

Intergenerational Income and Wealth Mobility in Spain: Geographic Analysis and Causal Place Effects

Javier Soria
Opportunity Insights, Harvard University

Octavio Medina
Google

September 9, 2026

Abstract

We provide the first estimates of intergenerational income and wealth mobility in Spain, using administrative tax records that link parents and children over more than two decades. Income persistence (a rank–rank slope of 0.274) and the probability of moving from the bottom to the top of the income distribution (10.4%) place Spain in the middle of the international ranking, but persistence at the top is exceptional: children born into the top 1% are roughly fifty times as likely to reach it in adulthood as children from the bottom decile. Linking these records to the universal registry of inheritances and gifts, we document that family transfers contribute substantially to this persistence at the top. At the scale at which families actually live, local samples tend to be too small for conventional estimation, a constraint that binds in most countries. We develop two Bayesian hierarchical estimators that separate signal from sampling noise by borrowing information across nested geographic levels, in proportion to how precisely each area is measured. Together they yield reliable mobility estimates at fine geographic levels and make it possible to recover causal place effects from a movers design at that same resolution. Upward mobility follows a marked northeast–southwest regional gradient, together with large variation across municipalities and neighborhoods in large cities. This local variation is most strongly associated with cross-class economic connectedness, local labor-market conditions and educational attainment, and 64% of the observed gap in upward mobility across municipalities is attributable to causal effects of childhood exposure to the places where children grow up.

Keywords: intergenerational mobility, inequality, causal place effects

JEL codes: E24, J16, J31, J61, J62, R1

This paper has benefited from helpful feedback and suggestions from Isabel Araújo, Jaime Arellano-Bover, Miguel Artola, Luis Bauluz, Raj Chetty, John Friedman, Rafael Frutos, Catherine Haeck, Ander Iraizoz, Drew Johnston, Kiko Llaneras, Rocío Martínez-Sampere, Clara Martínez-Toledano, Unai Oyón, Thomas Piketty, Hillel Rapoport, Jan Stuhler, and Tom Zohar. We are grateful to the Spanish Tax Agency (Agencia Estatal de Administración Tributaria, AEAT) for granting us access to the data. This research was funded by the Bank of Spain Graduate Fellowship, the Paris School of Economics, and a French government subsidy managed by the Agence Nationale de la Recherche (ANR) under the framework of the Investissements d’avenir programme, reference ANR-17-EURE-001. Himal Bamzai-Wokhlu, Hannah Guan, Shiloh Liu, and Lea Pacou provided fantastic part-time research assistance at different stages of the project. All errors remain our own.

1 Introduction

Equality of opportunity—the principle that a child’s background should not determine her economic future—matters both for fairness and for efficiency: societies in which prospects depend heavily on parental resources are judged less fair, and they fail to develop the talent of children born into disadvantaged families (Chancel et al., 2026; Piketty, 2020; Aghion et al., 2017; Bell et al., 2019). Traditionally, measuring how far the principle holds in practice long relied on household surveys, which observe parents and children over short portions of their working lives and leave estimates sensitive to transitory income and life-cycle bias (Solon, 1999, 2002; Black and Devereux, 2010; Blanden, 2013; van der Weide et al., 2021).

Linked administrative tax records transformed the intergenerational mobility literature, first by allowing persistence to be measured over long income windows and across the full distribution (Corak and Heisz, 1999; Österberg, 2000; Dahl and DeLeire, 2008), and then, with population-wide data, by moving beyond national averages altogether. Chetty et al. (2014b) documented that upward mobility varies sharply across the places where children grow up and identified the characteristics of the places where children from low-income families do best, and Chetty and Hendren (2018b,c) exploited variation in children’s age at move to show that much of this variation reflects the causal effect of childhood exposure to a place rather than the sorting of families across places. The framework has since been applied in a dozen countries, and, together with the growing availability of large-scale administrative data, it has reshaped what is known about the determinants of upward mobility, particularly how much the local environment in which a child grows up shapes her future adult outcomes.¹ Spain, one of the largest economies in Europe, has remained absent from it: a country with universal health coverage and near-free public higher education, and with income and wealth inequality close to the European average, yet with household wealth unusually concentrated in housing, a volatile labor market, long school-to-work transitions, and deep regional divides. (Eurostat, 2024; Chancel et al., 2026).

In this paper we provide the first estimates of intergenerational income and wealth mobility in Spain based on large-scale administrative data, at geographic levels ranging from the region down to the ZIP code. The data are a new panel, built from the records of the Spanish Tax Agency, that follows parents and children of the same families over 25 years.² Parents are linked to their children through the household composition of the 1998 income-tax filings, when the children were still dependents, and both generations are then observed annually through 2022. Income is measured comprehensively in each year for both generations—labor earnings, self-employment income, capital income, and public transfers—and the residence is recorded alongside it. We add two further sources covering the same individuals: wealth records from 2016 onwards, which

¹Canada (Connolly et al., 2019a), Italy (Acciari et al., 2022), Australia (Deutscher and Mazumder, 2020), Sweden (Heidrich, 2017), Switzerland (Chuard, 2021), Denmark (Eriksen and Munk, 2020), France (Kenedi and Sirugue, 2022), the Netherlands (Beekers, 2024), Ecuador (Pozo and Moreno, 2025) and Brazil (Britto et al., 2026), among others.

²An earlier version of this analysis, covering a shorter period, was first analyzed by Soria-Espin (2022). The sample excludes the Basque Country and Navarre, whose foral tax regimes are administered separately.

report the assets held by parents and by children, and the universal registry of the inheritance and gift tax, which records every transfer received between 2002 and 2022.

This combination—a universal registry of intergenerational transfers, covering inter vivos gifts as well as bequests, linked to individual income and wealth for both generations over a two-decade window—allows us to provide new evidence on the role of family transfers in intergenerational mobility; comparable data have so far been assembled only for Sweden, Norway, and the Netherlands ([Adermon et al., 2018](#); [Audoly et al., 2025](#); [Schulenberg et al., 2025](#)).³ Two features shape the transfers we observe. The first is timing: the parents of these cohorts are largely still alive, so the transfers that shape their children’s asset position are made during the parents’ lifetime rather than at its end, and the operative margin is the gift rather than the bequest. The second is the role of the family: in Spain, as in Southern Europe more broadly, the family rather than the state absorbs much of the risk faced by young adults. The quantitative role of private transfers in intergenerational persistence need not, therefore, resemble the Northern European benchmark.

We also develop a new approach to estimating intergenerational mobility when local samples are small, both for descriptive estimates at fine geographic levels and for the causal effect of childhood exposure to better places. One of the most active strands of this literature asks how the environments in which children grow up shape the prospects of those born into low-income families ([Chyn and Katz, 2021](#)). Answering that question requires upward mobility estimates for areas small enough to correspond to the environments children actually experience, and it is at that resolution that the data run short. Upward mobility is measured within cells defined by area, birth cohort and parental income, and the number of children in a cell falls quickly as areas are drawn more finely: in Spain, half of all municipalities contain only a handful of children per cohort.

Describing the geography of upward mobility area by area would then confound sampling noise with genuine spatial variation. Estimating the causal effect of childhood location is harder still. The standard design takes a measure of each area’s quality—the outcomes of the children of its permanent residents—as an input and regresses movers’ outcomes on it as if it were known, so the noise in those sparse cells passes directly into the exposure estimates and attenuates them toward zero ([Chetty and Hendren, 2018a](#)). This is innocuous where place quality is precisely measured, as it is at the geographies and in the population-wide samples for which the design was developed. It is not innocuous at finer resolutions, or in samples that are large but not population-wide, where the usual remedy—discarding cells that fall below a minimum size—removes most of the moves the design is estimated on. We address both problems with Bayesian hierarchical models that pool information across nested geographic levels in proportion to how precisely each area is measured. The first raises the reliability of local upward mobility estimates. The second estimates place quality jointly with the exposure coefficients, inside the movers design itself, which removes the attenuation the two-step procedure

³The Spanish registry records the occurrence, timing, type, and donor of each transfer, but not its amount. The evidence we provide on transfers therefore concerns the extensive margin of receipt, whereas the studies cited here observe transferred values.

cannot avoid without discarding any area, and which coincides with the standard estimator wherever place quality is precisely measured.

National estimates. The rank–rank slope is 0.274, meaning that a 10 percentile increase in parental income is associated with an average 2.7 percentile increase in the child’s adult income rank. A child whose parents sit at the 25th percentile reaches on average the 43rd percentile in adulthood, and 10.4% of children born to parents in the bottom income quintile reach the top quintile. These estimates place Spain in the middle of the international ranking: mobility is higher than in the United States ([Chetty et al., 2014b](#)) or Brazil ([Britto et al., 2026](#)), close to Italy ([Acciari et al., 2022](#)) and France ([Kenedi and Sirugue, 2022](#)), and well below Canada ([Connolly et al., 2019a](#)), Australia ([Deutscher and Mazumder, 2020](#)), Sweden ([Heidrich, 2017](#)) and Denmark ([Eriksen and Munk, 2020](#)). We show that these results are robust to the main measurement concerns in this literature, notably sample selection, life-cycle bias, and attenuation bias. Wealth is transmitted at least as strongly as income: ranking parents by their own wealth rather than their income raises the rank–rank slope from 0.274 to 0.313. Where Spain stands apart is at the very top of the distribution. A child born to parents in the top 1% is roughly fifty times as likely to reach the top 1% of adult income as a child from the bottom parental decile, and the corresponding ratio for wealth is higher again—to our knowledge the highest degree of top-tail persistence among countries with comparable administrative estimates.

Because we observe income, wealth and the universal registry of inheritances and gifts for the same individuals during two decades, we can document one of the channels through which this transmission operates, which the data available in most countries do not allow. Two facts organize what we find. Conditional on receiving, a transfer raises a child’s wealth rank by a similar amount at every parental origin, and a uniform shift of that kind does not by itself steepen the relationship between children’s ranks and their parents’. Access, by contrast, is far from uniform: gift receipt rises from 6% of children of bottom-half families to 42% of those of top-percentile families, while inheritance receipt is nearly flat. It is the second fact that feeds the gradient. Equalizing gift receipt across parental origins lowers the wealth rank–rank slope by 12.5% and the gradient in the probability of reaching the top percentile of wealth by 26.6%, while equalizing inheritance receipt changes neither; and among children of top-percentile parents, 5.1% of those with no observed transfer attain the top 1% of gross wealth against 21.8% of gift recipients. Hence, family transfers, in particular inter vivos gifts, operate as a vehicle for remaining at the top across generations rather than for reaching it from below ([?Adermon et al., 2018](#)).

The geography of economic opportunity. How much does a child’s chance of moving up depend on where she grows up? A great deal, and at a resolution well below the region. Upward mobility differs by more than 25 percentile points across Spanish municipalities, along a northeast-to-southwest gradient: it is high in Catalonia, the Madrid metropolitan area and several north-central provinces, and substantially lower

in Andalucía, Extremadura and the Canary Islands. This pattern aligns with a historical literature linking Spain's regional divide to the Reconquest and to lasting differences in land concentration, institutions, human capital and uneven industrialization (Oto-Peralías and Romero-Ávila, 2016; Oto-Peralías, 2020; Beltran Tapia and de Miguel Salanova, 2021; Tirado et al., 2016; Budí-Ors and Pijoan-Mas, 2022; Insa-Sánchez, 2024). These regional patterns coexist with wide differences across municipalities within the same province, and across neighborhoods within large cities. Measuring variation at this resolution is what the Bayesian hierarchical model makes possible, since it retains the small areas that area-by-area estimation would have to discard. We validate the model against well-powered comparisons across adjacent province borders and find that partial pooling attenuates genuine spatial differences rather than creating them, so the map is a conservative description of the heterogeneity in the data. Three characteristics of the places where children grow up are most strongly associated with upward mobility: bridging social capital, measured by cross-class economic connectedness in anonymized Facebook friendship networks (Chetty et al., 2022a,b; Johnston et al., 2026); the strength of the local labor market surrounding the parental generation; and local educational attainment. The ordering is in line with evidence from other countries, but the first two gradients are steeper in Spain than those reported elsewhere (Chetty et al., 2026c; Connolly et al., 2019a; Acciari et al., 2022; Britto et al., 2026).

Causal effects of childhood environments. These correlations describe where children from low-income families do better, but they cannot establish that places cause those differences rather than reflecting the families who sort into them. The distinction matters, because only causal differences can be acted upon by policies that target places. To separate the two we estimate childhood exposure effects, comparing children whose families make the same move at different ages (Chetty and Hendren, 2018a): if place matters causally, children who arrive earlier, and therefore spend more of childhood in the destination, should converge further toward the outcomes of the children raised there.

We estimate the design at the municipality, the same geography at which we measure the variation we are trying to explain. As set out above, the two-step procedure is not viable at that resolution: applying the sample filters used in the existing literature would discard the great majority of Spanish movers, and doing without them attenuates the exposure coefficients by roughly a third. A split-sample instrumental-variables estimator removes that attenuation with no Bayesian structure, by instrumenting one half-sample's estimate of the quality gap with the other's, but it halves an already thin first step and its variance cost is largest precisely where the correction is needed. Our joint model resolves this bias-variance trade-off. It attains the consistency of the instrumental-variables correction at a fraction of its variance, coincides with the two-step estimator wherever place quality is precisely measured, and passes a wide set of validation and robustness checks: against the instrumental-variables benchmark, against simulations calibrated to the observed distribution of municipality sizes, against the descriptive estimates obtained independently, and against alternative selection controls, outcome ages and sample definitions.

