

DISCUSSION PAPER SERIES

IZA DP No. 16633

The Heritability of Economic Preferences

Nathan Kettlewell
Agnieszka Tymula
Hong Il Yoo

NOVEMBER 2023

DISCUSSION PAPER SERIES

IZA DP No. 16633

The Heritability of Economic Preferences

Nathan Kettlewell

University of Technology Sydney and IZA

Agnieszka Tymula

University of Sydney

Hong Il Yoo

Loughborough University

NOVEMBER 2023

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Heritability of Economic Preferences*

We study the heritability of risk, uncertainty, and time preferences using a field experiment with a large sample of adult twins. We also offer a meta-analysis of existing findings. Our field study introduces a novel empirical approach that marries behavioral genetics with structural econometrics. This allows us to, for the first time, quantify the heritability of economic preference parameters directly without employing proxy measures. Our incentive-compatible experiment is the first twin study to elicit all three types of preferences for the same individual. Compared to previous studies, we find a greater role of genes in explaining risk and uncertainty preferences, and of the shared familial environment in explaining time preferences. Time preferences appear more important from policy and parenting perspectives since they exhibit limited genetic variation and are more than twice as sensitive to the familial environment as risk and uncertainty preferences.

JEL Classification: C93, D15, D81, D91, Z13

Keywords: risk preferences, ambiguity aversion, time preferences, twin study, genetics

Corresponding author:

Nathan Kettlewell
University of Technology Sydney
Economics Discipline Group
15 Broadway
Ultimo NSW 2007
Australia
E-mail: Nathan.Kettlewell@uts.edu.au

* We would like to thank Avery Liang for excellent research assistance on the meta-analysis included in the paper. This research was facilitated through access to Twins Research Australia, a national resource supported by the National Health & Medical Research Council (NHMRC) since 1981, administered by the University of Melbourne. We received funding support from the University of Technology Sydney and from the Australian Government through the Australian Research Council's Centre of Excellence for Children and Families over the Life Course (Project ID CE140100027).

1 Introduction

Risk and time preferences are fundamental to economic decision making. A risk lover may decide to start a new business, while a risk averse person chooses a secure job. A person with present bias may make a resolution to join a gym and never fulfill it. A more impatient investor may allocate more of their capital gains to consumption and less to savings. These choices, driven in part by our economic preferences, will have significant ramifications for life outcomes.¹

Our paper addresses a fundamental question on the nature of economic preferences concerning choice under risk and delay discounting. Where do they come from? To what extent are differences between people explained by their environments and their genes? This knowledge can help us to better understand the process through which variation in preferences may lead to inequalities in life outcomes. For example, a major social concern is the intergenerational persistence of socio-economic outcomes like income, wealth, education, occupation and health (see Blanden, 2013; Corak, 2013; Torche, 2015; Halliday et al., 2021), which are closely linked to economic preferences. How, and if, policymakers should intervene to address these inequalities depends on their underlying causes, including the role of genetics (Harden, 2021). Consider, for example, educational attainment, which is a highly heritable trait (Cesarini and Visscher, 2017) that is strongly negatively correlated with delay discounting (Golsteyn et al., 2014). Both our tolerance for and policy response to the link between discounting and education depend on how discounting behaviors are formed. If interpersonal differences in discounting are strongly explained by interpersonal genetic variation, we can debate the fairness of random genetic endowments playing a key role in educational inequalities through this channel. If discounting is strongly environmental, we should ask what aspects of the

¹Risk and time preferences have been shown to predict consequentially important choices such as risky health behaviors (Anderson and Mellor, 2008; Harrison et al., 2010; Courtemanche et al., 2015; Norrgren, 2022; Cheung et al., 2022), vaccinations (Lepinteur et al., 2023), migration (Jaeger et al., 2010; Ayhan et al., 2020), investment behavior (Bradford et al., 2017; Wong et al., 2019), altruistic behavior (Angerer et al., 2015), and job search and occupational selection (Fouarge et al., 2014; Cortés et al., 2021), as well as inequality in wealth at the societal level (Epper et al., 2020).

environment matter to evaluate whether this is fair or not. The estimates of heritability for behavioral traits are important measurements to contextualize the world we live in. Such information can help us to understand social outcomes, direct future research efforts, and inform social and economic policy decisions.

There have been isolated attempts to estimate the heritability of risk and time preferences using twin study designs. The results vary substantially across the studies, in part because of small samples in some papers, but more generally because of wide variation in the measurement approach used. To provide a structured summary, we conduct a meta-analysis of available literature (see Section 2.2). The existing studies define and measure preferences in a variety of ways, but many issues remain to be addressed. To our knowledge, only two studies (Cesarini et al., 2009; Zhong et al., 2009) use salient monetary incentives to elicit risk preferences. Crucially, there is no study of heritability which defines risk preferences in terms of structural parameters such as utility curvature and probability weighting. There are only two incentivized studies on time preferences (Anokhin et al., 2011; Sparks et al., 2014), both of them involving samples of adolescents. One study structurally estimates time preferences in the form of a discounting function (Anokhin et al., 2015), but they do not apply correction for utility curvature (Andersen et al., 2008). Perhaps more importantly, their analysis is based on an informal two-step procedure that initially estimates the discounting model without considering heritability, and subsequently uses predictions from this step as the dependent variable in a regression model of heritability which is otherwise unrelated to the discounting model. Despite the ongoing interest in uncertainty and ambiguity aversion—which concerns outcomes without precise probabilities and is thus distinguished from risk aversion—in the broad economics literature, we know nothing about heritability in this domain of preferences, except for findings reported by Cesarini et al. (2012). Finally, no study with salient incentives accounts for behavioral and measurement errors in observed responses, which can attenuate the estimates of genetic heritability.

Using a new dataset from an incentivized experiment with adult twins, we provide ev-

idence on the heritability of economic preferences. We decompose variation, in terms of nature and nurture, for many different types of preference—aversion to *probabilized* uncertainty (risk aversion), aversion to *non-probabilized* uncertainty (uncertainty aversion), long-run discounting, and present bias.² We develop a novel empirical approach that marries behavioral genetics with structural econometrics, and estimates structural parameters in economic theory jointly with heritability and errors in decision making. This allows us to, for the first time, identify heritability in measures of preferences which directly correspond to parameters in economic models of choice and are relevant to economic research. Specifically our approach extends the basic twin ACE genetic decomposition model to classes of structural decision models that are staples of economic decision theory, including expected utility theory (EUT), rank-dependent utility (RDU), exponential discounting, and quasi-hyperbolic discounting. We further consider the joint estimation of quasi-hyperbolic and RDU decision functions, accounting for the confoundedness of concave utility and delay discounting (Andersen et al., 2008). The existing studies define and measure preferences in a non-structural way, for example by equating risk preferences with the number of safe or sooner choices that an individual makes in a choice experiment. Compared to such approaches, the advantage of structural estimation is that it allows us to convert choices into quantitatively meaningful and theoretically grounded constructs that are primitives in many economic models.

Our paper also addresses another shortcoming of earlier research—measurement error. Measurement error is a serious concern in twin studies because estimates that do not correct for measurement error will underestimate the role of genes and overestimate the role of the unique environment. In the conventional twin study regression model, the unique

²In the literature on decision making under risk and uncertainty, the term *risk* typically refers to cases where decision makers make choices knowing probability distributions of outcomes. As Etner et al. (2012) point out, however, definitions of *uncertainty* and *ambiguity* vary from study to study in the current literature, and sometimes the two terms are used interchangeably. Similarly as Abdellaoui et al. (2011), we use *uncertainty* to describe cases where decision makers are making choices involving unknown probability distributions, and consider ambiguity aversion as a measure that relates to the difference between the same decision maker’s risk and uncertainty aversion. We acknowledge that studies such as Hey et al. (2010) and Stahl (2014) do not draw a distinction between uncertainty and ambiguity, and consider our notions of uncertainty and uncertainty aversion as ambiguity and ambiguity aversion, respectively.

environment is represented by the usual idiosyncratic disturbance term which captures genuine environmental effects as well as any measurement error. As a result, the presence of the latter understates the relative importance of genetic effects by construction. Existing studies are therefore potentially biased against finding a strong role for genes.³ When we estimate structural choice models, the variance of the error component is estimated and can therefore easily be deducted from the total variance of the elicited preference. For comparison, we also estimate preferences using non-structural models in line with the approaches in earlier research. Additionally, we expand on these earlier approaches by utilizing the fact that participants completed similar tasks multiple times. This allows for another method to deal with measurement error, in which we treat each task as a repeated measure and estimate the degree of variation due to noise which we purge from the decomposition equation (Ge et al., 2017).

In our most general structural models, we find that 36%-48% of the variation in risk aversion parameters is explained by genes, and in contrast to most other research, we also find a smaller but non-trivial role for the common family environment (8%-15%). In our meta-analysis of prior research, we find that 25% (95% CI [20%, 30%]) of variation in behaviorally elicited risk aversion can be explained by interpersonal differences in genes, with most of the remaining variation explained by unique environmental experiences. The reduced importance of unique environmental experiences in our study compared to the meta-analysis is consistent with behavioral and measurement errors inflating this factor. We also estimate a significant role of genes on parameters governing uncertainty preferences (20%-52%) and ambiguity aversion (26%-34%).

Results for time preferences are more nuanced. In our preferred specification, which jointly estimates delay discounting and present orientation along with an RDU decision function, we find no role for genes but 33% of variation is explained by the common sibling environment. However, when we estimate the same model without adjusting for risk aversion,

³Gillen et al. (2019) estimate that around 40% of variation in experimentally estimated risk preferences is due to noise.

the loading switches from the common sibling environment to genes. According to our meta-analysis of prior studies, the importance of genes in the case of delay discounting is quite high at 38%, but estimated quite imprecisely with a confidence interval [16%, 59%]. This suggests that the heritability in time preferences found in earlier studies may reflect the heritability of risk aversion rather than delay discounting per se because these earlier studies did not control for utility curvature in their estimates of time preference. Overall, our analysis also suggests that time preferences are quite malleable, and more so than risk aversion, which may have implications for parenting and education.

Finally, to further demonstrate that the way preferences are conceptualized and estimated matters for understanding how they form, we assess how our conclusions change under a variety of alternative modelling approaches. This includes estimating simple structural models that do not control for widely replicated behavioral phenomena (e.g., expected utility model without probability weighting) and non-structural regression models based on the raw counts of safe/sooner choices in the decisions tasks, which is conceptually closer to the majority of previous studies. Our estimates are highly sensitive to how preferences are treated, with genes having virtually no influence in some specifications. Consistently, the importance of unique environmental experiences is considerably larger in non-structural and parsimonious structural models. This is consistent with estimates of the variance share for unique environment masking unobserved behavioral and measurement errors in these models.

Our paper is organized as follows. In Section 2 we provide background information on the estimation of heritability and present the results of our meta-analyses. In Section 3 we discuss our data, including the behavioral decision tasks. In Section 4 we discuss how we estimate heritability, linking the classic twin study model with the estimation of heterogeneous structural parameters in economic experiments. In Section 5 we present our new results. Section 6 concludes.

2 Background

2.1 Methods for estimating heritability

The most popular approach to estimate heritability of economic preference is through the twin study method using the ACE decomposition model. The basic idea of the ACE decomposition model is that differences in the strength of correlation of a given phenotype (observable trait) between monozygotic (identical) versus dizygotic (fraternal) twins can be used to infer the degree to which that phenotype is formed by additive genetic effect (A), common sibling environment (C), and unique or individual-specific environment which is not shared between siblings (E).⁴ This model yields a tractable formula for measuring the degree of heritability, which can be interpreted as the share of the phenotype variation due to genes: $\sigma_A^2/(\sigma_A^2 + \sigma_C^2 + \sigma_E^2)$, where σ_k^2 denotes the variance in the phenotype that can be attributed to effect k . Section 4.1 summarizes a workhorse approach to estimating the three variance components.

The key statistical insight underpinning the twin study method is that monozygotic (MZ) twins share 100% of DNA, whereas dizygotic (DZ) twins share on average 50% of segregating genes, as with any sibling. The main assumption—known as the equal environments assumption—is that MZ and DZ twins are similarly exposed to shared environmental factors relevant to the phenotype. Under this assumption, MZ twins exhibit a stronger correlation in the phenotype than DZ twins due to a greater genetic similarity. Combined with the ACE model, this assumption allows one to rewrite the variance share of genetic effects as $2 \times (\rho_{MZ} - \rho_{DZ})$ (Falconer’s formula), where ρ_{MZ} and ρ_{DZ} refer to the phenotype correlation for the MZ and DZ twin pairs. In practice, however, researchers usually apply structural equations modelling (SEM) or mixed effects regression to estimate the variance components directly, in part to ensure that the finite sample estimate of the variance share lies in the

⁴Dominance genetic effects (D) – a genetic interaction effect between alleles at the same locus – may also be present but in the classic twin study cannot be identified without assuming no role for common environment, so are often assumed to be zero. This is often argued to be reasonable when the correlation in the phenotype for monozygotic twins is not more than double that of dizygotic twins.

$(0, 1)$ interval as theoretically required.

Genome-wide association studies (GWAS), which directly link genetic code to the phenotype, provide a main alternative to the twin study method.⁵ These studies require genetic sequencing of study participants which is typically done using single-nucleotide polymorphism (SNP) arrays, rather than the whole genome. SNPs are markers for variation in DNA that occur with high frequency throughout the genome. Variation in these SNPs is then linked to variation in the phenotype by, for example, running many hundreds of thousands of regressions, separately for each SNP. Adjustments are made to p-values for multiple hypothesis testing when identifying which of specific genetic variations are statistically significant. Heritability is estimated by methods that link variation in all SNPs to variation in the phenotype (see Yang et al., 2017).

Both twin and GWAS methods have their own relative strengths and weaknesses and can be thought of as complements (Friedman et al., 2021). However, when it comes to estimating the heritability of economic preferences, several considerations make the twin study design the more attractive option. While GWAS studies can identify specific genetic variants associated with the phenotype, to do so requires very large samples in the tens, if not hundreds, of thousands of participants. This is particularly the case for traits that are highly *polygenetic*, meaning that the phenotype is influenced by many genetic variants, each with small independent effects. Most complex human traits are known to be polygenetic (Chabris et al., 2015). Insufficient sample size is the main explanation for the so-called ‘missing heritability’ problem—the tendency for GWAS to estimate considerably smaller amounts of heritability than twin studies (Maher, 2008). Even with large samples, the presence of rare genetic variants and the fact that whole-genome sequencing is uncommon means that GWAS estimates for heritability are likely downward biased (see, for example, Wainschein et al., 2021).⁶

⁵See Bush and Moore (2012) and Uffelmann et al. (2021) for an overview.

⁶There is also the possibility that twin study estimates are upward biased, due to for example gene-gene or gene-environment interactions or violation of the equal environments assumptions. Conley and Fletcher (2017) discuss these and other possibilities in the context of behavioral genetics and conclude that such

The large sample requirement of GWAS methods is especially difficult to meet for studies that involve the elicitation of economic preferences, which requires participants to complete extensive decision tasks with meaningful incentives. The financial and practical challenges associated with administering such tasks across samples of tens or hundreds of thousands of people make GWAS a prohibitive option. The few GWAS studies that have been conducted had to rely on hypothetical decisions or attitudinal proxies (e.g., Benjamin et al., 2012; Sanchez-Roige et al., 2018; Linnér et al., 2019). Most of these studies have been insufficiently powered to identify specific genetic variants, and are likely to under-estimate heritability.⁷

Finally, economists are usually not interested in discovering which particular gene or interaction of genes is associated with a particular economic behavior. They are much more interested in broader issues rotating around the extent to which economic preferences are shaped by the collective effects of all genes rather than environmental factors, which can act as a guide to understanding the reach of parenting, education, and other potentially preference-shifting interventions. All these reasons make the twin study methodology particularly attractive to economists and social scientists in general.

2.2 Empirical evidence on the heritability of economic preferences

In Appendix B we summarize the existing studies that directly estimate the heritability of risk and time preferences (Tables B.1 and B.2).⁸ To provide the most comprehensive evidence

concerns are unlikely to explain much of the discrepancy. Wainschtein et al. (2021) find that in the case of BMI and height, once full genome sequencing is utilized, the gap between pedigree and GWAS estimates largely closes.

