson et al., 2010). Once a GWAS generates a series of coefficients associated with individual markers, these coefficients can be used to construct genetic indices called *polygenic scores*. These scores are typically linear combinations of individual markers. The Appendix provides considerably more detail on how this is done.

Our measure of genetic ability comes directly from a series of landmark GWAS discoveries that have identified some of the first direct associations between specific SNPs and educational attainment (Rietveld et al., 2013; Okbay et al., 2016). After documenting the first GWAS for education (Rietveld et al., 2013), the Social Science and Genetics Consortium recently extended their analysis to perform an educational attainment GWAS with an unprecedented sample size of 293,723 (Okbay et al., 2016). Our genetic measure is the polygenic score developed in this follow-up study, which combines all genotyped SNPs. We refer to this measure as the *EA Score*, indicating that "educational attainment" is the trait of interest. In recent work, Papageorge and Thom (2016) show that the *EA Score* predicts labor market outcomes independently of education, including wages and retirement.

We conclude this section by briefly discussing three important caveats and points of clarification. First, it is important to note that the genetic variants used in the construction of this genetic score are not located on sex chromosomes. For this reason, the distribution of these variants should be identical across men and women. Second, we do not claim to estimate *causal effects* of particular genetic variants. Any gene-outcome association that we observe in general reflects a combination of a direct effect and an indirect effect operating through the environments that parents make for their children. Parents with advantageous genetic endowments (which they pass on to their children) are more likely to have the resources or capacity to create better environments. Even so, an individual's genetic make-up is not *changed* by human capital investments. In contrast, IQ and other cognitive test scores are subject to the critique that they reflect environmental factors, such as earlier

human capital investments. For example, Bharadwaj, Løken, and Neilson (2013) find that variation in health care received by newborns has an impact on academic achievement years later.[8] Genetic indices are not subject to this critique because they are fixed at conception.

A final caveat to our use of genetic data is that it may misrepresent ability. By aggregating a number of genetic variants into a single score, we are implicitly assuming that these factors work together in determining a single scalar value. However, a growing literature suggests that ability is multi-dimensional. There may be distinct cognitive abilities (e.g., mathematical ability or facility with language), each possessing different associations with economic outcomes (Willis and Rosen, 1979; Heckman, 1995; Cawley et al., 1997). Ability may also encompass not just cognition, but non-cognitive factors as well (Heckman and Rubinstein, 2001).[9] On this point, we are tied to the state of the art in genetics. We do not yet have the tools to credibly determine whether the individual genetic markers that make up the score contribute to distinct abilities.

3 The HRS Sample, the Polygenic Score and Wealth

In this section, we introduce the data set we use to examine how genetic ability endowments relate to wealth. Section 3.1 provides basic summary statistics for our analytic sample. Section 3.2 provides details on our construction of household wealth.

3.1 Sample Construction

The Health and Retirement Study (HRS) is a longitudinal panel study that follows Americans over age 50 and their spouses. Surveys began in 1992 and occur every two years. The HRS collected genetic samples from 15,680 individuals over the course of three waves (2006, 2008, 2010). Genetic data from the 2010 wave have not yet been publicly released, so our sample only includes individuals genotyped in 2006 and 2008.[10] Individuals in the genotyped sample tend to be born in younger birth cohorts because survival until at least 2006 is required for

[8] Even birth weight, another proxy of innate endowments that has been used in prior literature is not immune to this critique as it reflects *in utero* investments, e.g., mothers' smoking behavior (Lien and Evans, 2005), exposure to pollutants (Currie, Neidell, and Schmieder, 2009) stress during pregnancy (Camacho, 2008; Currie and Rossin-Slater, 2013) or mothers' own health (Costa, 1998). See also Aizer and Currie (2014) for a recent discussion.

[9] On multidimensionality, Willis and Rosen (1979) emphasize manual skill, which they distinguish from academic skill.

[10] Release of genetic data from the 2010 wave is imminent. Thereafter, the genetic score variable must be re-computed for the larger sample. At that point, it should be straightforward to extend our econometric analysis to include the larger sample.

inclusion. Moreover, women and individuals with more education were more likely to agree to the collection of genetic data.

Our main analysis sample includes all genetically European financial respondents born before 1965 with non-missing genetic, education, and household wealth data.[11] We restrict the sample to European-Americans because the polygenic score we use was discovered in a sample of all genetic Europeans. We further restrict our sample to include only retired households in years 1996, 1998, and 2002-2010.[12] We limit our analysis to retired households because we have not yet included the *expected* defined benefit pension income for households still in the workforce.[13] We address defined benefit pension income for retired households by assuming that retired households that report receiving pension income are the only households that participated in defined benefit plans. This assumption would be incorrect if households report annuity or other income as pension income, or if households have delayed receiving benefits until a future date.

Our data restrictions are chosen to balance concerns about measurement error in wealth with concerns about selection biases that arise if we drop too many observations from our analysis. Our main analytic sample consists of observations where wealth information is not ambiguous. In robustness checks not included in this version of this paper, we demonstrate that our results are robust to (and in many cases stronger) using alternative data samples.[14] Therefore, we are confident that our main results relating ability to wealth are not driven by our choice of subsample.

The resulting analytic sample includes 3,990 financial respondents, with responses supplied for an average of seven waves. Table 1 provides basic descriptives on demographic and educational variables. The mean level of educational attainment is about 13 years, with 14% of the sample failing to graduate from high school or obtain a GED and about 23% of the sample earning at least a four year college degree. Roughly 47% of the sample is male. For financial respondents in our analytic sample, Figure 1 plots the sample (kernel-smoothed) density of the EA Score variable, the genetic index score we use for our analysis. Values of the score have been demeaned and re-scaled to measure standard deviations relative to the mean. Figure 1 suggests that the distribution of the EA Score is approximately normal.

[11] As part of the genetic data release, the HRS also released a file flagging 8,652 individuals as being of European descent based on their genes.

[12] By retired, we mean no member of the household is currently employed.

[13] In principle, expected defined benefit income is knowable in all years between 1992 and 2012 for working respondents, but only for the current job.

[14] Further details on wealth data, reasons for possible mis-measurement and possible alternative subsamples are found in Appendix B.

3.2 Household Wealth

The HRS contains rich and detailed information on household wealth. Unfortunately, data related to household retirement wealth and stock-market participation pose various challenges. Values of defined contribution plans from previous jobs are not asked in every wave; stock allocations in defined contribution plans are only asked in certain waves, and only for plans associated with the current employer; expected defined benefit pension income is also asked only of plans at the current employer. In some cases, such issues may be relatively unimportant. However, because this paper studies heterogeneity in wealth for elderly households, having a complete picture of the households' retirement assets is of fundamental importance. While some data issues have no hope of being resolved, our sample comprises households for whom wealth data is most likely to be both accurate and comprehensive.

Our measure of *total wealth* is designed to encompass all components of household wealth. Our data includes the present value of all pension, annuity, and social security income,[15] which come from the RAND HRS income files, as well as the net value of housing (including primary and secondary residences as well as investment property), the net value of private businesses, all financial assets including cash, checking accounts, savings accounts, CDs, stocks and stock mutual funds, bonds and bond mutual funds, trusts, and other financial assets, less the net value of non-housing debt, each of which we derive from the RAND HRS wealth files. Further, we include the account value of all defined-contribution retirement plans. We exclude from our wealth measure values of transportation and insurance.

We note that our measure of wealth includes both marketable securities, such as stocks which can be easily sold at publicly available prices, and non-marketable assets such as social security income. Our measure of wealth is therefore intended to capture the degree of financial security of the household, rather than the market value of household assets. For robustness, we will show results for household *financial wealth* (total wealth excluding retirement income and housing), which can be interpreted as the market value of households' salable financial assets.

Figure 2 shows the unconditional distribution of wealth for observations in our analytic sample. Notice that the distribution is right-skewed, which is consistent with a relatively small number of individuals who report very high levels of wealth. Figure 3 shows that the distribution of log wealth is somewhat more normally distributed.[16]

Table 2 shows the 10th, 25th, 50th 75th and 90th percentiles of our various wealth mea-

[15] We follow Yogo (2016) and assume a 1.5% guaranteed rate of return, discounted by the probability of death in each year conditional on age, cohort and gender as determined by the Social Security life tables.

[16] Recall from our discussion in Section 2 that the polygenic score does not reflect variants on sex chromosomes, so its distribution should be identical in men and women.

sures, as well as the average. The first row of Table 2 shows winzorised total wealth, which includes housing and pensions. The average for our sample is about $540,000. However, the median individual has total wealth of roughly $284,000, which is substantially lower. Again, this is due to high levels of wealth among individuals in the upper tail. For example, wealth at the 10th percentile is about $30,000; at the 90th percentile, wealth is a little over $1,230,000.

Rows two through four of Table 2 show wealth excluding the values of housing and pension income. A few interesting patterns emerge. First, housing makes up a larger portion of total wealth at the lower end of the distribution. For example, at the 10th percentile housing wealth is more than half of total wealth, whereas it accounts for only about one-fifth of total wealth at the 90th percentile. A similar result is found for pension wealth. In fact, for individuals at the 10th percentile, housing and pensions comprise the entirety of household wealth.

Table 3 shows the median, 75th and 90th percentiles of the individual components of total wealth. At the median, the table confirms that nearly all wealth is in the form of pensions and housing. The 75th percentile includes other sources of wealth, including IRA's, stock holdings, cash and CDs. At the 90th percentile, wealth is further diversified, including items such as secondary homes and real estate.