We find that each additional year of childhood exposure to a municipality one percentile higher in permanent-resident outcomes raises a child’s adult income rank by 0.020 percentile points, with a flat placebo profile for moves that occur after the outcome is measured. Extrapolating to a move at birth, 64% of the variation in upward mobility rates across Spanish municipalities is attributable to causal childhood place effects and the remaining third to the sorting of families across places, in the ballpark of the shares this design has recovered in other countries. Fitting the same model at province resolution yields an annual exposure effect of 0.036, close to the 0.040 that [Chetty and Hendren \(2018a\)](#) estimate across U.S. commuting zones, while our municipal estimate is close to their 0.022 for moves between counties within the same commuting zone. Selection behaves very differently across the two scales. Sorting accounts for 0.09 of the destination-origin gap between provinces but 0.32 between municipalities, so families who cross labor markets in Spain are barely selected on the quality of where they land, whereas those who move a short distance to a better municipality nearby are strongly selected. In income terms, a child from a low-income family who moves at age 10 from a municipality at the 10th to one at the 90th percentile of the quality distribution gains approximately 19% in adult income, net of selection. Because most of these gaps lie within provinces, much of that opportunity is available without long-distance migration.

1.1 Contributions to the literature

This paper contributes to three literatures. First, it extends the administrative-data literature on intergenerational mobility to a large European economy from which it has so far been absent. Building on the early linked tax-record studies of [Corak and Heisz \(1999\)](#) and [Österberg \(2000\)](#) and on the rank-based framework of [Chetty et al. \(2014b\)](#), we provide the first national estimates of income and wealth mobility for Spain and the first local estimates at municipality and ZIP-code resolution. Relative to existing country studies ([Connolly et al., 2019a](#); [Acciari et al., 2022](#); [Britto et al., 2026](#); [Deutscher and Mazumder, 2020](#); [Heidrich, 2017](#); [Chuard, 2021](#); [Eriksen and Munk, 2020](#); [Kenedi and Sirugue, 2022](#); [Beekers, 2024](#); [Pozo and Moreno, 2025](#)), the Spanish case combines average persistence close to the middle of the international distribution with an unusually strong concentration at the top—in income and in wealth alike—and a pronounced geography of economic opportunity.

Within this literature, the paper speaks to the small body of work that links intergenerational transfers to the income and wealth of the same parents and children ([Adermon et al., 2018](#); [Audoly et al., 2025](#); [Schulenberg et al., 2025](#); [Black et al., 2025](#)). Our contribution is to document where in the distribution these transfers operate. Consistent with [Adermon et al. \(2018\)](#) and [Audoly et al. \(2025\)](#), we find that family transfers are closely bound up with intergenerational persistence at the top; consistent with [Black et al. \(2025\)](#), we find that they are of limited relevance across the rest of the distribution. Two features of the Spanish case sharpen this asymmetry. Transfer receipt tracks children’s position in the wealth distribution and is essentially unrelated to

their position in the income distribution, so transfers sort children by assets rather than by earnings capacity. And at the top, the operative instrument is the inter vivos gift rather than the bequest: gift receipt is rare below the median and expands sharply in the upper tail, a pattern consistent with family transfers acting as a buffer that shields children born at the top from downward mobility rather than as a route by which children from low-income families move up.

Second, the paper contributes to the literature on the causal effect of childhood environments on economic opportunity ([Chetty and Hendren, 2018a](#); [Chyn and Katz, 2021](#); [Gaubert et al., 2025](#); [Deutscher, 2020](#); [Laliberté, 2021](#)). The childhood-exposure design was developed for, and has largely been implemented on, population-wide tax records covering tens of millions of children. That is what makes it possible to estimate each area's quality precisely in a first step and then treat those estimates as data in a second. Most countries do not have such records. Their administrative data are large but not population-wide, or population-wide but spread across a territory in which local cells are thin, and in either case the two-step procedure faces the same bind: retaining precision requires discarding most of the areas and most of the moves, and retaining the sample attenuates the exposure coefficients. Estimating place quality jointly with the exposure coefficients removes that bind, and the logic extends beyond this design to any two-step procedure in which a noisily estimated group-level quantity enters a second-stage regression.

The practical implication is that causal place effects can be estimated in a considerably wider set of countries than has so far been possible, and at the geographic scale at which neighborhoods, schools and local labor markets actually operate. That matters for what the literature can learn. Which features of a place raise the prospects of children growing up in disadvantaged areas is a question that can only be answered by comparing causal estimates across many settings and at a resolution fine enough for the candidate mechanisms to differ; the evidence available so far comes from a handful of countries and, for the most part, from geographies coarser than the mechanisms themselves. Our own estimates illustrate what this makes possible: fitting the design at two nested geographies on a single sample, we document how exposure and selection each vary with the scale at which place is defined—a comparison the existing literature does not report, and one that suggests international estimates computed at a single intermediate geography may average over a steep gradient in how much of the observed variation across places reflects sorting rather than causal effects. Together with [Beekers \(2024\)](#) for the Netherlands, this is one of only two applications of the childhood-exposure design to a European setting.

Third, the paper provides the most detailed evidence to date on intergenerational income and wealth persistence in Spain, and the first for recent cohorts. Historical and surname-based studies have documented declining occupational mobility in nineteenth-century Valencia, enduring class gaps in literacy, obstacles to secondary education for rural students, and status correlated across several generations ([Santiago-Caballero, 2021](#); [Beltran Tapia and de Miguel Salanova, 2021](#); [Insa-Sánchez, 2024](#); [Collado et al., 2013](#); [Güell et al., 2015](#));

several of these studies also provide the first evidence that opportunity in Spain has long varied across places, a regularity our estimates confirm for recent cohorts and at a far finer resolution. Survey-based work has traced an income elasticity that declined during the 1980s (Sánchez Hugalde, 2004), rising educational mobility unmatched by occupational mobility (De Pablos Escobar and Gil Izquierdo, 2016; Cervini-Plá, 2013), and a two-sample elasticity of roughly 0.40 to 0.42 that places Spain in an intermediate international position (Cervini-Plá, 2015; OECD, 2018a,b). What these sources could not jointly deliver, given small samples, imputed parental earnings, surname links, or geographically limited records, are internationally comparable national estimates, mobility measures at the level of municipalities and neighborhoods, and causal place effects. Comparability is itself a live concern in this literature, since differences in income definitions, measurement ages, and sample construction can move estimates enough to reconcile apparently conflicting findings, and we therefore report the sensitivity of our estimates to each of these choices.

The remainder of the paper proceeds as follows. Section 2 describes the administrative data, sample construction, and variable definitions. Section 3 reports national estimates of income and wealth mobility, including persistence at the top and the role of family transfers. Section 4 introduces the Bayesian small-area estimator, maps the geography of upward mobility, and studies its local correlates. Section 5 develops the joint movers design and estimates the causal effect of growing up in a higher-opportunity area. Section 6 concludes.

2 Data

2.1 Administrative Data

The backbone of our empirical analysis is a new intergenerational administrative panel—the Spanish Opportunity Atlas—constructed from individual-level tax records provided by the Spanish Tax Agency (*Agencia Estatal de Administración Tributaria*, AEAT)⁴. The panel covers 9.6 million unique individuals belonging to approximately 2 million households, tracked continuously over a 25-year period (1998–2022). By construction, the dataset targets birth cohorts 1980–1990 within the common fiscal territory⁵. To our knowledge, this constitutes the most comprehensive intergenerationally linked administrative database available for the study of income and wealth mobility in Spain.

⁴ An earlier version of this database, linking parents observed in 1998 to children observed in 2016, was first analyzed in Soria-Espin (2022). The panel used in this paper substantially improves that dataset in both temporal coverage and sample composition.

⁵ The common fiscal territory covers the 15 autonomous communities and 46 provinces under the general fiscal regime, representing approximately 93.6% of the Spanish population. The Basque Country and Navarra operate under special fiscal regimes (*régimen foral*) that grant them autonomy to collect most taxes independently; individual-level records from these regions are therefore not included in AEAT’s centralized files.

Intergenerational linkage. Children are matched to parents through household identifiers recorded in tax filings. The linkage is anchored in the 2003 tax records: the Tax Agency identified all children born between 1980 and 1990 claimed as dependents in 2003, when tax identification numbers are reliably recorded, and traced them back to their parents' 1998 income tax returns. The panel additionally includes spouses, siblings, and extended family members identified through a reclassification procedure described in Appendix G. From 1998 to 2007, the data are based on personal income tax returns (*IRPF*), which cover tax filers only; the 1998 exemption threshold was €7,212 of gross income (approximately €12,500 in 2022 prices), roughly 1.3 times the annual minimum wage, so non-filing is confined to households at the very bottom of the income distribution. Approximately 92% of parental returns in our sample are filed jointly, meaning that total household income—including the earnings of both spouses—is captured even when only one parent appears as the primary filer. From 2008 onward, the data source shifts to comprehensive fiscal census records that include non-filers, providing near-universal coverage of Atlas individuals for the period in which we measure children's adult income.

Two margins of non-coverage follow from this design, and Appendix G examines both in detail against census and tax benchmarks. The first margin is mechanical: children who had left the parental household by 2003 cannot be identified as dependents, so linkage rates vary with children's age in 2003 rather than with income. The Atlas links 2,996,295 children (target children plus siblings), 59.5% of the population of Spanish nationals born 1980–1990 in the common territory, with cohort-level coverage following an inverted-U from 40% for the 1980 cohort to 73% for 1987; coverage rises to 74.2% among children residing with both parents according to the 2001 Census, and roughly four in five tax-filing households with children in the target age range appear in the data. The second margin is income-related by construction: households in which no adult filed in 1998, necessarily earning below the filing threshold, are not observed. Appendix G shows that this group is small, drawing on the income composition of filers with dependent children and on parents' voluntary filing behavior around the 1999 reform that tripled the threshold. Crucially, the mobility estimates are invariant to both margins: cohorts covered at 40% and at 73% deliver the same estimates, reweighting the sample to population cohort and gender shares moves the rank-rank slope by less than 0.001, and simulations show that only implausibly extreme selective attrition could materially move the estimates, with the most plausible adverse scenarios attenuating rather than inflating measured persistence. Validation against the Spanish Survey on Living Conditions (*Encuesta de Condiciones de Vida*) further confirms that the income distribution of Atlas children closely tracks the population distribution of income across the same birth cohorts, with median differences below 10% and closely aligned zero-income shares (Appendix I).

Income and demographics. For each individual-year observation, we observe the full set of variables present in Spanish fiscal returns, including detailed income decompositions—wages, self-employment earnings, capital income, rental income, unemployment benefits, and social transfers—as well as demographic

characteristics such as age, sex, marital status, and household composition.

Wealth. From 2016, the panel is enriched with individual-level wealth data drawn from AEAT’s administrative wealth module, which reconstructs household balance sheets by combining Cadastre records, third-party filings from financial institutions, and wealth tax returns. The resulting data provide annual measures of real estate wealth, financial assets, bank deposits, and liabilities for every individual in the sample.

Inheritances and gifts. The Spanish Opportunity Atlas incorporates centralized records from the Spanish Tax Agency on all inheritance and gift events registered in Spain, appended at the individual level for each of the 9.6 million individuals in the dataset. These records are drawn from the centralized system through which the Tax Agency coordinates with regional tax agencies and provide a longitudinal record of every formal intergenerational wealth transfer received by individuals in our sample. For each event over 2002–2022, we observe the year of the transfer, its nature (gift or inheritance), the identity of the donor or decedent, and the kinship relationship between parties; we do not observe the transfer amount, associated tax benefits, or asset type. To our knowledge, no other dataset in Spain combines administrative income, wealth, and transfer records linked intergenerationally at this scale.

Geographic location and characteristics. Geographic location is observed annually at several levels: ZIP code, municipality, province, and autonomous community (region). We use these identifiers to assign each child a childhood location based on the parental ZIP code reported in 1998, and to characterize the local environments in which children grew up.⁶ For the geographic analysis, we merge the mobility estimates with a rich set of municipality-level characteristics drawn from three administrative sources: the 2001 Population and Housing Census (INE), the Digital Atlas of Urban Areas from the Ministry of Transport, and municipal income statistics from the Tax Agency. We complement these with a measure of bridging social capital, upward economic connectedness, following [Chetty et al. \(2022a,b\)](#) and using the cross-country release of [Johnston et al. \(2026\)](#), which constructs economic connectedness from anonymized Facebook friendship networks; in our setting it is measured at the ZIP-code level as the average socioeconomic status of the friends of low-SES individuals, where low SES corresponds to the 25th percentile of the national distribution.