⁷Linnér et al. (2019) is a notable exception with a sample of more than 900,000 individuals for the main measure of risk aversion (“Would you describe yourself as someone who is willing to take risks” [Yes/No]). They found 99 unique genetic loci associated with this variable (indicating it is highly polygenetic), although heritability was estimated to be only 5%. They obtained a higher estimate (16%) when using the first principal component from a set of questions about risky behaviors, which may reflect bias in their main estimate due to measurement error. Their study did not involve any quantitative measures of risk aversion comparable to those we analyze.

⁸We note that some studies decompose the variation in behaviors closely related to economic preferences, such as investment and saving decisions (e.g., Barnea et al., 2010; Cronqvist and Siegel, 2014, 2015). However, since these behaviors have determinants other than preferences, we do not include them in our review. We also acknowledge the literature that shows parents transmit their preferences to their children (see e.g., Dohmen et al., 2012; Alan et al., 2017; Heinrich and Shachat, 2020; Chowdury et al., 2022), but does not

to date on the heritability of risk and time preferences we meta-analyze the existing estimates for the share of variance explained by genes (A) and by the unique environment (E). The former share is based on the formula presented at the beginning of Section 2.1. The latter share is also based on the same formula, but places σ_E in the numerator instead of σ_A .

There are substantial differences across the studies in how results are reported. To standardize the approach in our meta-analysis, we adopt some rules. First, not all studies report estimates of all the components of the ACE model, and many report estimates from more than one model. In our meta-analysis, we use the estimates from the complete ACE model whenever they are reported. In the remaining cases, we use the estimates from the best-fitting model among those that were reported. A few studies estimated an ADE model which distinguishes between additive (A) and dominant (D) genetic effects; in slight abuse of notation, for such studies, what we refer to as A is in fact an aggregation of the two types of genetic effects. Second, many studies do not report standard errors. We used the reported standard error when available, and the inferred standard error based on the confidence interval in other cases.⁹ Third, studies use diverse methods to measure and estimate preferences, and many studies report heritability estimates for more than one measure of risk attitude. In Tables B.1 and B.2 we classify these measures into behavioral (incentivized, observational, and not incentivized) which are elicited from risky decisions with or without consequences, and stated measures which are self-reported responses to questions about one’s willingness to take risk and one’s tendency to be patient. We separately meta-analyze the heritability estimates of the behavioral and stated preference measures. If a study measured preferences using more than one instrument in the same sample, we use only one estimate. In the case of behavioral preferences, we prioritize estimates from incentivized elicitations and those that provide the most precise estimates.

Let j be a generic subscript for an existing estimate of the variance share due to genetic effects (A). To estimate the average role of each component in shaping preferences, we

distinguish between genetic and environmental transmission.

⁹Further details and sensitivity analyses are in Appendix B.

estimate the following random effects meta-analysis model

$$\Theta_j = \Theta_0 + \xi_j + \epsilon_j \quad (1)$$

where Θ_j is the j^{th} estimate of the variance share due to A. Under this model, each estimate is decomposed into the average effect parameter of interest Θ_0 , which is common to all existing estimates, and normally distributed sampling errors $\xi_j \sim N(0, \tau^2)$ and $\epsilon_j \sim N(0, \hat{\nu}_j^2)$, where between-study heterogeneity τ is an unknown parameter to be estimated alongside Θ_0 and $\hat{\nu}_j$ is the standard error of Θ_j that we obtain from the previous studies. We use the same model to estimate the average role of the unique environment (E), where Θ_j is then understood to be an existing estimate of the variance share due to E.

The random effects meta-analysis estimator of Θ_0 can be seen as a weighted average of all existing estimates in our data set:

$$\Theta_0^{RE} = \frac{\sum_j w_j \Theta_j}{\sum_j w_j} \quad (2)$$


where the weight for estimate j is the reciprocal of its total variance: $w_j = 1/(\hat{\nu}_j^2 + \tau_j^2)$. Thus, the estimates with higher precision (*i.e.*, smaller standard errors $\hat{\nu}_j$) receive larger weights.

2.2.1 Risk preferences

Nine studies estimated the heritability of risk preferences using twin methodology (see Table B.1). Three of these studies include preference measures obtained from incentive-compatible laboratory decision tasks, six include behavioral non-incentivized measures, and five include stated preferences measured as self-reported responses to questions about one's attitude to risk.

The meta-analytic average of heritability (*i.e.*, that of the variance share due to genes or A) in behaviorally measured risk attitude is 25% (95% CI [20%, 30%]) (see Figure 1). Common sibling environment (C) appears to be negligible. Most papers do not report this

Figure 1: Meta-analysis of the genetic role (A) in behavioral risk aversion




Note: Excluded estimates: Cesarini et al. (2009) behavioral not incentivized measure is not included because incentivized measure is available from the same sample. The second wave data (at 14 years old) from the longitudinal study by Anokhin et al. (2009) because the first wave dataset (at 12 years old) is larger due to attrition. Estimates obtained in the loss domain by Simonson and Sela (2010) because gain domain estimates are available. Lottery choice estimates from Nicolaou and Shane (2020) because investment task provides a more precise measure of risk preference. Harden et al. (2017) is excluded because the multivariate latent factor decomposition in their study is difficult to compare with the results from other studies where there is a clearly defined outcome variable.

estimate because models without the C component provide a better fit. The two studies that do estimate C find that it explains only 0% [0%, 37%] (Zhong et al., 2009) and 9% [1%, 22%] (Cesarini et al., 2009) of the variation. Finally, the most important factor in shaping preferences is the unique environment (E) explaining 68% of the variation in individual risk attitudes (95% CI [60%, 77%]) (see Figure 2).

The heterogeneity statistics offer a mixed picture. For heritability, $I^2 = 0$ implies that the between study variance accounts for none of the total variance (which follows from $\tau^2 = 0$), and $H^2 = 1$ means that the variance implied by the random effects meta analysis is identical to that of a fixed effects approach that assumes homogeneity. However, a formal test of study homogeneity rejects the null at the 5% level ($p = 0.041$) and the mixed results may reflect low power due to the small number of studies. To the extent there is heterogeneity, it is unclear

Figure 2: Meta-analysis of the unique environment role (E) in behavioral risk aversion



Note: Excluded estimates: Cesarini et al. (2009) behavioral not incentivized measure is not included because incentivized measure is available from the same sample. The second wave data (at 14 years old) from the longitudinal study by Anokhin et al. (2009) because the first wave dataset (at 12 years old) is larger due to attrition. Estimates obtained in the loss domain by Simonson and Sela (2010) because gain domain estimates are available. Lottery choice estimates from Nicolaou and Shane (2020) because investment task provides a more precise measure of risk preference. Harden et al. (2017) is excluded because the multivariate latent factor decomposition in their study is difficult to compare with the results from other studies where there is a clearly defined outcome variable.

whether this is due to genes being differently important in different populations, or due to differences in elicitation methods. The heterogeneity statistics for E are more equivocal and indicate differences across studies that are not due to sampling variation alone.

The results from the meta-analysis of the heritability of the stated risk preferences paint a similar picture with 24% of variation (95% CI [19%, 30%]) attributed to genetic factors (Figure B.1 in Appendix B). The common sibling environment is only estimated in two studies and is negligible (0% [0%, 10%] in Nicolaou and Shane (2020) and 5% [0%, 17%] in Cesarini et al. (2009)). The unique environment is estimated to account for 71% of variation in stated risk preferences (95% CI [64%, 78%]) (Figure B.2 in Appendix B).


2.3 Ambiguity preferences

A complete taxonomy on how the decision maker behaves under risk and uncertainty requires elicitation of preferences when the underlying probability distributions of outcomes are known to the decision maker (*risk*) as well as unknown to the decision maker (*uncertainty*). Since the seminal work due to Ellsberg (1961), several studies have reported that participants in experiments tend to make more conservative choices under uncertainty than risk, a phenomenon which is often referred to as ambiguity aversion. To date, only one study has estimated the heritability of ambiguity aversion (Cesarini et al., 2012) using a hypothetical choice between drawing a colored ball from an urn with known and unknown distribution of winning and losing balls. The choices in this study were not incentivized and did not impact participants' earnings. The heritability of selecting the risky urn in favor of the uncertain urn was estimated at 16% but is fairly imprecise (95% CI [0%, 29%]).

2.4 Time preferences

Table B.2 summarizes seven studies that estimated the heritability of time preferences. As with risk attitudes, there are substantial differences across the studies in how time preferences are elicited and measured. The only two studies that are incentive-compatible involved one-shot choices between \$7 now and \$10 in one week and collected data from adolescents only (Anokhin et al., 2011; Sparks et al., 2014). These studies do not allow for the estimation of the individual's discount rate. Four studies measured time preferences using hypothetical decisions and two included different self-reported questions or questionnaires about impatience or future orientation. Anokhin et al. (2015) is the only study that estimated preferences in a structural model, specifically a one-parameter hyperbolic discounting model with linear utility (they do not manipulate the front-end delay and hence cannot identify the quasi-hyperbolic model). Their analysis is based on an informal two-step procedure, which initially estimates time preferences without considering heritability, and subsequently uses individual-specific predictions from the first step as the dependent variable in a her-

Figure 3: Meta-analysis of the genetic role (A) in behavioral time preferences




Note: Excluded estimates: The second wave data (at 14 years old) from the longitudinal study by Anokhin et al. (2011) because the first wave dataset (at 12 years old) is larger due to attrition. From the longitudinal study by Anokhin et al. (2015) we exclude waves with fewer data points. Harden et al. (2017) is excluded because the multivariate latent factor decomposition in their study is difficult to compare with the results from other studies where there is a clearly defined outcome variable.

itability regression model. No studies distinguish between present bias and impatience or correct for utility curvature in their estimates of time preferences.

The meta-analytic average of heritability (*i.e.*, the variance share due to A) of time preferences is 38% with a relatively wide 95% confidence interval of [16%, 59%] (Figure 3). This is larger than for risk preferences but notably less precise. Like for risk preferences, estimates for the common sibling environment C are not always reported; the two studies that do report this find small and imprecise effects of 7% [0%, 26%] (Cesarini et al., 2012) and 12% [0%, 54%] (Sparks et al., 2014). Unique twin environment (E) accounts for 58% of variation in behavioral time preference (95% CI [40%, 76%]) (Figure 4). There is a large degree of heterogeneity across studies. For example, the I^2 implies that 80.4% of the total variance in the existing estimates is explained by between study variance τ .

Two studies measured heritability using self-reported responses to questionnaires. The multivariate latent factor decomposition in Harden et al. (2017) makes it difficult to compare their results to the traditional approaches with a clearly defined outcome variable. Hubler (2018) estimates of the genetic influence are substantially different from those obtained from

Figure 4: Meta-analysis of the unique environment role (E) in behavioral time preferences



Note: Excluded estimates: The second wave data (at 14 years old) from the longitudinal study by Anokhin et al. (2011) because the first wave dataset (at 12 years old) is larger due to attrition. From the longitudinal study by Anokhin et al. (2015) we exclude waves with fewer data points. Harden et al. (2017) is excluded because the multivariate latent factor decomposition in their study is difficult to compare with the results from other studies where there is a clearly defined outcome variable.

studies that use behaviorally measure time preference, with the A, C, and E components estimated at 0%, 23%, and 77%, respectively.

Our paper fills a gap in the literature by estimating the heritability of risk and time preferences of adults from incentive-compatible decisions, while carefully separating risk preferences from uncertainty attitudes and separating patience from future orientation thus providing more precise estimates of heritability.

3 Data

3.1 Recruitment and survey design

We use data from the Australian Twins Economic Preferences Survey (ATEPS) (Kettlewell and Tymula, 2021). ATEPS was created in cooperation with Twins Research Australia (TRA), the custodians for Australia's largest twin registry. TRA were involved with recruiting twins into the study, while administration of the online survey was handled by authors Nathan Kettlewell and Agnieszka Tymula.

At the time of the survey – September 2020-March 2021 – there were around 70,000 adult twin pairs in TRA’s registry. Participation in the registry is voluntary, and over its history TRA have used a variety of methods of recruitment including advertising with the Australian Multiple Birth Association, news releases, print media and other media engagements, word of mouth, the TRA website, newsletter and social media pages, and events such as information forums (Murphy et al., 2019).

TRA approached 6,848 adult twin pairs to take part in the study. These were adult twin pairs aged 18-65 years where both twins were active members of the registry with up-to-date contact details. Recruitment primarily occurred through email invites and SMS reminders. A limited number of phone calls were used to maximize the sample size. These calls were primarily used when one twin agreed to be contacted by the research team but the co-twin did not respond to the email or SMS invitation. Ultimately, 803 twin pairs agreed to be contacted by the research team, and 560 of those twin pairs fully completed the survey.

The survey was first emailed to participants on 8 September 2020 and then progressively sent to additional participants until 25 February 2021 before closing on 1 March 2021. Participants received unique survey links and were able to complete the survey at a time of their choosing. Due to the length of the survey, we allowed participants to pause at any point and recommence at a later time.

The survey began with general information about study protocols.¹⁰ To proceed past the first page, participants needed to correctly answer a series of yes/no questions to ensure they understood that (1) they could earn real money from their decisions, (2) that they would need to provide bank details at the end of the survey to facilitate this, (3) that both twins needed to complete the survey before payments could be made, and (4) that they should not discuss their answers with their co-twin until both had completed the survey. Participants answered a series of questions about twin similarity, which were used to determine zygosity

¹⁰The complete questionnaire is available at <https://dataverse.ada.edu.au/dataset.xhtml?persistentId=doi:10.26193/TTQEBQ>.

status.¹¹ They then completed a series of incentivized behavioral preference tasks. They were told that one task would be chosen at random at the end of the survey and that for all twin pairs who completed the survey, their choices in this task would determine their payment. In total, there were 15 separate behavioral tasks that could be selected for payment, with potential payoffs ranging from AUD\$0 to \$37.50 (approximately USD\$29), depending on their decisions and the outcomes of the games when played out for real. For the multiple price list tasks, participants were told that if this task is chosen for payment, one row will be selected at random and played out for real. All randomization was done using randomization branch logic through the Qualtrics survey software, although this was not explicitly described to participants. Payments were processed once a week via bank transfers and participants were told that payments would be received within 10 days of both twins completing the survey.¹² The average payment was AUD\$15. To emphasize the importance of participants' decisions, participants were only paid based on the behavioral tasks (no show-up fee). The median completion time for the survey was 55 minutes.

All protocols and procedures were approved by the University of Technology Sydney Human Research Ethics Committee (application numbers ETH19-4381 and ETH20-5410) and by TRA.

3.2 Descriptive statistics

Five hundred and sixty Australian twin pairs completed the online survey (401 MZ, 159 DZ). The average age in the sample is 45 years (range 18-66) and 81% of participants are female. Of the DZ pairs, 42 (26%) are mixed sex. Table 1 presents descriptive statistics for

¹¹Of the 518 same-sex twin pairs who fully completed the survey, the zygosity status of 184 (35%) is determined by self-reported genetic test results, 25 pairs (4.8%) are classified as dizygotic due to different blood type and the remainder are classified using the peas-in-a-pod questionnaire, which has been shown to predict zygosity with more than 90% accuracy (Ooki et al., 1990).

¹²Twenty-one participants completed the survey but were not paid, either because they declined the payment (e.g., by preferring the money be donated back to the research project) or provided incorrect bank details but could not be contacted to correct this. We retain these people in our sample because in such cases it is clear the participant had a strong intrinsic motivation to contribute to the research, and we have no reason to doubt the integrity of their responses.

