Family Background and Economic Mobility: Evidence from the US

Alejandro Puerta-Cuartas*

This paper examines how family background relates to economic mobility for disadvantaged children. We use data from the Panel Study of Income Dynamics for below-median income, multiple-child families. Using a novel approach combining family fixed effects, Empirical Bayes shrinkage, and SHapley Additive exPlanations, we identify which family characteristics most strongly predict children's economic outcomes relative to their parents, holding parental income constant. Our findings reveal that race and family structure are the primary predictors, accounting for 35% and 22.4% of the explained variation, respectively. While supporting the well-documented racial disparities in intergenerational mobility, our results suggest that the role of family structure in intergenerational mobility extends beyond the single- versus two-parent household distinction.

I. Introduction and related work

Children from underprivileged backgrounds are less likely to thrive. In a low-mobile society, the well-being of individuals is mainly determined by the socioeconomic level of their parents. This is especially detrimental for children of low-income families as it limits their chances of escaping the poverty trap. Studying the relationship between family background, particularly of disadvantaged children, and economic mobility offers a deeper understanding of the factors associated to the persistence of poverty.

Extensive research has documented childhood environment's lasting economic impacts. Seminal work by Corcoran et al. (1976, 1990) established that parental characteristics significantly influence children's educational attainment and earnings, while Björklund et al. (2007) emphasized the interplay of genetic and environmental factors. More recently, Bingley et al. (2017) identify families, rather than communities, as the dominant factor driving lifetime earnings inequality, while Chetty et al. (2014) link higher mobility to area traits including family stability, school quality, and neighborhood integration. These findings raise a central question for both research and policy: what family factors foster intergenerational mobility? Empirically, we investigate a more tractable, yet insightful question: which family factors best predict absolute mobility?

This paper provides a descriptive analysis of how family background influences the permanent income of children born in low-income households. While recent work has examined spatial heterogeneity in mobility (e.g., neighborhood (Chetty and Hendren, 2018; Cholli et al., 2024) or commuting zones (Chetty et al., 2014)), we instead analyze

* Department of Economics, Universidad Carlos III, Madrid, Spain, E-mail: alpuerta@eco.uc3m.es

heterogeneity at the family level. Drawing on data from the Panel Study of Income Dynamics (PSID) for multiple-child low-income families, we identify which family characteristics most strongly predict children's economic outcomes relative to their parents, holding parental income constant.

To identify key the family predictors of intergenerational mobility, we integrate shrinkage methods and interpretable machine learning. Our approach proceeds in three stages: first, we measure absolute mobility as the family fixed effect (FE) in a regression of children's permanent income on parental permanent income. This measure captures variations in children's economic outcomes relative to their parents, reflecting both upward and downward mobility. Specifically, the FE accounts for factors shared by siblings that influence their permanent income through channels other than parental income. As the FE increases, so does the predicted permanent income, suggesting that the FE can be interpreted as a family's ability to enhance their children's long-term economic outcomes, conditional on parental income. Next, we refine these estimates using Empirical Bayes shrinkage to mitigate noise from small family sizes. Finally, we apply SHapley Additive exPlanations (SHAP) values (Lundberg, 2017) to analyze which family factors best predict absolute mobility.

Our analysis reveals race and family structure as the strongest predictors of absolute mobility. SHAP analysis indicates that, conditional on parental income, race (white vs. non-white) explains 35% of the total variable importance among family background characteristics. Being raised in a nuclear family, defined as a household with a head, a spouse (if present), and children, accounts for 22.4%. Other relevant factors include the household's age at first birth (10.5%), number of children (9.4%), and education (9.3%). Because our methodology holds parental income constant, the estimated relative importance of these factors excludes income-mediated channels, including the monetary returns to education, income dilution in extended families, and racial disparities in parental income.

This paper contributes to different strands of the literature on intergenerational mobility. First, we provide a new perspective on the role of family structure on children's economic outcomes relative to their parents, advancing beyond the conventional focus on single- versus two-parent families. Kearney (2023) examines the widening economic disparities linked to family structure, emphasizing that children raised in stable two-parent households tend to achieve higher educational attainment and greater upward mobility. These findings align with Chetty et al. (2014), who finds that regions with weaker family structures exhibit significantly lower levels of intergenerational mobility. Similarly, Bloome (2017) shows that children raised outside stable two-parent families are less likely to attain high-income status in adulthood compared to their peers from stable two-parent households. By focusing on the broader distinction between nuclear and non-nuclear families, we document associations between family structure and economic mobility, offering a complementary perspective. In particular, our findings suggest that the role of family structure in intergenerational mobility extends beyond the single- versus two-parent household distinction.

Second, while our results supports the widely documented racial disparities in intergenerational mobility (Hertz, 2005; Bloome, 2014; Davis and Mazumder, 2018), we find that family structure emerges as the second most influential factor. This result gains further significance when considering that 68% of white children were consistently raised in nuclear families, compared to just 35% of non-white children in our sample. Together, these findings align with existing work on how family structure differences may contribute to persistent racial mobility gaps (McLanahan and Percheski, 2008), offering new descriptive evidence into this ongoing conversation.

Third, our study advances the empirical analysis of intergenerational mobility through several fronts. Our fixed effects methodology first residualizes children's income with respect to parental income, achieving two critical objectives: it prevents parental income from dominating subsequent analyses while capturing family background influences that are not transmitted via parental income. We further enhance estimation precision through Empirical Bayes shrinkage, leveraging information across families to overcome data limitations at the family level. The SHAP analysis then provides a sophisticated measure of relative variable importance that fully accounts for interactions and nonlinearities, offering substantial advantages over traditional correlation-based approaches. Finally, our simulation results indicate that integrating parametric Empirical Bayes estimation with SHAP yields superior performance compared to analyzing raw fixed effects. Together, this multi-faceted approach provides new tools for understanding mobility predictors while addressing fundamental challenges in family-level analysis.

The remainder of the paper is organized as follows: Section II provides an overview of the data. Section III presents the model, discusses identification, and describes the empirical procedure for estimating the fixed effects and analyzing which family factors best predict absolute mobility. Section IV presents simulation results, Section V describes the main findings, and Section VI concludes.

II. Data

We use data from the Panel Study of Income Dynamics (PSID), a commonly used dataset for analyzing intergenerational persistence in the U.S. Since our analysis examines mobility at the family level, we retain the Survey of Economic Opportunity (SEO), which is often excluded due to its overrepresentation of low-income populations.[1] Our sample, covering the period from 1968 to 2021, consists of 1088 multiple-child families whose family income in the father's generation was below the median of the income distribution for that generation. Following (Chetty et al., 2014), we focus on below-median income families for three key reasons : (i) their children face the most significant mobility barriers, (ii) understanding economic opportunity dynamics is most policy-relevant in this context, and (iii) this focus allows clearer examination of how family background relates to mobility in disadvantaged households.

We measure parental income as the three-year average of log family income when each child is between the ages of 13 and 21.[2] For 7.5% of fathers in our data, family income

[1]The SEO is typically excluded because its overrepresentation can influence estimates at the aggregate level, such as the intergenerational elasticity.

[2]This follows from Theorem 1's implications about within-family variation in parental income.

data during this period is insufficient. In these cases, we compute parental income by averaging log family income over the closest available three-year window to ages 13–21. In contrast, the child's permanent income is measured using a single realization of log family income between the ages of 30 and 40. Since we use family income, our analysis includes both male and female children. All dollar values are adjusted to 1968 dollars using the CPI. To address zero-income cases, we bottom-code family income at the 1st percentile, which applies to 0.18% of the observations in the raw PSID data.

Our dataset provides detailed information on family background characteristics. For education, the dataset records the highest level of education completed and whether the household head received additional training beyond standard school and college education. As regards family structure, the dataset records whether the household is nuclear (head, spouse (if any), and children) or extended (including other individuals). We utilize the available data to construct a variable representing the father's age at first birth and the proportion of male children. Furthermore, we use the retrospective inter-generation map of the PSID to obtain the number of children in the family.[3] The dataset further includes information on race (white or non-white), religion, and the family's ownership of housing and business.

To reflect the household conditions of children when they were growing up, we construct family background characteristics based on the period from the birth of the first child until the eldest child reaches 18 years of age. For time-invariant variables, we compute the mode, while for time-variant variables, we use the average. Accordingly, time-variant variables reflect the probability of being 1, e.g., the probability of house ownership when children were growing up. Summary statistics are reported in Table 1.

Table 1—: Summary Statistics for Below-Income Families in the Sample

	Mean	Median	Standard Deviation	Minimum	Maximum
Log of Father's Permanent Income	8.67	8.81	0.57	1.65	9.29
Log of Children's Permanent Income	8.69	8.85	0.94	-1.09	11.16
Education Level	11.19	12.00	3.20	0.00	17.00
Additional Training	0.18	0.00	0.39	0.00	1.00
Nuclear Household	0.86	1.00	0.23	0.00	1.00
Number of Children	3.74	3.00	1.98	2.00	15.00
Religion	0.90	1.00	0.30	0.00	1.00
Age at Birth of First Child	25.10	24.00	5.62	13.00	59.00
White	0.54	1.00	0.50	0.00	1.00
Proportion of Male Children	0.50	0.50	0.29	0.00	1.00
Business Ownership	0.10	0.00	0.22	0.00	1.00
House Ownership	0.64	0.79	0.38	0.00	1.00
Father's Birth Year	1941	1943	14.21	1892	1970
Children's Birth Year	1968	1966	12.69	1932	1991

[3]The number of children in the family may differ from the number of observed children in the sample due to some children not being observed between the ages of 30 and 40.