4 Genes, Wealth and Financial Decisions

In this section, we study the relationship between genetic labor market ability (as measured by the polygenic score), wealth and financial decisions. In Section 4.1 we establish that the relationship between the EA Score and household financial wealth is substantial. Controlling for factors such as education, bequests and household income reduces the size of the gradient by about half. In Section 4.2, we assess whether stock market participation can further explain the relationship between labor market genetic ability and wealth. Controlling for stock market participation accounts for an additional third of the size of the genetic gradient in wealth. In Section 4.3, we demonstrate that main results are robust when we incorporate spousal ability. Taken together, our results suggest a robust and economically meaningful relationship between household genetic ability and financial wealth. However, this link does not solely operate mechanically through higher earnings, but also appears to reflect how individuals make financial decisions. In Section 5 we examine preferences and expectations — individual characteristics that affect financial decision-making.

4.1 The Polygenic Score, Financial Wealth and Earnings

Figure 4 presents the stylized fact that motivates this paper. The nonparametric, unconditional relationship between the polygenic score of the financial respondent and household wealth, shown by a Lowess plot, is both positive and economically substantial. In Table 4 we regress wealth on the score and various sets of control variables. Unless otherwise indicated, all specifications throughout the paper include the following basic controls: the first ten principle components of the genetic data, a full set of birth year dummies, a full set of age dummies, a full set of calendar year dummies, a male dummy, and interactions between the male and age dummies and the male and birth year dummies. Panel A presents results for the log of wealth, and Panel B presents results for wealth in levels. The results in Column (1) indicate that a one standard deviation increase in genetic ability is associated with 24% higher total wealth (or about $117,000 in levels).

Since the score measures genetic endowments that predict educational attainment, it is natural to examine how much of this gradient can be explained by education. Column (2) adds controls for the financial respondent's years of schooling and degree, which reduces the coefficient by more than half; observed educational investments unsurprisingly play a large role in mediating this relationship. In Column (3), adding controls for parental education barely reduces the gradient, as might be expected given the strong inter-generational persistence in education. Regardless of education controls, the coefficients in the third column show that a strong and significant association remains. After controlling for own and parental education, a one standard deviation increase in the EA Score is associated with a 10% increase in wealth (or about $54,000).

We next explore mechanisms that might explain the strong positive association between the polygenic score and wealth, even after controlling for education. Since the score is directly related to human capital accumulation, perhaps the most obvious channel is performance in the labor market. Indeed, Papageorge and Thom (2016) demonstrate that the score is associated with higher wages after controlling for education and family background, though only for those with a college degree. Column (1) of Table 5 presents the association between the score and log wealth, now restricting to the sample with non-missing data on income and the other mechanism variables considered here (coeff. $= 0.12$).[17] In Column (2) of Table 5, we additionally control for the average log household income (averaged over years with non-missing household labor income). Though household income predicts higher household

[17]Note that this coefficient is slightly higher than the coefficient delivered by the specification in Column (3) of Table 4, which used exactly the same control set. This is due to the restriction in the sample to those with non-missing values of the income, inheritance, and stock market participation variables examined in this section.

wealth as expected, the coefficient on the EA Score remains nearly the same. This is not terribly surprising, since Papageorge and Thom (2016) find a modest average relationship between the score and personal income, with substantial gradients only appearing for the college educated evident among more recent birth cohorts.

Since individuals receive their genes from their parents, there will necessarily exist a high degree of correlation between their genetic score and the scores of their parents. High ability individuals with more education will tend to have more successful, high ability parents. This naturally suggests bequests and inheritance as a possible mechanism linking the polygenic score to household wealth. In Column (3) of 5, we add two separate controls for inheritances. First, we add the log of cumulative value of inheritances received in the current year (plus one). Second, we add a binary indicator for whether the individual ever receives inheritances in the HRS data. Adding these inheritance variables reduces coefficient on the EA Score only modestly from 0.12 to 0.11.

Together, our main findings from this section indicate that about half of the genetic gradient in wealth can be explained by education and earnings. In other words, the relationship between labor market ability and wealth is to some degree a mechanical consequence of labor market outcomes: people with higher earnings and access to bequests accumulate larger amounts of wealth. Nevertheless, much remains to be explained. We now ask if we can further explain the genetic gradient in wealth by considering financial decisions, in particular, stock market participation.

4.2 Examining the Role of Stock Market Participation

Besides influencing the earnings of an individual (and their parents), genetic ability could affect household wealth by altering the decisions that households make with their saved earnings and other financial resources. One of the most prominent financial decisions facing a household is how to allocate their wealth across different asset classes. Indeed, an enormous literature explores the portfolio choice problem and stock market participation in particular (Van Rooij, Lusardi, and Alessie, 2011). Returning to Table 5, in Column (4) we add an indicator variable for any stock market participation as a regressor in our specification. Accounting for stock market participation reduces the coefficient on the EA Score by about one-third, from 0.11 to 0.07. This suggests that portfolio choice decisions, and stock market participation in particular, may represent a critical channel linking genetic ability to household wealth.

To learn more about the relationship between genetic ability and stock market participation, Table 6 estimates specifications in which the dependent variable is a dummy for

whether or not the financial respondent's household owns any stocks. Column (1) indicates that a one standard deviation increase in the EA Score is associated with a 3.7 percentage point increase in the probability of stock market participation. Two explanations could rationalize this pattern. First, this could simply reflect the fact that individuals with higher polygenic scores are wealthier, and wealthier people tend to invest more of their wealth in risky assets such as stocks. Alternately, this genetic gradient could reflect differences in the portfolio choices that people make for a given level of wealth. In Column (2) 6, we add the lagged (last wave) values of the log of household wealth as a control variable. This is indeed strongly associated with stock market participation, and its inclusion reduces the coefficient on the EA Score to 0.028, cutting the association by about one fourth. When the average log household income is included in Column (3), the coefficient remains the same. Our conclusion is that the genetic score appears to be associated with stock market participation, even controlling for current wealth and the average of past wages.

In summary, our results up to this point demonstrate that genetic ability endowments predict wealth. A strong relationship remains even after controlling for a host of variables that capture direct channels related to income and bequests. Stock market participation appears to play a substantial mediating role in the genetic gradient, even after controlling for household wealth. This suggests financial decision-making likely plays a role. That is, the genetic gradient in wealth appears to be driven not only by the resources available to individuals due to higher earnings, but also how households manage their resources.

4.3 Non-Financial Respondent Ability and Household Wealth

Before turning to a more detailed discussion of labor market ability and financial decision-making, we briefly assess the robustness of our results if we incorporate the ability of other members of the household. Our analysis until now has only used the ability of the "financial respondent" (FR). Doing so ignores the possibility that household wealth may also be a function of the ability of the financial respondent's spouse, deemed the "non-financial respondent" (NFR). We demonstrate that NRF ability is predictive of financial wealth even after controlling for FR ability. However, we also demonstrate that our results are qualitatively similar if we instead consider maximum household ability.

In Panel A of Table 7, we show that the scores of both the NFR and FR independently predict financial wealth if we include both in a regression with log wealth as the outcome variable. In Column (1), we restrict the sample to the set of households where both FR and NFR EA Scores are available and we regress log wealth (plus 1) on the EA Score of the FR along with the standard set of controls and obtain a coefficient of 0.097. In Column (2),

if we include EA Score of the NFR, we find that the coefficient on the FR EA Score falls modestly to 0.092 and that the coefficient on NFR EA Score is 0.078. In Columns (3) and (4), we repeat the exercise with stock market participation as the outcome variable. If we do not include NFR ability, the coefficient on FR ability is 0.022. This falls to 0.019 once we control for NFR EA Score. Surprisingly, the coefficient on NFR labor market ability is larger, estimated at 0.034.

Together, the results in Panel A of Table 7 provide support for the idea that basic patterns we have shown until now remain intact if we include NFR EA Score as an additional variable. However, these results also seem to suggest that the FR EA Score may not sufficiently capture the ability endowments that are relevant for household wealth outcomes. That is, there may be alternative ways to incorporate the labor market ability of both spouses. This naturally raises the question of how the ability endowments of each spouse combine to produce joint household outcomes. The endowments of the FR and the NFR might substitute for one another if an individual's high ability can compensate for the low ability of their spouse. Alternately, spousal abilities might be complementary if high and low ability spouses have to reach compromise positions on financial decisions.

We defer a full analysis of this topic for future work. However, Panel B of Table 7 provides some suggestive evidence that a high ability spouse can compensate for the low ability of a financial respondent. Here we consider as a regressor not the FR EA Score, but the maximum of the FR and NFR scores. In Column (1), we replicate our baseline wealth regression, restricting to households with two genotyped spouses (EA Score coeff = 0.097). In Col (2) we instead measure household ability with the maximum score variable, and the estimated genetic gradient increases by over one third (coeff = 0.147). Similar results hold in Columns (3)-(4), where we repeat this exercise for stock market participation.