2.2 Sample Definition

The full Atlas contains 2,712,065 target children born between 1980 and 1990, distributed across 2,022,118 fiscal households. We define two nested samples by birth cohort. The *baseline sample* restricts to children born between 1980 and 1986, aged 35–42 when income is measured in 2021–2022, which minimizes lifecycle bias by

⁶From 2008 onward, location is also observed at the census-tract level, an administrative unit of approximately 1,500 residents.

observing adult income well into mid-career; it underlies all national-level results of Section 3. The *extended sample* includes all cohorts from 1980 to 1990 (aged 31–42), increasing statistical power for the geographic analyses of Section 4.2 at the cost of measuring income somewhat earlier in the lifecycle for the youngest cohorts. Section 3.6 shows that the main results are robust across both definitions. We apply three sequential filters to address measurement error and income misreporting, following standard practice in the literature (Chetty et al., 2014c; Britto et al., 2026). The first two act on parental income and the third on child income:

1. **Negative parental income.** We exclude parental households reporting negative income in any year of the 1998–2000 window used to compute average parental income, as negative incomes typically reflect business losses or capital deductions rather than persistent poverty.
2. **Parental income underreporting.** We remove households that simultaneously report imputed real estate income above the national 60th percentile and total income below the 10th percentile, a combination indicative of substantial unreported resources rather than genuine low income.
3. **Bottom of the child income distribution.** Within each birth cohort, we trim the bottom 5% of the child income distribution. This threshold is motivated by the empirical observation that the relationship between child income rank and mean parental income turns implausibly negative in the lowest percentiles: children with near-zero reported income are disproportionately drawn from non-poor households with significant unobserved income or transitory reporting anomalies. Appendix H documents this anomaly and develops a targeted alternative filter based on children’s own wealth records.

The first two filters, together with the requirement that parental income be observed over 1998–2000, yield an analysis dataset of 2,495,136 children in 1,754,723 households for the 1980–1990 cohorts, of which 1,580,994 belong to the baseline cohorts (1980–1986). The bottom trim then produces the estimation samples: 2,370,385 children in the extended sample and 1,501,948 in the baseline sample, the latter corresponding to 43.4% of its target population. Appendix Table A11 reports the full funnel from the census population to the estimation samples, by layer and cohort group. These choices do not drive the headline intergenerational mobility estimates. Appendix Table A17 reports the national rank-rank slope and absolute upward mobility under alternative bottom-trim thresholds and parental-income averaging windows, and Appendix H replaces the uniform trim with a targeted filter that removes only children with near-zero income but substantial personal wealth (0.6% of the sample): the rank-rank profile is unchanged and the top-tail persistence ratio of Section 3 is identical (51.4) under either sample. Appendix G additionally shows that the estimates are invariant to population reweighting and to the cohort variation in coverage.

2.3 Variable Definitions

Parental income. Parental economic status is measured as total household gross income—encompassing labor earnings, self-employment income, capital income, and taxable transfers—averaged over 1998–2000. For the baseline 1980–1986 cohorts, this window captures parental resources when children were aged 12–18. As we show in Section 3.6, averaging over three years is sufficient to eliminate attenuation bias from transitory income shocks. Parents are ranked nationally relative to other parents of children born in the same cohort.

Child income. Child economic status is measured as individual total gross income (same definition as parents) averaged over 2021–2022, when the baseline cohorts are aged 35–42. Measuring income at this stage of the life cycle minimizes lifecycle bias by capturing earnings well into mid-career, after the extended school-to-work transitions characteristic of Southern European labor markets. Section 3.6 confirms that our estimates are stable across alternative age windows. We focus on individual rather than household income because a substantial fraction of the target cohort remains unmarried or is not observed with a co-resident partner in tax records by the end of the panel, making household income an unreliable measure for this age range. Children are ranked nationally within their birth cohort. All monetary values are expressed in constant 2022 euros, deflated using the national Consumer Price Index ([Instituto Nacional de Estadística](#) , [INE](#)).

Gross wealth. From 2016 onward, we observe individual-level wealth stocks reconstructed from the Tax Agency’s wealth records, combining third-party reporting and cadastral data. Our main wealth measure is gross wealth, defined as the sum of real estate assets, financial assets, and bank deposits before subtracting liabilities.⁷ We use this variable both as an alternative outcome and to characterize the joint distribution of income and wealth across generations.

2.4 Descriptive Statistics

Table A1 reports summary statistics for the baseline sample of 1,580,994 parents. Mean parental household income (averaged over 1998–2000) is €29,247 in nominal terms, or approximately €42,400 in 2022 euros after CPI adjustment. Mean children’s individual income (averaged over 2021–2022) is €29,231, substantially lower than parental household income in real terms—consistent with measuring individual rather than household earnings and with the younger life-cycle position of children (ages 35–42 vs. parents at ages 48–52 during their income window). Both income distributions exhibit pronounced right-skewness and top-tail concentration

⁷We focus on gross rather than net wealth because housing is the dominant asset in household portfolios, and real estate values in the administrative wealth module are based on cadastral values rather than market values. Since cadastral values may understate the market value of housing and vary in how closely they track market prices across places, subtracting mortgage liabilities from these values can distort net-wealth rankings, especially for leveraged homeowners. Gross wealth therefore provides a more stable measure for ranking individuals in the joint distribution of income and wealth.

(Appendix Figures A1 and A2), motivating the use of rank-based measures throughout the analysis. Mean parental gross wealth in 2022 is €216,480, compared to €71,971 for children, a gap largely attributable to life-cycle wealth accumulation, as parents are observed at ages 68–72 and children at ages 35–42.

3 National Results

We begin by presenting national estimates of intergenerational income mobility for our baseline sample (Section 2.2): children born in 1980–1986, with parental household income averaged over 1998–2000 and child individual income averaged over 2021–2022, when children are observed between ages 35 and 42.

3.1 Baseline Intergenerational Income Mobility Estimates

The rank–rank specification. We measure mobility using the rank–rank specification of Chetty et al. (2014b). Within each birth cohort, we assign every child a percentile rank y_i in the distribution of child income and every parental household a percentile rank p_i in the distribution of parental income. We then estimate

$$y_i = \alpha + \beta p_i + \varepsilon_i, \quad (1)$$

where the slope β —the rank–rank slope (RRS)—measures *relative* mobility: how strongly a child’s position in the income distribution tracks that of their parents. *Absolute* upward mobility (AUM) is $\mu_{25} = \alpha + 25\beta$, the expected rank of a child from a below-median family, evaluated at the 25th percentile of the parental distribution.

We favor the rank–rank slope over the intergenerational income elasticity (IGE) for three reasons. First, the IGE is estimated in logs and is undefined for the zero and near-zero incomes that are common at the bottom of the distribution; the usual remedies—dropping these observations or adding an arbitrary constant—introduce sensitivity that is difficult to sign. Second, the IGE is acutely sensitive to lifecycle bias and to the ages at which parent and child incomes are measured, whereas ranks are considerably more stable. Third, the rank–rank relationship is close to linear over most of the distribution (Figure 1), so a single slope summarizes mobility across its full range and cleanly separates relative mobility (β) from absolute mobility (α)—a decomposition that is central to the geographic analysis in Section 4.

National estimates. Figure 1 summarizes the national rank–rank relationship. Panel A plots the mean child income percentile against the parental household income percentile in Spain. The rank–rank slope is 0.274, implying that a 10-percentile increase in parental income is associated with a 2.74-percentile increase in child income. The slope restricted to percentiles 10–90 is 0.253, confirming that the relationship is ap-

proximately linear through most of the distribution.⁸ Absolute upward mobility—the mean rank reached by children from a family at the 25th parental percentile—is 43, corresponding to roughly €21,080 in annual child gross individual income.


Panel B places the Spanish profile alongside the U.S. benchmark from [Chetty et al. \(2026a\)](#). The average slopes are strikingly similar—0.274 in Spain and 0.287 in the United States—despite large differences in income levels and institutions, so Spain does not look exceptional when summarized by a single slope. The contrast lies in the *shape*: the U.S. profile is close to linear throughout, whereas the Spanish profile bends upward above the 90th parental percentile, revealing stronger persistence at the very top.

Appendix Figure [A3](#) reports the intergenerational quintile transition matrix, where each cell gives the probability that a child born to parents in a given income quintile reaches each quintile of the adult distribution. The *rags-to-riches* probability—that a child from the bottom quintile reaches the top quintile—is only 10.4%. By contrast, 36.9% of children raised in the top quintile remain there and just 11.8% fall to the bottom, evidence of pronounced persistence at the top. Movement is more fluid among the middle three quintiles.⁹ The persistence visible in the bottom-left corner of the matrix—children of low-income parents who themselves remain at the bottom in adulthood—also surfaces in an independent administrative record of economic hardship. Receipt of the Minimum Basic Income (*Ingreso Mínimo Vital*), Spain’s national minimum-income scheme for households with insufficient resources, falls steeply with parental income: from roughly 4% of children in the lowest parental ventile to under 0.5% at the top (Appendix Figure [A5](#)). That a child’s exposure to a means-tested safety-net benefit in adulthood is so strongly predicted by parental income more than two decades earlier corroborates the persistence of disadvantage at the bottom of the distribution ([Badenes-Plá and Gambau-Suelves, 2020](#)).


⁸The rank–rank relationship is largely linear but shows pronounced upward curvature at the top, indicating extra persistence among the top 10%. This pattern is consistently observed in comparable studies ([Britto et al., 2026](#); [Acciari et al., 2022](#); [Deutscher and Mazumder, 2020](#); [Chuard, 2021](#); [Corak, 2020](#); [Kenedi and Sirugue, 2022](#)), with the exception of the United States ([Chetty et al., 2014b](#)), whose profile is exceptionally linear throughout.

⁹A comparable matrix using income deciles is provided in Appendix Figure [A4](#).

Figure 1: Mean Child Income Rank vs. Parent Income Rank



(a) Spain



(b) Spain and the United States

Notes: Panel A plots the mean child income percentile for each parental income percentile in Spain, estimated from Equation (1). The figure reports the rank–rank slope (RRS) for the full distribution and restricted to percentiles 10–90, together with national absolute upward mobility (AUM). Panel B overlays the Spanish profile with the U.S. profile from Chetty et al. (2026a). For Spain, child income is average individual income in 2021–2022 and parental income is average household income in 1998–2000. For the United States, child income is mean individual income in 2014–2015 and parental income is average family income over 1996–2000. Income definitions follow Section 2.3. Standard errors in parentheses.


Subgroup heterogeneity. Appendix Table A2 disaggregates the baseline estimates by child gender, birth order, parental marital status, and migration background. The largest heterogeneity is by gender.¹⁰ Sons face a flatter slope than daughters (0.260 vs. 0.285) yet enjoy 8.5 percentiles higher absolute upward mobility (47.8 vs. 39.3), an income gap of about €3,091. The gap is also visible in the rags-to-riches probability (12.7% for sons vs. 8.2% for daughters) and at the very top, where men are overrepresented by roughly 40% in the top percentile of child income relative to their population share (Appendix Figure A7). Remaining subgroup differences are small: first-born children have a slope of 0.277, children of married parents 0.280, and second-generation migrants 0.276.¹¹

Top-tail persistence. The upward curvature in Figure 1 implies that standard slopes understate persistence at the top. To capture it, we define the Top-Tail Relative Persistence Ratio (TTRPR) as the probability of reaching the top 1% given parents in the top 1%, divided by the same probability given parents in the bottom 10% (Appendix J details its construction and interpretation). Figure 3 plots the probability of reaching the top 1% of adult income by parental percentile; under perfect equality of opportunity this would equal 1% everywhere (dashed line). Instead, children from top-1% households have a 12.5% chance of reaching the top 1% themselves, against just 0.24% for children from bottom-decile households. The resulting TTRPR of 51.4 means a child born into the top 1% is over fifty times more likely to attain elite income status than a child from the bottom 10%.

¹⁰Appendix Figure A6 reports the rank–rank profile and quintile transition matrix separately by gender.

¹¹Boustan et al. (2025) provides a detailed comparative analysis of intergenerational mobility among second-generation migrants using the same Spanish data harmonized with other destination-country administrative datasets.

Figure 3: Conditional Probability of Reaching the Top 1% by Parental Income Percentile



Notes: Each point shows the probability that a child reaches the top 1% of adult income, conditional on parental household income percentile. Points are color-coded: red for the bottom decile and dark green for the top 1%. The dashed line at 1% denotes perfect equality of opportunity. Income definitions follow Section 2.3.


3.2 International Comparisons

Where does Spain stand internationally? Appendix Figure A9 ranks countries studied with comparable cohorts, the rank–rank framework, and administrative tax data.¹² By the rags-to-riches probability, Spain’s 10.4% lies above the United States (7.5%) and Brazil (2.5%) but below Australia (13.5%) and Switzerland (12.3%); by the rank–rank slope, Spain (0.274) is located well above the flattest Nordic economies (Denmark and Sweden, 0.14–0.20) and far below Brazil (0.55) and China (0.44). Overall, Spain sits in the lower-middle range of the international ranking of intergenerational income mobility. Spain’s distinctiveness emerges only at the top. Figure 4 reports the TTRPR across countries with comparable estimates: Spain’s ratio of 51.4 is the highest, exceeding Italy, Canada, and the United States. This reflects a distinctive combination—a low probability of moving from the bottom decile to the top 1%, similar to the United States, paired with unusually strong persistence at the top (Appendix Figure A10 reports the underlying conditional-probability profiles)¹³. The lesson mirrors the national rank–rank profile: Spain is not an outlier in average intergenerational mobility, but it is exceptional in top-percentile persistence.