The summary statistics highlight substantial heterogeneity in the background characteristics of below-median income families. As indicated by Table 1, children exhibit income levels similar to their parents, though with significantly greater variability in the children's generation. Most parents have completed at least high school, with educational attainment ranging from no schooling to graduate degrees. Additionally, 18% of parents received training beyond regular schooling. The majority of families are always nuclear (probability of family being nuclear equal to 1), consisting of the head, spouse (if any), and children, with three children per household. The gender ratio among children is balanced, with equal numbers of males and females in most families. Religiosity is prevalent, with 90% of parents reporting adherence to a religion. Homeownership is common, with 64% of families owning their dwelling, while business ownership is relatively rare at 10%. In terms of birth years, most parents in the sample were born between 1927 and 1957, while the majority of children were born between 1956 and 1978.

III. Empirical Framework

A. A Simple Model to Explain Absolute Mobility

Consider a random sample of independent and identically distributed observations $W = \left(\boldsymbol{Y}_j, \boldsymbol{Y}_j^P, \boldsymbol{X}_j\right)$, where $\boldsymbol{Y}_j := \left(Y_1, ..., Y_{n_j}\right)$ is a n_j-dimensional random vector representing the permanent income measures of the children of family j, $\boldsymbol{Y}_j^P := \left(Y_1^P, ..., Y_{n_j}^P\right)$ contains the corresponding parental permanent income measures for each child, and $\boldsymbol{X}_j$ are family background characteristics.

To analyze absolute mobility, we consider the following model:

$$\boldsymbol{Y}_j = \beta \boldsymbol{Y}_j^P + \mathbf{1}_{n_j} \alpha_j + \boldsymbol{\epsilon}_j, \quad j = 1, ..., N, \tag{1}$$

where β corresponds to the intergenerational elasticity, and $\mathbf{1}_{n_j}$ is the n_j-dimensional vector of ones. The fixed effects α_j encompass the influence of factors shared by siblings on their permanent income, holding parental income constant. These factors include both observed variables $\left(\boldsymbol{X}_j\right)$ and unobserved variables. Finally, $\boldsymbol{\epsilon}_j$ represents the n_j-dimensional random vector of idiosyncratic shocks.

Our measure of absolute mobility is represented by the family fixed effects (FEs) in equation (1).[4] This measure captures variations in children's economic outcomes relative to their parents, reflecting both upward and downward mobility. Specifically, the FE accounts for factors shared by siblings that influence their permanent income through channels other than parental income.

As the FE increases, so does the predicted permanent income, suggesting that FEs can be interpreted as a measure of a family's ability to enhance their children's long-term economic prospects, conditional on parental income. This interpretation forms the basis for our use of α_j as a measure of absolute mobility. This approach parallels teacher

[4] In this framework, the FEs exploit variation between sibling incomes to capture the impact of family background on children's income, which underscores the need to rely on multiple-child families.

value-added models (Kane and Staiger, 2008; Jacob and Lefgren, 2008; Gilraine et al., 2020), which isolate a teacher's contribution to student achievement (measured using fixed effects) while controlling for observable student and school characteristics. Similarly, our family fixed effects encompass the impact of family background on children's income, while holding parental income constant.

Unlike standard mobility measures, such as absolute and absolute upward mobility from Chetty et al. (2014), which capture overall income transmission, our measure isolates the influence of family background beyond parental income. By measuring absolute mobility using the family fixed effects in equation (1), we disentangle the influence of family background on children's outcomes beyond parental income. To illustrate this point, consider analyzing family structure and parental education. In a household with multiple children and extended relatives, family income needs to be allocated among many individuals, mechanically limiting the family's ability to boost their offspring's income. Thus, focusing on the FEs rather than the expected children's permanent income (or rank) enables us to assess whether family structure exerts an influence beyond income dilution.

In the case of parental education, one potential transmission channel is that higher education levels can lead to a higher lifetime income, which in turn allows parents to provide better opportunities for their children. Since our mobility measure holds parental permanent income constant, the effect of parental education on children's income rules out this channel. Thus, a positive relationship between parental education and α_j suggests that parental education can boost children's outcomes through channels other than income, such as cognitive and non-cognitive skill development or parental support and engagement.

Studying absolute mobility is especially relevant for children from disadvantaged backgrounds, as they often face the greatest challenges in achieving success. For this purpose, we restrict our analysis to children from below-median income families and specify our baseline model as

$$\boldsymbol{Y}_j = \beta \boldsymbol{Y}_j^P + \mathbf{1}_{n_j}\alpha_j + \boldsymbol{\epsilon}_j, \quad j \in \mathbb{S} \subseteq \{1, ..., N\}, \tag{2}$$

where $\mathbb{S}$ are the families whose income is below the median.

Our object of interest lies in estimating $\mathbb{E}\left[\alpha_j \mid \boldsymbol{X}_j, \mathbb{S}\right]$, which captures how family background characteristics $\boldsymbol{X}_j$. This conditional mean determines how observed family factors influence children's permanent income through mechanisms beyond parental income. Accordingly, it encapsulates the information regarding which family factors best predict economic mobility beyond parental income and how these factors influence children's permanent income. While we do not investigate specific causal mechanisms, analyzing these predictors provides valuable insights into family-level associations with intergenerational persistence.

We now study the identification of our object of interest. All proofs are provided in the Appendix.

B. Identification

To identify our parameter of interest, we begin by imposing mean independence between the idiosyncratic shocks and both parental permanent income and family fixed effects for below-median income families. This assumption is necessary, though not sufficient, for identifying the intergenerational elasticity and the conditional mean of the fixed effects.

ASSUMPTION 1 (Mean Independence):

$$\mathbb{E}\left[\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}_j^P, \alpha_j, \mathbb{S}\right] = 0.$$

Our identification strategy requires within-family variation in parental income, which merits careful discussion since permanent income is conceptually time-invariant. This mirrors the well-known non-identification of time-invariant covariates in fixed effects panel models. In our framework, if the same measure of parental income were used for all siblings, it would be perfectly collinear with the family fixed effects ($Y_j^P = \sum_{s=1}^N 1(s = j)Y_s^P$ for every $j \in \mathbb{S}$), undermining identification. We formalize this requirement in the Assumption stated below.

ASSUMPTION 2 (Within-Family Variation in Parental Income): *Parental permanent income varies within families almost surely; that is,*

$$\boldsymbol{Y}_j^P \notin span(\mathbf{1}_{n_j}) \quad \textit{almost surely}.$$

We incorporate within-family income variation by measuring parental income as the three-year average of log family income when each child is between the ages of 13 and 21. This approach captures economically meaningful differences in parental resources across children's upbringing periods, reflecting real-world dynamics. For example, career progression may increase earnings for later-born children, while economic shocks, such as job losses, may differentially impact siblings based on their birth timing. Crucially, this child-specific conceptualization of parental income directly aligns with our research objective: by allowing parental resources to vary across siblings, we can isolate how family background influences children's permanent income through channels beyond financial resources.

To identify the object of interest, $\mathbb{E}[\alpha_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}]$, it is essential to first disentangle the influence of parental permanent income from that of family fixed effects. The key insight is that by appropriately averaging outcomes net of the parental income component (weighted by β), we can isolate the family fixed effect. Therefore, the identification of β is necessary for identifying the distributional characteristics of α_j. Theorem 1 therefore establishes the identification of β, providing the foundation for our subsequent analysis of how family background relates to economic mobility for disadvantaged children.

THEOREM 1: *Under Assumptions 1 and 2 β is identified by*

$$\beta = \frac{\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j\right]}{\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j^P\right]},$$

where $\boldsymbol{Q}_j = \boldsymbol{I}_{n_j} - \frac{1}{n_j}\boldsymbol{1}_{n_j}\boldsymbol{1}_{n_j}'$ and the denominator is non-zero by Assumption 2.

A potential threat to the identification result in Theorem 1 is the weak identification of β. While Assumption 2 ensures that parental income is not identical across siblings (almost surely), it may still be highly similar, specially in families where children are born close together. To empirically assess this concern, we calculate the intraclass correlation coefficient (ICC), which measures the proportion of total variance in parental income that is attributable to differences between families. In our sample, the ICC is 0.689, indicating that 68.9% of parental income variance arises from between-family differences, while the remaining 31.1% reflects meaningful within-family variation. This result suggests that weak identification of β is unlikely to pose a significant issue in our analysis.

The following Corollary provides the theoretical foundation for identifying which family factors most strongly predict absolute mobility. A key requirement for this analysis is the identification of $\mathbb{E}[\alpha_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}]$, which is a necessary condition for (i) accurately recovering the predictive relationships between observable family characteristics and mobility outcomes, and (ii) determining which specific family factors most strongly influence absolute mobility.