Table 1: Descriptive statistics by zygosity

	Mean MZ	SD MZ	Mean DZ	SD DZ	Diff	P-val
Age	44.03	12.89	46.29	12.57	2.26	0.058
Male	0.14	0.35	0.25	0.43	0.10	0.002
Twin years together	22.38	5.40	20.40	3.91	-1.98	0.000
Born in Australia	0.87	0.34	0.90	0.30	0.03	0.260
Lives in a city	0.65	0.48	0.66	0.48	0.01	0.759
Couple	0.65	0.48	0.69	0.46	0.04	0.304
Household members	4.51	1.83	4.44	1.84	-0.08	0.577
Dependent children	1.89	1.42	2.02	1.52	0.13	0.289
University degree	0.59	0.49	0.60	0.49	0.01	0.816
Employed	0.86	0.35	0.84	0.37	-0.02	0.409
Retired	0.08	0.27	0.09	0.28	0.01	0.822
Income	1256.90	703.32	1321.91	719.97	65.01	0.246
Financially secure	3.16	0.76	3.18	0.75	0.02	0.748
Long-term health condition	0.22	0.41	0.19	0.39	-0.03	0.316
COVID-19 worry	2.82	2.70	2.81	2.66	-0.01	0.956
COVID-19 prob	10.92	15.54	9.57	13.30	-1.35	0.177
COVID-19 mort	14.04	20.34	13.51	19.07	-0.53	0.719
COVID-19 job loss	0.07	0.25	0.06	0.24	0.00	0.848
COVID-19 reduced income	0.13	0.33	0.14	0.34	0.01	0.660
COVID-19 work home	0.35	0.48	0.35	0.48	-0.01	0.806
COVID-19 reduced hours	0.14	0.34	0.13	0.34	-0.01	0.831
Num. COVID-19 positive friends	1.88	2.56	1.73	2.48	-0.15	0.371
Risk MPL1 num safe	8.03	3.29	8.32	3.35	0.29	0.228
Risk MPL2 num safe	5.17	2.11	5.34	2.13	0.17	0.238
Uncertainty MPL1 num safe	8.51	3.34	8.69	3.40	0.18	0.449
Uncertainty MPL2 num safe	5.41	2.19	5.52	2.27	0.11	0.488
Time MPL1 num sooner	3.82	3.31	3.48	3.24	-0.34	0.130
Time MPL2 num sooner	3.49	2.99	3.38	3.02	-0.11	0.605
Time MPL3 num sooner	4.00	3.36	3.88	3.39	-0.12	0.605
Time MPL4 num sooner	3.99	3.30	3.86	3.37	-0.13	0.569

Note: Descriptive statistics based on samples of 802 MZ twins (401 pairs) and 318 DZ twins (159 pairs). See Appendix Table C.1 for detailed variable definitions.

the analysis sample by zygosity status.

MZ and DZ twins are similar across most of the characteristics measured in our demographic and socioeconomic questionnaire. They are of similar age, are equally likely to be born in Australia, live in a city, and live as a couple. They have similar household characteristics (household size, number of children), education, work status, and COVID-related

experiences. DZ twins are more likely to be male (14% versus 25%) so we control for sex in our analysis. MZ twins on average report living longer with their co-twin than DZ twins (22.4 years versus 20.4 years), which is a potential threat to the equal environments assumption and could upward bias the estimates of heritability if MZ twins have more common experiences that influence economic preferences than DZ twins. We do not view this difference as particularly concerning because it has been widely replicated across behavioral genetics research that common environment plays a relatively small role in behavioral outcomes (Chabris et al., 2015). Indeed, also in our sample the correlations between years spent living together and the absolute difference in economic preferences are small, inconsistent in direction, and never statistically significant (Appendix Table C.2). MZ and DZ twins are generally very similar in their risk and time preferences, with no statistically significant differences. Altogether, our samples of MZ and DZ twins are highly similar, which is helpful in meeting the equal environments assumption embedded in our analysis.

3.3 Risk and uncertainty preference tasks

To measure risk preferences, we asked participants to choose between a sure payoff and a 50/50 lottery using a multiple price list (MPL) format. The lottery in each MPL remained fixed and the value of the sure payoff increased in each row (see Table 2).¹³ Participants completed two MPLs, which differed in their payoffs and number of rows, following advice in Gillen et al. (2019) in relation to identifying measurement error. The lottery was framed as an urn with black and red balls (there were 20 balls in MPL1 and 30 in MPL2, half of each color), out of which one ball would be picked and its color would determine if the participant won or lost the lottery.

To measure uncertainty preferences, we used the exact same MPLs with the same payoffs

¹³ATEPS also includes other measures of risk aversion: the Gneezy and Potters (1997) investment task, and Eckel and Grossman (2002) single lottery selection task. In this paper, we focus on the MPL for several reasons. MPL tasks are the most popular method to elicit preferences in field and lab research. We administered MPL also under uncertainty. It matches our time preference task in structure. Finally, it is ideal for structural estimation in the random utility framework as it involves many decisions from the same participant.

Table 2: Risk preference tasks

MPL1		MPL2	
Sure thing	50/50 chance	Sure thing	50/50 chance
\$2	\$30/\$0	\$2.50	\$25/\$0
\$4	\$30/\$0	\$5	\$25/\$0
\$6	\$30/\$0	\$7.50	\$25/\$0
\$8	\$30/\$0	\$10	\$25/\$0
\$10	\$30/\$0	\$12.50	\$25/\$0
\$12	\$30/\$0	\$15	\$25/\$0
\$14	\$30/\$0	\$17.50	\$25/\$0
\$16	\$30/\$0	\$20	\$25/\$0
\$18	\$30/\$0	\$22.50	\$25/\$0
\$20	\$30/\$0	\$25	\$25/\$0
\$22	\$30/\$0		
\$24	\$30/\$0		
\$26	\$30/\$0		
\$28	\$30/\$0		
\$30	\$30/\$0		

but the proportion of each ball color in the urn was unknown to participants. Participants could choose the ball color that would be associated with a win to ensure that they do not think that the number of balls in uncertain urns is skewed against them. The intuition for this task is that more uncertainty-averse participants will choose the sure payoff more often.

On average, our participants exhibited practically neutral attitudes to risk and uncertainty. They selected the sure payoff 53% of the time in the risk task and 56% in the uncertainty task, whereas a risk and uncertainty neutral person would choose the sure payoff 54% of the time in both types of tasks. While the respective sample frequencies are significantly smaller ($p = 0.011$) and greater ($p = 0.000$) than this benchmark, indicating risk seeking and uncertainty aversion, neither difference is large in magnitude.

3.4 Time preference tasks

To elicit time preferences, we asked participants to choose between sooner and smaller or later and larger payments, also using a MPL format. In total we used four MPLs. In the first

Table 3: Time preference tasks

MPL1 (MPL 3)		MPL2 (MPL4)	
Now (4 weeks)	8 weeks (12 weeks)	Now (12 weeks)	6 weeks (18 weeks)
\$15	\$15.50	\$13	\$13.50
\$15	\$16.50	\$13	\$15
\$15	\$17.50	\$13	\$16.50
\$15	\$18.50	\$13	\$18
\$15	\$19.50	\$13	\$19.50
\$15	\$20.50	\$13	\$21
\$15	\$21.50	\$13	\$22.50
\$15	\$22.50	\$13	\$24
\$15	\$23.50	\$13	\$25.50
\$15	\$24.50	\$13	\$27

MPL, participants chose between a payment ‘now’ versus in eight weeks and in the second MPL they chose between a payment ‘now’ and in 12 weeks. The third and fourth MPLs had the same payments and delay between the sooner and later payment, but the front-end delay (the delay to the sooner payment) was four and six weeks respectively (see Table 3 for the list of all decision scenarios). The addition of the longer front-end delay allows us to identify time preferences that are non-stationary (Halevy, 2015).¹⁴

On average, participants selected the sooner payment 36% of the time when its timing was framed as now, and 40% when the front-end payment was delayed. The significant difference between the two sample frequencies ($p = 0.000$) indicates that on average our participants have future-oriented, hence non-stationary, time preferences.

4 Econometric approach

We propose a new approach to measuring heritability in economic preferences. The ACE model has been widely used in twin studies to analyze the heritability of observable traits. However, the deep parameters of economic models, such as the Arrow-Pratt coefficient of

¹⁴We do not identify present bias because our “now” payments were processed within 10 days of both twins completing the survey due to practical constraints. Balakrishnan et al. (2020) empirically demonstrate the importance of administering immediate payments.

risk aversion and individual discount rates, are latent constructs that do not lend themselves to direct observation. We identify the latent preference parameters in the framework of behavioral econometrics and measure heritability in those parameters by integrating the ACE model with the random coefficient discrete choice model of interpersonal preference heterogeneity. In the remainder of this section, we describe the three pillars of our econometric approach—the ACE model, the behavioral econometric framework, and the integrative random coefficient model.

4.1 Behavioral genetics

The first building block of our econometric approach is the ACE model, which is the workhorse model for twin studies in the behavioral genetics literature. This model was introduced to twin research in economics by Cesarini et al. (2009). The acronym ACE refers to the key modeling assumption that variation in an observed trait of interest can be decomposed into three latent sources, namely the additive genetic factor (A), the common environment factor (C), and the unique environmental factor (E).

From an econometric perspective, the ACE model is similar to a two-way random effects model that considers twin siblings from the same family as two data points that make up the same panel unit. Let $n \in \{1, 2, \dots, N\}$ denote different twin pairs and $s \in \{1, 2\}$ be the index of individual siblings within a twin pair. Suppose that y_{sn} is an individual’s observed trait, such as a measure of an individual’s educational attainment. The ACE model of variation in this trait can be specified as

$$y_{sn} = \alpha + A_{sn} + C_n + E_{sn} \tag{3}$$

where α is the overall intercept; and A_{sn} , C_n and E_{sn} represent three independent error components. The additive genetic factor, $A_{sn} \sim N(0, \sigma_A^2)$, captures all genetic influences on the individual’s trait (assuming no dominance genetic effects). It is constant within

an MZ (that is, identical) twin pair ($A_{1n} = A_{2n}$) but varies between two siblings of a DZ (non-identical) twin pair ($A_{1n} \neq A_{2n}$). The common environmental factor, $C_n \sim N(0, \sigma_C^2)$, captures those environmental influences that are shared by both siblings regardless of whether they are MZ or DZ twins, such as their childhood neighborhood and home environment. Finally, the unique environmental factor, $E_{sn} \sim N(0, \sigma_E^2)$, captures the remaining influences that are unique to each sibling, such as individual-specific episodes of illnesses or other life events.

In this framework, heritability in trait y_{sn} refers to the proportion of total variation in this trait which is explained by variation in the additive genetic factor. Put another way, heritability is quantified as

$$\pi_A = \sigma_A^2 / (\sigma_A^2 + \sigma_C^2 + \sigma_E^2) \quad (4)$$

where variance parameters σ_A^2 , σ_C^2 and σ_E^2 are to be estimated from the data alongside the intercept α by applying, for instance, a method of maximum likelihood. To disentangle additive genetic factor A_{sn} and unique environmental factor E_{sn} , which share the same dimensions of variation as per the subscripts, structural constraints are placed on between-sibling correlations in these error components. Let $\text{cov}(X, Z)$ denote the covariance between random variables X and Z . In the ACE model, $\text{cov}(A_{1s}, A_{2s}) = \sigma_A^2$ for MZ twins who share the same genetic variation, and $\text{cov}(A_{1s}, A_{2s}) = 0.5\sigma_A^2$ for DZ twins who tend to share half the genetic variation in common cases where the two biological parents are genetically unrelated. By contrast, since the unique environmental factor accounts for individual-specific variation, $\text{cov}(E_{1s}, E_{2s}) = 0$ for any twin pair, regardless of whether they are MZ or DZ twins.

The ACE model in equation (3) does not distinguish E_{sn} from the usual regression disturbance term, meaning that σ_E^2 absorbs not only variation in the unique environment but also any other idiosyncratic variation such as deliberation errors in decision making and statistical measurement errors. This may lead to overestimation of the variance in the unique environmental factor, σ_E^2 , hence underestimation of heritability based on equation (4). When multiple measures of the same trait are observed within a short span of time

during which E_{sn} can be expected to remain invariant, one may exploit this invariance to improve the identification of σ_E^2 . Let $t = 1, 2, \dots, T$ be the index of different repetitions in the measurement of the same trait, and y_{snt} denote the measurement outcome in each repetition. Then, the ACE model may be generalized as

$$y_{snt} = \alpha + A_{sn} + C_n + E_{sn} + e_{snt} \quad (5)$$

to accommodate a dedicated idiosyncratic disturbance, $e_{snt} \sim N(0, \sigma_e^2)$, which is distinguished from E_{sn} by having an extra dimension of variation across repetition t . This specification, inspired by Ge et al. (2017), is effectively a multi-level model that sees twin pair n as an upper-level panel unit that nests each sibling s as a sub-unit which consists of T data points.

In the analysis of the heritability of economic preferences, a key limitation of the ACE model, be it in the form of equation (3) or (5), is its fundamental assumption that the trait in question is observable, univariate, and measurable without the knowledge of other traits. The structural preference parameters of interest in economic forecasting and welfare evaluation (known as the “deep” parameters, especially in the macroeconomics literature) do not fit into this framework. In most cases, these parameters (e.g., measures of utility curvature and long-run discount rates) are latent constructs that make up a theorized decision making process rather than observable characteristics of an individual.¹⁵ Moreover, many behavioral models characterize a given domain of preferences in terms of multiple parameters—consider, for instance, non-expected utility models that attribute risk preferences to utility curvature and probability weighting—which defy a univariate measurement approach. Finally, the logic of an assumed behavioral model often makes it desirable to estimate several structural parameters jointly, even when they are referring to seemingly disjoint domains of preferences.

¹⁵Observed choice patterns in a preference elicitation experiment, summarized by descriptive measures such as the number of safe choices or the switching point in an MPL, often provide a useful insight into individual choice behavior. However, it is difficult to relate the observed choice patterns directly to latent preference parameters, unless one assumes away deliberation errors in the subject’s responses.

Perhaps the most well-known example is due to Andersen et al. (2008) who show that one can draw more plausible inferences about the discount rates elicited in a time preference experiment by conditioning such inferences on utility curvature that is identified with the aid of a risk preference experiment.

To address this limitation, we propose a new modeling strategy that integrates the ACE model with the structural parameters in decision-theoretic models of economic preferences.

4.2 Behavioral econometrics

The second building block of our econometric approach is behavioral econometric modeling, which acts as a bridge between economic theory at one end and experimental data at the other end (Harrison, 2019). In each of our risk, uncertainty, and time preference tasks, an individual sibling in a twin pair provided a series of binary choice responses. Behavioral econometrics offers a logically coherent framework that enables us to take each choice response directly as a data point and model these data points as realizations of a stochastic process that incorporates theoretical predictions based on structural parameters, as well as errors in decision-making that can induce potential violations of such predictions.

From an econometric perspective, the behavioral econometric model for each type of experiment can be represented as a latent dependent variable model with a non-linear index function. We first lay out a generic notation that applies to all three experiments. Let subscript snt identify the data point on sibling $s \in \{1, 2\}$ from twin pair $n \in \{1, 2, \dots, N\}$ in binary choice task $t \in \{1, 2, \dots, T\}$. Let c_{snt} indicate an observed binary indicator that equals one if the observed choice is Option A and zero if Option B. Consider latent dependent variable, c_{snt}^* , which is specified as

$$c_{snt}^* = V_{A,snt}[\boldsymbol{\theta}_{sn}] - V_{B,snt}[\boldsymbol{\theta}_{sn}] + \epsilon_{snt} \quad (6)$$

where $V_{j,snt}[\boldsymbol{\theta}_{sn}]$ is the index component that represents the individual sibling's preference

for Option $j \in \{A, B\}$ under an assumed theory of decision making; $\boldsymbol{\theta}_{sn}$ denotes a vector of relevant structural parameters in that theory, and ϵ_{snt} is an idiosyncratic disturbance term which accounts for behavioral errors. We assume that ϵ_{snt} follows a logistic distribution with mean zero and potentially heteroskedastic scale factor μ_{snt} , where the latter can be seen as the behavioral noise parameter. Typically, economic theory makes deterministic predictions such as the sibling chooses Option A if $V_{A,snt}[\cdot]$ is greater than $V_{B,snt}[\cdot]$ and Option B if the inequality is reversed. By applying the observation rule $c_{snt} = 1[c_{snt}^* > 0]$ where $1[\cdot]$ denotes an indicator function, we allow for violations of these deterministic predictions due to behavioral errors. Technically, each term in equation (6) can also carry superscript $type \in \{risk, unc, time\}$ to clarify the type of experiment being studied because the assumed theory, relevant structural parameters and the extent of behavioral noises may vary from experiment to experiment. We omit this extra superscript for now to avoid notational cluttering.