¹²See Appendix Table A3 for the full set of estimates and sources.

¹³Appendix Figure A58 shows that this higher top-1% persistence is robust to differences in sample construction

Figure 4: Top-Tail Relative Persistence Ratio (TTRPR) Across Countries



Notes: The TTRPR is defined as the ratio of (i) the probability that a child reaches the top 1% given parents in the top 1%, to (ii) the probability of reaching the top 1% given parents in the bottom 10%. Country estimates from Connolly et al. (2019a) for Canada, Acciari et al. (2022) for Italy, and Chetty et al. (2014b) for the United States.

3.3 Intergenerational Wealth Mobility

Measurement and sample. We now replicate the national income analysis for wealth, providing the first evidence on intergenerational wealth persistence in Spain from linked parent–child administrative records. All estimates use the baseline estimation sample of Section 3.1, so that differences between the income and wealth results reflect the outcome under study rather than sample composition. Child gross wealth is the sum of real-estate holdings, financial assets, and bank deposits, averaged over 2021–2022, when children are aged 35–42, the same window used for child income. An advantage of the Spanish registries is that wealth is observed from administrative property and financial records rather than from wealth-tax filings: coverage therefore does not select on wealth-tax liability, and no child is lost in moving from income to wealth.¹⁴ Parents are ranked by their 1998–2000 household income throughout. Parental income is the only measure of parental resources that is pre-determined with respect to child outcomes: no parental wealth measure exists for the late 1990s, and parental wealth measured at the end of the panel is mechanically depressed for parents who have already given assets away, so that ranking parents by contemporaneous wealth would classify the most generous donors as poor (Appendix Figure A11).

Baseline results: income vs. wealth mobility. Appendix Figure A12 plots mean child income and gross wealth percentiles against the parental income percentile. The two profiles have the same qualitative shape—

¹⁴Real estate enters at cadastral values, which average roughly 60–70% of market value. Because mortgage debt is concentrated among younger households, netting liabilities out of below-market asset values distorts the bottom of the net-wealth distribution; gross wealth is therefore our primary measure, with net wealth retained as a robustness check. Percentile ranks are recomputed within the estimation sample: within birth cohort for children, and pooled at the household level for parents. Appendix Table A4 summarizes the sample and measurement windows.

approximate linearity through most of the distribution and pronounced curvature at the top—and the wealth gradient is the flatter of the two, at 0.226 against 0.274, with slightly higher absolute upward mobility (44.7 against 43.4). The quintile transition matrices deliver a consistent picture (Appendix Figure A13). A flatter slope with respect to parental income should not, however, be read as weaker persistence in the wealth dimension itself. Parental income measures the resources relevant to wealth accumulation with error—parental income and wealth ranks correlate at only 0.55, and child income and wealth ranks at 0.53 (Appendix Figure A14)—and this imperfect correspondence attenuates the wealth gradient estimated against parental income. Ranking parents by their own wealth reverses the ordering: the wealth-on-wealth slope of 0.313 exceeds the income-on-income slope of 0.274, and each child outcome is best predicted by the same parental dimension (Appendix Table A5).¹⁵ Wealth is therefore transmitted at least as strongly as income across Spanish generations, with a wealth-on-wealth slope in the range documented for Scandinavia and the United States (Boserup et al., 2013; Adermon et al., 2018; Fagereng et al., 2021; Audoly et al., 2025).

Intergenerational mobility by family transfer status. One channel through which this persistence may operate is direct intergenerational transfers, which we observe in the universal records of the Spanish inheritance and gift tax: every gift and every inheritance registered between 2002 and 2022 for the children in our sample. Using the stock of transfers received over this window, we classify children into four mutually exclusive categories—no observed transfer, an inheritance only, an inter vivos gift only, or both. A quarter of children (25.3%) have received something: 15.7% an inheritance alone, 7.6% a gift alone, and 2.0% both.¹⁶


Figure 5 decomposes both rank–rank relationships by this classification, on a common vertical scale. In the wealth dimension the four categories form distinct layers that are roughly parallel across the parental distribution: at any given parental origin, gift recipients sit 19.2 percentiles above non-recipients, inheritance recipients 8.1, and those receiving both 23.5. In income the same layers are three to four times closer together and, more tellingly, asymmetric: gift recipients sit 5.3 percentiles above non-recipients and those receiving both 5.6, while inheritance recipients are indistinguishable from them (−0.2). Two readings follow. First, on both dimensions the largest gaps belong to gifts and to the combination of gift and inheritance, not to inheritance alone. Second, the asymmetry is what one would expect from the nature of the two margins: an inheritance arrives when a parent dies, an event largely unrelated to the child’s own earnings capacity, so a large wealth gap combined with no income gap is the signature of an unselected transfer, whereas an inter vivos gift is an active decision by families with assets to spare and its recipients differ from non-recipients on both dimensions. The margin that matters for children’s outcomes is therefore the discretionary one: not what parents leave behind, but what they choose to hand over while their children are still building their own

¹⁵Because end-of-panel parental wealth is itself depressed by prior giving, 0.313 is a lower bound: measuring parental wealth in 2016 raises the slope to 0.338 while leaving the parent-wealth-to-child-income slope essentially unchanged (Appendix Table A6), the pattern implied by transfer-driven attenuation and the opposite of what classical measurement error would produce.

¹⁶The registry records the occurrence, timing, and type of each transfer, but not its amount, so all results concern the extensive margin of receipt. Children were 16–22 years old in 2002, so transfers predating the window are likely rare.

position.¹⁷

Figure 5: Mean Child Gross Wealth and Income Ranks by Parental Income and Transfer Status



Notes: Panel (A) plots the mean child gross wealth percentile and Panel (B) the mean child income percentile against the parental household income percentile (two-percentile bins), separately for four mutually exclusive transfer categories—no observed transfer, only gifts, only inheritances, and both—as recorded in the Spanish inheritance- and gift-tax registry over 2002–2022. The black series reproduces the unconditional profile. Both panels use the same vertical scale, so the magnitude of the transfer-associated gaps is directly comparable across outcomes. Averaged over the parental distribution, the wealth gaps relative to no-transfer children of the same origin are +19.2 (only gift), +8.1 (only inheritance) and +23.5 (both) percentiles; the corresponding income gaps are +5.3, –0.2 and +5.6, and are reported bin by bin in Appendix Figure A16. The asymmetry between the two panels is informative about selection: inheritance receipt shifts the wealth rank but not the income rank, as an event triggered by parental mortality should, whereas gift receipt shifts both. Child gross wealth is the 2021–2022 average of real-estate (cadastral value), financial-asset and deposit holdings; child income is average individual gross income in 2021–2022; parental income is average household income in 1998–2000. Ranks are computed within the baseline estimation sample ($N = 1,501,948$). Transfer amounts are not observed; the categories capture the extensive margin of receipt only, and all differences are descriptive associations rather than effects.

Access to family transfers and the intergenerational wealth gradient. Two features of Figure 5 determine how transfers bear on intergenerational mobility. Conditional on receiving, the jump in a child’s wealth and income rank is broadly the same at every parental origin: the layers are close to parallel, widening only within the top decile.¹⁸ A jump of the same size everywhere moves recipients vertically without steepening the relationship between their rank and their parents’, and so is close to mobility-neutral in the rank–rank sense. But receipt itself is strongly graded by parental income. Panel A of Figure 6 shows gift receipt rising from 4.8% of children of bottom-half families to 32.7% of those of top-percentile families, while inheritance

¹⁷Regression versions with cohort, gender, and childhood-region fixed effects confirm the pattern (Appendix Figure A17). Parental mortality is itself correlated with socioeconomic status, so the inheritance comparison bounds rather than eliminates selection. Part of the wealth gap is definitional, since a transferred asset enters gross wealth on receipt, but transfers are typically received well before wealth is measured (Appendix Figure A19) and the gap does not grow with time since receipt (Appendix Figure A18), so it reflects assets retained rather than differential accumulation.

¹⁸The gift gap stands at 17.4 ranks at the tenth parental percentile and 21.3 at the ninetieth. The widening at the top is the intensive margin—conditional on receiving, children of richer parents appear to receive more—which we cannot observe and which makes the estimates below lower bounds.


receipt stays nearly flat.¹⁹ The intergenerational gradient should therefore be fed less by the size of the jump than by its incidence—by who receives at all. How much recipients receive plausibly matters too, and the widening of the gaps at the top suggests it does, but amounts are not recorded in the registry; the exercise that follows quantifies only the part of the gradient that runs through the extensive margin of receipt, and is a lower bound for that reason.

Panel B asks how much of the observed gradient survives if access to transfers is equalized (the extensive margin). Writing the observed profile of mean child wealth rank as a weighted average of the category-specific profiles, we hold receipt rates fixed at their population level across all parental origins while leaving conditional outcomes as observed, which isolates the composition margin. Equalizing access to gifts lowers the wealth rank–rank slope from 0.226 to 0.198, a fall of 12.5%; equalizing access to inheritances leaves it unchanged, and equalizing both delivers essentially the gift result. The same exercise on child income moves the gradient by only 3.0%.²⁰ Family transfers therefore contribute to intergenerational persistence in Spain not because they are large, but because the discretionary ones are unequally distributed: inheritances reach children of all origins at similar rates and are correspondingly neutral for relative mobility, while gifts, which parents choose to make and only wealthy parents can afford, reach the children who are already advantaged. Making that redistribution of parental assets a matter of family choice rather than of mortality is what turns a level shift into a gradient. The top of the wealth distribution is where the two margins meet.

¹⁹The categories are mutually exclusive, so high-origin children receiving both transfer types are counted in the combined category rather than the inheritance-only one, which makes inheritance-only receipt appear to decline with parental income. On the binary margins, gift receipt rises from 6.2% to 41.8% between the same two origin groups while inheritance receipt rises only from 17.6% to 22.5% (Appendix Figure A20); the decompositions use these margins.

²⁰Seen from the other side of the distribution, transfer receipt rises from about 11% of children at the bottom of the wealth distribution to 68% in the top percentile and continues to rise within it, while the corresponding gradient across the income distribution is far flatter (Appendix Figures A22, A23 and A24).

Figure 6: Access to Family Transfers and the Parental Income Gradient in Child Wealth



Notes: Panel (A) plots the share of children receiving each type of transfer by parental household income percentile, for the three receiving categories of Figure 5 (children with no observed transfer are the omitted complement). Gift receipt rises from 4.8% of children of bottom-half families to 32.7% of children of top-percentile families and expands sharply within the top percentiles; inheritance-only receipt is far flatter, and its mild decline partly reflects high-origin children being reclassified into the combined category, since the categories are mutually exclusive. Panel (B) plots the observed profile of mean child gross wealth rank against parental income rank together with three counterfactual profiles. Writing that profile as a weighted average of the category-specific profiles, $E[y | p] = \sum_k \pi_k(p) \mu_k(p)$, where $\pi_k(p)$ is the share of children of parental rank p in transfer category k and $\mu_k(p)$ their mean child rank, each counterfactual replaces $\pi_k(p)$ by its population value $\bar{\pi}_k$ while leaving the conditional means $\mu_k(p)$ as observed. Each line therefore shows the gradient that would obtain if every parental origin faced the same probability of receiving the stated type of transfer, holding fixed what recipients do with what they receive; the in-figure box reports each slope and its change relative to the observed one. Equalizing gift access lowers the slope from 0.226 to 0.198, equalizing inheritance access leaves it unchanged, and equalizing both delivers essentially the gift result. Because only the extensive margin of receipt is observed, transfer amounts are unavailable, and most bequests are not yet realized, these are lower bounds on the role of transfers in the gradient. Sample and variable definitions as in Figure 5.

3.4 Persistence at the Top of the Wealth Distribution

The averages just described conceal what happens in the upper tail, which is where the Spanish case departs most sharply from the international pattern. The probability of reaching the top percentile of the child wealth distribution is essentially flat through the lower nine deciles of the parental distribution and rises steeply within the top percentiles, more sharply than for income: the top-tail relative persistence ratio is 62.8 for wealth against 49.2 for income on the same sample (Appendix Figure A15).²¹ Persistence at the top is thus stronger in assets than in earnings, and the transfer records allow us to ask how the top of the wealth distribution is reached.


Family transfer receipt and the probability of reaching the top Figure 7 plots the probability of entering the top percentile of child wealth (Panel A) and of child income (Panel B) against parental income,

²¹Both figures are computed on the baseline estimation sample and are therefore internally comparable; the income ratio differs slightly from the one reported in Section 3.1, which uses the full 1980–1990 analysis sample to be comparable across countries.

separately by transfer status. Parental origin confers a substantial advantage on its own: among children of top-percentile parents who received no observed transfer, 5.1% reach the top percentile of wealth and 8.8% the top percentile of income, five and nine times the benchmark under independence. Receiving a transfer multiplies that advantage rather than substituting for it, and it does so far more strongly in wealth: from the same origin, the probability of reaching the top of the wealth distribution roughly quadruples, to 18.3% for inheritance recipients and 21.8% for gift recipients, while in income it rises much more modestly, to 11.2% and 14.9%.