COROLLARY 1: *Under Assumptions 1–2 and $\mathbb{E}\left[\epsilon_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] = 0$, for a fixed β, the least squares estimator $\hat{\alpha}_j^{OLS} := \frac{1}{n_j}\boldsymbol{1}_{n_j}'(\boldsymbol{Y}_j - \beta\boldsymbol{Y}_j^P)$ satisfies:*

$$\mathbb{E}\left[\hat{\alpha}_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] = \mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right].$$

The standard sibling earnings models (corresponding to Equation 2) imposes the orthogonality assumption $\mathbb{E}\left[\alpha_j\boldsymbol{Y}_j^P\right] = 0$ to estimate the sibling correlation. In contrast, the identification result in Corollary 1 relaxes this assumption, incorporating that family characteristics (as captured by α_j) are inherently correlated with parental income. By explicitly conditioning on parental income while allowing for this correlation, we better isolate the influence of family background. This approach thus provides a more nuanced understanding of which family factors most strongly predict absolute mobility.

While Corollary 1 holds in the population, in finite samples the OLS estimates of the family fixed effects, though unbiased, are noisy, particularly for families with few children. As our simulations illustrate (see Section IV), shrinking offers substantial gains when analyzing the conditional mean of α_j conditional on $\boldsymbol{X}_j$, especially in a noisy and heteroskedastic setting. In account of this, we now introduce our EB shrinkage approach to estimates the family FEs.

C. Empirical Bayes Estimation of Absolute Mobility

Limited within-family observations make fixed effects estimation particularly challenging, highlighting the need to refine raw OLS estimates. Although the ordinary least squares (OLS) estimates for the FEs are unbiased, they are noisy (Gilraine et al., 2020), especially for units (families) with few observations (children). This issue is particularly relevant in our sample, where most families have between 2 and 5 children. A common approach that substantially improves upon the raw fixed effects estimator is based on EB shrinkage (Armstrong et al., 2022), where the preliminary estimate is shrunken towards a baseline value. The strength of Empirical Bayes in our framework stems from considering the estimation of the FEs as a compound decision problem (Efron, 2012), where many estimates of similar quantities are required, but little information for each is available. Thus, we exploit information about other families' transmission processes to estimate the individual FEs.[5]

We adopt an empirical Bayes approach for both theoretical and empirical reasons. Theoretically, EB is the optimal estimator under the mean squared error (MSE) criterion, as it minimizes Bayes risk and provides the best predictor of unobserved heterogeneity given the available data (Gaillac, 2023). Moreover, it remains optimal for estimating averages with respect to the distribution of unobservables even under local deviations from normality (Bonhomme and Weidner, 2022) – a crucial property given our focus on estimating $\mathbb{E}\left[\alpha_j \mid \boldsymbol{X}_j, \mathbb{S}\right]$.

Empirically, shrinkage estimation, the key feature of EB, has demonstrated its utility across diverse fields, including income dynamics (Gu and Koenker, 2017), healthcare (van Houwelingen et al., 2020), and genetics (Endelman and Jannink, 2012). Notably, it has been particularly effective in enhancing predictive accuracy when estimating intergenerational mobility (Chetty and Hendren, 2018). Additionally, EB estimates are well-justified predictors of the quality (as measured by fixed effects) of individual teachers (Bonhomme and Weidner, 2022), further supporting their application in estimating family fixed effects.

Standard parametric empirical Bayes (EB) shrinks preliminary estimates toward the grand mean, but performs poorly for family fixed effects due to high estimation noise from small family sizes. We address this limitation by implementing the robust EB procedure of Armstrong et al. (2022),[6] which improves upon conventional EB in two key ways. First, rather than shrinking toward a common mean, it shrinks the initial OLS estimates $\hat{\alpha}_j^{OLS}$ toward family-specific predictions $\boldsymbol{X}_j'\hat{\gamma}$ from observable characteristics. Second, it determines optimal shrinkage intensity based on the relative variance of sampling errors $\left(\sigma^2_{OLS,j}\right)$ versus effect heterogeneity $\left(\sigma^2_{\alpha,j}\right)$. Formally, following Armstrong

[5]Another notable feature of the EB approach is its adaptive shrinkage, which depends on the available information. Families with fewer children entail noisier estimates, whereas those with more children provide more accurate estimates. Accordingly, the EB estimates experience greater shrinkage when less information is available.

[6]While their method was originally designed for constructing robust EB confidence intervals, we adapt their flat prior limited information Bayes approach to ensure positive shrinkage weights in our empirical Bayes estimator.

et al. (2022), we assume

$$\hat{\alpha}_j^{OLS} \mid \alpha_j, \boldsymbol{X}_j, \sigma_{OLS,j} \sim \mathcal{N}\left(\alpha_j, \sigma^2_{OLS,j}\right), \tag{3}$$

along with the working assumption that the sampling distribution of the α_j is conditionally normal:

$$\alpha_j \mid \boldsymbol{X}_j, \sigma_{OLS,j} \sim \mathcal{N}\left(\boldsymbol{X}_j'\gamma, \sigma^2_{\alpha,j}\right). \tag{4}$$

Thus, according to equation (3) and (4), the EB estimate of α_j is given by

(5)

$$\hat{\alpha}_j^{EB} = \boldsymbol{X}_j'\hat{\gamma} + \frac{\hat{\sigma}^2_{\alpha,j}}{\hat{\sigma}^2_{\alpha,j} + \hat{\sigma}^2_{OLS,j}} \left(\hat{\alpha}_j^{OLS} - \boldsymbol{X}_j'\hat{\gamma}\right),$$

$$\hat{\sigma}^2_{\alpha,j} = \max\left\{\frac{\sum_{j=1}^{N} \omega_j \left(\hat{u}_j^2 - \hat{\sigma}^2_{OLS,j}\right)}{\sum_{j=1}^{N} \omega_j}, \frac{2\sum_{j=1}^{N} \omega_j^2 \hat{\sigma}^4_{OLS,j}}{\sum_{j=1}^{N} \omega_j \sum_{j=1}^{N} \omega_j \hat{\sigma}^2_{OLS,j}}\right\}, \quad \hat{u}_j := \hat{\alpha}_j^{OLS} - X_j'\hat{\gamma},$$

where the weights ω_j are usually set to $\hat{\sigma}^{-2}_{OLS,j}$ or $1/n$.

We refer to the normality assumption in equation (4) as a "working assumption" because the EB estimate in equation (5) neither depends on normality nor requires a Bayesian interpretation of α_j. The EB estimator has lower MSE, averaged across units, than the unshrunk unbiased estimators, even when the individual effects are treated as nonrandom (James and Stein, 1992).

We now describe our empirical approach to identify which family factors best predict absolute mobility.

D. *Analyzing the Relationship Between Family Background and Absolute Mobility*

To identify which family factors best predict absolute mobility, we build upon the empirical approach of Chetty et al. (2014) and Chetty and Hendren (2018). In the former study, the authors correlate a measure of absolute mobility with local area characteristics to examine why some CZ areas are more (upwardly) mobile than others. In the latter, the causal effect of growing up in a particular commuting zone on children's adult salaries is estimated using a parametric shrinkage estimator. Subsequently, the authors identify the characteristics of places that produce high levels of upward mobility by regressing the causal estimates on observables.

We analyze family background's influence on absolute mobility by relating the EB estimates to family characteristics. To this end, we utilize the SHapley Additive exPlanations (SHAP) values from the explainable artificial intelligence literature (Lundberg, 2017). This approach builds upon the shapley values (Shapley, 1953), which measure the individual contribution of players in a cooperative game. By considering variables as players and models as coalitions in a predictive framework, SHAP values assess the influence of independent variables on the dependent variable. Consequently, the SHAP values in our framework allow us to pinpoint the most relevant family factors to explain absolute mobility and analyze their impact on children's permanent income.

Shapley values (Shapley, 1953) allocate the surplus in a cooperative game among players. Specifically, in a coalition C of N agents, the j-th shapley value represents the fair share of the coalition's value (V), that agent j should receive. It is defined as

$$\phi_j(V) = \frac{1}{N}\sum_S \frac{[V(S \cup \{j\}) - V(S)]}{\binom{N-1}{k_s}},$$

where the summation is over all the subsets S, of the team $T = \{1, ..., N\}$, that one can construct after excluding j, k_s is the number of agents in the coalition S, $V(S)$ is the value achieved by subteam S, and $V(S \cup \{m\})$ is the value after j joins S. Accordingly, the shapley value measures the average contribution of j, corresponding to her fair share.

The shapley values for regression measure the contribution of each explanatory variable in predicting the outcome for each observation. In this context, agents are the explanatory variables, and coalitions correspond to subsets of variables used in the model. Intuitively, shapley values for regression are calculated by comparing the model's predictions with and without each regressor to assess its impact. Due to the large number of possible combinations, the SHAP values rely on an approximation (see Lundberg (2017) for further details).

A key advantage of this approach is its ability to capture non-linear relationships through variable importance measures. For instance, if the benefits of higher parental education are magnified for nuclear families (a non-linear interaction), such patterns would remain opaque to linear models. SHAP values naturally quantify these complex effects by evaluating feature contributions across different variable combinations.

A conventional SHAP analysis of child income that jointly considers family characteristics and parental income would identify composite predictors of economic outcomes. However, our approach first residualizes child income with respect to parental income—a strategy that yields two critical advantages. First, it prevents parental income from dominating the explanatory power of the SHAP analysis. Second, it captures family background influences that are not transmitted via parental income.

By analyzing $m\left(\boldsymbol{X}_j\right) := \mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right]$, we aim to identify the key elements $\boldsymbol{Z}_j \subset \boldsymbol{X}_j$ to predict absolute mobility. To this end, we estimate the SHAP values from the machine learning (ML) regression of $\hat{\alpha}_j^{EB}$ from equation (5) on $\boldsymbol{X}_j$. These SHAP values then allow us to quantify feature importance.