Our analysis of the risk preference task is based on the Rank-Dependent Utility model (RDU) due to Quiggin (1982). Let $U[M; r_{sn}]$ denote a utility function for monetary outcome M

$$U[M; r_{sn}] = M^{1-r_{sn}} \quad (7)$$

where r_{sn} is a measure of concavity in this function, which would be equivalent to the Arrow-Pratt coefficient of relative risk aversion had we assumed Expected Utility Theory (EUT). Given RDU and our experimental design that prompts a choice between a 50:50 prospect of $M_{A,snt}$ or nothing and a sure payoff of $M_{B,snt}$, the preference index for each option j can be specified as a function of two preference parameters, $\boldsymbol{\theta}_{sn} = \{r_{sn}, \omega_{sn}\}$, as follows

$$\begin{aligned} V_{A,snt}[\boldsymbol{\theta}_{sn}] &= \omega_{sn} U[M_{A,snt}; r_{sn}] - (1 - \omega_{sn}) U[0; r_{sn}] \text{ and} \\ V_{B,snt}[\boldsymbol{\theta}_{sn}] &= U[M_{B,snt}; r_{sn}] \end{aligned} \quad (8)$$

where the new parameter ω_{sn} is a decision weight which results from probability weight-

ing.¹⁶ If the decision weight coincides with the objective probability of $M_{A,snt}$ (that is, if $\omega_{sn} = 0.5$), RDU simplifies to EUT. If the sibling displays pessimism ($\omega_{sn} < 0.5$), however, probability weighting enhances risk aversion by reducing the certainty equivalent of Option A relative to what EUT predicts. Similarly, if the sibling displays optimism ($\omega_{sn} > 0.5$), probability weighting enhances risk-seeking by increasing the certainty equivalent. We complete the model specification by adopting the Contextual Utility model of behavioral errors due to Wilcox (2011), which posits that the behavioral noise parameter is proportional to the maximal utility difference in the choice task: $\mu_{snt} = (U[M_{A,snt}; r_{sn}] - U[0; r_{sn}])\mu$, where μ denotes the baseline noise parameter. When each risky choice task involves three outcomes ($0 < M_{B,snt} \leq M_{A,snt}$) as with our design, this heteroskedastic specification ensures that the implied probability of choosing the sure payoff is monotone increasing in the parameter r_{sn} that measures how risk-averse the sibling is in terms of the utility function.¹⁷

Our analysis of the uncertainty preference task is based on equations (7) and (8), but we allow the preference parameters and the baseline scale of behavioral errors (r_{sn} , ω_{sn} and μ) to take different values from their risk preference counterparts. The resulting model specification is consistent with RDU as well as several other models of uncertainty preferences cataloged by Hey et al. (2010) and Kothiyal et al. (2014). For instance, given our experimental design that prompts a choice between an uncertain prospect of $M_{A,snt}$ or nothing and a sure payoff of $M_{B,snt}$, decision weight ω_{sn} can be seen directly as the subjective probability of $M_{A,snt}$ in Subjective Expected Utility (SEU). It can be also seen as a reduced-form parameter which jointly accounts for the effects of the alpha weight and relevant priors in

¹⁶With an experimental design that considers two or more interior probabilities, one may identify a probability weighting function (PWF), $w[P]$, and use this function to derive a decision weight that substitutes for any objective probability P . Since there is only one interior probability in our design ($P \in \{0, 0.5, 1\}$) and $w[0] = 0$ and $w[1] = 1$ by definition, we can identify $w[0.5] = \omega$, the decision weight at $P = 0.5$, but cannot trace how $w[P]$ varies over P to identify the PWF itself. While it is straightforward to choose any of one-parameter PWFs and solve the equation $w[0.5] = \omega$ for the value of a shape parameter that produces the target decision weight, it remains the case that we cannot tell apart those alternative PWFs. For this reason, we directly consider the decision weight as a structural parameter.

¹⁷As shown by Wilcox (2011) and re-iterated by Apesteguia and Ballester (2018), the Fechner model of behavioral errors ($\mu_{snt} = \mu$ for all data points) violates this type of monotonicity.

Alpha MaxMin Expected Utility (Alpha-MEU).¹⁸ Alternatively, one may drop the assumption of probabilistic sophistication shared by these models and consider ω_{sn} as a capacity weight in Choquet Expected Utility. Regardless of its behavioral content, a decrease in ω_{sn} enhances uncertainty aversion by reducing the certainty equivalent of the uncertain prospect. For convenience of reference, we will call parameter ω_{sn} the decision weight even when it pertains to choice under uncertainty. At a substantive level, we remain agnostic about the exact interpretation of ω_{sn} since our study is not intended to evaluate alternative theories of uncertainty preferences.

Finally, our analysis of the time preference task is based on a discounted utility model with the quasi-hyperbolic (QH) discounting function due to Phelps and Pollak (1968). Also known as the β - δ discounting function, the QH discounting function assumes that a sure payment to be received in q periods from now is discounted by a factor of

$$\begin{aligned} D[q; \beta_{sn}, \delta_{sn}] &= 1 && \text{if } q = 0 \\ &= \beta_{sn} \delta_{sn}^q && \text{if } q > 0 \end{aligned} \tag{9}$$

where β_{sn} is a measure of non-stationary time preferences that relates to the notion of present bias, and δ_{sn} is a baseline discount factor that can be inverted to obtain a long-run discount rate.¹⁹ We code time horizon q in weeks so that equation (9) defines a weekly discount factor. If preferences are stationary ($\beta_{sn} = 1$), this function simplifies to the exponential discounting function with a constant discount factor of δ_{sn} . Typically the decision maker is said to exhibit present bias if $\beta < 1$, and future bias if $\beta > 1$.²⁰ As summarized in Section 3.4, our payment protocol introduces some discrepancy between the “now” ($q = 0$) frame

¹⁸Stahl (2014) provides a related discussion in the context of his two-parameter stochastic choice model, which is consistent with multiple theories of uncertainty preferences including Alpha-MEU.

¹⁹The baseline discount factor δ_{sn} can be seen as a measure of the sibling’s long-run delay aversion. Let d denote the implied discount rate which solves $1/(1+d)^q = \beta\delta^q$. As q tends to infinity, d tends to $1/\delta - 1$ and the effect of β vanishes.

²⁰Consider a choice between a smaller-sooner payment (SS) and a larger-later payment (LL). A present-biased individual may prefer SS to LL if SS is immediately available, but prefer LL to SS if both payments are scheduled for future dates. A future-biased individual may display the reverse choice pattern.

and the actual date of payment. To avoid making claims to incentivized identification of present or future bias, we will use “present-orientation” or “future-orientation” to describe the respective cases. Given the QH function and our experimental design that prompts a choice between a smaller payment of $M_{A,snt}$ in q_{snt} weeks and a larger payment of $M_{B,snt}$ in $q_{snt} + p_{snt}$ weeks, the preference index for each option j can be specified as a function of three preference parameters, $\theta_{sn} = \{\beta_{sn}, \delta_{sn}, r_{sn}\}$, as follows

$$\begin{aligned} V_{A,snt}[\theta_{sn}] &= D[q_{snt}; \beta_{sn}, \delta_{sn}]U[M_{A,snt}; r_{sn}] \text{ and} \\ V_{B,snt}[\theta_{sn}] &= D[q_{snt} + p_{snt}; \beta_{sn}, \delta_{sn}]U[M_{B,snt}; r_{sn}] \end{aligned} \tag{10}$$

where utility function $U[\cdot]$ uses the same curvature parameter r_{sn} as the RDU model of risk preferences. Intuitively, it is difficult to disentangle the effects of QH discounting and of utility curvature from the discounting choice tasks alone because both heavier delay discounting and more concave utility can make the smaller-sooner payment more attractive. The joint estimation of risk and time preferences, proposed by Andersen et al. (2008), helps us distinguish discounting function $D[\cdot]$ from utility function $U[\cdot]$. We complete the model specification by adopting the Fechner model of behavioral errors ($\mu_{snt} = \mu$ for all data points in the time preference experiment) and allow behavioral noise parameter μ to take a different value from the baseline noise parameters in the Contextual Utility models of risk and uncertainty preferences.

4.3 Structural estimation of heritability

The third and final building block of our econometric approach is the random coefficient model of unobserved interpersonal heterogeneity in preferences (McFadden and Train, 2000). In the random coefficient framework, each individual sibling’s preference vector θ_{sn} is modeled as a draw from a common statistical distribution, which represents the distribution of preferences across individuals in the population of twin siblings. To bring together the behavioral genetic models of heritability and behavioral econometric models of decision-making,

we first specify the structural parameters in the latter models as random coefficients. Then, we assume that each random coefficient is distributed in the population according to an ACE model, and estimate that ACE model to infer heritability in the relevant structural parameter.

Our behavioral econometric models are characterized by six structural parameters—namely r_{sn}^{risk} , ω_{sn}^{risk} , r_{sn}^{unc} , ω_{sn}^{unc} , β_{sn}^{time} and δ_{sn}^{time} —where the superscripts have been added to clarify the domains of preferences. We specify the population distribution of each parameter as an S_B distribution due to Johnson (1949). Algebraically, an S_B distribution is derived by applying a logit transformation to a normally distributed random variable. Similarly as a beta distribution, it allows the estimation results to capture a wide range of distributional features (*e.g.*, uniformity, unimodality, bimodality, and left and right skewness) without requiring us to impose any particular shape restriction *a priori* (Harrison et al., 2023). For our purposes, its connection to a normal distribution is another attractive aspect because this provides a convenient pathway to integrating the ACE model into the structural parameters.

We assume that variation in the structural parameters across individual siblings can be decomposed into variation in the genetic (A), common (C), and unique (E) environmental factors. To operationalize this assumption, we apply the ACE model in equation (3) to the primitive normal variate for each structural parameter’s population distribution. For example, consider the RDU decision weight ω_{sn}^{risk} which must fall into the unit interval. In our analysis, this parameter is specified as

$$\omega_{sn}^{risk} = \Lambda[\alpha\{\omega^{risk}\} + A_{sn}\{\omega^{risk}\} + C_n\{\omega^{risk}\} + E_{sn}\{\omega^{risk}\}] \quad (11)$$

where $\Lambda[z] = \exp[z]/(1+\exp[z])$ is the standard logistic distribution function; each summand in its argument is defined analogously as the corresponding term in equation (3); and suffix $\{.\}$ emphasizes that the values of these terms vary from structural parameter to structural parameter even within the same individual. Then, heritability in the decision weight can be

measured by evaluating the earlier formula in equation (4) at the variances of $A_{sn}\{\omega^{risk}\}$, $C_{sn}\{\omega^{risk}\}$, and $E_{sn}\{\omega^{risk}\}$, the three independent error components which make up the primitive normal variate.

More generally, consider generic structural parameter θ_{sn} which lies in interval (κ_L, κ_U) . This parameter can be specified as $\theta_{sn} = \kappa_L + (\kappa_U - \kappa_L)\Lambda[y_{sn}^*]$, where y_{sn}^* denotes a primitive normal variate which follows an ACE model specific to that parameter. We assume that the decision weight in the uncertainty experiment (ω_{sn}^{unc}) belongs to the unit interval, just as its risk counterpart (ω_{sn}^{risk}). We also place the baseline discount factor (δ_{sn}^{time}) in the unit interval, following the conventional assumption that long-run discount rates are positive. Each of utility curvature parameters r_{sn}^{risk} and r_{sn}^{unc} is assumed to fall into interval $(-10, 1)$, based on the algebraic structure of our choice tasks.²¹ Finally, the present bias parameter (β_{sn}^{time}) is assumed to fall into interval $(0.75, 1.25)$, based on parametric estimates reported by previous studies (Andersen et al., 2014; Augenblick et al., 2015).

As discussed, the canonical ACE regression model in equation (3) potentially confounds the unique environmental factor E with the idiosyncratic variation that accounts for behavioral errors. The panel ACE regression in equation (5) introduces a distinction between the two types of unobservables by assuming that the E factor remains constant within an individual across decision tasks, whereas the idiosyncratic variation is task-specific. Our econometric approach, based on equations (6) through (11), maintains this statistical distinction, and also adds a more visible distinction in terms of functional specification: Since the E factor operates via economic preference parameters that enter the model as power or multiplicative coefficients, it has non-linear effects on the latent dependent variable c_{snt}^* in equation (6), in contrast to the idiosyncratic variation that has linear effects.

Our econometric approach involves a total of 27 primitive parameters to be estimated.

²¹Consider a participant who acts on EUT preferences without making errors in decision-making. If this subject's switching points are located in the interior of our risk MPLs, we can infer that r_{sn}^{risk} belongs to a sub-interval of $(-9.05, 0.74)$. If a participant acts similarly on SEU preferences and under the belief that each ambiguous outcome is equally probable, interior switching points in our uncertainty MPLs allow us to locate r_{sn}^{unc} in the same interval.

Each ACE decomposition of a structural parameter like equation (11) entails estimation of four primitive parameters—the intercept and the variances of the three normal error components—and we consider six structural parameters. In addition, we include a distinct behavioral noise parameter (μ) for each of the risk, uncertainty, and time preference experiments. We compute the maximum simulated likelihood (MSL) estimates of the primitive parameters which maximize a simulated analogue to a sample likelihood function that pools choice observations across individuals and over choice tasks. As documented in Appendix A, the likelihood function resembles that of a multi-level mixed effects model, once choice tasks, individual siblings and twin pairs are seen as three successive levels. We adjust all standard errors and test statistics for clustering at the twin pair level, thereby addressing the highest panel dimension of our data at both the inferential and the modeling stages.²²

5 Results

5.1 Risk preferences

Our main estimates for the risk and uncertainty preference tasks are reported in Table 4. For each structural parameter of the RDU decision function, we report the shares of variance explained by the A, C, and E components as specified in equation (4). We also report estimates that constrain the common environment to be zero (the AE model) and the genetic heritability to be zero (the CE model). By definition, the A, C, and E shares fall into the $[0, 1]$ interval. We calculate 95% confidence intervals which satisfy this boundary condition by applying 10,000 repetitions of a parametric bootstrapping procedure due to Krinsky and Robb (1986). We compare the fit of the alternative specifications using the Akaike and Bayesian Information Criteria (AIC and BIC) statistics.

²²Similarly as the random intercept in the random effects probit model induces panel correlation within an individual, the ACE components of our random coefficient model induce correlation across repeated observations on the same individual as well as the same twin pair. The remark on the modeling stage refers to this aspect of our econometric approach.

Table 4: ACE share decomposition results: structural risk preference parameters

A	C	E	LL	AIC	BIC
RDU Risk preferences					
γ^{risk}					
0.36 (0.34, 0.39)	0.15 (0.13, 0.18)	0.48 (0.45, 0.52)	-7706.91	15431.83	15470.78
0.43 (0.37, 0.49)	—	0.57 (0.51, 0.63)	-7721.67	15457.34	15487.64
—	0.53 (0.39, 0.67)	0.47 (0.33, 0.61)	-7721.91	15457.83	15488.12
w^{risk}					
0.48 (0.42, 0.54)	0.08 (0.06, 0.11)	0.44 (0.39, 0.48)			
0.46 (0.36, 0.57)	—	0.54 (0.43, 0.64)			
—	0.44 (0.31, 0.58)	0.56 (0.42, 0.69)			
RDU Uncertainty preferences					
γ^{unc}					
0.52 (0.41, 0.63)	0.02 (0.00, 0.06)	0.46 (0.32, 0.58)	-7190.85	14399.71	14438.66
0.51 (0.44, 0.57)	—	0.49 (0.43, 0.56)	-7204.32	14422.63	14452.93
—	0.48 (0.43, 0.54)	0.52 (0.46, 0.57)	-7191.73	14397.46	14427.76
w^{unc}					
0.20 (0.10, 0.31)	0.26 (0.02, 0.57)	0.54 (0.32, 0.70)			
0.49 (0.38, 0.60)	—	0.51 (0.40, 0.62)			
—	0.43 (0.39, 0.48)	0.57 (0.52, 0.61)			
RDU Ambiguity preferences					
$\gamma^{unc} - \gamma^{risk}$					
0.26 (0.21, 0.30)	0.11 (0.08, 0.13)	0.64 (0.57, 0.70)	-14897.76	29831.54	29909.44
0.30 (0.25, 0.35)	—	0.70 (0.65, 0.75)	-14925.99	29879.97	29940.57
—	0.30 (0.21, 0.39)	0.70 (0.61, 0.79)	-14913.64	29855.29	29915.88
$w^{unc} - w^{risk}$					
0.34 (0.29, 0.40)	0.06 (0.04, 0.08)	0.60 (0.55, 0.65)			
0.32 (0.24, 0.43)	—	0.68 (0.57, 0.76)			
—	0.37 (0.26, 0.48)	0.63 (0.52, 0.74)			

Note: The estimation sample comprises 560 twin pairs (401 MZ and 156 DZ). Variance shares are derived from MSL regression estimates from structural choice models (see eq. 8) using choices in the decision tasks outlined in Table 2. Column A is the estimated fraction of variance explained by additive genetic effects. Column C is the estimated fraction of variance explained common environment. Column E is the estimated fraction of variance explained by unique environment. Krinsky and Robb (1986) confidence intervals in parenthesis.