Two features of this pattern are worth noting. First, within each outcome gifts and inheritances deliver similar probabilities: conditional on having received something from a top-percentile family, it matters little which of the two it was. What separates the two margins at the population level is therefore access rather than what receipt delivers to those who obtain it. Second, the categories diverge sharply only for children who received both, whose probability of reaching the top percentile of wealth is 43.1% against 19.3% for income—though this is a small group, 2.0% of children overall and 9.1% of those raised in top-percentile families, so the contrast should not be over-read. In both outcomes and for every category, the profile is flat and close to the independence benchmark throughout the lower nine deciles of parental income.

Figure 7: Probability of Reaching the Top 1% of Gross Wealth and of Income, by Parental Income and Transfer Status




Notes: Panel (A) plots the probability that a child reaches the top 1% of the child gross wealth distribution and Panel (B) the probability of reaching the top 1% of the child income distribution, in both cases conditional on the parental household income percentile and separately by the four mutually exclusive transfer categories of Figure 5. Both panels use the same vertical scale, so the magnitude of the transfer-associated differences is directly comparable across outcomes; the dashed line at 1% is the benchmark under independence. At the top parental percentile the probability of reaching the top 1% of wealth is 5.1% for children with no observed transfer, 18.3% for inheritance-only recipients, 21.8% for gift-only recipients and 43.1% for those receiving both; the corresponding probabilities for income are 8.8%, 11.2%, 14.9% and 19.3%. The implied top-tail relative persistence ratios—the probability at the top parental percentile divided by the mean probability in the bottom parental decile—are 33.4, 44.0, 24.7 and 26.6 for wealth against 62.8 pooled, and 43.5, 43.6, 21.5 and 17.6 for income against 49.2 pooled. In both outcomes every within-category ratio lies below the pooled one, so part of what the pooled ratio measures is the greater prevalence of transfers among children of high-income parents rather than a larger advantage within any group. Transfer status is post-determined, so the panels describe who reaches the top rather than the effect of transfers on doing so. Sample and variable definitions as in Figure 5.

The composition of transfers among top-wealth children. The previous figure asked, for children of a given parental origin, how likely they were to reach the top. We now reverse the question and look only at the children who did reach it, asking where they came from and what role transfers played on the way. Figure 8 takes the children in the top percentile of gross wealth, groups them by the parental income origin they were raised in, and decomposes each group by transfer status. Three things are visible. Most of these children come from the upper half of the parental distribution, and only a small minority from bottom-half families. Within every origin group, transfer receipt is common, but it rises steadily with origin, from 53% among those raised in bottom-half families to 85% among those raised in top-percentile families. And the composition of that receipt rotates as origin rises: among those who came from below, the modal observed route into the top percentile is an inheritance (27.7% against 15.9% for gifts alone), while among those who came from the top it is an inter vivos gift (45.0% against 15.6%)²². Read alongside the low incidence of gifts elsewhere in the distribution, this pattern is consistent with inter vivos giving acting as a channel of persistence at the top rather than of upward mobility: children who reach the top of the wealth distribution from below are considerably more likely to have inherited, while gifts are concentrated among those whose parents were already there.

²²The same decomposition for the top percentile of child *income* (Appendix Figure A26) shows the same pattern but lower receipt at every origin: 67% of top-earning children of top-percentile parents have received a transfer against 85% of their top-wealth counterparts, and 30% against 53% from bottom-half origins.

Figure 8: Transfer Composition of Top-1% Gross Wealth Children, by Parental Income Group



Notes: Children in the top 1% of child gross wealth ($N = 15,023$), grouped by parental household income origin: Bottom 50% (percentiles 1–50), Middle 40% (51–90), Top 10% (91–99, excluding the top percentile) and Top 1% (100). Bars decompose each origin group by transfer category; the bold label above each bar reports the share of all top-1%-wealth children contributed by that origin, and the annotation above it the within-group receipt rate and cell size. Receipt rises monotonically with origin, from 53% among children of bottom-half families to 85% among children of top-percentile families, and the mix rotates: inheritance is the modal observed route from below (27.7% versus 15.9% gift-only), gifts from above (45.0% versus 15.6%). The counterpart for the top percentile of child income is reported in Appendix Figure A26, and the corresponding distribution of parental origins in Appendix Figure A25. Transfer status is post-determined; shares describe composition, not effects.

A final exercise puts a magnitude on the pattern. Repeating the equal-access counterfactual of Section 3.3 for this outcome, holding gift receipt rates fixed at their population level across all parental origins lowers the gradient in the probability of reaching the top percentile of wealth by 26.6%, against 12.5% for the average wealth rank; equalizing inheritance receipt again leaves it close to unchanged (Appendix Figure A27). The accounting contribution of family transfers to intergenerational persistence in Spain is thus concentrated at the upper end of the wealth distribution, where both the incidence of gifts and the advantage they accompany are largest. Three qualifications bound how far this can be taken. Transfer status is post-determined and only its extensive margin is recorded, so the counterfactuals describe how much of the observed gradient is accounted for by differences in who receives, not how much of it transfers cause; amounts are unobserved and most bequests are not yet realized, so the magnitudes are lower bounds on the accounting contribution itself; and transfers are one channel among several through which parental advantage may be transmitted, alongside education, occupational networks and assortative matching, which these data do not separate.

3.5 Evolution of Intergenerational Mobility Across Cohorts

Finally, we ask whether measured persistence differs across the cohorts we observe. Because the youngest cohorts are followed only until 2022, comparing cohorts requires fixing the age at which child income is measured; we set this age at 31, the oldest age common to all eleven cohorts (1980–1990), and average income over a two-year window centered on it.²³ Appendix Figure A31 plots the resulting cohort-specific rank–rank slope. It traces a shallow U: persistence falls from 0.214 for the 1980 cohort to a minimum of 0.186 for the 1983 cohort, then rises steadily to 0.265 for the 1990 cohort. Measured persistence is thus higher for cohorts reaching adulthood after 2008 than for those reaching it before. This shape does not depend on the particular age at which we fix measurement: as Appendix Figure A32 shows, the same U-shaped profile appears when child income is measured at ages 28, 29, or 30 rather than 31, with only the level shifting.²⁴ Appendix Figure A33 shows that the rise in persistence for the later cohorts reflects movement at both ends of the parental distribution: a decline in the mean adult rank of children from families at the 25th percentile and an increase for children from families at the 75th percentile. Taken together, this descriptive analysis suggests that intergenerational income mobility has worsened for young adults born between 1980 and 1990 in Spain.

Several explanations could lie behind this pattern. The later cohorts entered the labor market during or after the Great Recession, and adverse conditions at the education-to-work transition are known to have persistent effects on the earnings of Spanish workers, particularly those from less advantaged backgrounds (Arellano-Bover, 2022a; Gurkova, 2025); a scarring channel of this kind would raise measured persistence for exactly the cohorts in which we observe it. A second possibility is compositional: the boom that preceded the recession expanded low-skill employment in construction and services, segments into which children from lower-income families disproportionately sort, so the early cohorts may have enjoyed temporarily compressed persistence that later unwound (Bernardi and Gil-Hernández, 2021). A third is that rank–rank slopes estimated in recession years differ mechanically from those estimated in expansion years, independently of any change in underlying mobility, since the calendar year of measurement moves one-for-one with birth cohort once age is held fixed. Broader structural changes over this period—the collapse of the housing market, austerity-driven cuts to public services, and rising inequality—could each contribute as well.²⁵ Because these channels move together across cohorts, our data cannot separate them, and disentangling cohort, age, and period effects in intergenerational trends is known to be difficult even with richer variation (Nyblom and Stuhler, 2024). We

²³This differs from the baseline sample of Section 3.1, where child income is measured in 2021–2022 at ages 36–42. Here income is measured at a common *age* but in a different calendar *year* for each cohort, ranging from 2011 for the 1980 cohort to 2021 for the 1990 cohort.

²⁴Levels shift downward at younger ages because rank–rank slopes are attenuated earlier in the life cycle (Appendix Figure A28); the cross-cohort *shape*, rather than its level, is the object of interest.

²⁵Immigrant children, who might otherwise raise composition concerns across cohorts, represent less than 3% of each cohort and cannot enter the panel if they arrive after the 1998–2000 parental income window; they are therefore unlikely to drive the profile. See Boustan et al. (2025) for more information on immigration and intergenerational mobility in Spain and many other countries in comparative perspective

therefore read Appendix Figure A31 as a descriptive feature of the cohorts we observe rather than as evidence on the mechanism generating it.

3.6 Robustness

Life-cycle bias. A standard concern in intergenerational mobility estimation is that measuring child income too early in adulthood understates the influence of parental background, because individuals with high lifetime income often experience steeper early-career earnings growth (Haider and Solon, 2006; Solon, 2002; Böhlmark et al., 2006). This concern is especially relevant in Spain, where school-to-work transitions are prolonged (Salvà-Mut et al., 2016). Appendix Figure A28 shows that the estimated rank–rank slope varies substantially when child income is measured before age 35 but converges tightly to the baseline value of 0.274 from age 36 onward. Appendix Figure A29 further shows that the estimate is robust to alternative child-income averaging windows and cohort definitions.

Attenuation bias. A second concern is that using too few years of parental income introduces measurement error and attenuates rank–rank slope estimates (Nyblom et al., 2018). Appendix Figure A30 reports the slope as a function of the number of years used to average parental income. The one-year estimate is 0.261, the three-year estimate used in our baseline is 0.274, and extending the window to six years yields 0.280. The modest increase from three to six years suggests that our baseline window captures nearly all of the permanent-income signal.

4 The Geography of Intergenerational Income Mobility in Spain

Does economic opportunity depend on where in Spain a child grows up? In this section we estimate the rank–rank slope and absolute upward mobility statistics introduced in Section 3 at every level of Spain’s administrative geography—autonomous communities (regions), provinces, municipalities, and ZIP codes—following the approach developed by Chetty et al. (2014c) for the United States and since applied to Italy, Brazil, and France (Acciari et al., 2022; Britto et al., 2026; Kenedi and Sirugue, 2022). We assign each child to the ZIP code where they were first claimed as a dependent on their parents’ 1998 income tax return, and measure adult outcomes at the same ages used in the national analysis (children born between 1980 and 1990, observed in 2022 at ages 32 to 42). For each geographic unit g we estimate

$$y_i = \alpha_g + \beta_g p_i + \varepsilon_i,$$

where y_i is child i 's national income rank in adulthood and p_i is the national rank of their parents' household income. The slope β_g measures relative mobility within unit g , and $\mu_{25,g} = \alpha_g + 25\beta_g$ is the predicted adult rank of a child raised at the 25th national percentile (*absolute upward mobility*). The answer to the opening question is affirmative, and strikingly so: we document a northeast–southwest gradient of roughly 25 percentile points in μ_{25} between the highest- and lowest-mobility municipalities—a gap comparable to the entire cross-country range documented in Section 3—that persists from the 1980 to the 1990 birth cohort.

4.1 Bayesian Hierarchical Estimation of Intergenerational Mobility at the Local Level

Spain's population is spread across its territory with extreme unevenness. The country is divided into 50 provinces, 8,131 municipalities, and some 11,752 ZIP codes, but these units differ in size by orders of magnitude: a handful of metropolitan areas hold a large share of the population, while much of the interior consists of municipalities with a few hundred inhabitants.²⁶ Because a local upward mobility estimate rests only on the children observed in that locality, this unevenness translates directly into unevenness in what can be measured. Sparse areas are not a peripheral nuisance here but the modal case: half of Spanish municipalities, and four in ten ZIP codes, contain fewer than twenty-five children in our cohort window.²⁷

The consequence is that cell-by-cell estimation fails over most of the country, and fails in a way that is easy to overlook. Fitting a rank-rank regression separately in each area yields values that are not merely imprecise but frequently impossible—predicted ranks below zero or above one hundred, negative slopes—and, more importantly, error of the same order of magnitude as the geographic variation we set out to describe. A Monte Carlo exercise on our own panel makes this concrete in the units of the finding itself: at twenty-five children, the sample size of the modal Spanish area, unpooled OLS recovers absolute upward mobility with a root mean squared error of roughly nine percentile points, against a cross-area range of about twenty-five points. More than a third of the range that constitutes our substantive result would therefore be sampling noise.

Two responses are available, and neither is satisfactory on its own. Restricting attention to areas with enough data preserves the properties of the estimator but discards roughly three quarters of Spanish municipalities and ZIP codes—precisely the small and rural places whose mobility prospects are of independent policy interest, and whose exclusion would leave the geography of opportunity described only where population is dense. Estimating everywhere without adjustment retains the full map but confounds sampling error with genuine spatial variation, so that the apparent differences between two small municipalities would mostly reflect how many children happen to be observed in each.

²⁶Maps of each administrative level appear in Appendix Figures A34 and A35.

²⁷Roughly three quarters of municipalities and of ZIP codes contain fewer than one hundred children. The median ZIP code contains 36, as noted above.