In the context of ML, SHAP provides a unified measure of feature importance by computing Shapley values from cooperative game theory. The importance of feature j, denoted ϕ_j, is defined as:

$$\phi_j = \sum_{S \subseteq F \setminus \{j\}} \frac{|S|!(|F| - |S| - 1)!}{|F|!} \left[f_{S\cup\{j\}}(x_{S\cup\{j\}}) - f_S(x_S)\right]$$

where F is the set of all features, S is a subset of features excluding j, $|S|$ is the size of subset S, $|F|$ is the total number of features, f_S represents the model's prediction using feature subset S, and x_S are the values of the features in subset S. The SHAP value ϕ_j represents the average marginal contribution of feature j across different feature combinations. For our analysis, we normalize variable importance scores so that they sum to 100.

Our empirical approach can be summarized in the following three-step procedure.

1) Estimate by OLS the equation

$$\boldsymbol{Y}_j - \hat{\beta}\boldsymbol{Y}_j^P = \mathbf{1}_{n_j}\alpha_j + \boldsymbol{\epsilon}_j, \quad j \in \mathbb{S} \subseteq \{1, ..., N\},$$

to compute $\hat{\alpha}_j^{OLS}$ and $\hat{\sigma}_{OLS,j}$, where $\hat{\beta}$ is the pooled OLS estimate from the regression of $\boldsymbol{Y}_j$ on $\boldsymbol{Y}_j^P$.

2) Construct the empirical Bayes estimator by:

   a) Estimating the best linear predictor $\boldsymbol{X}_j'\hat{\gamma}$ via regression of $\hat{\alpha}_j^{OLS}$ on $\boldsymbol{X}_j$

   b) Computing shrinkage weights:

$$\hat{\sigma}^2_{\alpha,j} = \max\left\{\frac{\sum \omega_j(\hat{u}_j^2 - \hat{\sigma}^2_{OLS,j})}{\sum \omega_j}, \frac{2\sum \omega_j^2 \hat{\sigma}^4_{OLS,j}}{\sum \omega_j \sum \omega_j \hat{\sigma}^2_{OLS,j}}\right\}$$

   where $\hat{u}_j = \hat{\alpha}_j^{OLS} - \boldsymbol{X}_j'\hat{\gamma}$, and $\omega_j = \hat{\sigma}^{-2}_{OLS,j}$.

   c) Computing the EB estimate:

$$\hat{\alpha}_j^{EB} = \boldsymbol{X}_j'\hat{\gamma} + \frac{\hat{\sigma}^2_{\alpha,j}}{\hat{\sigma}^2_{\alpha,j} + \hat{\sigma}^2_{OLS,j}}(\hat{\alpha}_j^{OLS} - \boldsymbol{X}_j'\hat{\gamma}).$$

3) Obtain the SHAP values from the ML regression of $\hat{\alpha}_j^{EB}$ on $\boldsymbol{X}_j$ to compute feature importance and estimate the effect of the most relevant family characteristics on the predictions made by the model.

Incorporating ML into our methodology requires hyperparameter tuning. To this end, we utilize k-fold cross-validation combined with Bayesian Grid Search (Shcherbatyi et al., 2024), using the MSE as scoring metric.[7] The primary advantage of Bayesian optimization lies in its probabilistic modeling of the objective function, which allows

[7] In our application we consider two learners: Random Forest and XGBoost. For the Random Forest regressor, the parameters for the grid search include the number of trees in the forest, the maximum depth of the trees, the minimum number of samples required to split an internal node, the minimum number of samples required to be at a leaf node, and the maximum number of features to consider for splitting a node. For XGBoost, we optimize nine key hyperparameters: the learning rate, number of boosting rounds, maximum tree depth, minimum child weight, minimum loss reduction required for node splitting, L1 and L2 regularization, subsample ratio of training instances, and feature subsampling fraction per tree.

it to determine the most promising hyperparameters based on previous results. This approach significantly reduces the number of evaluations required, as it avoids the exhaustive search of every possible combination in a pre-defined grid.

IV. Simulations

This section presents a simulation study evaluating whether relative variable importance can be accurately recovered when fixed effects must be estimated, and how shrinkage methods mitigate estimation noise. Our analysis reveals that our three-step procedure successfully recovers both the correct ranking and relative importance of family background characteristics, as illustrated through comparison with an Oracle benchmark using true fixed effects. Notably, shrinkage, particularly toward the best linear predictor, yields substantial accuracy gains, reducing mean-squared error by adaptively mitigating noise in the fixed effect estimates.

We examine three empirical Bayes estimators. The first (BLP), shrinks $\hat{\alpha}_j^{OLS}$ towards its best linear predictor given $\boldsymbol{X}_j$, the second (RF) towards a non-parametric estimate from the Random Forest regression of $\hat{\alpha}_j^{OLS}$ on $\boldsymbol{X}_j$, and the third (PEB) towards the sample mean.

As a baseline, we also consider the raw OLS estimator, which applies no shrinkage (NS). This comparison enables us to assess the extent to which shrinkage mitigates the three main sources of noise in OLS estimates within our framework:(i) idiosyncratic errors $\boldsymbol{\epsilon}_j$, (ii) family-level shocks v_j, and (iii) family size n_j. First, $\hat{\alpha}_j^{OLS}$ is computed as the simple average of $\boldsymbol{Y}_j - \hat{\beta}\boldsymbol{Y}_j^P$, making it sensitive to within-family income variation. Since this variation is captured by the idiosyncratic error term ϵ_{ij}, extreme child-specific earnings can lead to its overestimation or underestimation. Second, the variance of v_j may generate outliers that distort covariate importance in SHAP analyses. Third, the small number of observed children per family $\left(n_j\right)$, means that each fixed effect estimate $\hat{\alpha}_j^{OLS}$ is based on very limited data.

We simulate data using the observed covariates $\boldsymbol{X}_j$ and parental income Y_j^P, generating children's permanent income Y_j setting $\beta = 0.5$ according to the following data generating process (DGP):

$$
\begin{aligned}
&\boldsymbol{Y}_j = \beta\boldsymbol{Y}_j^P + \mathbf{1}_{n_j}\alpha_j + \boldsymbol{\epsilon}_j, \quad j \in \mathbb{S} \subseteq \{1, ..., N\}, \quad \boldsymbol{\epsilon}_j \mid \boldsymbol{Y}_j^P, \alpha_j \sim \left(0, \sigma_\epsilon^2 \boldsymbol{I}_{n_j}\right), \\
&\alpha_j = m\left(\boldsymbol{X}_j\right) + v_j, \quad v_j \mid \boldsymbol{X}_j \sim \left(0, \sigma_{v_j}^2\right).
\end{aligned}
$$

We consider three distinct scenarios:

(i) **Gaussian, homoskedastic, linear, low-noise**:

$$
\begin{aligned}
\epsilon_{ij} \,|\boldsymbol{Y}_j^P, \alpha_j &\sim \mathcal{N}(0, 0.1) \\
v_j \,|\boldsymbol{X}_j &\sim \mathcal{N}(0, 0.1) \\
m\left(\boldsymbol{X}_j\right) &= 1.5 \times \text{race} + 2.5 \times \text{nuclear} + 0.07 \times \text{educ}.
\end{aligned}
$$

(ii) **Gaussian, heteroskedastic, linear, noisy**:

$$\begin{aligned}\epsilon_{ij} \,|\boldsymbol{Y}_j^P, \alpha_j &\sim \mathcal{N}(0,5)\\ v_j \,|\boldsymbol{X}_j &\sim \mathcal{N}\left(0, 1/n_{cj}\right)\\ m\left(\boldsymbol{X}_j\right) &= 1.5 \times \text{race} + 2.5 \times \text{nuclear} + 0.07 \times \text{educ},\end{aligned}$$

where n_{cj} is the number of children in the family according to the retrospective inter-generation map, which might differ from n_j, the number of children in family j in the sample.

(iii) **Non-Gaussian, non-linear, noisy**:

$$\begin{aligned}\epsilon_{ij} \,|\boldsymbol{Y}_j^P, \alpha_j &\sim t_7\left(0, \sqrt{25/7}\right)\\ v_j \,|\boldsymbol{X}_j &\sim t_7\left(0, \sqrt{25/7}\right)\\ m\left(\boldsymbol{X}_j\right) &= 1.5 \times \text{race} + 3.5 \times \text{nuclear} + 1.8 \times \text{nuclear}^2\\ &\quad + 0.1 \times \exp\{\text{educ}/5\} + 0.2 \times \text{educ} \times \text{race}.\end{aligned}$$

Here, the scaled t-distribution with $\nu = 7$ degrees of freedom and scale parameter $\sqrt{25/7}$ ensures the variance matches scenario (ii) $\left(\sigma_\epsilon^2 = 5\right)$.