The top panel of Table 4 reports the results for choice under *risk*, that is when the probability distribution of outcomes is known to the decision maker. We find that the

variance in risk preferences is explained in large part by additive genetic effects. The results indicate that genes explain 36% of the variation in the utility curvature parameter r and 48% in the probability weighting parameter ω . We also find a smaller, but non-trivial role for common family environment (15% and 8% respectively). These results imply only around 45 to 50% of variation in risk preferences is due to unique environmental experiences – substantially less than the 68% estimate from the meta-analysis (Figure 1).

The second panel reports the corresponding results for choice under *uncertainty*, that is when the probability distribution of outcomes is unknown to the decision maker. Our estimates show that the unique environment accounts for around 50% of the variance in uncertainty preferences, and this share is similar across ACE, AE, and CE models. Although the best fitting model constrains the role of genes to be zero, the amount of improvement in model fit that the resulting CE model offers over the unconstrained ACE model is relatively small.²³ We therefore focus on the ACE model, which avoids the strong assumption that genes explain none of the variation. We find that 52% of variation in the utility curvature parameter and 20% of the variation in the weighting parameter is explained by genes.

Our last set of estimates is based on the notion of ambiguity attitude similar to Abdellaoui et al. (2011). We measure ambiguity attitude in terms of the difference between structural parameters estimated for the risk and uncertainty preference tasks. As discussed in the previous section, the interpretation of such parameters depends on the theoretical model one chooses to adopt. In our case, ambiguity attitude pertains to either greater curvature of the utility function under uncertainty compared to risk ($r^{unc} - r^{risk}$) or higher weight placed on the best outcome under uncertainty compared to risk ($\omega^{unc} - \omega^{risk}$). Our estimates imply that 26% of the variation in the difference in curvature is explained by genes, and 11% by common environment. 34% of the variation in the difference in the weighting parameter is explained by genes, and 6% by common environment. This leaves a slightly larger role for unique environment in the formation of ambiguity preferences than risk preferences – the

²³According to Burnham and Anderson (2002), a difference in AIC statistics ($AIC_j - AIC_{min}$) of at least 4 is required to strongly support the alternative model (this value is only 2.25 in our setting).

Table 5: ACE share decomposition results: time preferences

A	C	E	LL	AIC	BIC
Quasi-hyperbolic discounting with RDU risk					
δ					
0.00 (0.00, 0.00)	0.33 (0.32, 0.34)	0.67 (0.66, 0.68)	-18940.12	37916.23	37994.14
0.23 (0.21, 0.25)	–	0.77 (0.75, 0.79)	-19006.79	38041.59	38102.18
–	0.33 (0.32, 0.35)	0.67 (0.65, 0.68)	-18947.80	37923.60	37984.19
β					
0.00 (0.00, 0.01)	0.33 (0.21, 0.44)	0.67 (0.55, 0.78)	-18940.12	37916.23	37994.14
0.42 (0.22, 0.63)	–	0.58 (0.37, 0.78)	-19006.79	38041.59	38102.18
–	0.31 (0.17, 0.43)	0.69 (0.57, 0.83)	-18947.80	37923.60	37984.19

Note: The estimation sample comprises 560 twin pairs (401 MZ and 156 DZ). Variance shares are derived from MSL regression estimates from structural choice models (see eq. 9) using choices in the decision tasks outlined in Table 3. Column A is the estimated fraction of variance explained by additive genetic effects. Column C is the estimated fraction of variance explained common environment. Column E is the estimated fraction of variance explained by unique environment. Krinsky and Robb (1986) confidence intervals in parenthesis.

E component accounts for 60 ($\omega^{unc} - \omega^{risk}$) to 64% ($r^{unc} - r^{risk}$) of ambiguity preferences compared to 44 (ω^{risk}) to 48% (r^{risk}) of risk preferences.

5.2 Time preferences

Our main results for time preferences are presented in Table 5. Here we see that under the best fitting models (CE and ACE depending on the criteria), the variation due to genes is estimated to be zero for both long-run discounting factor δ and future orientation β , while the common environment explains as much as 1/3 of the variation in both parameters.

The lack of genetic heritability in time preferences that we have found contrasts with our meta-analysis of the existing studies. As reported in Section 2, the previous estimates suggest that genetics accounts for 38% of variation in choices made in delay discounting (Figure 3). The discount utility model, such as equation (10), offers a structural explanation that may reconcile these two sets of findings: The choice between a smaller-sooner reward and a larger-later reward depends not only on the discounting function that captures time

preferences, but also on the utility function that transforms the raw reward amounts into undiscounted utility flows. Thus, as long as the utility function exhibits genetic variation, so will the observed choices between the two rewards, regardless of whether the discounting function exhibits genetic variation. By contrast, since our joint estimation approach formally distinguishes between the discounting function and the utility function, the results in Table 5 only capture genetic variation in the discounting function. We return to this issue when we discuss related findings in Figure 8.

5.3 Alternative measurement models

In the previous section, we presented results using our most general behavioral choice models, which allowed for utility curvature, probability weighting, delay aversion, and non-stationary time preferences. In this section, we consider how our results change when we adopt simpler behavioral models—such as expected utility and exponential discounting models—that omit one or more of these features. We also report estimates using a standard regression approach comparable to the earlier literature, where risk and time preferences are equated with the count of safe and sooner choices in raw data.

Figure 5 summarizes the ACE decompositions of risk preferences using alternative estimation approaches.²⁴ For comparability, we only present results from the full ACE specifications. The estimates ‘MPL 1’ and ‘MPL 2’ are based on the number of safe choices in the first and second MPL respectively (see Table 2). For those outcomes, we estimate the ACE shares using the standard ACE regression model in equation (3) that does not draw a distinction between the unique environment and the idiosyncratic behavioral/measurement errors. The estimates ‘Panel’ refer to the panel ACE regression model in equation (5), where the two counts of safe choices are treated as repeated measures and the unique environment is distinguished from the idiosyncratic errors by the usual random effects assumption. Since the two MPLs are not exactly the same, the panel specification includes MPL fixed effects to


²⁴The estimates are also reported in Table C.3.

improve comparability between the two count measures. Comparing the panel estimates to the MPL-by-MPL estimates allows us to gauge the importance of measurement error when the ACE decomposition is obtained from a non-structural regression model. The ‘RDU’ estimates are our structural estimates discussed in Sections 5.1 and 5.2, and the ‘EUT’ estimates are obtained by estimating a special case of the RDU model that assumes no probability distortion.

The results in Figure 5 call for caution against a naive dichotomy between regression-based and structural estimation of heritability. The results of a restricted structural model which neglects the joint presence of alternative behavioral phenomena (here EUT which assumes no probability distortion and equates risk preferences with utility curvature) are more comparable to the standard and/or panel ACE regression models than to our general structural model (RDU which accommodates the joint effects of utility curvature and probability weighting on choice under risk). None of our standard regression specifications indicates a significant role for genes in explaining risk preferences; instead, 85-90% of variation is attributed to unique environment and a small share to common environment. The panel regression model, which accounts for measurement error, slightly reduces the role of the unique environment, as expected. We obtain similar results in the EUT specification. However, a more general decision function that allows for probability weighting overturns these findings; indicating a large role for genes and the unique environment explaining less than 50% of the variation. Compared to EUT, the RDU model has considerably better fit (the log-likelihood increases from -9965.98 to -7706.91).


The corresponding decompositions of uncertainty preferences yield a similar pattern of findings (see Figure 6). Across standard and panel regression specifications, we estimate heritability between 10-20%, with around 70 to 80% of the variance explained by unique environment. The EUT estimates agree that the unique environment explains around 80% of the variance, though it attributes most of the remaining variation to the common environment instead of genes. In contrast, the RDU estimates indicate that the unique environment

Figure 5: Alternative estimates of ACE shares: Risk preferences



Note: The estimation sample comprises 560 twin pairs (401 MZ and 156 DZ). Variance shares are derived from regression estimates using choices in the decision tasks outlined in Table 2 where probability distributions are known. A is the estimated fraction of variance explained by additive genetic effects. C is the estimated fraction of variance explained common environment. E is the estimated fraction of variance explained by unique environment. *MPL 1-MPL 2*: estimates from the classical ACE model with the number of ‘safe’ choices in the respective MPL used as the dependent variable. *Panel*: estimates from the classical ACE model with repeat measures with the number of ‘safe’ choices in MPL tasks used as the dependent variable. The panel regression equation includes a separate intercept for each task and individual intercepts modeled as a normally distributed random parameter. The variance of this parameter is interpreted as measurement error and subtracted from the overall variance when estimating the A, C and E shares. *EUT*: MSL regression estimates from a structural choice model using choices in both MPL tasks assuming a single parameter expected utility theory decision function with constant relative risk aversion (see eq. 7). *RDU_r* and *RDU_w*: MSL regression estimates from a structural choice model using choices in both MPL tasks assuming a two parameter rank-dependent utility theory decision function (see eq. 8). *r* is utility curvature and *w* is probability weighting. Error bars are Krinsky and Robb (1986) confidence intervals.

Figure 6: Alternative estimates of ACE shares: Uncertainty preferences




Note: See Figure 5. The only difference is that here we use the MPL tasks with unknown (rather than known) probability distribution of outcomes.

accounts for less than 50% of the variance in either uncertainty preference parameter, with genes explaining 52% and 20% of variations in utility curvature and probability weighting, respectively.

We next examine ambiguity preferences in Figure 7. The standard ACE regression measures ambiguity attitude in terms of sample choice frequencies by taking the difference in the number of safe choices between risk and uncertainty versions of MPL 1 or MPL 2. The results indicate a small role for genes (around 5%), with the rest of the variation attributed to the unique environment. In this case, using the panel ACE regression to deal with measurement error has a marked impact, increasing the role of genes to 24% and reducing the role of the unique environment by the same amount. The EUT model agrees with the standard ACE regression that genes play virtually no role. By contrast the more general RDU model

Figure 7: Alternative estimates of ACE shares: Ambiguity preferences




Note: See Figure 5. The only difference is that here as the dependent variable in *MPL 1-MPL 2* and *Panel* we use the difference between the number of ‘safe’ choices in the respective MPL when probabilities are unknown versus known.

indicates that genes explain as much as 34% of the variation in ambiguity preferences.

Figure 8 displays alternative decompositions of time preferences. In the standard ACE regression (MPL 1 to MPL 4), this is measured by the number of sooner choices in each MPL. As earlier, the panel ACE regression (‘Panel’) allows for MPL fixed effects but otherwise treats the four counts as repeat measures. We present estimates from four different structural models: ‘EXP’ assumes a simple exponential discounting function and a linear utility function; ‘EUTEXP’ assumes a simple exponential discounting function but estimates this jointly with risk preferences under EUT to control for utility curvature; ‘QH’ assumes quasi-hyperbolic discounting and linear utility; and finally ‘QHRDU’ (our most general specification) jointly estimates the quasi-hyperbolic discounting function along with

Figure 8: Alternative estimates of ACE shares: Time preferences (Delay aversion)




Note: Variance shares are derived from regression estimates using choices in the decision tasks outlined in Table 3. *MPL 1-MPL 4*: estimates from the classical ACE model with the number of ‘sooner’ choices in the respective MPL used as the dependent variable. *Panel*: estimates from the classical ACE model with repeat measures with the number of ‘sooner’ choices in MPL tasks used as the dependent variable. The regression equation includes a separate intercept for each task and individual intercepts modelled as a normally distributed random parameter. The variance of this parameter is interpreted as measurement error and subtracted from the overall variance when estimating the A, C and E shares. *EXP*: MSL regression estimates from a structural choice model using choices in all MPL tasks assuming an exponential discount function (i.e., eq. 9 with $\beta = 1$) with linear utility. *EUTEXP*: MSL regression estimates from a structural choice model using choices in all MPL tasks assuming an exponential discount function (i.e., eq. 9 with $\beta = 1$) with constant relative risk aversion utility (i.e., eq. 7) estimated from risky choice MPLs. *QH*: MSL regression estimates from a structural choice model using choices in all MPL tasks assuming a quasi-hyperbolic discount function (i.e., eq. 9) with linear utility. *RDUQU*: MSL regression estimates from a structural choice model using choices in all MPL tasks assuming a quasi-hyperbolic discount function (i.e., eq. 9) with rank dependent utility (see eq. 8) estimated from risky choice MPLs. Error bars are Krinsky and Robb (1986) confidence intervals. See Figure 5 for further details.

utility curvature and probability weighting under RDU.

We continue to observe the importance of adopting a general model specification which recognizes that an individual’s response to a decision task is jointly influenced by several latent behavioral phenomena. The standard and panel ACE regression models attribute around 20% of variation to genes and 80% to unique environment. However, in the structural specifications the loading moves from genes to the common environment. Interestingly, this is not the case for the quasi-hyperbolic decision function when we *do not* control for utility curvature (‘QH’). One interpretation of this is that genes do matter for the choice between a smaller-sooner reward and a later-larger reward, but they operate through the utility function (which is shared between risk and time preferences) instead of the discounting function (which is a direct measure of time preferences in relation to delay discounting). We also see that the role of the unique environment is smallest when we use our most general structural specification (‘RDUQH’).

Figure 9 reports estimates for the measures of non-stationary time preferences. In the context of the standard and panel ACE regression models, this measure refers to the difference in the number of sooner choices between MPLs with and without front-end delays. The estimates ‘QH’ and ‘RDUQH’ refer to the β parameter in the quasi-hyperbolic discounting function—also known as the β - δ discounting function—based on the specifications described above. Again, the results caution against a naive regression versus structural dichotomy. The standard ACE regression indicates the dominant role of the unique environment, which accounts for more than 90% of the variance. Both the panel ACE regression and the QH structural model indicate a somewhat diminished role of the unique environment at 80%, with the remaining variation split between genes and the common environment. Once we control for utility curvature by adopting the general RDUQH specification, we find yet a smaller role for the unique environment (67%) and more a prominent role for the common environment (33%).

Figure 9: Alternative estimates of ACE shares: Time preferences (Non-stationarity)



Note: Variance shares are derived from regression estimates using choices in the decision tasks outlined in Table 3. *MPL i - MPL j* : estimates from the classical ACE model with the difference in the number of ‘sooner’ choices between MPL i and j used as the dependent variable. *Panel*: estimates from the classical ACE model with repeat measures with the difference in the number of ‘sooner’ choices between MPL i and j used as the dependent variable. The regression equation includes a separate intercept for each task and individual intercepts modelled as a normally distributed random parameter. The variance of this parameter is interpreted as measurement error and subtracted from the overall variance when estimating the A, C and E shares. *QH* and *RDUQH*: see Figure 8 (estimates correspond to the future orientation parameters). Error bars are Krinsky and Robb (1986) confidence intervals. See Figure 5 for further details.