5 Genes, Financial Literacy and Expectations

The preceding results suggest that i) there is a substantial relationship between the EA Score and household wealth, ii) this is not entirely or even primarily related to higher income, and iii) this appears to be at least partially mediated by stock market participation, which is a financial choice variable. This leads us to investigate whether individuals with different levels of the polygenic score differ in terms of how they think about financial decisions. This section is divided into two main parts. In Section 5.1, we establish that the polygenic score does indeed predict differences in how people think about financial-related factors. We first document a relationship between genetic ability and financial literacy. Next, we examine genetic ability and subjective beliefs about macroeconomic outcomes. Individuals with higher

values of the EA score report subjective expectations that are closer to objective probabilities, and are less likely to report "extreme" beliefs that take some outcomes as certainties. The second part of this section (Section 5.2) demonstrates that these self-reported beliefs, though not incentivized, are indeed linked to wealth and financial outcomes. Section 5.3 documents that reported beliefs matter in that they predict financial behavior. For example, individuals with lower genetic scores are more likely to predict excessively high probabilities of low stock market returns and are subsequently more likely to be observed avoiding the stock market.

5.1 The Polygenic Score and Financial Literacy

An obvious theoretical starting point is financial literacy (Delavande, Rohwedder, and Willis, 2008). Do individuals with higher EA Scores exhibit greater sophistication in their understanding of financial choices? Fortunately, the HRS data contain a number of questions that directly assess an individual's financial literacy. Unfortunately, these are asked in a small module in the 2010 wave, which leaves us with less than 700 genotyped respondents for these questions. The 2010 module asks three basic financial literacy questions:

- **Compounding Interest:** "First, suppose you had $100 in a savings account and the interest rate was 2% per year. After 5 years, how much do you think you would have in the account if you left the money to grow – more than $102, exactly $102, or less than $102?"
- **Real Interest Rate:** "Imagine that the interest rate on your savings account was 1% per year and inflation was 2% per year. After 1 year, would you be able to buy more than today, exactly the same as today, or less than today with the money in this account?"
- **Diversify Stocks:** "Do you think that the following statement is true or false: buying a single company stock usually provides a safer return than a stock mutual fund?"

Columns (1)-(3) of Table 8 present linear probability models explaining whether respondents correctly answered these questions as function of the EA Score and our standard set of controls.[18] The score is positively related to correctly answering the Real Interest Rate and Diversification questions, but only the coefficient for the Real Interest Rate questions is statistically significant (p-value < 0.05). In Column (4), the dependent variable is an indicator for whether the individual correctly answered all three questions. We find that

[18]Those responding that they "Don't Know" were coded as not responding correctly.

a one standard deviation increase in the EA Score is associated with a 4 percentage point increase in the probability of correctly answering all questions (p-value <0.1). The financial literacy module also asks a separate question on whether creditors or debtors would be helped by inflation. This question is not asked to the individuals who answered the three questions listed above. Column (5) present estimates related to a correct answer on this question. The genetic association for this question is statistically significant and economically substantial: a one standard deviation increase in the score is associated with an 8 percentage point increase in the probability of a correct answer (p-value < 0.01).

Taken together, the results in Table 8 provide some evidence that individuals with higher values of the genetic score tend to be more financially literate. However, we again reiterate that these results should be interpreted cautiously since the financial literacy questions are available only for a small subset of the individuals in our main analysis sample. In particular, we may be underpowered to detect true effects with such relatively small samples.

5.2 The Polygenic Score and Financial Expectations

An important element of financial decision-making is an assessment of the risks and uncertainties associated with the macroeconomy and the payoffs to different possible financial choices. Here we investigate whether the EA Score is associated with differences in the beliefs and expectations about objective macro events that are relevant for financial choices. The HRS data are uniquely well-suited for this analysis, since most respondents are repeatedly asked to provide subjective probabilities on a range of events. Individuals are asked to provide a probability on a scale of 0 to 100, for the following three events:

- **Stock Market Goes Up**: "By next year at this time, what is the percent chance that mutual fund shares invested in blue chip stocks like those in the Dow Jones Industrial Average will be worth more than they are today?"
- **Economic Depression**: "What do you think are the chances that the U.S.economy will experience a major depression sometime during the next 10 years or so?"
- **Double Digit Inflation**: "And how about the chances that the U.S. economy will experience double- digit inflation sometime during the next 10 years or so?"

The panels of Figure 5 present histograms of the responses, pooling all person-year observations across waves. Across all three questions, we see evidence of heaping at focal probabilities. Specifically, there are pronounced spikes at answers of 0, 50, and 100. It is essential to note that respondents are given specific instructions about the meaning of a

response of 0 or 100. That is, they are told to supply these answers if they believe that there is "absolutely no chance" that the event will happen, or if it is "absolutely certain" to happen. All three of these macroeconomic events can be associated with objective probabilities that should be common knowledge in an economy with fully informed agents. Our objective benchmark probability for the stock market going up in a single year is 71 percent, which is the average number of years of positive returns for the S&P 500 over the period 1992-2015. There is no common definition of an economic depression, but clearly this refers to an unusually severe period of economic contraction. We use data from the Federal Reserve Bank of Saint Louis on annual real GDP growth over the period 1948-2016, and define an unusually severe contraction as a year with growth less than or equal to -0.73 percent, which is the 25th percentile of the distribution of growth rates for negative-growth years. Based on this metric, the unconditional probability of a severe contraction is 4.4 percent per year, which implies a 36 percent probability for such an event over a 10 year period. Finally, we note that the Bureau of Labor Statistics reports two years with double digit inflation (1980, 1981) over the period 1958-2015. This implies an approximate probability of 3.4 percent for double digit inflation in any year, or about a 29 percent chance for double digit inflation over a 10 year period.

Table 9 provides estimates of the association between the EA Score and individual beliefs about the probabilities of these macroeconomic events. Each panel presents results on a different expectation (Panel A: probability the stock market goes up, Panel B: probability of a depression, Panel C: probability of double digit inflation). Our first measure related to these expectations variables is the absolute value of the deviation between the respondent's probability and the objective probability. We regress this deviation on our standard controls and the EA Score in Column (1). For all three events, higher values of the polygenic score are associated with a statistically significant reduction in the deviation between the respondent's subjective probability and the objective probability. For example, in Column (1) of Panel A, the coefficient estimate of -0.453 suggests that a one standard deviation increase in the score is associated with a reduction in the deviation of about one half of a percentage point.

Columns (2)-(4) of Table examine binary outcomes indicating whether respondents answered with specific focal probabilities (0, 50, and 100, respectively). For all three events, we observe the same pattern of association: The EA Score is negatively associated with providing a subjective probability indicating complete certainty (0 or 100), and is largely uncorrelated with providing a focal probability of 50 percent. It should be noted that the magnitudes of these associations are substantial. For example, Column (2) of Panel B suggests that a one standard deviation increase in the EA Score is associated with a 0.5 percentage point reduction in the probability of believing there is 0 chance that the economy

will suffer a major depression in the next 10 years. For the sake of comparison, 6.7 percent of individuals responded with a 0 belief for this item.

5.3 Linking Financial Beliefs to Wealth and Financial Decisions

The results of the preceding section suggest that individuals with lower genetic scores appear to report beliefs that are at odds with objective probabilities and, moreover, tend to heap on "focal" beliefs. One possibility is that these reported beliefs are not related to individual behavior in a meaningful way. For example, it may be that individuals with lower polygenic scores simply report numbers that are not reflective of their beliefs. This would suggest that their answers are more prone to measurement error, but do not necessarily imply greater difficulties with financial decisions.

In Table 10, we investigate how beliefs relate to behavior and wealth. In each specification, we control for the EA Score. In Column (1), we show that individuals who ever predict a 0 belief that the stock market will go up are significantly less wealthy by about 22 percentage points. Similar magnitudes are evident for individuals who report that they are certain there will be a recession or double-digit inflation. Interestingly, individuals who report they are certain that the stock market will go down report gains to their wealth of about 26%. In Column (2), we add deviations discussed in the previous section and find that it is not only focal beliefs, but also deviations from the objective probabilities that matter, at least for the stock market and depression events. To investigate these patterns a bit further, Columns (3) and (4) relate beliefs to whether or not the individual owns stocks. Not surprisingly, individuals who are pessimistic are less likely to participate in the stock market, while individuals who are certain that the market will go up are more likely to participate in the stock market. This table suggests that expectations about the economy (include focal beliefs and beliefs that deviate from objective probabilities) predict individual financial decision-making. This is consistent with a world in which beliefs matter because they are directly informing decisions. Notice, in particular, that individuals who are certain that the stock market will go up are substantially more likely to hold stocks (by about 10 percentage points). This indicates that such extreme beliefs are indeed associated with financial choices consistent with these beliefs.

6 Evidence on Policy Implications

In this section, we turn to some policy implications of our results. The results of the preceding sections suggest that individuals with different levels of labor market ability differ in terms

of financial literacy and the accuracy of their beliefs. Since these are important inputs for effective financial decision-making, individuals with lower levels of labor market ability might benefit if at least some of their financial decisions are managed by a third party. Defined benefit pensions (employer based pensions) offer one arrangement in which savings are effectively managed by an individual's employer (or a third party), and an individual receives a guaranteed stream of income without having to make investment decisions over the life-cycle. Of course, pension participation is not necessarily randomly assigned. In table 11, we regress an indicator for defined benefit pension holdings onto the EA Score. We show that after including our standard battery of controls, there is no economically meaningful or statistically significant relationship between the EA Score and pension holdings. In Column (2), we assess the association between pensions and wealth and show that holding a pension is associated with a 37 percent increase in wealth. Finally, in Column (3) of Table 11, we ask if pension holdings moderate the relationship between genes and wealth. In particular, aside from our general set of controls, the EA Score and a dummy for holding a pension, we also include an interaction between the EA Score and the pension dummy. The coefficient on the dummy is negative and significant, suggesting that pensions, which reduce financial decision-making autonomy, seem to moderate the relationship between genetic ability and wealth. These findings suggest that wealth inequality could be mitigated by public pension schemes that reduce autonomy of financial decision-making.