For each scenario, we simulate data according to the corresponding DGP, and perform 500 bootstrap iterations by resampling $(\alpha_j, \hat{\alpha}_j^{OLS}, \hat{\sigma}_{OLS,j}, \boldsymbol{X}_j)$ N times with replacement, denoting $\left(\alpha_j^{(b)}, \hat{\alpha}_j^{OLS,(b)}, \hat{\sigma}_{OLS,j}^{(b)}, \boldsymbol{X}_j^{(b)}\right)$ the resampled dataset for iteration b. At each iteration b:

1) **EB estimators (BLP/RF/PEB)**: Implement steps 2–3 of our procedure using $(\alpha_j^{(b)}, \hat{\alpha}_j^{OLS,(b)}, \hat{\sigma}_{OLS,j}^{(b)}, \boldsymbol{X}_j^{(b)})$, differing only in shrinkage targets, with Random Forest regression for the final step.
2) **NS estimator**: Compute SHAP values directly from Random Forest regression of $\hat{\alpha}_j^{OLS,(b)}$ on $\boldsymbol{X}_j^{(b)}$.
3) **Oracle**: Compute SHAP values from Random Forest regression of $\alpha_j^{(b)}$ on $\boldsymbol{X}_j^{(b)}$, bypassing estimation steps.

Although only three covariates enter the DGP, we include all ten available covariates in our analysis to reflect realistic conditions.

In the first (low-noise) scenario, we anticipate minimal differences between $\hat{\alpha}_j^{EB}$ and $\hat{\alpha}_j^{OLS}$ for SHAP analysis, as both $\boldsymbol{\epsilon}_j$ and v_j exhibit low variance. In the second scenario, increased noise in $\boldsymbol{\epsilon}_j$ and heteroskedasticity in v_j reduce the accuracy of $\hat{\alpha}_j^{OLS}$. The simple average of $\boldsymbol{Y}_j - \hat{\beta}\boldsymbol{Y}_j^P$ becomes noisier, while heteroskedasticity causes families with more children to have more stable fixed effects estimates by design, as they exhibit lower variance. Since empirical Bayes adaptively shrinks estimates based on n_j, we expect it to outperform raw OLS in this case.

The third scenario introduces heavy-tailed shocks via a t-distribution. The empirical Bayes estimate relies on two normality assumptions: (1) the sampling distribution of $\hat{\alpha}_j^{OLS}$ and (2) the prior distribution of α_j. Neither holds here: n_j is too small to rely on the Central Limit Theorem, and we explicitly violate normality with t_7 errors. Moreover, simulating α_j from the same heavy-tailed distribution creates additional challenges for reliable inference. Furthermore, simulating α_j from a heavy-tailed $t-$distribution, causes both the likelihood and prior components of the empirical Bayes estimator deviate from their assumed normal forms, creating additional challenges for reliable inference.

We incorporate nonlinearities in $m(\boldsymbol{X}_j)$ to show that our procedure can reliably estimate variable importance despite nonlinear relationships. However, we note that shrinking $\hat{\alpha}_j^{OLS}$ toward a non-parametric estimate in step 2 does not improve upon shrinking toward the best linear predictor. While machine learning in step 3 captures nonlinearities, preliminary non-parametric shrinkage may induce regularization and model selection bias (Chernozhukov et al., 2022) both in $\hat{\alpha}_j^{EB}$ and for the SHAP analysis.

We assess estimator performance through two key metrics. First, we measure ranking accuracy—the frequency with which each method correctly orders variables by their relative importance across bootstrap replicates. Second, we compute the mean squared error (MSE) of the importance estimates, which captures both bias and precision in quantifying each variable's contribution. Both metrics use the Oracle estimates as ground truth, where the ranking reflects the true relative variable importance observed in each bootstrap sample.

Table 2—: Percentage of Correct Variable Rankings by Feature and Estimator, Across Scenarios

	Oracle	BLP	RF	PEB	NS
Scenario (i): Gaussian, homoskedastic, linear, low-noise					
Race	100	100	100	100	100
Family Structure	100	100	100	100	100
Education	100	100	100	100	100
Scenario (ii): Gaussian, heteroskedastic, linear, noisy					
Race	100	93.6	93.8	91.0	94.4
Family Structure	100	93.6	92.6	91.0	93.8
Education	100	83.0	76.0	78.8	78.8
Scenario (iii): Non-Gaussian, non-linear, noisy					
Race	100	100	100	100	100
Family Structure	100	100	100	99.8	99.8
Education	100	100	100	99.8	99.8

Table 2 presents the ranking accuracy results, revealing distinct patterns across simulation scenarios. Under the favorable conditions of Scenario (i), all estimators—including raw OLS—perfectly recover the true variable rankings (100% accuracy), indicating that shrinkage offers no improvement in this dimension when the setting is well-behaved. In Scenario (ii), ranking accuracy remains high across all estimators. The two most influential variables—Race and Family Structure—are correctly ranked in over 91% of bootstrap samples by every estimator. Accuracy for the less informative Education variable drops more substantially, ranging from 76.0% to 83.0%. The BLP estimator achieves slightly higher accuracy for this variable, but the differences across estimators are small. These results indicate that even in more challenging settings, the most important variables tend to be reliably identified, and any gains from shrinkage in terms of ranking are present but subtle.

Scenario (iii) introduces both non-Gaussian errors and nonlinearities, yet all estimators continue to perform remarkably well. BLP maintains perfect ranking accuracy for all variables, while PEB and NS only misclassify Family Structure and Education in one bootstrap sample out of 500. Overall, the results suggest that ranking performance is generally resilient to noise, non-Gaussianity, heteroskedasticity and non-linearities, with only marginal gains from shrinkage in more complex settings. While the correct ranking is unsurprising given that the three scenarios only includes three variables, the consistent performance across challenging conditions remains noteworthy.

Table 3—: Oracle Mean, Bias, Variance, and MSE by Feature and Estimator, Across Scenarios

		Bias				**Variance**				**MSE**			
Feature	Oracle Mean	BLP	RF	PEB	NS	BLP	RF	PEB	NS	BLP	RF	PEB	NS
Scenario (i): Gaussian, homoskedastic, linear, low-noise													
Race	51.75	-1.45	-1.01	-2.04	-1.99	2.29	2.72	2.37	2.37	4.38	3.75	6.52	6.33
Family Structure	34.32	0.94	0.72	0.07	0.10	3.60	3.78	3.60	3.61	4.48	4.30	3.61	3.62
Education	11.83	-1.66	-2.08	-2.05	-2.11	1.69	2.21	1.90	1.94	4.44	6.54	6.13	6.37
Scenario (ii): Gaussian, heteroskedastic, linear, noisy													
Race	52.07	-5.18	-5.40	-12.02	-10.72	15.49	31.75	21.43	21.37	42.28	60.89	165.94	136.34
Family Structure	33.70	-1.12	-5.94	-7.07	-8.37	21.64	36.75	23.30	24.89	22.88	72.08	73.34	94.91
Education	11.46	-0.67	-2.26	-1.32	-2.15	14.63	12.94	6.98	7.84	15.07	18.05	8.72	12.49
Scenario (iii): Non-Gaussian, non-linear, noisy													
Race	55.18	-2.22	-2.84	-4.17	-4.23	4.68	5.55	4.93	4.88	9.62	13.63	22.29	22.77
Family Structure	27.65	-1.04	-1.19	-1.81	-1.91	5.53	6.46	5.96	5.95	6.61	7.88	9.24	9.59
Education	12.46	-0.07	-0.06	-0.51	-0.19	4.61	5.15	4.21	4.12	4.62	5.16	4.47	4.16

Table 3 presents the bias, variance, and MSE of estimated relative variable importance across bootstrap iterations, broken down by feature and estimator. All results use the Oracle relative importance at each bootstrap iteration as the ground truth across our three scenarios.

In the low-noise setting of Scenario (i), all estimators closely approximate the Oracle benchmark, with only minor differences in variance. For example, Race importance estimates range narrowly from 49.72 (PEB) to 50.74 (RF), with consistently low variances.

As a result, MSE differences are modest—RF achieves the lowest Race MSE (3.75), edging out BLP (4.38) by just 0.63; PEB performs slightly better in Family Structure (3.61 vs. BLP's 4.48); and BLP holds a modest advantage in Education (4.44 vs. PEB's 6.13). These findings align with theoretical expectations: when estimation noise is negligible, the unshrunk OLS estimator performs comparably to shrinkage-based alternatives, offering little room for improvement through regularization. However, the more realistic Scenarios (ii) and (iii) reveal BLP's consistent superiority.

While BLP underestimates the most important variable in Scenario (ii), it exhibits a remarkable precision-accuracy tradeoff. For the two most important variables, Race and Family Structure, it achieves both the lowest variance (15.49 and 21.64) and mean estimates closest to the Oracle (46.89 vs. 52.07; 32.59 vs. 33.70). Although its variance for Education is slightly higher than RF's (14.63 vs. 12.94), BLP offers better accuracy (10.79 vs. 11.46). In contrast, PEB and NS substantially underestimate key predictors, with larger variances and pronounced bias, e.g., Race importance drops to 40.05 (PEB) and 41.35 (NS). These advantages translate into substantially lower MSEs for BLP: for Race, its MSE (42.28) is nearly four times smaller than PEB's (165.94) and three times smaller than NS's (136.34); for Family Structure, BLP outperforms NS by a factor of four (22.88 vs. 94.91). Although RF's flexibility may seem beneficial in theory, in practice it results in overfitting, with MSEs 45% higher for Race (60.89 vs. 42.28) and over three times higher for Family Structure (72.08 vs. 22.88). Overall, BLP's parametric shrinkage provides robust performance, maintaining both stability and accuracy in a challenging setting.