We therefore estimate local intergenerational mobility with a Bayesian hierarchical model that partially pools information across nested geographic levels (Gelman et al., 2020; McElreath, 2020). Each area’s estimate is a compromise between what its own children’s outcomes imply and the pattern of its containing province, with the weight on the former determined by how precisely the area is measured: sparse cells are pulled toward their province, data-rich cells are left essentially untouched, and the degree of pooling is estimated from how much areas within provinces actually differ rather than imposed. Every area in the country therefore receives an estimate, and each carries an explicit measure of its own reliability. For absolute upward mobility and the rank-rank slope, we fit

$$y_i \sim \text{Normal}(\mu_i, \sigma) \quad (2)$$

$$\mu_i = (\alpha + \alpha_p + \alpha_m) + (\beta + \beta_p + \beta_m) x_i \quad (3)$$

$$\alpha_p \sim \text{Normal}(0, \tau_{\alpha p}), \quad \alpha_m \sim \text{Normal}(0, \tau_{\alpha m}) \quad (4)$$

$$\beta_p \sim \text{Normal}(0, \tau_{\beta p}), \quad \beta_m \sim \text{Normal}(0, \tau_{\beta m}) \quad (5)$$

where x_i is the parental income rank of child i and p and m index her province and municipality. The same hierarchy applied at the province-by-ZIP level yields ZIP-resolution estimates²⁸.

We validate the estimator against a ground truth constructed from the panel itself. Within ZIP codes large enough for their full-sample fit to be treated as the local truth, we draw subsamples of the sizes actually encountered across Spain, refit unpooled OLS and the hierarchical model on each draw, and compare both to that truth. Partial pooling cuts the error of the modal area by more than half, from roughly nine percentile points to four, and unpooled OLS remains two to three times less accurate until sample sizes that the large majority of Spanish municipalities and ZIP codes never reach. Adding the province level to the hierarchy improves accuracy further at small samples, by borrowing information from neighboring areas within the same province. Crucially for interpretation, the shrinkage is conservative in the direction that matters: because sparse areas are pulled toward their province, partial pooling attenuates genuine spatial differences rather than manufacturing them, so the northeast–southwest gradient we document is a lower bound on the dispersion in the underlying data.²⁹

For *rags-to-riches* probabilities $P(Q_5 \mid Q_1)$, sparse-cell estimation is more demanding still, since each entry of the 5×5 transition matrix requires its own observation count. We estimate these with a hierarchical ordered logistic model with province-varying cutpoints and monotonic effects of parental quintile, the discrimination parameter varying by province and quintile; full specifications for both models are given in Appendix K.2. The causal analysis of Section 5 rests on the same partial-pooling logic, applied to the place quality of permanent residents at this same municipal resolution, with one difference in architecture and one in specification.

²⁸We do not add an autonomous-community level, because several autonomous communities consist of a single province and the two effects would not be separately identified.

²⁹Full simulation design and per-parameter results (intercept, slope, μ_{25} , μ_{50} , μ_{100}) appear in Appendix K and Figure A61.

4.2 Main Geographic Patterns

Figure 9 maps absolute upward mobility across Spanish municipalities. Panel (a) shows direct OLS estimates of μ_{25} , the expected mean child income rank for children whose parents are at the 25th percentile of the national parental income distribution. These estimates reveal substantial local variation but are necessarily noisy in small municipalities. Panel (b) reports Bayesian hierarchical estimates that partially pool information across geographic units, reducing small-area noise and allowing a more complete picture of local mobility across the entire municipal map.


The dominant pattern is a clear North-East/South-West divide. The highest levels of upward mobility are concentrated in Catalonia, parts of Aragón, the Madrid metropolitan area, and several north-central provinces such as Segovia, Soria, and Guadalajara. The lowest levels are found in Andalucía, Extremadura, the Canary Islands, and parts of the South-West. The gap between the highest- and lowest-ranked municipalities exceeds 25 percentile points, and rags-to-riches probabilities display a similar geography: children from the bottom quintile in Barcelona or Lleida are over three times more likely to reach the top quintile than comparable children in Tenerife or Sevilla (Appendix Figure A36). Crucially, the Bayesian estimates in Panel (b) preserve the same broad regional geography visible in Panel (a) while smoothing the patchiness generated by sparse local samples, confirming that the North-East/South-West divide is not an artefact of the hierarchical model.

The origins of this stark spatial divide date back centuries and still cast a long shadow on observed disparities in economic opportunity. A large body of economic history research traces persistent regional disparities in Spain to the medieval Reconquest and its legacy of unequal landownership, institutional development, and social capital (Oto-Peralías and Romero-Ávila, 2016; Oto-Peralías, 2020); to the geographic concentration of industrialization and market access (Budí-Ors and Pijoan-Mas, 2022; Beltran Tapia and de Miguel Salanova, 2021; Beltrán Tapia et al., 2021); and to the resulting spatial inequalities in human capital accumulation and educational infrastructure (González-Val et al., 2026; Insa-Sánchez, 2024). Our spatial mobility estimates suggest that these historical forces continue to shape the geography of economic opportunity in present-day Spain. At the same time, the municipal estimates reveal a more nuanced geography beyond broad regional divides. These patterns are robust across geographic scales: they persist at the urban area, province, and autonomous community levels (Appendix Figures A37–A39) and remain visible at the finer ZIP-code level (Appendix Figure A40).³⁰


³⁰Outside of urban areas, the vast majority of Spanish municipalities correspond to a single ZIP code, so that ZIP-code and municipality-level estimates are often identical in rural areas.

Figure 9: Geographic Variation in Absolute Upward Mobility in Spain

(a) OLS estimates



(b) Bayesian hierarchical estimates



Notes: This figure maps absolute upward mobility (μ_{25}) across Spanish municipalities. μ_{25} is the expected mean child income percentile for children whose parents are at the 25th percentile of the national parental income distribution. Panel (a) presents direct municipality-level estimates. Panel (b) presents estimates from the Bayesian hierarchical model described in Section 4.1. The direct estimates are noisier, especially in municipalities with smaller samples, while the Bayesian model partially pools information across geographic units and provides a smoother representation of the same broad regional pattern. The color scales represent mean child income percentiles.

Robustness. The province random effect raises a natural concern about the municipal map of Figure 9(b): partial pooling could either manufacture cross-province discontinuities or attenuate genuine local differences. We test this where the worry is sharpest, at province borders. For each pair of adjacent provinces we compare the mobility gap between municipalities on either side of the boundary, as estimated by the hierarchical model, with the same gap estimated by a single OLS regression pooling all border municipalities on each side—an aggregate with thousands of children per side and no pooling at all. The two agree closely (correlation 0.84–0.88 across alternative definitions of a border municipality), and where they differ the hierarchical estimates are always the smaller ones. Partial pooling attenuates border gaps rather than creating them, and the attenuation shrinks monotonically with municipal sample size, effectively vanishing in the largest municipalities. Appendix L.1 reports the full exercise.

4.3 Correlates of Upward Mobility

What local environments are associated with higher upward mobility for children from low-income families? To answer this question, we combine the mobility estimates constructed in the previous sections with a broad set of municipality-level characteristics measured around the period in which the children in our sample were growing up. These variables come from three administrative sources. First, we use the 2001 Population and Housing Census from the National Statistics Institute (INE), which provides information on educational attainment, occupational structure, household composition, housing tenure, and commuting patterns. Second, we use the Digital Atlas of Urban Areas in Spain, produced by the Ministry of Transport, Mobility and Urban Agenda, which provides indicators on labor-market conditions, land use, demographic structure, and urban amenities, mostly measured between 2001 and 2007. Third, we use municipal income statistics from the Spanish Tax Agency, including mean income and inequality measures. Appendix N provides detailed variable definitions and sources.

To complement this administrative characterization of local environments, we incorporate a new measure of bridging social capital: upward economic connectedness. In their pioneering work, Chetty et al. (2022a,b) construct measures of economic connectedness using billions of Facebook friendship links in the United States to capture the extent to which individuals from low- and high-socioeconomic-status backgrounds are socially connected. They show that this measure is one of the strongest predictors of upward mobility across local areas. At the global scale, Johnston et al. (2026) extend this approach by estimating economic connectedness and related social-capital measures using data on 595 billion social connections among 2.1 billion Facebook users from 183 countries. We use their finest geographic estimates for Spain: upward economic connectedness at the ZIP-code level, defined as the average socioeconomic status of the friends of low-SES individuals, where low SES corresponds to the 25th percentile of the national SES distribution.³¹ To our knowledge, this is the

³¹See the Data and Methodology section in Johnston et al. (2026) for details on the construction of the social-capital measures.

first fine-grained measure of bridging social capital in Spain based on large-scale friendship networks across socioeconomic groups.

Main correlates of upward mobility. We begin by estimating univariate regressions of upward mobility on each standardized municipal characteristic³². Figure 10 reports the strongest correlate within each thematic category, comparing results based on the two municipality-level upward mobility estimates introduced in Figure 9: the OLS estimates shown in Panel A of that figure and the Bayesian shrinkage estimates shown in Panel B.³³ The OLS correlations are estimated on 3,067 municipalities, while the Bayesian correlations are estimated on the extended set of 5,676 municipalities³⁴. Economic connectedness is the only variable with a different coverage structure: the raw EC data contain 3,308 Spanish ZIP codes and, after excluding the Basque Country and Navarra, 2,984 ZIP codes are matched to the mobility data. Aggregating these ZIP-code measures to the municipality level yields 1,900 municipalities with EC and Bayesian upward mobility estimates, and 1,812 municipalities with EC and OLS-based upward mobility estimates. These municipalities cover 88.6% of the population outside the Basque Country and Navarra. Coverage is nearly complete among large and medium-sized municipalities, but more limited among small rural municipalities; Appendix Figure A41 documents this coverage pattern.³⁵


³²See Appendix O for further details

³³See Appendix N for a detailed description of the variables included in each thematic category and their sources.

³⁴The choice between OLS and Bayesian estimates involves a trade-off. OLS estimates provide the cleanest basis for correlational analysis because they are computed independently for each municipality, without borrowing information from neighboring areas or covariates. Bayesian shrinkage estimates, by contrast, partially pool information across geographic units toward group means, which improves precision for ranking individual places—particularly small municipalities—but can mechanically strengthen correlations with variables that covary at the province or regional level. For this reason, OLS estimates are the preferred input for the correlational exercise. We display both sets of estimates in Figure 10 to show that the Bayesian estimates preserve the same qualitative ranking and quantitative structure of correlates, indicating that the patterns documented here are not driven by measurement error in small municipalities.

³⁵The top three correlates within each category are reported in Appendix Figure A43.

Figure 10: Main Correlates of Upward Mobility in Spain at the Municipality Level



Notes: Each point represents the coefficient from a univariate regression of upward mobility (μ_{25}) on one standardized covariate. Horizontal lines denote 95 percent confidence intervals. Variables are grouped into thematic categories; the figure displays the strongest predictor from each group. OLS and Bayesian shrinkage estimates of upward mobility are shown side by side, corresponding to Panels A and B of Figure 9, respectively. The OLS correlations are estimated on 3,067 municipalities, while the Bayesian correlations are estimated on 5,676 municipalities. Economic connectedness is estimated on the subset of 1,812 municipalities with matched EC and OLS-based mobility data. The sign of the association is indicated by (+) or (–) next to each variable name. See Appendix Figure A43 for the top three correlates per category.

The most striking pattern in Figure 10 is the dominance of economic connectedness. Despite the rich set of local covariates available, the extent to which low- and high-income individuals befriend each other within a community emerges as the single strongest predictor of upward mobility. The mechanism emphasized in the social capital literature is that cross-class friendships provide low-income children with access to information, professional networks, and aspirational norms that would otherwise remain out of reach (Chetty et al., 2022b; Putnam, 2000; Bourdieu, 1986, 1987). This relationship appears especially pronounced in Spain: Johnston et al. (2026) show that the economic connectedness and upward mobility gradient in Spain is the second steepest among all countries for which geographically disaggregated estimates of intergenerational mobility based on administrative data are available, and the finding is consistent with the U.S. evidence in Chetty et al. (2022a), who identify this form of social capital as the strongest predictor of upward mobility across U.S. counties.

Beyond social capital, the remaining administrative correlates align with patterns documented internationally (Britto et al., 2026; Kenedi and Sirugue, 2022; Connolly et al., 2019a; Chetty et al., 2014b). Local labor-market conditions and human capital emerge as the next most powerful correlates. Municipalities with a higher share of households affected by unemployment exhibit substantially lower upward mobility, reinforcing a key economic mechanism in Spain: access to stable employment during formative years shapes long-run outcomes both across cohorts and across places. This is consistent with a broad literature documenting the persistent scarring effects of adverse labor-market conditions (Kahn, 2010; Oreopoulos et al.,

2012; Fernández-Kranz and Rodríguez-Planas, 2018; Arellano-Bover, 2022b), and with recent evidence that parental job loss itself has lasting consequences for children’s economic trajectories (Chetty et al., 2026c; Oreopoulos et al., 2008; Hilger, 2016). Municipalities with higher educational attainment among young adults tend to have greater upward mobility, likely through both direct returns to schooling and indirect spillovers via the local skill composition, better schools, stronger peer effects, and a wider range of accessible professional occupations (Moretti, 2004b,a; Sacerdote, 2001; Chetty et al., 2014a). Conversely, a larger share of agricultural employment, which tends to offer lower wages, seasonal contracts, and limited opportunities for occupational upgrading (Blázquez Cuesta, 2005), predicts lower mobility. The negative association with youth dependency ratios complements this picture: municipalities with a larger share of young dependents face greater pressure on household resources and local public services, constraining the investments that support children’s long-run outcomes (Van Bavel et al., 2011; Geys et al., 2008). The remaining correlates are weaker but consistent with this characterization. Car-dependent commuting patterns, which proxy for spatial isolation from employment centers, are negatively associated with mobility. Mean municipal income and the share of foreign-born residents display positive but modest associations, suggesting that aggregate income levels matter less than how local economies translate resources into opportunity (Chetty et al., 2014b).