We also ask if high childhood SES mitigates the likelihood that individual report extreme beliefs. Some individuals with disadvantageous ability endowments, but who are born into high-SES environments, may have access to resources and investments that would improve their ability to process information. We consider four retrospective childhood SES measures: whether or not the individual grew up in poverty; average income of father's occupation; whether the family ever moved or asked for health due to financial difficulties and whether the individual's father was ever unemployed for long periods of time.[19] In Table 12, we relate extreme beliefs to "high SES" for each measure. In Panel A, we consider the belief that the start market will go up with 100% probability. High SES (not growing up in poverty) is associated with a lower probability of reporting an extreme belief. If the family moved or asked for help due to financial troubles, the individual is more likely to report an extreme belief. We ask whether there is an interaction between SES and a higher genetic score and find little evidence of such an effect. In other words, there is little evidence that the high SES attenuates the relationship between a low polygenic score and extreme beliefs. We find similar results for extreme beliefs except for beliefs about double digit inflation and the

[19]These four childhood SES measures are discussed at length and used in conjunction with the polygenic score in Papageorge and Thom (2016).

income measures of childhood SES.[20]

The results on education provide some evidence that education many not be a solution to the problem that some ability endowments are associated with difficulties in probabilistic thinking. Of course, much research would need to be done to assess whether this is truly the case. Still, our results provide additional support for the idea that public pension schemes could mitigate inequality due to poor financial decision-making. These types of schemes have come under fire precisely because of the sound argument that individuals, acting on their own, could make better portfolio decisions than a public pension scheme, for example. The opposing argument is that poverty among the elderly beyond being undesirable on ethical grounds also imposes costs onto society. If so, then allowing people to make autonomous financial decisions that lead to poverty creates an externality and it may be welfare enhancing to reduce autonomy.

7 Conclusion

This paper shows that the same genetic endowments that predict educational attainment and earnings are also associated with higher wealth. This could arise purely from an association between ability and earnings, as high earnings will mechanically generate high wealth. We show that controlling for education and earnings does indeed attenuate the genetic gradient in wealth, but only accounts for roughly one half of the association. Stock market participation accounts for an additional one third of the genetic gradient in wealth, which suggests financial decision-making as another mechanism linking labor market ability and wealth. We show that those with a higher polygenic score perform better on standard financial literacy questions, and are less likely to report extreme beliefs about the economy, including the likelihood of stock market appreciation or a severe recession.

We also show evidence that childhood SES does not appear to modify the relationship between the polygenic score and beliefs. This is troubling as it suggests that reallocating resources for educational investments is not an easy solution for overcoming difficulties in financial decision-making. On the other hand, we show that the genetic gradient in wealth is weaker among individuals who have less autonomy in their financial decisions due to participation in traditional pension plans, while participation in pensions is not itself predicted by the polygenic score. Our findings suggest that policies that reduce autonomy in financial decision-making, such as public pension schemes, might play an important role in reducing wealth inequality. This is particularly important given our findings that the same ability en-

[20]Results remain the same under various sets of controls, i.e., whether or not we control for parents' or own education.

dowments that predict low earnings also predict disadvantageous financial decision-making, which could further exacerbate wealth inequality among the elderly.

References

Aizer, Anna and Janet Currie. 2014. "The Intergenerational Transmission of Inequality: Maternal Disadvantage and Health at Birth." *Science* 344 (6186):856–861.

Bach, Stefan, Andreas Thiemann, and Aline Zucco. 2015. "The Top Tail of the Wealth Distribution in Germany, France, Spain, and Greece." DIW Berlin Discussion Paper.

Benhabib, Jess and Alberto Bisin. 2016. "Skewed Wealth Distributions: Theory and Empirics." NBER Working Paper.

Benhabib, Jess, Alberto Bisin, and Shenghao Zhu. 2011. "The Distribution of Wealth and Fiscal Policy in Economies with Finitely Lived Agents." *Econometrica* 79 (1):123–157.

Benjamin, Daniel J, David Cesarini, Christopher F Chabris, Edward L Glaeser, David I Laibson, Vilmundur Guðnason, Tamara B Harris, Lenore J Launer, Shaun Purcell, Albert Vernon Smith et al. 2012. "The Promises and Pitfalls of Genoeconomics." *Annual Review of Economics* 4:627–662.

Bharadwaj, Prashant, Katrine Vellesen Løken, and Christopher Neilson. 2013. "Early Life Health Interventions and Academic Achievement." *American Economic Review* 103 (5):1862–1891.

Bierut, Laura Jean. 2010. "Convergence of Genetic Findings for Nicotine Dependence and Smoking Related Diseases with Chromosome 15q24-25." *Trends in Pharmacological Sciences* 31 (1):46–51.

Camacho, Adriana. 2008. "Stress and Birth Weight: Evidence from Terrorist Attacks." *American Economic Review* :511–515.

Cawley, John, Karen Conneely, James Heckman, and Edward Vytlacil. 1997. "Cognitive Ability, Wages, and Meritocracy." In *Intelligence, Genes, and Success: Scientists Respond to The Bell Curve*, edited by Bernie Devlin, StephenE. Fienberg, Daniel P. Resnick, and Kathryn Roeder. Springer New York, 179–192.

Cesarini, David, Christopher T Dawes, Magnus Johannesson, Paul Lichtenstein, and Björn Wallace. 2009. "Genetic Variation in Preferences for Giving and Risk Taking." *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.

Costa, Dora L. 1998. "Unequal at Birth: A Long-Term Comparison of Income and Birth Weight." *The Journal of Economic History* 58:987–1009.

Cronqvist, Henrik and Stephan Siegel. 2015. "The Origins of Savings Behavior." *Journal of Political Economy* 123 (1):123–169.

Currie, Janet, Matthew Neidell, and Johannes F Schmieder. 2009. "Air Pollution and Infant Health: Lessons from New Jersey." *Journal of Health Economics* 28 (3):688–703.

Currie, Janet and Maya Rossin-Slater. 2013. "Weathering the Storm: Hurricanes and Birth Outcomes." *Journal of Health Economics* 32 (3):487–503.

Delavande, Adeline, Michael Perry, Robert Willis et al. 2006. "Probabilistic Thinking and Early Social Security Claiming." Michigan Retirement Research Center Working Paper.

Delavande, Adeline, Susann Rohwedder, and Robert J Willis. 2008. "Preparation for Retirement, Financial Literacy and Cognitive Resources." Michigan Retirement Research Center Research Paper.

Dominitz, Jeff and Charles F Manski. 2007. "Expected Equity Returns and Portfolio Choice: Evidence from the Health and Retirement Study." *Journal of the European Economic Association* 5 (2-3):369–379.

Fagereng, Andreas, Luigi Guiso, Davide Malacrino, and Luigi Pistaferri. 2016. "Heterogeneity in Returns to Wealth and the Measurement of Wealth Inequality." *American Economic Review* 106 (5):651–655.

Fletcher, Jason M. 2012. "Why Have Tobacco Control Policies Stalled? Using Genetic Moderation to Examine Policy Impacts." *PLoS ONE* 7 (12):1–6.

Fletcher, Jason M and Steven F Lehrer. 2011. "Genetic Lotteries within Families." *Journal of Health Economics* 30 (4):647–659.

Heckman, James J. 1995. "Lessons from the Bell Curve." *Journal of Political Economy* :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.

Hewitt, John K. 2012. "Editorial Policy on Candidate Gene Association and Candidate Gene-by-Environment Interaction Studies of Complex Traits." *Behavior Genetics* 42 (1):1–2.

Hurd, Michael D. 2009. "Subjective Probabilities in Household Surveys." *Annual Review of Economics* 1:543.

Jones, Charles I. 2015. "Pareto and Piketty: The Macroeconomics of Top Income and Wealth Inequality." *The Journal of Economic Perspectives* 29 (1):29–46.

Kézdi, Gábor and Robert J Willis. 2003. "Who Becomes a Stockholder? Expectations, Subjective Uncertainty, and Asset Allocation." Unpublished manuscript, University of Michigan.

———. 2009. "Stock Market Expectations and Portfolio Choice of American Households." Unpublished manuscript, University of Michigan.

———. 2013. "Expectations, Aging and Cognitive Decline." In *Discoveries in the Economics of Aging*. University of Chicago Press, 305–337.

Kindermann, Fabian and Dirk Krueger. 2014. "High Marginal Tax Rates on the Top 1%? Lessons from a Life Cycle Model with Idiosyncratic Income Risk." NBER Working Paper.

Lien, Diana S and William N Evans. 2005. "Estimating the Impact of Large Cigarette Tax Hikes: The Case of Maternal Smoking and Infant Birth Weight." *Journal of Human Resources* 40 (2):373–392.

Lillard, Lee and Robert J Willis. 2001. "Cognition and Wealth: The Importance of Probabilistic Thinking." *Michigan Retirement Research Center Working Paper* .

Okbay, Aysu, Jonathan P Beauchamp, Mark Alan Fontana, James J Lee, Tune H Pers, Cornelius A Rietveld, Patrick Turley, Guo-Bo Chen, Valur Emilsson, S Fleur W Meddens 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. 2016. "Genes, Education and Labor Outcomes: Evidence from the Health and Retirement Study." IZA Discussion Paper 10200.