Scenario (iii) provides compelling evidence for the robustness of shrinkage estimation. Despite non-Gaussian error distributions, BLP maintains superior performance, with mean estimates closest to the Oracle benchmark across all variables[8] and consistently achieves the lowest variance among all estimators (except in Education, where PEB and NS perform slightly better). The MSE results are particularly striking: BLP attains 9.62 for Race (versus NS's 22.77), 6.61 for Family Structure (versus 9.59), and 4.62 for Education (versus 4.16), demonstrating remarkable resilience to violated assumptions. While RF shows modest improvements over PEB and NS, its higher variance (5.55 versus BLP's 4.68 for Race) highlights the inherent accuracy-precision tradeoff. Crucially, BLP's advantage is most pronounced for the most influential variables, where it achieves substantially lower MSE. These findings indicate that shrinkage preserves significant advantages over unshrunk OLS even under non-Gaussian conditions, though the benefits are somewhat less dramatic than in Scenario (ii)'s heteroskedastic setting.

This last setting illustrates that nonparametric shrinkage offers no advantage over linear shrinkage, despite the true nonlinear DGP. While machine learning in step 3 captures nonlinear patterns, preliminary nonparametric shrinkage in step 2 introduces regularization and model selection bias (Chernozhukov et al., 2022) both in $\hat{\alpha}_j^{EB}$ and for the SHAP analysis. This result suggests that shrinking toward the best linear predictor yields more reliable estimates, even when the true relationship is nonlinear. Moreover, BLP's lower variance relative to RF across the three scenarios reflects the fundamental trade-off be-

[8]The only exception is Education, where BLP and RF show negligible differences.

tween parametric models and flexible ML methods (James et al., 2021).

Our simulation exercise highlights the robustness and stability of the BLP estimator across a range of data-generating environments. While gains are modest under ideal conditions, BLP consistently outperforms alternative estimators as noise increases or classical assumptions break down. Its ability to balance bias and variance makes it particularly well suited for applied settings where data are limited and estimation error is a concern. The results underscore the practical value of shrinking towards the best linear predictor, offering stable and interpretable estimates of variable importance even in complex, noisy, or misspecified environments.

The results in Table 3 are consistent with theoretical expectations: EB estimation minimizes Bayes risk and achieves MSE-optimality, yielding accurate predictions of unobserved heterogeneity (Gaillac, 2023). However, this superior predictive performance does not automatically translate to better variable importance estimation. SHAP analysis evaluates importance through mean absolute SHAP values, a distinct metric that may not directly reflect MSE optimization. While EB improves global prediction accuracy, its shrinkage properties may inadvertently affect the relative importance of individual features in finite samples.

The key finding from our simulations is that EB estimation offers meaningful advantages in our framework—particularly in noisy and heteroskedastic settings. These benefits are most pronounced when shrinking toward the best linear predictor, an approach that combines theoretical rigor with operational stability. Although the relationship between MSE and SHAP values requires nuanced interpretation, the EB approach delivers more stable estimates of relative variable importance while being theoretically grounded. In summary, these results support our proposed procedure for identifying the family characteristics that best predict absolute mobility, highlighting its potential for empirical research on intergenerational mobility.

V. Results

Our analysis reveals race and family structure as the strongest predictors of absolute mobility. As depicted in Figure 1 (a), SHAP analysis indicates that, conditional on parental income, race (white vs. non-white) explains 35% of the total variable importance among family background characteristics. Being raised in a nuclear family, defined as a household with a head, a spouse (if present), and children, accounts for 22.4%. Other relevant factors include the household's age at first birth (10.5%), number of children (9.4%), and education (9.3%). Other factors show negligible predictive power, collectively accounting for less than 14% of the explained variation. The results regarding family structure gain additional significance when contextualized within the dramatic 19-percentage-point decline in two-parent households between 1960 (87.7%) and 2016 (68.7%) documented by U.S. Census Bureau (2017).

Figure 1 (b) displays results that additionally control for birth cohort and regional fixed effects. As a robustness check, we account for generational differences (birth cohorts: pre-1940 [reference], 1940–49, 1950–59, 1960+) and geographic variation (U.S. regions: North Central, Northeast, West, with South as reference). These controls help isolate the

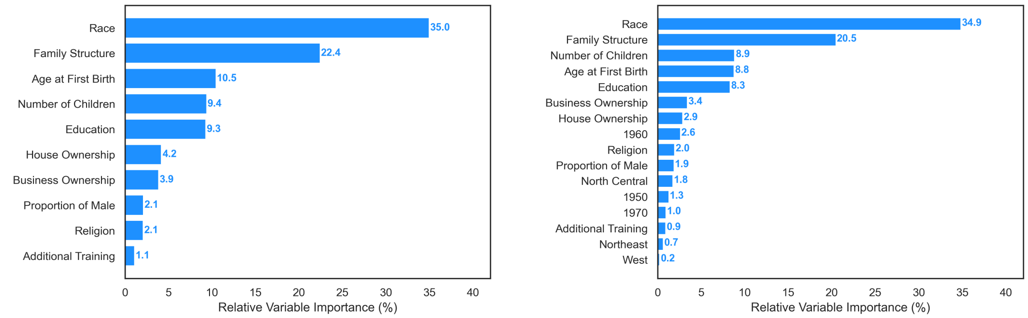

(a) Baseline

(b) Controlling for Cohort and Regional Fixed Effects

Figure 1. : Relevance of Family Background Characteristics in Shaping Absolute Mobility

The importance of explaining absolute mobility is computed according to the average SHAP values (in absolute value), and normalized so that the sum of variable importance adds up to 100.

relationship between family characteristics and absolute mobility from temporal trends or regional disparities. The variables ranked third to fifth show some reordering due to their nearly equivalent contributions. However, the top two predictors (race and family structure) maintain their rank order. These results indicate that temporal and regional factors contribute relatively little to explaining absolute mobility after accounting for family background characteristics.

Previous research has highlighted the importance of family structure in shaping intergenerational mobility. Kearney (2023) examines the widening economic disparities linked to family structure, emphasizing that children raised in stable two-parent households tend to achieve higher educational attainment and greater upward mobility. While past research has emphasized the challenges faced by children of single mothers, Kearney (2023) underscores how the decline in stable marriages has deepened class divisions, noting that the economic and social advantages of two-parent families are now increasingly concentrated among higher-income households, thereby exacerbating inequality.

These findings align with McLanahan and Percheski (2008), who highlights the role of family structure in perpetuating class, race, and gender inequalities, and with Chetty et al. (2014), who finds that regions with weaker family structures exhibit significantly lower levels of intergenerational mobility. Similarly, Bloome (2017) underscores the long-term economic consequences of family instability, showing that children raised outside stable two-parent families are less likely to attain high-income status in adulthood compared to their peers from stable two-parent households.

Our descriptive analysis offers a distinct perspective on these well-documented relationships. Whereas much existing work compares single- versus two-parent households, our methodology captures the relative importance of nuclear family structure within a broader set of mobility determinants. The observed patterns suggest that the documented decline in two-parent households may relate to mobility outcomes through multiple compositional changes, not only through increases in single-parent families but potentially through other family structures as well. In particular, our findings suggest that the role of family structure in intergenerational mobility extends beyond the single- versus two-parent household distinction.

The relevance of family structure stands out, particularly since it is conditional on parental income. In the absence of this adjustment, one might attribute the association to income dilution, where resources are spread more thinly in certain household structures.However, by excluding this pathway, the findings emphasize the relevance of family structure through non-monetary channels.

Figure 1 reinforces well-documented racial disparities in intergenerational mobility. Hertz (2005) highlights the low rate of upward mobility for black families at the bottom of the income distribution, which is further corroborated by Bloome (2014), who shows that black children are more likely than white children to remain in the lowest income quintile. Moreover, Chetty et al. (2014) finds that both blacks and whites living in areas with large African American populations, such as the Southeast of the U.S., experience lower rates of upward mobility. Segregation is identified as a key mechanism behind this trend, with more racially segregated areas exhibiting lower mobility and underscoring the importance of racial context in shaping mobility outcomes. Building on this, Davis and Mazumder (2018) explores the geographic dimension of mobility, revealing that the low mobility in the Southeast, as documented by Chetty et al. (2014), is primarily driven by the low mobility of whites. In contrast, blacks raised in the Southeast experience higher mobility than those raised in the Northeast and Midwest.

Beyond supporting the well-documented racial gaps in intergenerational mobility, our results underscore the crucial role of family structure, which emerges as the second strongest predictor—a dynamic that is particularly striking in light of racial disparities in family composition. Notably, 68% of white children in our sample were consistently raised in nuclear families, compared to just 35% of non-white children (see Figure 3). Together, these findings align with existing work on how family structure differences may contribute to persistent racial mobility gaps (McLanahan and Percheski, 2008), offering new descriptive evidence into this ongoing conversation.

Figure 2 presents SHAP scatter plots for the five most influential predictors of absolute mobility. Each plot illustrates the marginal relationship between a given variable and its impact on mobility, quantified by SHAP values. Individual observations (families) are represented as points, where the x-axis indicates the observed value of the predictor variable and the y-axis shows the corresponding SHAP value, a measure of the variable's directional effect on mobility. Positive SHAP values reflect positive shifts in absolute mobility outcomes, while negative values the opposite. The dispersion of points reveals heterogeneity in how these key factors influence outcomes across families.