6 Discussion

Risk and time preferences are ubiquitous in economic research. In economic theory, structural parameters representing these preferences are fundamental primitives that drive individual decision making and influence the welfare consequences of changes in the state of the world. Macroeconomists employ the estimates of these preference parameters as inputs to

operationalize their computational models of the economy. Empirical economists use the structural and proxy measures of risk and time preferences as key control variables to model insurance, occupational, and health-related choices. Behavioral economists believe that understanding the determinants of risk and time preferences is necessary for effective policy that needs to take into account how malleable preferences are.

We quantify the influence of nature and nurture on economic preferences by developing a novel econometric approach to measuring heritability. Ours is the first paper to marry structural estimation of economic preference parameters with the canonical twin study method that decomposes variation in individual traits into genetic and environmental sources (the ACE model). Structural estimation has been widely embraced in the economic analysis of decision-making, and the structural parameter estimates have many desirable properties, including direct mapping to economic theory. Importantly, these parameters have a clear quantitative interpretation and are comparable across studies.

Our analysis highlights three major limitations in the existing approaches to estimating heritability in economic preferences. First, accurate estimation of heritability needs to properly account for the behavioral and measurement errors which can be substantial in preference elicitation tasks (Hey and Orme, 1994; Gillen et al., 2019). Importantly, these errors are indistinguishable from the notion of unique environment in the standard ACE regression model that is frequently used to establish heritability. This may lead to over-estimation of the role of unique environment. Structural estimation deals with these concerns by explicitly considering choice errors as part of the data generating process, so that they are effectively controlled for when preferences are estimated. Second, accurate estimation of the heritability needs to use preference measures that are as close to their intended economic interpretations and uses as possible. Equating raw survey or experimental response measures with risk or time preferences is unsatisfactory in this regard since it is often difficult, if not impossible, to infer the precise values of relevant preference parameters from such measures. Finally, the regression analysis with just a single response variable, by construction, makes

it difficult to accommodate the joint effects of well-documented behavioral phenomena, such as probability distortion, utility curvature, and non-stationarity of time preferences, which simultaneously shape the individual’s risk and time preferences.

Indeed, in our preferred structural model of preferences, we find that the importance of genes to risk attitudes is substantially greater, and that of the unique environment smaller, than what we observe in our meta-analysis of the standard ACE estimates. When we apply a panel data extension of the ACE regression to account for measurement errors without a further departure from the standard approach, we also find a diminished role of the unique environment but the difference from the standard estimates is rather small. This suggests that controlling for measurement errors from a statistical perspective and modeling a structural decision making process from an economic perspective both make contributions to accurate estimation of heritability. The second channel is particularly important to the identification of time preferences (Andersen et al., 2008), and we speculate that the relatively high meta-analytic estimate of the heritability of time preferences (38% with a wide 95% CI of [16%,59%]) is driven by the heritability of risk preferences, which affects raw responses in time preference tasks via the utility function. In accordance with this view, when we control for utility curvature, our structural estimates of the genetic influence on delay discounting are practically 0%. When we estimate a restricted structural model that does not control for the utility curvature, our estimate is in the 27% range, closer to the meta-analytic estimate.

At first glance, our findings suggest that risk and time preferences are influenced more strongly by each individual’s own interaction with a wider society than by their genetic lotteries or upbringing. Across all of our empirical estimates and also in our meta-analysis, the unique environment—which captures individual-specific variation that is not shared between twin siblings—is always the most influential source of variation in people’s preferences. For delay discounting, our preferred estimates even indicate no role of genes. Nevertheless, on top of finding greater heritability in risk preferences compared to existing studies, our structural analysis suggests that time preferences may be more malleable than risk preferences,

an important insight for policy making. When reviewing previous literature, we have found that the common twin environment—which captures non-genetic variation that is shared between twin siblings—has negligible effects on shaping individual preferences. Is the home environment indeed irrelevant and can we ditch all of our parenting books? Our preferred structural models, accounting for non-expected utility and non-stationary discounting behavior, paint a slightly different picture than the previous studies. We estimate the role of common environment to be 15% for risk preferences and 33% for time preferences, both of which are much larger than previously estimated. The fact that the common environment has twice as large of an effect on shaping time preferences as risk preferences offers another important policy and parenting insight.

From an empirical perspective, our paper provides new benchmark evidence on the formation of economic preferences. From a methodological perspective, our paper demonstrates the importance of how preferences are defined and measured. An understanding of the anatomy of risk and time preferences has clear value for applied and theoretical research, considering that most economic studies relate to the allocation of resources over time under risk and uncertainty. Admittedly the twin study method, on its own, does not identify specific drivers of preference heterogeneity—say, particular genetic variants or transformative experiences. It does, nevertheless, provide bounds on how much variation those drivers are likely to explain, which can help to direct and frame future research efforts. There is an undeniable human curiosity in understanding why we are the way that we are. This may even affect how we interpret others’ behaviors, and interact with them. For example, it might be easier to empathize with gambling addiction as a chronic health condition if we view that behavior in light of genetic effects on preferences. Studies like ours can help to contextualize the social worlds we exist in.

References

- Abdellaoui, M., A. Baillon, L. Placido, and P. P. Wakker (2011). The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review* 101(2), 695–723.
- Alan, S., N. Baydar, T. Boneva, T. F. Crossley, and S. Ertac (2017). Transmission of risk preferences from mothers to daughters. *Journal of Economic Behavior & Organization* 134, 60–77.
- Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2008). Eliciting risk and time preferences. *Econometrica* 76(3), 583–618.
- Andersen, S., G. W. Harrison, M. I. Lau, and E. E. Rutström (2014). Discounting behavior: A reconsideration. *European Economic Review* 71, 15–33.
- Anderson, L. R. and J. M. Mellor (2008). Predicting health behaviors with an experimental measure of risk preference. *Journal of Health Economics* 27(5), 1260–1274.
- Angerer, S., D. Glätzle-Rützler, P. Lergetporer, and M. Sutter (2015). Donations, risk attitudes and time preferences: A study on altruism in primary school children. *Journal of Economic Behavior & Organization* 115, 67–74.
- Anokhin, A. P., S. Golosheykin, J. D. Grant, and A. C. Heath (2009). Heritability of risk-taking in adolescence: a longitudinal twin study. *Twin Research and Human Genetics* 12(4), 366–371.
- Anokhin, A. P., S. Golosheykin, J. D. Grant, and A. C. Heath (2011). Heritability of delay discounting in adolescence: a longitudinal twin study. *Behavior Genetics* 41(2), 175–183.
- Anokhin, A. P., J. D. Grant, R. C. Mulligan, and A. C. Heath (2015). The genetics of impulsivity: evidence for the heritability of delay discounting. *Biological Psychiatry* 77(10), 887–894.
- Apesteguia, J. and M. A. Ballester (2018). Monotone stochastic choice models: The case of risk and time preferences. *Journal of Political Economy* 126(1), 74–106.
- Augenblick, N., M. Niederle, and C. Sprenger (2015). Working over time: Dynamic inconsistency in real effort tasks. *The Quarterly Journal of Economics* 130(3), 1067–1115.
- Ayhan, S. H., K. Gatskova, and H. Lehmann (2020). The impact of non-cognitive skills and risk preferences on rural-to-urban migration in ukraine. *Journal of Comparative Economics* 48(1), 144–162.
- Balakrishnan, U., J. Haushofer, and P. Jakiela (2020). Thow soon is now? evidence of present bias from convex time budget experiments. *Experimental Economics* 23, 294–321.
- Barnea, A., H. Cronqvist, and S. Siegel (2010). Nature or nurture: What determines investor behavior? *Journal of Financial Economics* 98(3), 583–604.

- Benjamin, D. J., D. Cesarini, M. J. Van Der Loos, C. T. Dawes, P. D. Koellinger, P. K. Magnusson, C. F. Chabris, D. Conley, D. Laibson, M. Johannesson, et al. (2012). The genetic architecture of economic and political preferences. *Proceedings of the National Academy of Sciences* 109(21), 8026–8031.
- Blanden, J. (2013). Cross-country rankings in intergenerational mobility: A comparison of approaches from economics and sociology. *Journal of Economic Surveys* 27(1), 38–73.
- Bradford, D., C. Courtemanche, G. Heutel, P. McAlvanah, and C. Ruhm (2017). Time preferences and consumer behavior. *Journal of Risk and Uncertainty* 55, 119–145.
- Burnham, K. P. and D. R. Anderson (2002). *Model selection and multimodel inference : a practical information-theoretic approach*. Springer New York.
- Bush, W. S. and J. H. Moore (2012). Chapter 11: Genome-wide association studies. *PLoS Computational Biology* 8(12), e1002822.
- Cesarini, D., C. T. Dawes, M. Johannesson, P. Lichtenstein, and B. Wallace (2009, 05). Genetic Variation in Preferences for Giving and Risk Taking. *The Quarterly Journal of Economics* 124(2), 809–842.
- Cesarini, D., M. Johannesson, P. K. E. Magnusson, and B. Wallace (2012). The behavioral genetics of behavioral anomalies. *Management Science* 58(1), 21–34.
- Cesarini, D. and P. M. Visscher (2017). Genetics and educational attainment. *npj Science of Learning* 2(1), 1–7.
- Chabris, C. F., J. J. Lee, D. Cesarini, D. J. Benjamin, and D. I. Laibson (2015). The fourth law of behavior genetics. *Current Directions in Psychological Science* 24(4), 304–312.
- Cheung, S., A. Tymula, and X. Wang (2022). Present bias for monetary and dietary rewards: Evidence from chinese teenagers. *Experimental Economics* 25, 1202–1233.
- Chowdury, S., M. Sutter, and K. F. Zimmerman (2022). Economic Preferences across Generations and Family Clusters: A Large-Scale Experiment in a Developing Country. *Journal of Political Economy* 130, 2361–2410.
- Conley, D. and J. Fletcher (2017). *The Genome Factor*. Princeton University Press.
- Corak, M. (2013). Income inequality, equality of opportunity, and intergenerational mobility. *Journal of Economic Perspectives* 27(3), 79–102.
- Cortés, P., J. Pan, E. Reuben, L. Pilossoph, and B. Zafar (2021, May). Gender differences in job search and the earnings gap: Evidence from the field and lab. Working Paper 28820, National Bureau of Economic Research.
- Courtemanche, C., G. Heutel, and P. McAlvanah (2015). Impatience, incentives and obesity. *The Economic Journal* 125(582), 1–31.

- Cronqvist, H. and S. Siegel (2014). The genetics of investment biases. *Journal of Financial Economics* 113(2), 215–234.
- Cronqvist, H. and S. Siegel (2015). The origins of savings behavior. *Journal of Political Economy* 123(1), 123–169.
- Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2012). The intergenerational transmission of risk and trust attitudes. *The Review of Economic Studies* 79(2), 645–677.
- Eckel, C. C. and P. J. Grossman (2002). Sex differences and statistical stereotyping in attitudes toward financial risk. *Evolution and Human Behavior* 23(4), 281–295.
- Ellsberg, D. (1961). Risk, ambiguity, and the Savage axioms. *The Quarterly Journal of Economics*, 643–669.
- Epper, T., E. Fehr, H. Fehr-Duda, C. T. Kreiner, D. D. Lassen, S. Leth-Petersen, and G. N. Rasmussen (2020). Time discounting and wealth inequality. *American Economic Review* 110(4), 1177–1205.
- Etner, J., M. Jeleva, and J.-M. Tallon (2012). Decision theory under ambiguity. *Journal of Economic Surveys* 26(2), 234–270.
- Fouarge, D., B. Kriechel, and T. Dohmen (2014). Occupational sorting of school graduates: The role of economic preferences. *Journal of Economic Behavior & Organization* 106, 335–351.
- Friedman, N. P., M. T. Banich, and M. C. Keller (2021). Twin studies to GWAS: there and back again. *Trends in Cognitive Sciences* 25(10), 855–869.
- Ge, T., A. J. Holmes, R. L. Buckner, J. W. Smoller, and M. R. Sabuncu (2017). Heritability analysis with repeat measurements and its application to resting-state functional connectivity. *Proceedings of the National Academy of Sciences* 114(21), 5521–5526.
- Gillen, B., E. Snowberg, and L. Yariv (2019). Experimenting with measurement error: Techniques with applications to the caltech cohort study. *Journal of Political Economy* 127(4), 1826–1863.
- Gneezy, U. and J. Potters (1997). An experiment on risk taking and evaluation periods. *The Quarterly Journal of Economics* 112(2), 631–645.
- Golsteyn, B. H., H. Grönqvist, and L. Lindahl (2014). Adolescent time preferences predict lifetime outcomes. *The Economic Journal* 124(580), F739–F761.
- Halevy, Y. (2015). Time consistency: stationarity and time invariance. *Econometrica* 83(1), 335–352.
- Halliday, T., B. Mazumder, and A. Wong (2021). Intergenerational mobility in self-reported health status in the US. *Journal of Public Economics* 193, 104307.

- Harden, K. P. (2021). *The Genetic Lottery: Why DNA Matters for Social Equality*. Princeton University Press.
- Harden, K. P., N. Kretsch, F. D. Mann, K. Herzhoff, J. L. Tackett, L. Steinberg, and E. M. Tucker-Drob (2017). Beyond dual systems: A genetically-informed, latent factor model of behavioral and self-report measures related to adolescent risk-taking. *Developmental Cognitive Neuroscience* 25, 221–234. Sensitive periods across development.
- Harrison, G. W. (2019). The methodologies of behavioral econometrics. *Contemporary Philosophy and Social Science: An Interdisciplinary Dialogue*, 107.
- Harrison, G. W., M. I. Lau, and E. E. Rutström (2010). Individual discount rates and smoking: Evidence from a field experiment in denmark. *Journal of Health Economics* 29(5), 708–717.
- Harrison, G. W., M. I. Lau, and H. I. Yoo (2023). Constant discounting, temporal instability and dynamic inconsistency in denmark: A longitudinal field experiment. CEAR Working Paper 2020-09.
- Heinrich, T. and J. Shachat (2020). The development of risk aversion and prudence in Chinese children and adolescents. *Journal of Risk and Uncertainty* 61(3), 263–287.
- Hey, J. D., G. Lotito, and A. Maffioletti (2010). The descriptive and predictive adequacy of theories of decision making under uncertainty/ambiguity. *Journal of Risk and Uncertainty* 41(2), 81–111.
- Hey, J. D. and C. Orme (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica: Journal of the Econometric Society* 62, 1291–1326.
- Hubler, P. M. A. (2018). Heritability of time preference: Evidence from german twin data. *Kyklos* 71(3), 433–455.
- Jaeger, D. A., T. Dohmen, A. Falk, D. Huffman, U. Sunde, and H. Bonin (2010). Direct evidence on risk attitudes and migration. *The Review of Economics and Statistics* 92(3), 684–689.
- Johnson, N. L. (1949). Systems of frequency curves generated by methods of translation. *Biometrika* 36(1/2), 149–176.
- Kettlewell, N. and A. Tymula (2021). The Australian Twins Economic Preferences Survey. *Forthcoming*, 1–6.
- Kothiyal, A., V. Spinu, and P. P. Wakker (2014). An experimental test of prospect theory for predicting choice under ambiguity. *Journal of Risk and Uncertainty* 48(1), 1–17.
- Krinsky, I. and A. L. Robb (1986). On approximating the statistical properties of elasticities. *Review of Economics and Statistics* 68(4), 715–719.
- Lepinteur, A., L. G. Borga, A. E. Clark, C. Vögele, and C. D’Ambrosio (2023). Risk aversion and covid-19 vaccine hesitancy. *Health Economics Forthcoming*.