Poterba, James M and David A Wise. 1998. "Individual Financial Decisions in Retirement Saving Plans and the Provision of Resources for Retirement." In *Privatizing Social Security*. University of Chicago Press, 363–401.

Rietveld, Cornelius A, Sarah E Medland, Jaime Derringer, Jian Yang, Tõnu Esko, Nicolas W Martin, Harm-Jan Westra, Konstantin Shakhbazov, Abdel Abdellaoui, Arpana Agrawal et al. 2013. "GWAS of 126,559 Individuals Identifies Genetic Variants Associated with Educational Attainment." *Science* 340 (6139):1467–1471.

Rohwedder, Susann and Robert J Willis. 2010. "Mental Retirement." *The Journal of Economic Perspectives* 24 (1):119–138.

Saez, Emmanuel and Gabriel Zucman. 2014. "Wealth Inequality in the United States since 1913: Evidence from Capitalized Income Tax Data."

Thompson, Owen. 2014. "Economic Background and Educational Attainment: The Role of Gene-Environment Interactions." *Journal of Human Resources* 49 (2):263–294.

Thorgeirsson, Thorgeir E, Daniel F Gudbjartsson, Ida Surakka, Jacqueline M Vink, Najaf Amin, Frank Geller, Patrick Sulem, Thorunn Rafnar, Tõnu Esko, Stefan Walter et al. 2010. "Sequence Variants at CHRNB3-CHRNA6 and CYP2A6 Affect Smoking Behavior." *Nature Genetics* 42 (5):448–453.

Van Rooij, Maarten, Annamaria Lusardi, and Rob Alessie. 2011. "Financial Literacy and Stock Market Participation." *Journal of Financial Economics* 101 (2):449–472.

Vilhjalmsson, et al., Bjarni J. 2015. "Modeling Linkage Disequilibrium Increases Accuracy of Polygenic Risk Scores." *The American Journal of Human Genetics* 87:576–592.

Vissing-Jørgensen, Annette. 2002. "Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures." NBER Working Paper.

Willis, Robert J. and Sherwin Rosen. 1979. "Education and Self-Selection." *Journal of Political Economy* 87 (5):S7–S36.

Yogo, Motohiro. 2016. "Portfolio choice in retirement: Health risk and the demand for annuities, housing, and risky assets." *Journal of Monetary Economics* 80:17–34.

8 Tables and Figures

Table 1: Summary statistics

Variable	Mean	Std. Dev.	N
Male	0.42	0.49	4433
< 1930	0.32	0.47	4433
1930-1934	0.19	0.39	4433
1935-1939	0.21	0.41	4433
1940-1944	0.15	0.36	4433
1945-1949	0.07	0.26	4433
1950-1954	0.05	0.22	4433
Education (Years)	12.88	2.57	4433
None	0.16	0.37	4429
GED	0.05	0.21	4429
High School	0.53	0.5	4429
College (2 year)	0.04	0.21	4429
College (4 year)	0.13	0.33	4429
Masters	0.07	0.25	4429
Advanced	0.02	0.14	4429
Yrs. of ed: Father	9.76	3.61	3193
Yrs. of ed: Mother	10.26	3.11	3329

Notes: Summary statistics for our main sample of individuals.

Table 2: Wealth Distribution

	p10	p25	p50	p75	p90	Mean	St Dev
Wealth (Winz)	29,550	105,089	284,176	631,047	1,235,471	540,640	826,944
Wealth (No Housing)	12,861	45,000	146,686	413,621	919,036	411,180	1,355,300
Wealth (No Pensions)	1,212	66,144	224,520	549,822	1,124,892	495,831	1,083,248
Wealth (No H or P)	0	5,469	85,422	327,315	810,000	330,375	950,170

Notes: Wealth mean and distribution (10th, 25th, 50th, 75th and 90th percentiles) for total wealth, non-housing wealth, non-pension wealth and wealth that includes neither pensions nor housing.

Table 3: Components of Wealth

	p50	p75	p90	Mean	Med Share	Mean Share
Ret Plans (Employer)	0	0	0	22,343	0 %	4 %
Ret Inc (PV)	33,641	74,736	136,550	58,461	21 %	9 %
Real Estate	0	0	57,717	46,671	0 %	7 %
Business	0	0	0	36,742	0 %	6 %
IRAs	0	55,055	202,557	72,207	0 %	12 %
Stocks	0	28,000	214,708	89,883	0 %	14 %
Cash Equiv	8,102	30,000	81,023	33,169	5 %	5 %
CDs	0	5,408	60,000	22,980	0 %	4 %
Bonds	0	0	0	14,667	0 %	2 %
Other Assets	0	0	16,224	14,382	0 %	2 %
Other Debts	0	0	5,000	2,706	0 %	0 %
Trusts	0	0	0	2,380	0 %	0 %
Home Value	118,131	208,100	354,475	164,562	74 %	26 %
Mortgage	0	0	58,742	16,499	0 %	3 %
Home Loan	0	0	0	2,315	0 %	0 %
Second Home	0	0	27,704	21,964	0 %	4 %
Second Morgt.	0	0	0	1,450	0 %	0 %

Notes: Different sources of wealth (mean and distribution, including the 50th, 75th and 90th percentiles). Columns 5 and 6 are median and mean share, respectively, of each component in total wealth.

Table 4: The Polygenic Score and Wealth

	Log Wealth	**Log Wealth**	**Log Wealth**
EA2Score	0.235***	0.103***	0.099***
	(0.022)	(0.022)	(0.022)
Resp Education	No	Yes	Yes
Parental Education	No	No	Yes
Obs.	15202	15202	15202
R^2	0.151	0.255	0.259
	Wealth (level)	**Wealth (level)**	**Wealth (level)**
EA2Score	117,301***	57,475***	54,520***
	(12,938)	(12,260)	(12,290)
Resp Education	No	Yes	Yes
Parental Education	No	No	Yes
Obs.	15517	15517	15517
R^2	0.089	0.169	0.172

Notes: Relating the genetic score to wealth (logs and levels) with varying sets of controls.

Table 5: Accounting for Income, Inheritances, and Stocks

	Total Wealth	**Total Wealth**	**Total Wealth**	**Total Wealth**
EA Score	0.120***	0.121***	0.108***	0.069***
	(0.027)	(0.026)	(0.026)	(0.023)
Avg log HH Inc		0.347***	0.331***	0.257***
		(0.038)	(0.038)	(0.032)
Log Run Sum Inher			0.018***	0.014***
			(0.005)	(0.004)
Ever Rec Inher			0.247***	0.152***
			(0.059)	(0.051)
Owns Stocks				1.043***
				(0.041)
Obs.	6943	6943	6943	6943
R^2	0.291	0.330	0.346	0.449

Notes: Relating the genetic score to wealth controlling for household income, bequests and stock ownership.

Table 6: Polygenic Score and Stock Ownership

	Own Stocks	**Own Stocks**	**Own Stocks**
EA2Score	0.037***	0.027***	0.028***
	(0.009)	(0.010)	(0.010)
Lag of log Wealth		0.156***	0.152***
		(0.008)	(0.009)
Avg log HH Inc			0.025**
			(0.012)
Obs.	7893	4938	4938
R^2	0.221	0.374	0.376

Notes: Relating the genetic score to stock ownership.

Table 7: Non-Financial Respondent Score and Financial Wealth

Panel A: Spouse's Score				
EA Score	0.097***	0.092***	0.022*	0.019*
	(0.028)	(0.028)	(0.011)	(0.011)
EA Spouse		0.078***		0.034***
		(0.026)		(0.011)
Obs.	5166	5166	5182	5182
R^2	0.314	0.318	0.228	0.232
Panel B: Max Score				
EA Score	0.097***		0.022*	
	(0.028)		(0.011)	
Max EA Score		0.136***		0.042***
		(0.032)		(0.014)
Obs.	5166	5166	5182	5182
R^2	0.314	0.316	0.228	0.231

Notes: Relating the genetic score to wealth with and without NFR score.

Table 8: *EA Score* and Financial Literacy

	(1)	(2)	(3)	(4)	(5)
	Compound Interest	**Real Interest**	**Diversify**	**All Correct (1)-(3)**	**Inflation and Lending**
EA2Score	-0.001	0.032**	0.027	0.040*	0.080***
	(0.018)	(0.016)	(0.021)	(0.021)	(0.020)
Obs.	666	667	667	663	674
R^2	0.245	0.203	0.239	0.294	0.270

Notes: Relating genetic ability to financial literacy.

Table 9: *EA Score* and Beliefs

	Dev. from Objective	**0% Prob**	**50% Prob**	**100% Prob**
Market Up				
EA2Score	-0.453***	-0.004***	-0.002	-0.003**
	(0.141)	(0.001)	(0.003)	(0.001)
Obs.	39626	39626	39626	39626
R^2	0.071	0.031	0.011	0.021
U.S. Depression				
EA2Score	-0.343***	-0.005***	-0.002	-0.004**
	(0.125)	(0.002)	(0.003)	(0.002)
Obs.	32971	32971	32971	32971
R^2	0.064	0.030	0.022	0.044
Double Digit Inf				
EA2Score	-0.612***	-0.007***	0.000	-0.007***
	(0.184)	(0.002)	(0.004)	(0.002)
Obs.	19541	19541	19541	19541
R^2	0.050	0.027	0.023	0.040

Notes: Relating genetic ability to "heaping" on focal beliefs.