The top-left panel in Figure 2 reveals stark racial patterns in absolute mobility. Non-white individuals cluster in the negative SHAP value range, reflecting systematically

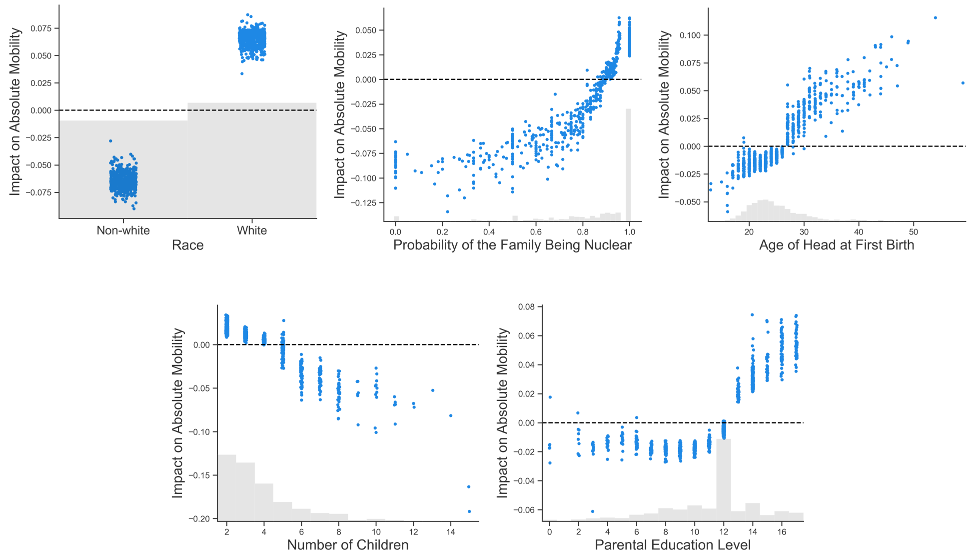

Figure 2. : SHAP scatter plots for the five most influential variables in predicting absolute mobility.

Note: the impact on absolute mobility is measured by the SHAP value for each of the five variables in the figure. The dots in each figure depict the SHAP value of the corresponding variable for an individual in the sample, highlighting that the effect of each characteristic on absolute mobility is heterogeneous across individuals.

lower mobility outcomes compared to their white counterparts, who show consistently positive SHAP values. In the top-center panel, the probability of being raised a nuclear family shows a nonlinear relationship with mobility. Below a 0.9 probability, SHAP values are predominantly negative. Beyond this threshold, the trend reverses sharply: households with higher probabilities of being nuclear exhibit positive SHAP values, suggesting that a nuclear family structure is associated with mobility advantages.

The top-right panel tracks household head age at first birth, revealing a gradual transition from negative to positive SHAP values. Early parenthood (before age 27) aligns with worse mobility outcomes, while delayed childbearing correlates with progressively higher SHAP values. These findings align with established literature documenting that children of younger parents face greater risks of health, social, and economic disadvantages compared to those raised by older parents (Hofferth and Reid, 2002).

The bottom-left panel in Figure 2 examines the relationship between the number of children in a family and absolute mobility. The results indicate a nuanced pattern: families with 2-4 children show a modest positive association with mobility, though the effect size is relatively small. In contrast, families with more than four children show a clear negative association with mobility. Similar to our findings for nuclear families, these results operate independently of income effects, indicating that the observed patterns reflect more than just income dilution.

Finally, the bottom-left panel analyzes parental education level. SHAP values remain mostly negative for individuals whose parents have fewer than 12 years of education, highlighting that lower parental education is associated with lower mobility. However, once parental education reaches 12 years (high school completion), SHAP values become positive, and they continue to rise as education increases. Given that parental income is held constant in our analysis, these results suggest the observed positive association likely captures non-monetary advantages associated with education (Vila, 2005; Heckman et al., 2018).

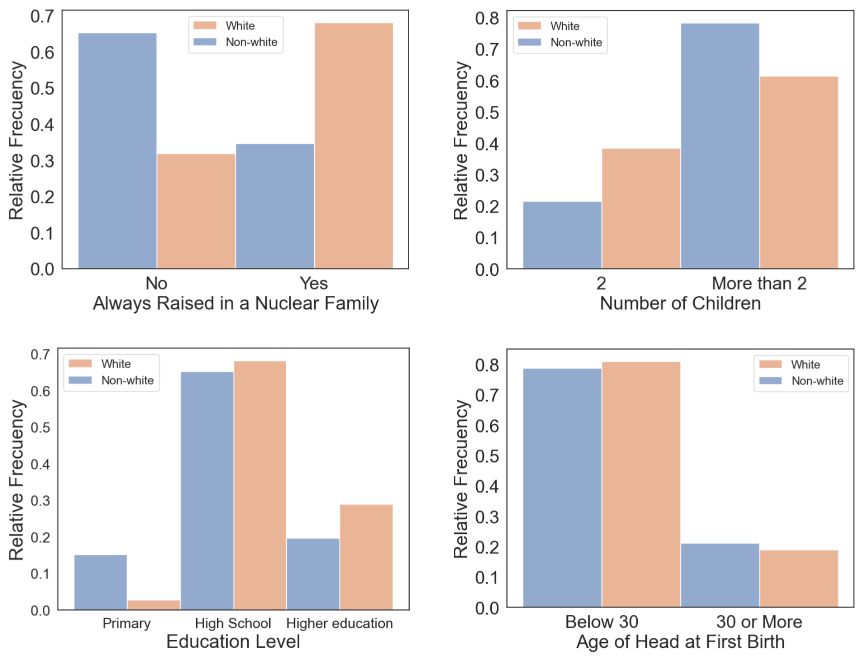

Figure 3. : Histogram for the Most Relevant Characteristics by Race

Note: Primary includes primary school or less, High School covers more than primary but no higher than high school, and Higher Education refers to any education beyond high school.

The data reveal pronounced racial disparities in family structure, educational attainment, and the number of children, though we find no systematic differences in the head's age at first birth. As shown in Figure 3, white children are far more likely to be raised in nuclear families (68%) compared to non-white children (35%). Additionally, while 38% of white parents have two children, only 22% of non-white parents do. Educational attainment also varies by race. Although the proportion of parents with more than primary but no higher than a high school education is similar across groups (65% for whites vs. 68% for non-whites), disparities emerge in other categories. Non-white households are overrepresented in primary education (15% vs. 3% for white households), whereas white households are more likely to have higher education (29% vs. 20%). These racial disparities in family structure and educational attainment may jointly contribute to widening gaps in absolute mobility.

VI. Conclusions

This paper provides a descriptive analysis of how family background influences the permanent income of children born in low-income households. Our proposed measure of mobility corresponds to the family fixed effect in the regression of children's permanent income on parental permanent income. This object captures the families' ability to boost their children's permanent income, conditional on parental income. To address estimation challenges arising from limited data for families, we employ an Empirical Bayes approach that improves the accuracy of our fixed effect estimates. We then analyze the relative importance of different family background characteristics using SHapley Additive exPlanations (SHAP), a unified framework for interpreting predictive models that allows us to quantify the contribution of each factor while accounting for potential nonlinearities and interactions.

Our analysis suggests that racial disparities in mobility persist, with family structure playing a significant role. In particular, we identify race and family structure as the most powerful predictors of absolute mobility, with race explaining 35% of predictive importance and nuclear family structure accounting for 22.4%, conditional on parental income. Notably, the difference in nuclear family prevalence between white (68%) and non-white (35%) children sheds light on the connection between family structure and racial mobility gaps. Other influential factors include the household's age at first birth, the number of children and parental education. Importantly, our methodology holds parental income constant, isolating the effects of family background and excluding income-mediated channels.

Given the descriptive nature of our findings, future research should investigate the causal mechanisms linking family background to absolute mobility. The framework developed by Gaillac (2023) provides a promising approach, as it enables estimation of individual-level causal effects in settings like ours. Applying this approach could yield important insights regarding two key questions: (1) how the decline of two-parent households in the U.S. has affected intergenerational mobility, and (2) the extent to which family structure differences contribute to persistent racial disparities in mobility outcomes. Such analysis would provide a deeper understanding of how family background shapes economic opportunity across generations.

References

Armstrong, T. B., Kolesár, M., and Plagborg-Møller, M. (2022). Robust empirical bayes confidence intervals. *Econometrica*, 90(6):2567–2602.

Bingley, P., Cappellari, L., and Tatsiramos, K. (2017). Family, community and life-cycle earnings: Evidence from siblings and youth peers. Technical Report 6743, Center for Economic Studies & Ifo Institute.

Björklund, A., Jäntti, M., and Solon, G. (2007). Nature and nurture in the intergenerational transmission of socioeconomic status: Evidence from swedish children and their biological and rearing parents. *The BE Journal of Economic Analysis & Policy*, 7(2).

Bloome, D. (2014). Racial inequality trends and the intergenerational persistence of income and family structure. *American sociological review*, 79(6):1196–1225.

Bloome, D. (2017). Childhood family structure and intergenerational income mobility in the united states. *Demography*, 54(2):541–569.

Bonhomme, S. and Weidner, M. (2022). Posterior average effects. *Journal of Business & Economic Statistics*, 40(4):1849–1862.