Local income inequality is a natural candidate correlate of intergenerational persistence, given the well-documented “Great Gatsby Curve” across countries (Corak, 2013). In Spain, however, this relationship is only detectable among large cities (Appendix Figure A44). Among the 50 largest municipalities, the correlation between the Gini coefficient (restricted to the bottom 99% of the distribution) and the rank-rank slope is approximately 0.53. The pattern does not extend to smaller municipalities because Gini coefficients estimated from small populations are noisy, and because smaller municipalities display limited dispersion in measured inequality (Appendix Figure A45). The Great Gatsby Curve in Spain is therefore primarily an urban phenomenon, emerging where income segregation across neighborhoods is most pronounced.

Economic connectedness and the geography of opportunity. The prominence of economic connectedness motivates a closer examination of its spatial relationship with upward mobility. Figure 11 plots upward mobility against upward economic connectedness across the 1,807 municipalities for which both measures are available.³⁶ The association is striking: municipalities where low-SES individuals are more likely to befriend higher-SES peers display substantially higher rates of upward mobility. In the population-weighted regression, a one-standard-deviation increase in economic connectedness is associated with 0.77 standard

³⁶Unlike the other local characteristics we consider, economic connectedness is not measured while the children in our sample were growing up: it is constructed from Facebook friendship networks observed in recent years, when they are already adults. The association is therefore open to reverse causality, since more upwardly mobile children may themselves form more cross-class ties in adulthood, or may sort into places where such ties are common. Two considerations bound this concern. Economic connectedness is a feature of places rather than of individuals, and the children in our sample are a small share of any municipality’s adult population, so their own friendships can move the local measure only marginally. And the evidence for the United States indicates that connectedness reflects stable features of local social structure—the composition of schools, workplaces and neighborhoods, and the exposure and friending behavior these generate—rather than short-run outcomes, so that a recent measure is informative about the environment of earlier cohorts (Chetty et al., 2022a,b). We nonetheless read this association as descriptive.

deviations higher upward mobility ($SE = 0.052$), while the unweighted coefficient is 0.65, indicating that the relationship is not driven solely by Spain's largest cities.

Figure 11: Economic Connectedness and Upward Mobility in Spain at the Municipality Level



Notes: Each point represents a municipality. The horizontal axis measures upward economic connectedness and the vertical axis measures upward mobility based on OLS estimates. Point size is proportional to 2001 municipal population. Blue points denote municipalities above the national median of upward mobility; red points denote those below. The fitted line is from a population-weighted OLS regression. Selected municipalities are labeled for reference. The sample includes municipalities with available mobility and EC data.

A natural concern is that the strong association between economic connectedness and upward mobility simply reflects broad regional income differences in Spain. Both upward mobility and economic connectedness are higher in many northern and north-eastern areas, and lower in parts of southern and western Spain. To assess this concern, Appendix Table A8 sequentially controls for municipal income, municipal inequality, and geographic fixed effects. Two results emerge. First, the association between economic connectedness and upward mobility is not explained by mean municipal income or local income inequality. Adding these controls barely changes the coefficient on economic connectedness in the unweighted specifications, and only moderately attenuates it in the population-weighted specifications. This suggests that economic connectedness captures a dimension of local social structure that is distinct from the aggregate level and distribution of income. Second, adding regional or province fixed effects substantially reduces the coefficient on economic connectedness. This attenuation is informative. It indicates that a large part of the relationship operates across

broad regions, consistent with the spatial gradients visible in the maps (Figure 9). At the same time, the coefficient remains positive and statistically significant even after controlling for province fixed effects. Economic connectedness is therefore not merely a proxy for the North–South income divide: within the same broad regional environments, municipalities where low-SES individuals are more connected to higher-SES peers also tend to exhibit higher upward mobility. The sharp increase in explanatory power after adding geographic fixed effects underscores that opportunity in Spain is deeply territorial, but the remaining coefficient shows that social connectedness captures variation within that geography.

Multivariate analysis. To assess whether economic connectedness retains predictive power after conditioning on the full set of observable local characteristics, Appendix Table A9 reports a more demanding horse-race specification. Starting from economic connectedness, we sequentially add the leading individual-variable correlates: unemployment, education, youth dependency, and agricultural employment. Economic connectedness remains positive and statistically significant after adding these controls, though its coefficient is attenuated. This pattern suggests that economic connectedness is partly a summary statistic for a broader bundle of high-opportunity local conditions, especially stronger labor markets and more favorable demographic and occupational structures. When geographic fixed effects are added to this richer specification, the remaining economic connectedness coefficient is further reduced, confirming that much of the predictive power of economic connectedness reflects broad regional variation.

Taken together, the correlational evidence suggests that the geography of opportunity in Spain is shaped by broad regional-level, structural differences in labor-market conditions, human capital, occupational composition, and in particular, economic connectedness, generating the pronounced North–South gradient that dominates the spatial distribution of upward mobility. Importantly, a methodological implication of the correlational analysis is that the expanded Bayesian estimates appear to be a reliable basis for extending geographic analysis to smaller local areas. As shown in Figure 10, both the direction and the magnitude of the univariate correlations are closely aligned across the OLS and Bayesian specifications. This suggests that the hierarchical shrinkage procedure preserves the underlying correlational structure of upward mobility rather than distorting it, lending confidence to the use of Bayesian estimates for spatial analyses in municipalities where direct OLS-based upward mobility estimates would be imprecise.

4.4 Evolution of Upward Mobility Across Cohorts in Large Municipalities

The decline in intergenerational mobility documented at the national level in Figure A31 was widespread across the territory, but territorially uneven and especially pronounced in areas that already displayed relatively low upward mobility. Appendix Figure A46 compares upward mobility for children from low-income families across Spain’s 50 largest municipalities for the 1980 and 1990 birth cohorts, measuring outcomes at

the same age of 32. The central pattern is clear: almost all large municipalities experienced a decline in upward mobility between cohorts. Appendix Figure A47, which extends the comparison to all municipalities through two map snapshots, shows the same broad deterioration and the expansion of lower-mobility areas, particularly in southern and central Spain.

The magnitude of the decline, however, varied across local areas. Some suburban municipalities around Barcelona, such as Terrassa and L'Hospitalet de Llobregat, remained among the highest-mobility places for the 1990 cohort, indicating relative resilience within the broader national decline. By contrast, several suburban municipalities around Madrid, including Móstoles and Getafe, experienced sharper declines. These cohort differences should be interpreted descriptively rather than as causal estimates of place effects, since they combine national macroeconomic conditions at labor-market entry, local labor-market dynamics, family sorting, and changes in local opportunity structures. Nevertheless, the spatial pattern is informative: the fall in mobility prospects for children from low-income families was not a uniform national shift, but a geographically uneven process that reshaped the map of opportunity across cohorts. This interpretation is consistent with recent evidence that economic opportunity can change substantially across adjacent birth cohorts as local childhood environments and social communities evolve (Chetty et al., 2026c).

5 The Causal Effects of Childhood Environments on Upward Mobility

The previous section documents large differences in upward mobility across Spanish municipalities and identifies the local characteristics—economic connectedness, labor-market conditions, educational attainment—most strongly associated with them. Those associations describe where children do better; they cannot establish whether the gaps between places reflect the causal effect of growing up in one rather than another, or the sorting of families with different unobserved characteristics across them. The distinction has direct policy implications, since place-based interventions can raise upward mobility only to the extent that cross-area gaps are causal (Chyn and Katz, 2021; Gaubert et al., 2025).

To separate the two, we implement the childhood-exposure movers design of Chetty and Hendren (2018a), subsequently applied by Deutscher (2020), Laliberté (2021), Beekers (2024), and Britto et al. (2026), at the municipality—the same geography at which the variation and its correlates are measured, so that the causal estimates speak directly to the map of Section 4. The design compares children whose families make the same move—from the same origin, to a higher-mobility destination, from the same parental income stratum—at different ages. If exposure to place causally affects adult income, children who move earlier, and therefore spend more of childhood in the destination, should converge more fully toward the outcomes of children raised there; if the cross-area gaps reflect sorting alone, the age at which the family moves should be irrelevant.

The canonical implementation of the movers design proceeds in two steps: it first estimates each area's

quality from the outcomes of its permanent residents, and then regresses movers' outcomes on those first-stage estimates as if they were data. This is a classical errors-in-variables problem, and it is benign only when the first stage is precise. Precision, however, is governed by the number of permanent residents per area, cohort, and income stratum—and at any resolution fine enough to correspond to a neighborhood or a school catchment, that number is small in most countries, not only in Spain. The design then faces a trade-off with no acceptable resolution inside the two-step framework. Restricting the sample to areas whose quality is precisely estimated preserves the properties of the estimator, but at fine resolution such filters retain only a small and non-random fraction of moves, since the areas they exclude are the small ones and the families they exclude are disproportionately those whose children the design is meant to study. Dropping the filter retains the sample, but the first-stage noise then attenuates the exposure coefficients toward zero, and the resulting estimates understate the causal role of place. The standard econometric correction does not escape the trade-off either: a split-sample instrumental-variables estimator restores consistency by instrumenting one half-sample's quality gap with the other's, but it does so by halving an already-thin first stage, and its variance cost is largest precisely where the correction is needed most. Consistency is thus attainable, but not with the precision required to estimate an age profile.

5.1 Bayesian movers design specification

Conceptual framework. We resolve this trade-off by never estimating place quality separately: we propose a Bayesian movers design estimated in a single step, in which place quality is not an estimate to be plugged in but a latent quantity identified jointly with the exposure coefficients, from the permanent residents and the movers at once. Two properties follow. First, nothing needs to be filtered: the permanent residents enter through the exact sufficient statistics of their likelihood, so the model compresses the *data*, which is lossless, rather than an *estimate*, which is not, and every area contributes exactly its information content however small it is.³⁷ Second, the uncertainty about the quality of sparsely populated areas is carried into the exposure coefficients rather than discarded, so the attenuation that the two-step estimator cannot avoid does not arise. The model has the same two ingredients as the canonical design—a measurement of place quality from permanent residents, and a regression of movers' outcomes on the quality gap they experience—but the two are estimated simultaneously rather than sequentially, and share the same parameters rather than passing an estimate from one to the other. We set out the model below, then its relation to the alternatives, the estimation sample, and the parameterizations we report.

Place quality and the childhood exposure regression. Following Chetty and Hendren (2018a), we measure the quality of each place by the outcomes of the children of its *permanent residents*: $\theta_m(c, x)$ denotes the

³⁷The group-level likelihood depends on the data only through six sums, so the collapse is exact rather than approximate (Appendix P.2).

expected adult income rank of a child of cohort c raised in municipality m by parents at income rank x , whose family never moves. A place is good if the children who grow up there reach high adult ranks given their parents' rank, measured on the families that stay. This definition, and the interpretation of the exposure coefficients below, are inherited unchanged from the canonical design; the estimator departs from it in three places, which we flag as they arise.

The first is how θ is obtained. Because most Spanish municipalities contain few permanent residents per cohort, we do not estimate θ cell by cell in a first stage, but model it as a latent hierarchical surface,

$$\theta_m(c, x) = \underbrace{\alpha_{p(m),c} + \beta_{p(m),c} x}_{\text{province-by-cohort rank-rank line}} + \underbrace{\text{dev}_m}_{\text{municipal shift}}. \quad (6)$$

Within each province and cohort, children's adult ranks rise linearly with their parents' rank, and every municipality in that province sits on the same line displaced up or down: $\alpha_{p(m),c}$ is the expected adult rank of a child from the bottom of the parental distribution, $\beta_{p(m),c}$ the local persistence of parental advantage, and dev_m the municipality's own contribution to upward mobility, by which its children exceed or fall short of what their province and their parents' rank predict. A municipality with $\text{dev}_m = 2$ raises the expected adult rank of every child raised in it by two percentiles. Place quality is therefore defined for every municipality in the country and carries an explicit measure of its own precision, where the canonical design has a point estimate treated thereafter as known, and nothing at all below its size filter.