Table 10: Beliefs and Financial Wealth

	log Wealth	**log Wealth**	**Owns Stocks**	**Owns Stocks**
Ever Pr Mrkt Up 0%	-0.219***	-0.125**	-0.100***	-0.058***
	(0.055)	(0.064)	(0.016)	(0.020)
Ever Pr Mrkt Up 100%	0.261***	0.261***	0.100***	0.098***
	(0.063)	(0.063)	(0.021)	(0.021)
Ever Pr Rec 0%	0.008	0.041	-0.018	-0.012
	(0.055)	(0.056)	(0.017)	(0.018)
Ever Pr Rec 100%	-0.250***	0.002	-0.057***	0.002
	(0.059)	(0.079)	(0.018)	(0.025)
Ever $>$ 10% Inf 0%	-0.077	-0.082	0.002	0.002
	(0.067)	(0.067)	(0.021)	(0.021)
Ever $>$ 10% Inf 100%	-0.199***	-0.131	-0.088***	-0.067**
	(0.068)	(0.092)	(0.020)	(0.028)
MaxDev PMUP		-0.003**		-0.002***
		(0.001)		(0.000)
MaxDev Prec		-0.008***		-0.002***
		(0.002)		(0.001)
MaxDev DD Inf		-0.001		-0.000
		(0.002)		(0.001)
EA2Score	0.088***	0.086***	0.028***	0.028***
	(0.022)	(0.022)	(0.007)	(0.007)
Obs.	14045	14045	14510	14510
R^2	0.269	0.274	0.204	0.208

Notes: Relating "heaping" on focal beliefs to wealth and stock ownership.

Table 11: Pensions and Financial Wealth

	Has Pension	**log Wealth**	**log Wealth**
EA Score	-0.003	0.100***	0.141***
	(0.006)	(0.022)	(0.033)
DB Pension		0.367***	2.039*
		(0.037)	(1.056)
EA Score x DB			-0.077**
			(0.037)
Obs.	23156	15202	15202
R^2	0.110	0.272	0.276

Notes: Relating pensions to the genetic ability-wealth gradient.

Table 12: Beliefs and Childhood SES

	Not Poverty	**Income**	**Move or Help**	**Father Unemployed**
Panel A: Stock Market Up				
High SES	-13.868*	3.028	16.837**	5.196
	(7.621)	(8.036)	(8.059)	(7.769)
EA2 × High SES	-0.305	-0.381	-0.352	-0.387
	(0.323)	(0.320)	(0.324)	(0.321)
Obs.	39035	30533	38987	39145
R^2	0.059	0.057	0.059	0.057
Panel B: Recession				
High SES	-4.048	-0.810	8.171	-0.081
	(6.594)	(6.868)	(7.137)	(6.854)
EA2 × High SES	0.179	-0.092	0.134	-0.031
	(0.277)	(0.273)	(0.282)	(0.282)
Obs.	32509	26164	32419	32569
R^2	0.060	0.060	0.061	0.059
Panel C: Double Digit Inflation				
High SES	-12.452	15.732	-3.961	-1.200
	(9.971)	(9.721)	(11.293)	(10.121)
EA2 × High SES	-0.508	-1.215***	-0.136	-0.391
	(0.416)	(0.375)	(0.431)	(0.413)
Obs.	19290	17454	19254	19312
R^2	0.043	0.045	0.043	0.042

Notes: Relating beliefs to different measures of high childhood SES.

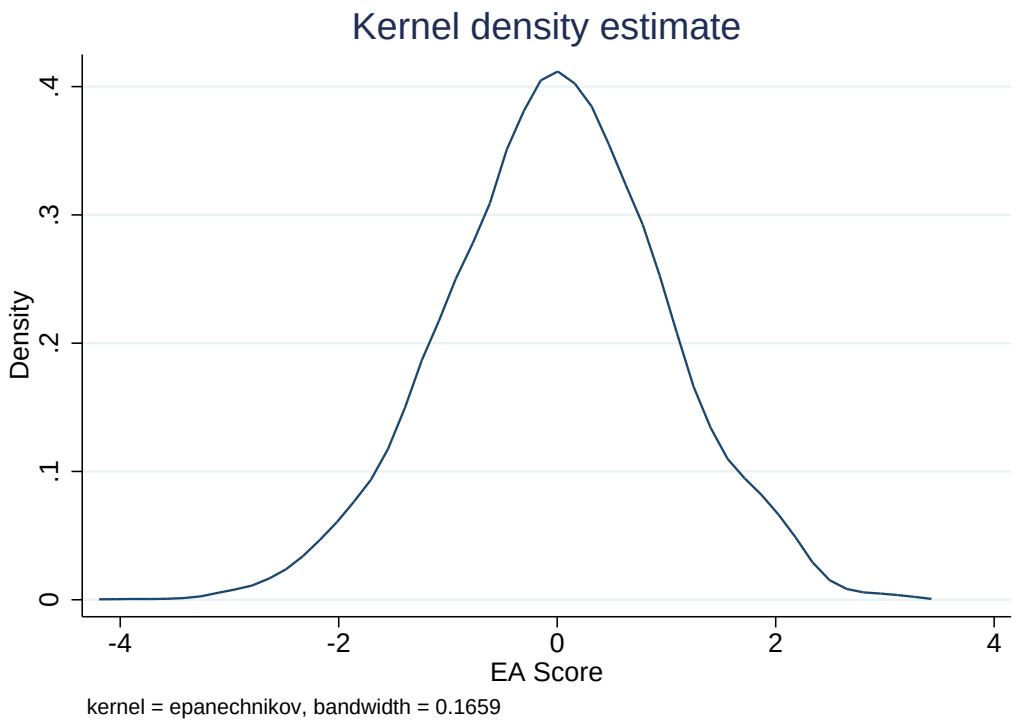


Figure 1: *Notes*: EA2 Score Distribution among HRS Individuals.

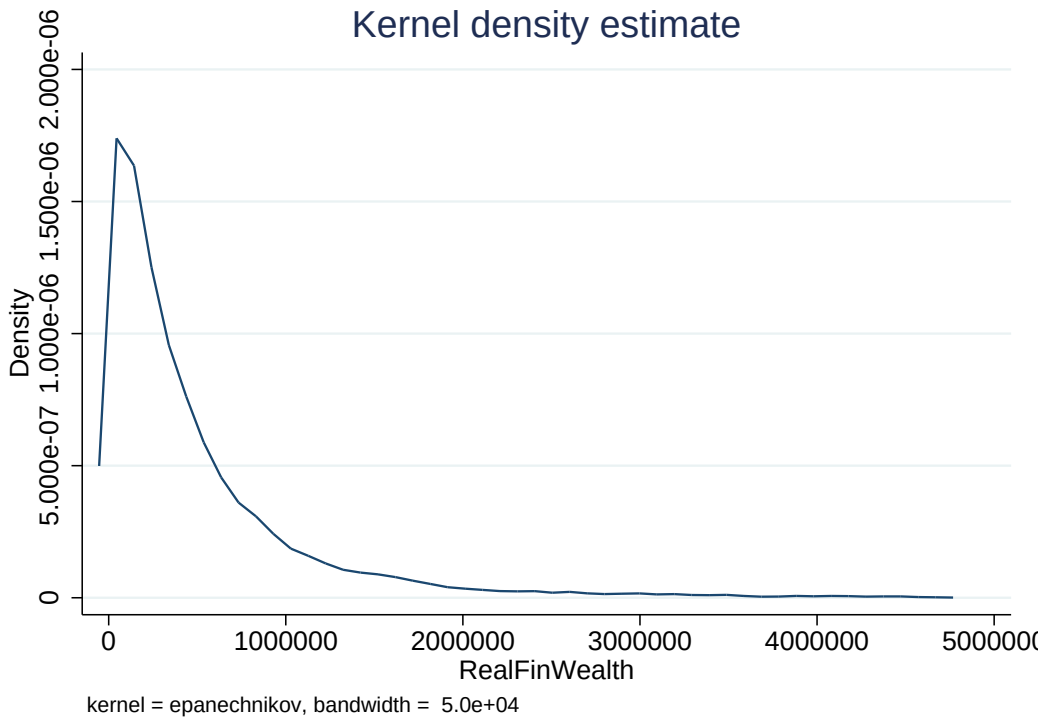


Figure 2: *Notes*: Wealth Distribution among HRS Individuals.

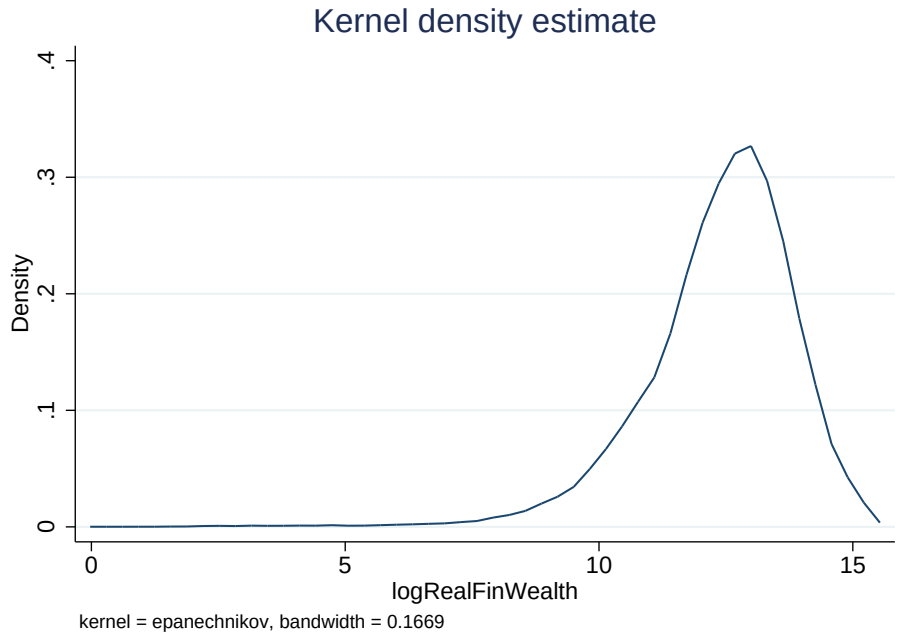


Figure 3: *Notes*: Log Wealth Distribution among HRS Individuals.