Chernozhukov, V., Escanciano, J. C., Ichimura, H., Newey, W. K., and Robins, J. M. (2022). Locally robust semiparametric estimation. *Econometrica*, 90(4):1501–1535.

Chetty, R. and Hendren, N. (2018). The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228.

Chetty, R., Hendren, N., Kline, P., and Saez, E. (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics*, 129(4):1553–1623.

Cholli, N. A., Durlauf, S. N., Landersø, R., and Navarro, S. (2024). Understanding the heterogeneity of intergenerational mobility across neighborhoods. Technical report, National Bureau of Economic Research.

Corcoran, M., Gordon, R. H., Laren, D., and Solon, G. (1990). Effects of family and community background on economic status. *The American Economic Review*, 80(2):362–366.

Corcoran, M., Jencks, C., and Olneck, M. (1976). The effects of family background on earnings. *The American Economic Review*, 66(2):430–435.

Davis, J. and Mazumder, B. (2018). Racial and ethnic differences in the geography of intergenerational mobility. *Available at SSRN 3138979*.

Efron, B. (2012). *Large-scale inference: empirical Bayes methods for estimation, testing, and prediction*, volume 1. Cambridge University Press.

Endelman, J. B. and Jannink, J.-L. (2012). Shrinkage estimation of the realized relationship matrix. *G3: Genes— genomes— genetics*, 2(11):1405–1413.

Gaillac, C. (2023). Predicting unobserved individual-level causal effects.

Gilraine, M., Gu, J., and McMillan, R. (2020). A new method for estimating teacher value-added. Technical report, National Bureau of Economic Research.

Gu, J. and Koenker, R. (2017). Unobserved heterogeneity in income dynamics: An empirical bayes perspective. *Journal of Business & Economic Statistics*, 35(1):1–16.

Heckman, J. J., Humphries, J. E., and Veramendi, G. (2018). The nonmarket benefits of education and ability. *Journal of human capital*, 12(2):282–304.

Hertz, T. (2005). Rags, riches, and race. *Unequal chances: Family background and economic success*, 165:165–91.

Hofferth, S. L. and Reid, L. (2002). Early childbearing and children's achievement and behavior over time. *Perspectives on Sexual and Reproductive Health*, pages 41–49.

Jacob, B. A. and Lefgren, L. (2008). Can principals identify effective teachers? evidence on subjective performance evaluation in education. *Journal of labor Economics*, 26(1):101–136.

James, G., Witten, D., Hastie, T., and Tibshirani, R. (2021). *An Introduction to Statistical Learning with Applications in R*. Springer Texts in Statistics. Springer, 2nd edition. ISBN 978-1-0716-1417-4.

James, W. and Stein, C. (1992). Estimation with quadratic loss. In *Breakthroughs in statistics: Foundations and basic theory*, pages 443–460. Springer.

Kane, T. J. and Staiger, D. O. (2008). Estimating teacher impacts on student achievement: An experimental evaluation. Technical report, National Bureau of Economic Research.

Kearney, M. S. (2023). The two-parent privilege: How americans stopped getting married and started falling behind. In *The Two-Parent Privilege*. University of Chicago Press.

Lundberg, S. (2017). A unified approach to interpreting model predictions. *arXiv preprint arXiv:1705.07874*.

McLanahan, S. and Percheski, C. (2008). Family structure and the reproduction of inequalities. *Annu. Rev. Sociol*, 34(1):257–276.

Shapley, L. S. (1953). A value for n-person games. *Contribution to the Theory of Games*, 2.

Shcherbatyi, I., Louppe, G., and Head, T. (2024). Scikit-learn hyperparameter search wrapper. https://scikit-optimize.github.io/stable/auto_examples/sklearn-gridsearchcv-replacement.html. Accessed: 2025-02-27.

U.S. Census Bureau (2017). Living arrangements of children under 18 years old: 1960 to present. Retrieved from U.S. Census Bureau website.

van Houwelingen, H. C., Brand, R., and Louis, T. A. (2020). Empirical bayes methods for monitoring health care quality. *arXiv preprint arXiv:2009.03058*.

Vila, L. E. (2005). The outcomes of investment in education and people's well-being. *European journal of education*, 40(1):3–11.

Appendix

Proof of Theorem 1

To identify the expectation of the family fixed effect for below-median income families $\mathbb{E}\left[\alpha_j \mid \mathbb{S}\right]$, we start by defining the following two objects:

$$\boldsymbol{Q}_j = \boldsymbol{I}_{n_j} - \frac{1}{n_j}\mathbf{1}_{n_j}\mathbf{1}'_{n_j} \quad \text{(within-family centering matrix)}$$
$$\boldsymbol{H}_j = \frac{1}{n_j}\mathbf{1}'_{n_j} \quad \text{(family averaging operator).}$$

The row vector $\boldsymbol{H}_j$ allows us to isolate family effects by exploiting the conditional mean independence assumption:

$$\begin{aligned}\mathbb{E}\left[\boldsymbol{H}_j(\boldsymbol{Y}_j - \beta\boldsymbol{Y}^P_j) \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] &= \mathbb{E}\left[\boldsymbol{H}_j\mathbf{1}_{n_j}\alpha_j + \boldsymbol{H}_j\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] \\ &= \mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] + \mathbb{E}\left[\boldsymbol{H}_j\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] \\ &= \mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right],\end{aligned}$$

where the last equality follows because:

$$\begin{aligned}\mathbb{E}\left[\boldsymbol{H}_j\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] &= \boldsymbol{H}_j\mathbb{E}\left[\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] \\ &= \mathbf{0} \quad \text{(by Assumption 2).}\end{aligned}$$

We now integrate out the parental income variation to obtain the conditional expectation:

$$\begin{aligned}\mathbb{E}\left[\mathbb{E}\left[\boldsymbol{H}_j(\boldsymbol{Y}_j - \beta\boldsymbol{Y}^P_j) \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] \mid \mathbb{S}\right] &= \mathbb{E}\left[\mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}^P_j, \mathbb{S}\right] \mid \mathbb{S}\right] \\ &= \mathbb{E}\left[\alpha_j \mid \mathbb{S}\right].\end{aligned}$$

The key insight is that by appropriately averaging outcomes net of the parental income component (weighted by β), we can isolate the family fixed effect. Thus, identification of β is necessary for identifying the distributional characteristics of α_j.

The centering matrix $\boldsymbol{Q}_j$ identifies β through the moment condition:

$$\begin{aligned}\mathbb{E}\left[(\boldsymbol{Y}^P_j)'\boldsymbol{Q}_j\boldsymbol{\epsilon}_j\right] &= 0 \quad \text{(by Assumption 2)} \\ \mathbb{E}\left[(\boldsymbol{Y}^P_j)'\boldsymbol{Q}_j(\boldsymbol{Y}_j - \mathbf{1}_{n_j}\alpha_j - \boldsymbol{Y}^P_j\beta)\right] &= 0\end{aligned}$$

By construction, $\boldsymbol{Q}_j\mathbf{1}_{n_j} = \mathbf{0}$, so the last display boils down to:

$$\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j\right] = \beta\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j^P\right],$$

yielding the population estimand for β:

$$\beta = \frac{\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j\right]}{\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j^P\right]}$$

Thus, for identification of β, we require:

$$\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j^P\right] \neq 0.$$

This condition is guaranteed by Assumption 2, which states that that $\boldsymbol{Y}j^P \notin \text{span}(\mathbf{1}n_j)$ almost surely—meaning parental income must exhibit within-family variation.

This condition has two immediate implications. First, it requires that families have multiple children ($n_j > 1$), since when $n_j = 1$, $\boldsymbol{Q}_j = \mathbf{0}$ and within-family variation cannot be assessed. Second, since the centering matrix $\boldsymbol{Q}_j$ projects any vector onto the subspace orthogonal to $\mathbf{1}_{n_j}$, when $\boldsymbol{Y}_j^P$ lies entirely within the span of $\mathbf{1}_{n_j}$ — that is, if $\boldsymbol{Y}_j^P = \mathbf{1}_{n_j} Y_j^P$ for some scalar Y_j^P — then

$$\mathbb{E}\left[(\boldsymbol{Y}_j^P)'\boldsymbol{Q}_j\boldsymbol{Y}_j^P\right] = 0.$$

Thus, when parental permanent income is constant across siblings β is not identified in equation (1). As a result, the distributional characteristics of α_j are not identified either.

Proof of Corollary 1

By the conditional mean independence $\mathbb{E}\left[\epsilon_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] = 0$ we have that

$$\begin{aligned}
\mathbb{E}\left[\hat{\alpha}_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] &= \mathbb{E}\left[\frac{1}{n_j}\mathbf{1}'_{n_j}\left(\boldsymbol{Y}_j - \beta\boldsymbol{Y}_j^P\right) \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] \\
&= \mathbb{E}\left[\frac{1}{n_j}\mathbf{1}'_{n_j}\left(\mathbf{1}'_{n_j}\alpha_j + \boldsymbol{H}_j\boldsymbol{\epsilon}_j\right) \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] \\
&= \mathbb{E}\left[\alpha_j + \frac{1}{n_j}\mathbf{1}'_{n_j}\boldsymbol{\epsilon}_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right] \\
&= \mathbb{E}\left[\alpha_j \mid \boldsymbol{Y}_j^P, \boldsymbol{X}_j, \mathbb{S}\right]. Q.E.D.
\end{aligned}$$