- Linnér, R. K., P. Biroli, E. Kong, S. F. W. Meddens, R. Wedow, M. A. Fontana, M. Lebreton, S. P. Tino, A. Abdellaoui, A. R. Hammerschlag, et al. (2019). Genome-wide association analyses of risk tolerance and risky behaviors in over 1 million individuals identify hundreds of loci and shared genetic influences. *Nature Genetics* 51(2), 245–257.
- Maher, B. (2008). Personal genomes: The case of the missing heritability. *Nature* 456(7218), 18–22.
- McFadden, D. and K. Train (2000). Mixed mnl models for discrete response. *Journal of Applied Econometrics* 15(5), 447–470.
- Murphy, K., J. Lam, T. Cutler, J. Tyler, L. Calais-Ferreira, S. Li, C. Little, P. Ferreira, J. M. Craig, K. J. Scurrah, and et al. (2019). Twins Research Australia: A new paradigm for driving twin research. *Twin Research and Human Genetics* 22(6), 438–445.
- Neale, M. C. and M. B. Miller (1997). The use of likelihood-based confidence intervals in genetic models. *Behavior genetics* 27, 113–120.
- Nicolaou, N. and S. Shane (2020). Common genetic effects on risk-taking preferences and choices. *Journal of Risk and Uncertainty* 59(3), 261–279.
- Norrgrén, L. (2022). Time preference, illness, and death. *Journal of Health Economics* 86, 102692.
- Ooki, S., K. Yamada, A. Asaka, and K. Hayakawa (1990). Zygosity diagnosis of twins by questionnaire. *Acta Geneticae Medicae et Gemellologiae: Twin Research* 39(1), 109–115.
- Phelps, E. S. and R. A. Pollak (1968). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies* 35(2), 185–199.
- Polderman, T. J., B. Benjamin, C. A. De Leeuw, P. F. Sullivan, A. Van Bochoven, P. M. Visscher, and D. Posthuma (2015). Meta-analysis of the heritability of human traits based on fifty years of twin studies. *Nature Genetics* 47(7), 702–709.
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization* 3(4), 323–343.
- Sanchez-Roige, S., P. Fontanillas, S. L. Elson, 23andMe Research Team, A. Pandit, E. M. Schmidt, J. R. Foerster, G. R. Abecasis, J. C. Gray, H. de Wit, et al. (2018). Genome-wide association study of delay discounting in 23,217 adult research participants of European ancestry. *Nature Neuroscience* 21(1), 16–18.
- Simonson, I. and A. Sela (2010). On the Heritability of Consumer Decision Making: An Exploratory Approach for Studying Genetic Effects on Judgment and Choice. *Journal of Consumer Research* 37(6), 951–966.
- Sparks, J. C., J. D. Isen, and W. G. Iacono (2014). Preference on cash-choice task predicts externalizing outcomes in 17-year-olds. *Behavior Genetics* 44(2), 102–112.

- Stahl, D. O. (2014). Heterogeneity of ambiguity preferences. *Review of Economics and Statistics* 96(4), 609–617.
- Torche, F. (2015). Analyses of intergenerational mobility: An interdisciplinary review. *The ANNALS of the American Academy of Political and Social Science* 657(1), 37–62.
- Uffelmann, E., Q. Q. Huang, N. S. Munung, J. de Vries, Y. Okada, A. R. Martin, H. C. Martin, T. Lappalainen, and D. Posthuma (2021). Genome-wide association studies. *Nature Reviews Methods Primers* 1(1), 1–21.
- Wainschein, P., D. Jain, Z. Zheng, T. A. W. Group, N. T.-O. for Precision Medicine (TOPMed) Consortium, L. A. Cupples, A. H. Shadyab, B. McKnight, B. M. Shoemaker, B. D. Mitchell, B. M. Psaty, C. Kooperberg, C.-T. Liu, C. M. Albert, D. Roden, D. I. Chasman, D. Darbar, D. M. Lloyd-Jones, D. K. Arnett, E. A. Regan, E. Boerwinkle, J. I. Rotter, J. R. O’Connell, L. R. Yanek, M. de Andrade, M. A. Allison, M.-L. N. McDonald, M. K. Chung, M. Fornage, N. Chami, N. L. Smith, P. T. Ellinor, R. S. Vasan, R. A. Mathias, R. J. Loos, S. S. Rich, S. A. Lubitz, S. R. Heckbert, S. Redline, X. Guo, Y.-D. I. Chen, C. A. Laurie, R. D. Hernandez, S. T. McGarvey, M. E. Goddard, C. C. Laurie, K. E. North, L. A. Lange, B. S. Weir, L. Yengo, J. Yang, and P. M. Visscher (2021). Recovery of trait heritability from whole genome sequence data. *bioRxiv*.
- Wilcox, N. T. (2011). ‘stochastically more risk averse:’ a contextual theory of stochastic discrete choice under risk. *Journal of Econometrics* 162(1), 89–104.
- Wong, R. S., M. Osman, W. H. Wong, Y. Lin, and K. Ho (2019). Saving for a better retirement: How risk attitudes affect choice of retirement scheme. *Psychological Reports* 122(1), 305–322.
- Yang, J., J. Zeng, M. E. Goddard, N. R. Wray, and P. M. Visscher (2017). Concepts, estimation and interpretation of SNP-based heritability. *Nature Genetics* 49(9), 1304–1310.
- Zhong, S., S. H. Chew, E. Set, J. Zhang, H. Xue, P. C. Sham, R. P. Ebstein, and S. Israel (2009). The heritability of attitude toward economic risk. *Twin Research and Human Genetics* 12(1), 103–107.
- Zyphur, M. J., J. Narayanan, R. D. Arvey, and G. J. Alexander (2009). The genetics of economic risk preferences. *Journal of Behavioral Decision Making* 22(4), 367–377.

Appendix A. Sample likelihood function

This appendix summarizes the algebraic structure of our sample likelihood functions. To convey the main ideas without excessive notational cluttering, we focus on an illustrative example that considers structural estimation of utility curvature parameter r_{sn} under EUT. Continuing with the notation defined in Section 4 of the main text, this theory of risk preferences is derived by assuming that decision weight ω_{sn} in the risk preference experiment equals the objective probability of 0.5 for all sibling $s \in \{1, 2\}$, twin pair $n \in \{1, 2, \dots, N\}$ and choice task $t \in \{1, 2, \dots, T\}$.

Let c_{snt} denote a binary indicator that equals 1 if sibling s of pair n selected Option A in task t and 0 if Option B. Conditional on structural parameter r_{sn} and behavioral noise parameter μ , the likelihood of this choice observation can be specified as

$$P_{snt}[r_{sn}, \mu] = \Lambda[\Delta V_{snt}[r_{sn}]/\mu_{snt}]^{c_{snt}} \Lambda[-\Delta V_{snt}[r_{sn}]/\mu_{snt}]^{(1-c_{snt})} \text{ and} \quad (\text{A1})$$

$$\mu_{snt} = (U[M_{A,snt}; r_{sn}] - U[0; r_{sn}])\mu$$

where $\Delta V_{snt}[\cdot] = V_{A,snt}[\cdot] - V_{B,snt}[\cdot]$ denotes an expected utility difference, which is derived by evaluating the RDU index functions in equation (8) at $\omega_{sn} = 0.5$, and other notations are defined in Section 4.

Our econometric approach considers structural parameter r_{sn} as a draw from an S_B distribution, which in turn is generated from an ACE model. We can represent this chain of modeling assumptions by writing $r_{sn} = r[\alpha, A_{sn}, C_n, E_{sn}]$, where function $r[\cdot]$ denotes a mapping of the ACE components to the structural parameter.²⁵ Recall that unique environmental factor E_{sn} is uncorrelated between two siblings of the same twin pair; common environmental factor C_n displays perfect positive correlation; and additive genetic factor A_{sn} displays perfect or imperfect positive correlation, depending on whether the twins are identical (MZ) or non-identical (DZ). Accordingly, we initially write out a sibling-level like-

²⁵To state $r[\cdot]$ explicitly, suppose that r_{sn} falls into interval (κ_L, κ_U) . Then, $r[\alpha, A_{sn}, C_n, E_{sn}] = \kappa_L + (\kappa_U - \kappa_L)\Lambda[\alpha + A_{sn} + C_n + E_{sn}]$.

likelihood function which is conditional on the knowledge of error components shared by both siblings— C_n and part of A_{sn} that is correlated between the siblings—before progressing to integrate out these two error components to derive a twin pair-level likelihood function.

Consider first MZ twins who share the same genetic factor so that $A_{1s} = A_{2s} = A_n$. Conditional on shared error components A_n and C_n , the joint likelihood of T choice observations on sibling s of MZ pair n is given by

$$l_{sn}^{MZ}[\alpha, \sigma_E^2, \mu; A_n, C_n] = \int \prod_{t=1}^T P_{snt}[r[\alpha, A_n, C_n, E_{sn}], \mu] f[E_{sn}|0, \sigma_E^2] dE_{sn} \quad (\text{A2})$$

where $f[\cdot|0, \sigma_E^2]$ denotes the density function for a normally distributed random variable with mean 0 and variance σ_E^2 . Apart from the conditioning on A_n and C_n , equation (A2) represents what an individual-level likelihood function for the random coefficient EUT model looks like in typical applications that do not consider twins data.²⁶ We can derive the joint likelihood of $2 \times T$ choice observations on pair n as follows

$$L_s^{MZ}[\alpha, \sigma_A^2, \sigma_C^2, \sigma_E^2, \mu] = \int \int l_{1s}^{MZ}[\alpha, \sigma_E^2, \mu; A_n, C_n] l_{2s}^{MZ}[\alpha, \sigma_E^2, \mu; A_n, C_n] \times f[A_n|0, \sigma_A^2] f[C_n|0, \sigma_C^2] dA_n dC_n \quad (\text{A3})$$

by integrating out the shared error components.

Consider DZ twins who only share half the genetic variation, in the sense of $\text{cov}[A_{1s}, A_{2s}] = \sigma_A^2/2$. To operationalize the implied imperfect correlation, we specify A_{sn} as a sum of two independent normal error components: $A_{sn} = a_{sn} + b_n$, where $a_{sn} \sim N(0, \sigma_A^2/2)$ account for genetic influences unique to sibling s and $b_n \sim N(0, \sigma_A^2/2)$ account for genetic influences which are shared by both siblings. This decomposition does not change the overall distribution of A_{sn} , which is normal with mean zero and variance σ_A^2 as it is for MZ twins. Conditional on shared error components b_n and C_n , the joint likelihood of T choice observa-

²⁶Put another way, suppose that both A_n and C_n take values of zero. Then, equation (A2) coincides with an individual-level likelihood function in such applications.

tions on sibling s of DZ pair n is given by

$$l_{sn}^{DZ}[\alpha, \sigma_a^2, \sigma_E^2, \mu; b_n, C_n] = \int \int \prod_{t=1}^T P_{snt}[r[\alpha, (a_{sn} + b_n), C_n, E_{sn}], \mu] \times \\ f[a_{sn}|0, \sigma_a^2] f[E_{sn}|0, \sigma_E^2] da_{sn} dE_{sn} \quad (\text{A4})$$

where $\sigma_a^2 = \sigma_A^2/2$. Equation (A4) is identical to equation (A2), apart from that A_s has been replaced with $(a_{sn} + b_n)$ thereby adding one more unique error component (namely a_{ns}) to be integrated out at the sibling level. Similarly as with the MZ case, we can derive the joint likelihood of $2 \times T$ choice observations on pair n as follows

$$L_s^{DZ}[\alpha, \sigma_A^2, \sigma_C^2, \sigma_E^2, \mu] = \int \int l_{1s}^{DZ}[\alpha, \sigma_a^2, \sigma_E^2, \mu; b_n, C_n] l_{2s}^{DZ}[\alpha, \sigma_a^2, \sigma_E^2, \mu; b_n, C_n] \times \\ f[b_n|0, \sigma_A^2/2] f[C_n|0, \sigma_C^2] db_n dC_n \quad (\text{A5})$$

by integrating out the remaining, shared error components.

To estimate primitive parameters in the argument list of $L_s^{MZ}[\cdot]$ and $L_s^{DZ}[\cdot]$, we maximize a sample log-likelihood function which adds up the natural log of either likelihood function as appropriate across all twin pairs in the sample. As this sample log-likelihood function does not have an analytic expression, in practice we maximize its simulated analogue instead. Each pair-level likelihood function can be simulated in the same sequence as our discussion has progressed. In the first step, a sibling-level likelihood function is simulated by making repeated draws of the unique error components, holding fixed a particular set of draws of the shared error components. In the second step, the product of the two sibling-level simulated likelihood functions is averaged across repeated draws of the shared error components to simulate a pair-level likelihood function. Our likelihood evaluator uses 100 Halton draws per each sibling-level error component which carries subscript sn , and another set of 100 Halton draws per each pair-level error component which carries index n .

The logic of this illustrative example extends to other models of economic preferences. In a nutshell, for each structural parameter in the analysis, one must include a distinct set of

ACE parameters $\alpha, \sigma_A^2, \sigma_C^2$, and σ_E^2 . Thus, the number of integrals in each equation above is doubled if one is to estimate the RDU model of risk preferences that we present in the main text; and quadrupled if one is to jointly estimate the RDU model of risk preferences with the QH discounting model of time preferences. Allowing behavioral noise parameter μ to vary from model to model does not affect the number of integrals because this parameter is a primitive parameter to be estimated directly rather than a random coefficient which is to be integrated out.

Appendix B. Meta analysis

B.1 Identifying studies for inclusion

To identify studies for inclusion, we searched Google Scholar using the following search terms: “risk attitudes/preferences”, “risk aversion”, “time preferences”, “intertemporal choice”, “temporal discounting”, “delay discounting”, “present bias” + “heritability”, “genetic basis”, “twin studies”. We supplemented the results from this search with by adding additional studies that we were aware of. We also searched additional databases such as PsycINFO but did not find any further studies to include. The studies we identified are summarized in Tables B.1 and B.2.

Table B.1: Literature summary of the heritability of risk attitudes

Study	Measurement method	MZ/DZ	F/M	Registry	Age
Cesarini et al. (2009)	B (I) Certainty equivalent determined by six choices between a 50/50 gamble for SEK100 and varying sure payoffs	307/135	20/80	Swedish Registry	Twin
	B (NI) Proportion of £1m invested (options: 0, 200k, 400k, 600k, 800k, 1m)	319/139	20/80		
	S Dohmen et al. (2005) (10-point scale)	317/139	20/80		
Anokhin et al. (2009)	B (I) Number of pumps in BART	82/71	0/100	US	12.5 (0.21)
	B (I) Number of pumps in BART	87/49	100/0		12.5 (0.21)
	B (I) Number of pumps in BART	41/41	0/100		14.6 (0.64)
	B (I) Number of pumps in BART	57/24	100/0		14.6 (0.64)
Zhong et al. (2009)	B (I) Ranking of 3 options: (1) 50/50 of getting 40 or 0, (2) 20 for sure, (3) 15 for sure (high risk taking if (1) was ranked highest, medium if (2), low if (3))	158/62	50/50	China	30 (14.63)
Zyphur et al. (2009)	B (NI) Composite score based on: (1) choice between \$2k for sure, 50% chance of \$5k, and 20% chance of \$15k, (2) choice between 3 retirement funds ranging from safe to risky, (3) hypothetical trade of salary for company stock (3 options, more risky=more stocks).	111/89	0/100	Minnesota Registry	Twin 36.7 (1.12)
Simonson & Sela (2011)	B (NI) % of safe choices in lottery choice (gain domain)	110/70	78/22	North Carolina Twin Registry	MZ: 46.6 DZ:49
	B (NI) % of safe choices in lottery choice (loss domain)	110/70	78/22		
Harden et al. (2017)	B (NI) Count of plays in Iowa Gambling task	153/284	48/52	Texas Project	Twin 15.9 (1.4)
	B (NI) Number of pumps in BART	153/284	48/52		
	B (NI) % failed stops at crossroads in Stoplight Task	153/284	48/52		
	S Risk perceptions (28 items on 7 activities)	153/284	48/52		

Beauchamp et al. (2017)	B (NI)	Four ordinal categories based on a hypothetical choice between a sure payoff and a 50/50 gamble to decide the lifetime monthly salary	1000/1083	54/46	Swedish Registry	Twin	R=52-67
	S	“Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?” (10-point scale)	1128/1200	54/46			
	S	11 response categories (10-point scale)	1134/1212	54/46	TwinsUK	Reg- istry	
	B (NI)	Proportion of £100k invested (options: 0, 20k, 40k, 60k, 80k, 100k)	1898/1344	92/8			
	B (NI)	Choice between 3 lotteries (Zyphur et al., 2009)	1898/1344	92/8			
Nicolaou & Shane (2020)	S	“Are you generally a person who is fully prepared to take risks or do you try to avoid taking risks?” (10-point scale)	1898/1344	92/8			
	S	Domain-specific: financial matters, car driving, leisure and sports, health, career (10-point scale)	1898/1344	92/8			
	S	“How much risk are you willing to tolerate when deciding how to invest your money?” (10-point scale)	867/1008	57/43			
	S	“How much risk are you willing to tolerate when deciding how to invest your money?” (10-point scale)	867/1008	57/43			
Le et al. (2010)	S	“How much risk are you willing to tolerate when deciding how to invest your money?” (10-point scale)	867/1008	57/43	Australian Twin Younger Cohort	Registry	37.7