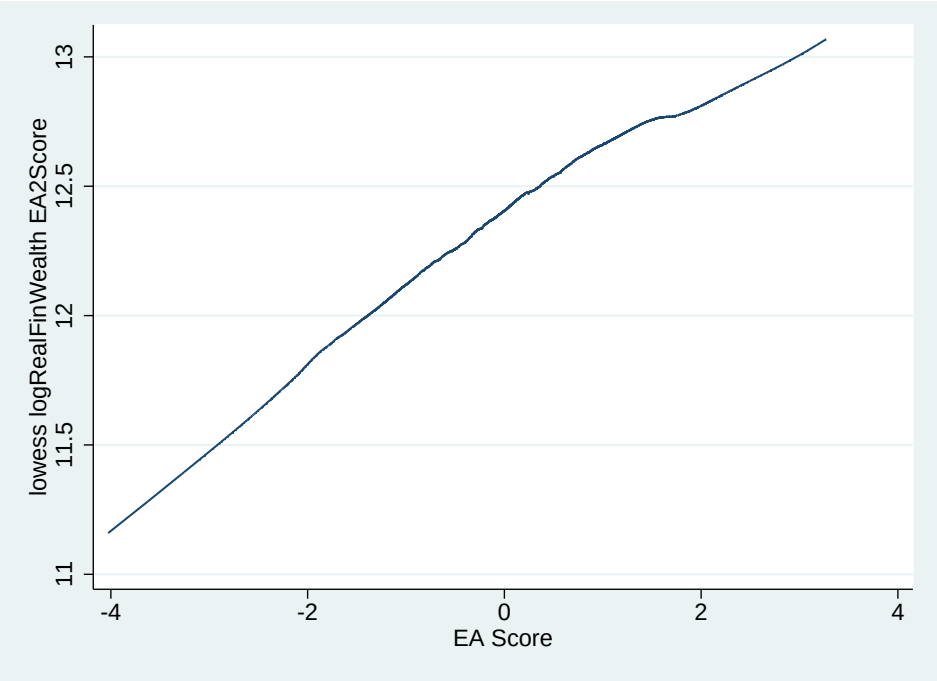


Figure 4: *Notes*: EA2 Score and Wealth among HRS Individuals.

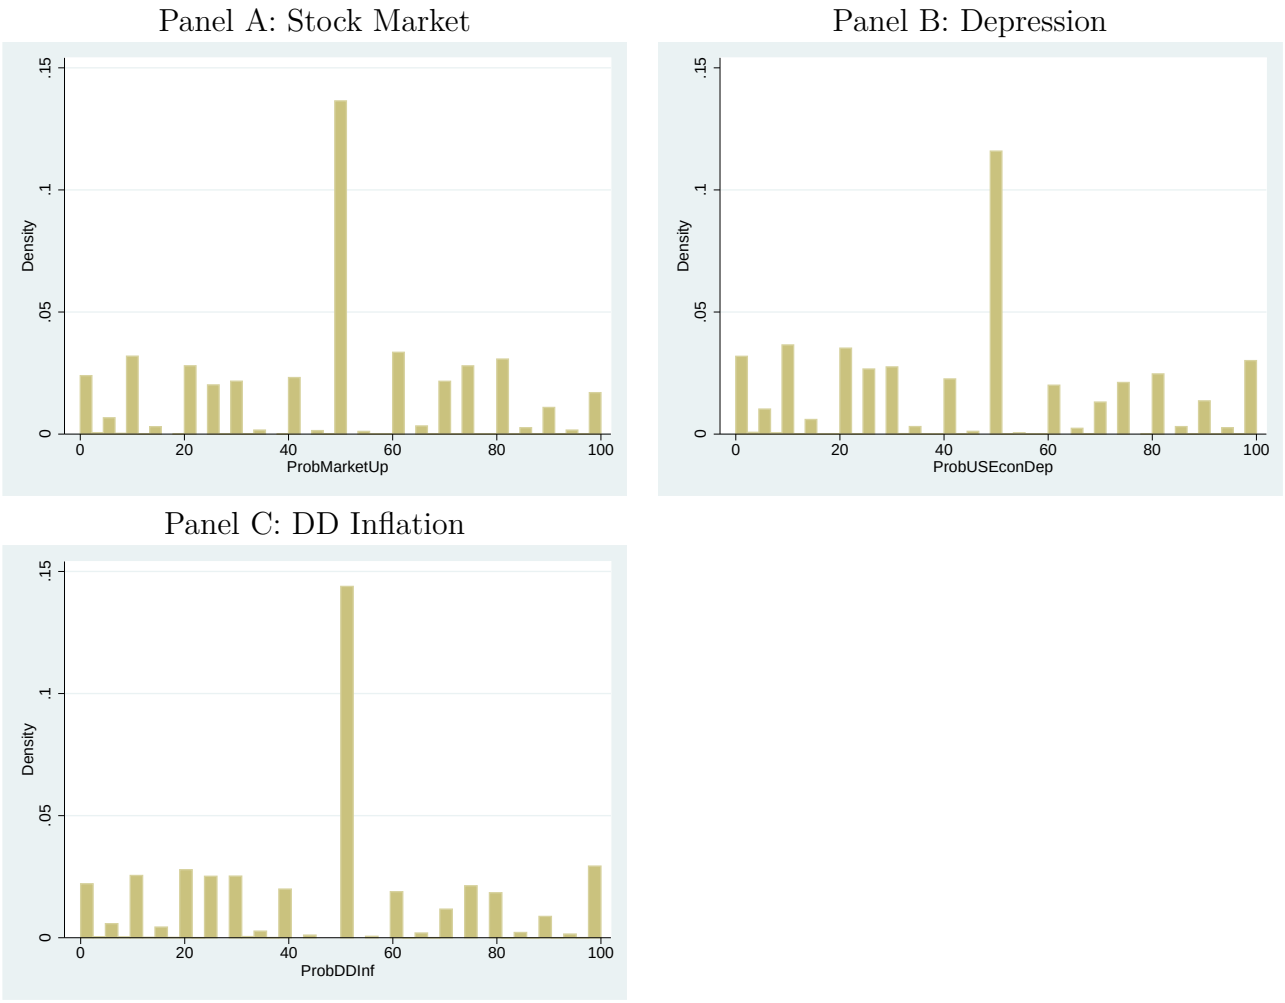


Figure 5: *Notes*: Distributions of Subjective Beliefs.

Online Appendix:

"Genes, Expectations and Financial Wealth:
Evidence from the Health and Retirement Study"

By: Danny Barth, Nicholas W. Papageorge and Kevin Thom

A Additional Details on GWAS and Construction of the EA2 Score

What follows is nearly identical to the genetic data appendix in an earlier paper, Papageorge and Thom (2016), which uses the same genetic score used in the current paper. We reprint the appendix here solely for the reader's convenience In this appendix, we provide a brief introduction to molecular genetics and the kinds of genetic data that we use in this study. We repeat some portions of Section 2 so that this appendix can provide a self-contained introduction to GWAS and the EA2 Score used in our analysis. First, we describe some basic features of the human genome. Next, we discuss how statistical gene-discovery projects can produce scores that are useful for the prediction of economic outcomes such as educational attainment. We highlight how recent advances permit credible and replicable inference.

The human genome consists of approximately 3 billion nucleotide base pairs spread out over 23 chromosomes. Each individual possesses two copies of each chromosome, one from each parent. A gene is a subsequence of base pairs within a chromosome. On average, each gene is made up of over 100,000 base pairs. Each base pair can either be an adenine-thymine (AT) pair, or a guanine-cytosine (GC) pair. Thus, the human genome can be thought of as a series of 3 billion genetic addresses, each of which can contain one of two nucleotide pairs.

A particular location in the genome can be referred to by a name (e.g. rs7937), which indicates its position in the genome. At the vast majority of such locations (about 99%), there is no variation in the observed nucleotide pair, both across humans and across chromosomes within a human. A single-nucleotide polymorphism (SNP) exists when there are differences in the nucleotide pair present at a particular location on the genome. An allele refers to one of the variants that may be present at a particular SNP. If AT is more commonly found at a particular SNP, it is referred to as the major allele, and then GC is referred to as the minor allele.

A traditional approach to the discovery of gene-behavior associations rests on examining *candidate genes*. Under this paradigm, researchers use some knowledge of the relevant biological processes to suggest places in the genome that might contain SNPs associated with a particular outcome. Unfortunately, this approach to identifying gene-economic outcomes has

also generated a large number of reported associations that have failed to replicate outside of their discovery samples. This problem has been so widespread that an editorial statement from the journal *Behavior Genetics* stated that "The literature on candidate gene associations is full of reports that have not stood up to rigorous replication," and that "it now seems likely that many of the published findings of the last decade are wrong or misleading and have not contributed to real advances in knowledge," (Hewitt, 2012). This pattern has emerged, in part, because traditional candidate gene studies have been severely underpowered to detect real genetic effects. Sample sizes in general have been too small relative to the true effect sizes of individual SNPs, making it likely that statistically significant associations are the result of chance. This problem is exacerbated when studies search over many candidate genes, creating a multiple hypothesis testing problem that increases the likelihood of finding false positive results (Benjamin et al., 2012).

An alternative to candidate genes is an approach called a genome-wide association study (GWAS). Under the GWAS methodology, researchers scan the entire genome for SNPs that are associated with a particular phenotype (trait or outcome), but adopt strong measures to deal with multiple hypothesis-testing. For a particular outcome of interest, y_i, and for a set of observed SNPs, $\{SNP_{ij}\}_{j=1}^{N^J}$, a GWAS study proceeds by obtaining estimates of N^J separate regressions of the form:

$$y_i = \mu X_i' + \beta_j SNP_{ij} + \epsilon_{ij} \tag{1}$$

Here SNP_{ij} measures the number of copies of a reference allele possessed by individual i for SNP j. For example, if the reference allele at SNP j is AT, then SNP_{ij} could take the values 0, 1, or 2. The maximum value of 2 reflects the fact that an individual can have at most two copies of the reference allele - one on each inherited chromatid. Additionally, X_i is a vector of controls, including principal components of the genetic variables $\{SNP_{ij}\}_{j=1}^{N^J}$. Principal components of the genetic data are added to control for population stratification. For example, it could be that SNP_{ij} is correlated with a particular ethnicity or ancestry group. Failure to control for the principal components could generate observed SNP-phenotype relationships that reflect the influence of broader ethnic differences rather than the influence of a particular genetic marker.

After obtaining estimates for all N^J versions of equation (1), those estimated coefficients $\widehat{\beta}_j$ with sufficiently small p-values are said to reflect relationships that are genome-wide significant. Given the huge number of regressions run under this methodology, the significance thresholds in modern GWAS are typically very strict. A conventional threshold is 5×10^{-8}. This approach has become popular and as a consequence of its stringency requirements,

has led to the discovery of a number of credible genetic associations. For example, the well-known FTO gene for obesity was discovered through a GWAS, despite the lack of any existing biology that would have suggested it as a candidate gene (Benjamin et al., 2012).

Existing work has demonstrated the importance of credibly identified SNPs for several economic outcomes. These SNPs either directly emerged from a GWAS, or were candidate genes that were validated by later GWAS results. An established literature documents a number of credible genetic associations with smoking behaviors (Bierut, 2010; Thorgeirsson et al., 2010). Fletcher (2012) demonstrates that a SNP associated with smoking intensity also appears to moderate the effect of tobacco taxes. More closely related to our work, another set of studies suggests indirect linkages between genetic variants and human capital. For example, Fletcher and Lehrer (2011) use a set of SNPs associated with health outcomes to provide exogenous within-family variation to estimate a causal relationship between health and education. Finally, Thompson (2014) shows that a variant associated with the MAOA gene appears to moderate the relationship between income and education.

Recent work using GWAS has discovered some of the first direct associations between specific SNPs and education. Rietveld et al. (2013) identified three SNPs (rs9320913, rs11584700, rs4851266) attaining genome-wide significance in a GWAS for educational attainment. Follow-up work by the same team (the Social Science and Genetics Consortium) has recently extended the Rietveld et al. (2013) study to perform an educational attainment GWAS with a sample size of 293,723. This follow-up study, Okbay et al. (2016), has discovered 74 SNPs that attain genome-wide significance. We build our analysis here on the gene-education associations found in this follow-up study.

One common technique adopted in the GWAS literature is to take observed SNPs and the estimated GWAS coefficients (the $\widehat{\beta}_j$) and aggregate them into a polygenic score that can be used for prediction. Typically these scores take the following form:

$$PGS_i = \sum_j \widetilde{\beta}_j SNP_{ij} \tag{2}$$

where $\widetilde{\beta}_j$ is some transformation of the underlying GWAS coefficients. The $\widehat{\beta}_j$ estimates are typically corrected to account for correlation between SNPs and prevent over or under prediction. The follow up study Okbay et al. (2016) combines all genotyped SNPs into a polygenic score that attains a predictive power of up to 3.85% of the variation in educational attainment.[21] In our study, we use SNP weights $\widehat{\beta}_j$ that have been adjusted using a technique

[21]We note as well that the polygenic score that we use in this study combines all SNPs analyzed in Okbay et al. (2016), not just those reaching genome-wide significance. As noted in Okbay et al. (2016), this maximizes the predictive power out of sample.

called LD Pred (Vilhjalmsson, 2015), and applied to the genetic data in the HRS.[22] We refer to the polygenic score created using these weights as the *EA Score*, where "EA" stands for "educational attainment". We refer to it this way since other polygenic scores exist, which capture genetic variation explaining different outcomes.

B Data Issues

This appendix provides details on the construction of our wealth data and our measurement of stock-market participation. Our data are largely constructed from the RAND wealth and income files. The RAND files are carefully cleaned and consistently coded by RAND Corportation and are available for public use. The RAND files have been used in both academic and industry publications, and ensure comparability and consistency across HRS waves and research projects. We refer the reader to the RAND codebook and documentation for further details.

One important shortcoming of the RAND wealth files is the exclusion of employer-sponsored retirement plan account balances. While the RAND wealth files do include the balances of IRAs and other non-employer-sponsored plans, wealth accumulated in employer-sponsored 401k, 403(b), and other such accounts are not included. For households at or near retirement, such accounts can be a significant source of wealth. Further, such accounts may be the only vehicles through which households invest in the stock market, and measures of stock-market participation will understate true participation if these plans are not considered.

Unfortunately, data on employer-sponsored retirement plans are not asked in every wave, and are sometimes inconsistently coded across waves. The remainder of this section focuses on our methodology for coding retirement account balances and stock-market participation inferred from those accounts.

B.1 Wealth in Retirement Accounts

Broadly speaking, there are two types of retirement plans: defined benefit plans, such as traditional pensions (which the HRS calls type A plans), and defined contribution plans, such as 401k and 403(b) plans (which the HRS calls type B plans). We discuss each type of plan in turn.

[22] We would like to especially thank Aysu Okbay, a member of the Social Science and Genetics Consortium, for graciously generating and sharing this score with us.

B.1.1 Defined Benefit Plans

To deal with issues arising from type A style retirement plans, our sample includes only households fully in retirement (households in which no member of the household is currently working). We exclude working households because expected benefits from defined-benefit pension plans are likely to be both an important source of financial wealth and noisily measured. For retired households, our assumption is that those who report receiving pension income were included in defined-benefit pension plans at some point during their working lives, and those who do not receive pension income in retirement were not included in such plans. To the extent that households misreport pension income, for example if income from an annuity converted from a 401k plan is reported as pension income, or if households have delayed receiving pension benefits until some future date, our assignment of households participating in type A plans will be biased. Further, because the household earns a guaranteed stream of income regardless of the underlying investments that support that income (and because we do not observe these underlying investments), we do not consider a household's participation in type A pension plans to be participation in the stock-market.

We include retirement income in our household wealth measure by calculating the price of an actuarily fair annuity based on the entirety of household retirement income, which includes pension income, annuity income, and income from social security. We follow Yogo (2016) by calculating the present discounted value of this income based on a 1.5% annual risk-free rate of return, and discount income in each year by the probability of the recipient surviving until that year.[23] Specifically, we calculate the present value of retirement income, P_t, as:

$$P_t = Y_t \sum_{s=1}^{T-t-1} \frac{\prod_{u=1}^{s} p_{t+u}}{R}, \tag{3}$$

where Y_t is total retirement income, p_t is the recipient's survival probability in period t and is a function of gender, birth cohort, and age, and $R = 1.015$ is the annual risk-free rate of return.

B.1.2 Defined Contribution Plans

Wealth in defined-contribution style plans is a bit trickier. Households may have plans associated with multiple previous employers. To calculate comprehensive measures of wealth and stock market participation, we would like to know both the balances and asset allocations

[23] We differ from Yogo (2016) in that we use the probability of death of the individual receiving the income, rather than of the female spouse.

of all employer-sponsored type B plans from all previous jobs. Unfortunately, this is not always possible.

In years 1996, 1998, and 2002-2010 (comprising even-numbered years), we have the highest quality data on total balances in employer sponsored type B retirement plans.[24] In these years, our wealth data include balances of employer-sponsored plans that are still maintained through that employer, and have not been converted to annuities or rolled-over into IRAs. The HRS refers to such plans as *dormant plans*. Unfortunately, the value of dormant plans at employers prior to retirement are not asked in 1992, 1994, and 2000.

Dormant plans also present problems for measurement of stock market participation. While in years 2002-2010 the stock allocation within a respondent's retirement plan at the current employer is observable for working households, the stock allocation in dormant plans for retired households is not. This means our stock market participation variable does not include stock ownership in dormant plans. In future iterations of this paper, we will develop a dummy variable for whether the household has *ever* owned stocks through an employer-sponsored retirement plan, which will allow us to utilize the information in the current retirement plan section reported in previous waves. As of now, the stock market participation variable is determined only by information in the assets and income section of the data, which comprises only stock and stock mutual funds as well as the stock allocation in IRA and Keogh accounts.

[24] In 2012, the pension data were changed to an entirely new format. We will include this data shortly.