Table B.2: Literature summary of the heritability of impatience

Study	Measurement method		MZ/DZ	F/M	Registry		Age
Anokhin et al. (2011)	B (I)	\$7 today or \$10 in 7 days (categorical)	169/203	48/52	US		12.5 (0.21)
Sparks et al. (2014)	B (I)	\$7 today or \$10 in 7 days (categorical)	N=606	48/52	Minnesota Twin Registry	Twin	14.6 (0.24)
	B (I)	\$7 today or \$10 in 7 days	239/159	50/50			$\sim 17y.o.$
Cesarini et al. (2012)	B (NI)	Number of times immediate option was chosen (money today or more money in a week, 3 items)	1150/2362	54/46	Swedish Registry	Twin	R=51-66
Isen et al. (2014)	B (NI)	\$0.5-\$10 now or \$10 after a delay of 1, 2, 10, 30, 180 or 365 days (AUC)	148/97	45/55	Minnesota Registry	Twin	15.1 (0.55)
Anokhin et al. (2015)	B (NI)	138 choices between \$X now or \$100 in Y days (AUC)	50/39	51/49	US		16.6 (0.26)
	B (NI)	138 choices between \$X now or \$100 in Y days (AUC)	113/125	51/49			18.5 (0.21)
	B (NI)	138 choices between \$X now or \$100 in Y days (hyperbolic k)	50/38	51/49			16.6 (0.26)
	B (NI)	138 choices between \$X now or \$100 in Y days (hyperbolic k)	108/124	51/49			18.5 (0.21)
Harden et al. (2017)	B (NI)	Indifference point in smaller sooner and larger later choices (staircase approach)	153/284	48/52	Texas Project	Twin	15.9 (1.4)
	S	UPPS Impulsivity Scale (45 items, 4 dimensions)					
	S	Future Orientation Scale (15 items, 3 dimensions)					
Hubler (2018)	S	“Would you describe yourself as a patient or impatient person?” (11-point scale)	703/775	56/44	German Life	Twin-	17.1 (5.1)

B.2 Effect sizes and standard errors

If a study measured preferences using more than one instrument in the same sample, we use only one estimate. In the case of behavioral preferences, we prioritize estimates from incentivized elicitations and those that provide the most precise estimates.

Two major challenges arise when summarising estimates from AC(D)E models. These are differences in the models that are estimated, and inconsistent reporting of standard errors.

Some papers only report estimates from the best fitting model (often an AE model, with C constrained to zero), while others estimate the full ACE specification. Our preferred approach is to use the full ACE estimates to avoid constraining C, so we use the estimates from the complete ACE model whenever they are reported. In the remaining cases, we use the estimates of from the best-fitting model among those that were reported. In the few papers that estimated an ADE model, in slight abuse of notation, what we refer to as A is in fact an aggregation of the additive (A) and dominant (D) genetic effects. One concern with our adopted approach is that studies that only report AE estimates may get higher weight than they should, since constraining C will often lower the standard error for A%. We address this concern with sensitivity analysis below.

For conventional meta-analysis we need the standard error for each study; however, it is a common convention in twin studies to only report confidence intervals. In situations where the confidence interval is symmetric, the standard error can be recovered if we assume that it was constructed as $\pm 1.96 \times SE$. But in some studies the confidence intervals are constructed to respect the fact that variance shares are bounded $[0,1]$, so are not symmetric. Details on how confidence intervals were constructed were frequently missing in the studies we reviewed. We suspect that in many cases likelihood-based intervals (see Neale and Miller, 1997) were used, since this is the default in popular open-source R programs used for twin studies.

Table B.3 shows, for each study included in our analyses, the point estimate, the standard error we used, and how that standard error was obtained.

Table B.3: Estimates and standard errors used in meta-analyses

Study	A% (SE)	E% (SE)	Details
Behavioral risk preference studies			
Cesarini et al. (2009)	0.16 (0.077)	0.75 (0.056)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% credible intervals. The intervals were approximately symmetric.
Anokhin et al. (2009)*, M 12 y.o.	0.28 (0.071)	0.72 (0.071)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were symmetric.
Anokhin et al. (2009)*, F 12 y.o.	0.17 (0.087)	0.83 (0.087)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric (slight asymmetry possibly due to rounding).
Zhong et al. (2009)	0.54 (0.226)	0.46 (0.111)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were not symmetric.
Zyphur et al. (2009)#	0.63 (0.25)	0.37 (0.290)	Reported SE for A%. SE not reported for the E variance share, but was reported for the E path coefficient (e). Since $E=e^2$, we used the delta method ($SE_e \times f'(E)$).
Simonson & Sela (2011)*	0.33 (0.158)	0.67 (0.158)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric, but not close enough to be explained by rounding.
Beauchamp et al. (2017)#	0.42 (0.226)	0.58 (0.036)	Reported SE for E. For A%, because an ADE model was estimated this was A+D. SE is calculated as $\sqrt{SE_A^2 + SE_D^2}$ using the reported SE_A and SE_D .
Nicolaou & Shane (2020)*	0.25 (0.031)	0.75 (0.031)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric, but not close enough to be explained by rounding.
Stated risk preference studies			
Cesarini et al. (2009)	0.29 (0.077)	0.65 (0.051)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% credible intervals. The intervals were approximately symmetric.
Beauchamp et al. (2017)#	0.36 (0.176)	0.65 (0.031)	Reported SE for E. For A%, because an ADE model was estimated this was A+D. SE is calculated as $\sqrt{SE_A^2 + SE_D^2}$ using the reported SE_A and SE_D .

Nicolaou & Shane (2020)	0.22 (0.066)	0.78 (0.046)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The A% intervals were approximately symmetric, but not close enough to be explained by rounding. The E interval was symmetric.
Le et al. (2010)*	0.23 (0.033)	0.77 (0.033)	SE for A% obtained by dividing the estimate by the reported t-statistic. The regression approach used in this paper did not directly estimate E, so we set its SE equal to the SE for A%.

Behavioral time preference studies

Anokhin et al. (2011)*	0.30 (0.097)	0.70 (0.097)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric, but not close enough to be explained by rounding.
Sparks et al. (2014)	0.37 (0.189)	0.51 (0.107)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric, but not close enough to be explained by rounding.
Cesarini et al. (2012)	0.18 (0.092)	0.75 (0.056)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric but not by construction (obtained using likelihood approach).
Anokhin et al. (2015)*	0.62 (0.056)	0.38 (0.056)	$SE = \max[UL - Est., Est. - LL]/1.96$ using the reported 95% confidence intervals. The intervals were approximately symmetric, but not close enough to be explained by rounding.

Note: *Estimates from an AE model. #Estimates from an ADE model.

An alternative to using the available information about parameter uncertainty reported in the papers is to approximate standard errors using a standardized approach. For example, in Polderman et al. (2015) the authors transform variance shares into Fisher Z-values with the approximate standard errors depending only on the sample size, conduct the meta-analysis and then back-transform the estimates at the end. This approach can also be used on the MZ/DZ correlations with heritability estimated using Falconer's formula, rather than the usual SEM or multi-level regression approach; however, we do not do this since the Falconer

approach does not compare neatly to our own analysis of deep parameters.

We explored the possibility of using transformed Z-values for our study but ultimately rejected this due to unsatisfactory weights given to the different studies. Using the Fisher approach, our meta-estimate for behavioral risk preference is 36% (95% CI [23%,48%]), which is larger than our preferred estimate of 25% [20%,30%]. This difference is driven by the fact that the studies received nearly identical weights using the Fisher approach, even though some studies were much less precise than others. For example, the study by Zyphur et al. (2009) (A%=63% with 95% CI [15%,100%] in the original study) received a weight of 0.96% in our preferred approach, reflecting the large degree of uncertainty, compared to 12.16% in the Fisher approach. In contrast, Nicolaou and Shane (2020) (A%=30% with 95% CI [19%,30%] in the original study) received a weight of 64.21% in our preferred approach, reflecting its far greater precision than other studies, compared to 13.67% in the Fisher approach – a weighting only slightly larger than that given to Zyphur et al. (2009).

B.3 Sensitivity

One concern with using the reported parameter uncertainty details is that A% estimates will be more precise in studies that only report AE model estimates because constraining the C component reduces uncertainty. To gauge whether this is likely to meaningfully affect our results in practice, we conducted a bounding exercise by scaling the standard errors in studies that use the AE model by a (relatively large) factor of 1.5. For behavioural risk preference, our meta-estimate is identical but with a slightly larger confidence interval as expected (25% [18%,32%]). Our estimates are also nearly identical for stated risk (25% [18%,31%]) and for behavioral time preference (37% [16%,59%]), differing by only one percentage point from what we report in the paper. We therefore conclude that any bias from using the AE estimates is likely to be small.

B.4 Additional meta-analysis results

Figure B.1: Meta-analysis of the genetic role (A) in stated risk aversion





Figure B.2: Meta-analysis of the unique environment (E) role in stated risk aversion



Appendix C. Additional results

Table C.1: Variable descriptions

Variable	Definition	MZ obs.	DZ obs.
Age	Age at last birthday	802	318
Male	= 1 if male	802	318
Twin years together	How many years (including your childhood) lived with your twin	802	318
Australia born	= 1 if born in Australia	802	318
Lives in a city	= 1 if currently live in a major city (Sydney, Melbourne, Brisbane, Adelaide, Perth, Canberra)	795	317
Couple	= 1 if married or in a defacto relationship	800	314
Household members	How many people live in your household	799	315
Dependent children	Number of dependent children	766	310
University degree	= 1 if highest level of education obtained is a university degree	802	318
Employed	= 1 if worked any time in the last 7 days or if had a job but did not work in the last 7 days due to holidays, sickness or any other reason	802	318
Retired	= 1 if currently retired from the workforce	802	318
Income	Average usual weekly own income in the last month using midpoint value for the following categories: \$1-\$149, \$150-\$299, \$300-\$399, \$400-\$499, \$500-\$649, \$650-\$799, \$800-\$999, \$1,000-\$1,249, \$1,250-\$1,499, \$1,500-\$1,749, \$1,750-\$1,999, \$2,000-\$2,999, \$3,000 or more (coded as \$3000). Negative or nil coded as missing.	692	275
Financially secure	Given your current needs and financial responsibility, would you say that you and your family are: = 1 if Poor, = 2 if Just getting along, = 3 if Comfortable, = 4 if Very comfortable, = 5 if Prosperous.	802	318
Long-term health condition	= 1 if has a long-term health condition, impairment or disability that has lasted more than 6 months	800	318

COVID-19 worry	Worry or concern about contracting COVID-19 on a scale of 1 to 10	798	318
COVID-19 prob	Probability participant believes they will get COVID-19 in the next 3 months	796	315
COVID-19 mort	If you do get COVID-19, what is the percent chance you will die from it?	795	317
COVID-19 job loss	= 1 if experienced job loss due to COVID-19	802	318
COVID-19 reduced income	= 1 if experienced reduction in income due to COVID-19	802	318
COVID-19 work home	= 1 if experienced working from home due to COVID-19	802	318
COVID-19 reduced hours	= 1 if experienced a reduction in working hours due to COVID-19	802	318
Num. COVID-19 positive friends	How many relatives or close friends have tested positive for COVID-19	799	318
Risk MPL1 num safe	Number of safe choices in MPL task 1 (known probabilities)	802	318
Risk MPL2 num safe	Number of safe choices in MPL task 2 (known probabilities)	802	318
Uncertainty MPL1 num safe	Number of safe choices in MPL task 1 (unknown probabilities)	802	318
Uncertainty MPL2 num safe	Number of safe choices in MPL task 2 (unknown probabilities)	802	318
Time MPL1 num sooner	Number of sooner choices in MPL task 1	802	318
Time MPL2 num sooner	Number of sooner choices in MPL task 2	802	318
Time MPL3 num sooner	Number of sooner choices in MPL task 3	802	318
Time MPL4 num sooner	Number of sooner choices in MPL task 4	802	318

Table C.2: Correlations between number of years living with twin and absolute difference between twin pairs in economic preference variables

	Correlation	P-val
Risk MPL1 num safe	-0.063	0.054
Risk MPL2 num safe	-0.049	0.113
Uncertainty MPL1 num safe	-0.021	0.514
Uncertainty MPL2 num safe	-0.029	0.364
Time MPL1 num sooner	-0.026	0.376
Time MPL2 num sooner	-0.05	0.104
Time MPL3 num sooner	-0.005	0.889
Time MPL4 num sooner	-0.024	0.432

Table C.3: ACE share decomposition results: additional models

Model	A	C	E
Risk preferences			
MPL1	0.01 (0.00, 0.18)	0.10 (0.02, 0.22)	0.89 (0.73, 0.95)
MPL2	0.04 (0.00, 0.19)	0.05 (0.00, 0.13)	0.91 (0.75, 0.98)
Panel	0.01 (0.00, 0.10)	0.14 (0.05, 0.25)	0.86 (0.73, 0.93)
EUT	0.01 (0.00, 0.14)	0.09 (0.03, 0.17)	0.90 (0.75, 0.96)
Uncertainty preferences			
MPL1	0.11 (0.01, 0.30)	0.06 (0.00, 0.24)	0.83 (0.67, 0.91)
MPL2	0.23 (0.10, 0.27)	0.05 (0.00, 0.18)	0.73 (0.60, 0.81)
Panel	0.21 (0.06, 0.40)	0.04 (0.00, 0.21)	0.75 (0.59, 0.85)
EUT	0.00 (0.00, 0.13)	0.19 (0.10, 0.28)	0.81 (0.67, 0.88)
Ambiguity preferences			
MPL1	0.05 (0.00, 0.16)	0.00 (0.00, 0.03)	0.95 (0.83, 1.00)
MPL2	0.05 (0.00, 0.19)	0.00 (0.00, 0.02)	0.95 (0.80, 1.00)
Panel	0.24 (0.01, 0.52)	0.07 (0.00, 0.46)	0.69 (0.42, 0.77)
EUT	0.00 (0.00, 0.08)	0.05 (0.02, 0.08)	0.95 (0.86, 0.98)
Time preferences (Delay aversion)			
MPL1	0.15 (0.07, 0.25)	0.00 (0.00, 0.02)	0.85 (0.75, 0.93)
MPL2	0.16 (0.08, 0.26)	0.01 (0.00, 0.03)	0.83 (0.73, 0.92)
MPL3	0.22 (0.14, 0.32)	0.00 (0.00, 0.01)	0.78 (0.68, 0.86)
MPL4	0.24 (0.15, 0.34)	0.00 (0.00, 0.05)	0.76 (0.65, 0.85)
Panel	0.21 (0.11, 0.33)	0.00 (0.00, 0.03)	0.79 (0.66, 0.89)
EXP	0.00 (0.00, 0.01)	0.20 (0.17, 0.22)	0.80 (0.77, 0.82)
EUTEXP	0.01 (0.00, 0.03)	0.24 (0.22, 0.26)	0.75 (0.73, 0.76)
QH	0.27 (0.23, 0.31)	0.01 (0.00, 0.02)	0.72 (0.68, 0.75)
Time preferences (Non-stationary)			
MPL1-MPL3	0.04 (0.00, 0.12)	0.01 (0.00, 0.35)	0.95 (0.62, 0.99)
MPL2-MPL4	0.08 (0.02, 0.16)	0.00 (0.00, 0.16)	0.91 (0.76, 0.97)
Panel	0.14 (0.00, 0.48)	0.09 (0.00, 0.46)	0.77 (0.46, 0.85)
QH	0.10 (0.02, 0.23)	0.10 (0.03, 0.18)	0.81 (0.71, 0.87)

Note: See Figure 5–9 for further details.