The treatment relevant for a mover is the gap in this surface between destination and origin, evaluated at her own cohort and parental rank, $\Delta_i = \theta_{d(i)}(c_i, x_i) - \theta_{o(i)}(c_i, x_i)$: how much higher an adult rank the destination delivers, for a child with her parents' rank, than the origin she left. Her own adult income rank y_i , measured at the same age as that of the permanent residents to whom she is compared, loads on that gap through

$$y_i = \alpha_{\text{cell}(i)} + \lambda_{c_i} + h \text{dev}_{o(i)} + b_{a_i} \Delta_i + \varepsilon_i. \quad (7)$$

Two terms are standard. $\alpha_{\text{cell}(i)}$ is a fixed effect for the mover's origin-province $\times$ parental-quintile $\times$ age-at-move cell, absorbing any difference in average outcomes across families who relocate at a given age, from a given province, out of a given income stratum; as in [Chetty and Hendren \(2018a\)](#) and [Britto et al. \(2026\)](#), b_a is therefore identified within such cells—among families alike in origin, timing, and parental income, some move to better destinations than others—and from how that covariation changes with age at move.³⁸ The cohort intercepts λ_c absorb national differences in labor-market entry conditions. The term $h \text{dev}_{o(i)}$ is the second departure. Recall that dev_m is not a raw municipal estimate but a partially pooled one: the model weighs each municipality's own permanent residents against the rank-rank line of its province in proportion to how precisely that municipality is measured, so a municipality with a handful of resident children is pulled

³⁸ A Mundlak-style control function in place of these cells leaves the exposure slope unchanged (Appendix P.6).

close to its provincial line while one with several thousand is left essentially where its own data place it. This is the shrinkage logic of the small-area estimator of Section 4.1, applied here inside the causal design. The term itself is absent from implementations in which the selection cells and the quality gap share a geography. Ours do not: the cells are provincial while Δ_i varies across municipalities within them, so two families in the same cell may leave origins of very different quality. Because origin quality enters Δ_i with a negative sign, a family leaving a low-mobility municipality mechanically faces a larger gap; omitting the term would let that mechanical relation bias the exposure coefficients downward, most severely for movers who stay within their province.³⁹ The coefficient of interest, b_a , retains its usual meaning: the share of the destination-vs-origin quality gap realized by children who relocate at age a , with $b_a = 1$ implying full convergence to the outcomes of the destination’s permanent residents and $b_a = 0$ none. The third departure is that equations (6) and (7) are estimated jointly, on the movers together with the sufficient statistics of the 1.32 million children of permanent residents, rather than sequentially: Δ_i is never treated as known, and the two blocks share parameters instead of one passing an estimate to the other.⁴⁰

Relation to two-step and instrumental-variables estimators. Three properties of the joint estimator are worth stating before turning to the sample. First, it is not an alternative to the canonical design but an extension of it to a regime the canonical design cannot reach: where place quality is precisely measured the two coincide, and they separate only where measurement error is large.⁴¹ Second, the Bayesian structure is not doing the identification. The joint model attains the consistency of the split-sample instrumental-variables correction at a fraction of its variance cost: the two estimators agree on the exposure profile, and only one of them is precise enough to report it.⁴² Third, the correction does not flatter the causal reading: attenuation biases the selection parameter as much as the exposure parameter, so correcting it *raises* the share of the geographic variation attributed to the sorting of families rather than to place.⁴³

Sample and outcome definitions. A child’s adult outcome is her income rank at age 24, computed within cohort from 2008 onward (Section 2.1). As in Chetty and Hendren (2018a) and Britto et al. (2026), we measure income at the earliest age observable for every cohort: this maximizes the range of observable ages at move and turns later moves into a mechanical placebo. It also fixes the cohort window at 1984–1994, the cohorts whose age-24 incomes fall within the near-universal coverage regime.⁴⁴ Movers are identified from the parents’ residential histories over 1998–2022 — the same location panel that underlies the geographic analysis of

³⁹Restricting dev_m to a cohort-invariant intercept, rather than allowing municipality-specific rank-rank slopes or cohort-varying municipal effects, is data-driven: neither carries identifiable signal in these data (Appendix P.2, which also reports the fitted loading h).

⁴⁰The resulting latent quality surface reproduces the independent municipal estimates of Section 4 (correlation 0.93; Appendix P.7). Appendix P.2 reports the full specification, priors, and convergence diagnostics.

⁴¹At province resolution, where the quality gaps are precisely measured, the two-step and joint estimates agree (Appendix P.4).

⁴²The instrumental-variables estimator is unbiased throughout but far noisier at the reliability of our data (Appendices P.4 and P.5).

⁴³A two-step implementation places the selection level at 0.19 rather than 0.32 (Appendix P.4).

⁴⁴Coverage for the 1980–83 cohorts is incomplete and income-selective at age 24; they enter the age-28 robustness series instead (Appendices P.3 and P.6).

Section 4 — under a single classification rule: the origin is the first observed municipality, and a move requires a change of municipality followed by three years of stability. Extending these histories back to 1998 requires recovering family moves over 1999–2005 from municipal codes embedded in the tax records themselves; these contribute 38% of the mover sample and essentially all moves observed at the youngest ages.⁴⁵ The estimation window covers moves at ages 10–30, extended to 33 in the placebo region. Moves at ages 24–25 are displayed but excluded from estimation, a convention more conservative than the $a > 24$ placebo rule of Chetty and Hendren (2018a) and Britto et al. (2026).⁴⁶ The estimation sample comprises 1,320,889 children of permanent residents across 5,507 municipalities and 192,436 movers.

Reported parameterizations. As in the previous literature we summarize the exposure curve two ways. The *semiparametric* specification estimates each b_a freely, producing the per-age curve of Figure 12. The *linear-hinge* specification of Chetty and Hendren (2018a) collapses the curve into interpretable scalars: over the causal window ($a \leq 23$) a line with level b_{23} at age 23 and slope $-\gamma_{\text{exp}}$, so each additional year of childhood exposure closes γ_{exp} of the destination-vs-origin gap; over the placebo region ($a \geq 26$) a level δ_{sel} and slope δ_{post} , which capture pure selection, since moves at these ages postdate the measured outcome. The flatness of the placebo segment ($\delta_{\text{post}} \approx 0$) is the design’s internal validity test, and it is also what makes the two levels comparable: both b_{23} and δ_{sel} are evaluated at age 23, so their difference isolates the causal component of a childhood move only if the placebo profile is genuinely flat.

5.2 Main results

Figure 12 reports the childhood exposure effects: for each age at move, the share of the destination-vs-origin gap realized by children who moved at that age. The figure displays the canonical pattern of the movers-design literature (Appendix Figure A67): $\hat{b}_a$ declines from 0.68 for moves at age 10 to 0.52 at age 23 (the fitted hinge line runs from 0.76 to 0.50 over the same ages), and is flat at 0.32 throughout the placebo region. The placebo test is passed: the post-outcome profile is flat, as required by the identifying assumption that selection does not vary with age at move⁴⁷.


Exposure and selection effects and the causal share of geographic variation. The linear-hinge fit gives $\hat{\gamma}_{\text{exp}} = 0.020$: each additional year of childhood exposure to a municipality one percentile higher in permanent-resident outcomes raises the child’s adult rank by 0.020 percentile points—equivalently, each

⁴⁵Appendix P.1 documents the classification rules and audits the recovered moves against an independent location source.

⁴⁶Recorded move years can trail true moves by up to one year, so these ages mix partial exposure, disruption, and selection (Appendix P.3).

⁴⁷The estimates at ages 24–25 fall between the causal segment and the placebo level ($\hat{b}_{24} = 0.38$, $\hat{b}_{25} = 0.30$), consistent with partial same-year exposure and the one-year lag in recorded move dates; as noted in Section 5.1, we display these two ages but exclude them from estimation.

Figure 12: Childhood Exposure Effects on Income Ranks in Adulthood



Notes: This figure shows the per-age-at-move exposure effects of moving to a higher-opportunity municipality. Each point is the posterior mean of b_a from the semiparametric specification of the joint model (equation (7); see Appendix P); vertical bars are 90% credible intervals. The solid line is the exposure segment of the linear-hinge fit, with slope $-\hat{\gamma}_{\text{exp}} = -0.020$ over $a \leq 23$; the dotted horizontal line marks the selection level $\hat{\delta}_{\text{sel}} = 0.32$ ($\hat{\delta}_{\text{post}} = -0.000$). The vertical dashed line marks age 24, the outcome-measurement age above which moves can no longer plausibly affect the child's adult income. Estimates at ages 24–25 are displayed but excluded from the hinge fit and from $\hat{\delta}_{\text{sel}}$, because moves recorded in the outcome year or the year after mix partial exposure, disruption, and selection (Section 5.1). Outcome: income rank at age 24; cohorts 1984–1994; moves at ages 10–33 under the unified 1998–2022 residential classification.

year of childhood spent in the destination closes 2% of the destination-vs-origin gap. Extrapolating the fitted profile to a move at birth answers the question that motivates this section. A child raised in the destination from birth accumulates the exposure effect over all 23 pre-outcome years ($23 \times 0.020 = 0.46$ of the gap), and the fitted causal segment sits $\hat{b}_{23} - \hat{\delta}_{\text{sel}} = 0.18$ above the selection level even for the latest childhood moves; adding the two, and netting out selection, she realizes 0.64 of the destination-vs-origin gap in permanent-resident outcomes. Roughly two-thirds (64%) of the upward mobility variation across Spanish municipalities therefore reflects the causal effect of growing up in a place, and the remaining third the sorting of families across places.⁴⁸ That remaining third is what the placebo level measures: at $\hat{\delta}_{\text{sel}} = 0.32$, families who move to better municipalities differ in unobservables worth about a third of the gap they cross.

Where that sorting occurs, however, is not where one might expect. Refitting the same model on the same sample at province resolution gives a larger exposure effect, $\hat{\gamma}_{\text{exp}} = 0.036$, and a selection level of 0.07, less than a quarter of the municipal figure. Sorting on place quality in Spain therefore does not operate through selective long-range migration: families who cross labor markets are close to unselected on the quality of where they land. It operates within them, among families who move a short distance to a better municipality nearby. This qualifies the map of Section 4 in a specific way: cross-municipality gaps within a province embed substantially more sorting than cross-province gaps of the same size, so the finer the geography at which local mobility differences are read, the larger the share of them that reflects who lives there rather than what the place does.⁴⁹

Comparison with the literature. Appendix Table A24 places our estimates, at both geographic resolutions, alongside the closest international evidence based on a movers design. Such comparisons must hold resolution fixed, because exposure effects tend to fall—and selection levels to rise—as places are defined more finely. On exposure, the Spanish estimates match the U.S. evidence scale for scale: 0.036 at province level against 0.037–0.040 for U.S. moves crossing a commuting zone, and 0.020 at municipal level against 0.022 for county moves within one (Chetty and Hendren, 2018a). The Brazilian (0.024) and Dutch (0.021) estimates fall between our two figures, consistent with intermediate geographic units (Britto et al., 2026; Beekers, 2024). Growing up in a better place matters about as much in Spain as elsewhere, at either scale. Selection is where the Spanish estimates look different, and in both directions: the province-scale figure of 0.07 lies below the reported international range and the municipal figure of 0.32 at its upper end. Whether this gradient is peculiar to Spain cannot currently be tested, since no study reports selection at two resolutions—Chetty and Hendren (2018a) estimate county-level selection coefficients but do not report them. Our estimates suggest the margin may be substantial, and that international selection statistics, computed at a single intermediate

⁴⁸The net causal share of a move at age a is $\hat{b}_{23} + \hat{\gamma}_{\text{exp}}(23 - a) - \hat{\delta}_{\text{sel}}$; the extrapolation below age 10 follows Chetty and Hendren (2018a) and Britto et al. (2026). Since $\hat{\delta}_{\text{sel}}$ enters negatively, 0.64 is the most conservative figure across our specifications (Appendix P).

⁴⁹A two-step implementation places the municipal selection level at 0.19 rather than 0.32; the difference is the measurement attenuation of Appendix P.4.

geography, may average over a sharp gradient in how much sorting the observed place gaps contain.⁵⁰

Robustness. Each ingredient of the design—the architecture of the selection controls, the residential histories recovered from the tax records, the age at which the outcome is measured—is a decision that could in principle drive the exposure profile. None of them does. Replacing the hard cells of equation (7) with a Mundlak-style control function leaves the exposure slope at 0.021 and the selection level at 0.29. Excluding the 1999-dated recovered moves, the one dating margin the audit of the residential panel leaves open, gives $\hat{\gamma}_{\text{exp}} = 0.021$; shifting all code-dated moves by one year, the offset between the two location sources, changes no conclusion; and excluding the 2020 move spike, which falls entirely in the placebo region, leaves the estimates unchanged. Refitting the model with income measured at age 28—which adds the 1980–83 cohorts and relocates the mechanical placebo boundary to 28, as the design requires—reproduces the curve, with a per-age slope of 0.019 and the same selection level of 0.32; this last check carries particular weight, since early-career and mature place quality differ within provinces, and the invariance of the estimates to the outcome age is precisely the immunity the same-age comparison is meant to deliver. Appendix P reports all of these checks in full, together with the simulation evidence and the design-based corroborations of Section 5.1.

5.3 Heterogeneity in causal effects by age and type of move


How much better off would a low-income child be if she moved, at different ages, to a municipality with greater opportunities? Figure 13 translates the exposure estimates of Section 5.2 into income terms, for two moves of the kind Spanish families actually make: one between municipalities at the 25th and 75th percentiles of the quality distribution—a gap of 8.2 percentile points of expected adult rank—and a wider P90–P10 move worth 16.5 points. The share of that gap the child causally realizes falls with the age at which she arrives, from 0.44 at age 10 to 0.18 at the last pre-outcome age, and the resulting rank gains are converted into euros at the mature age-28 income distribution.⁵¹ The gains are large and strongly age-dependent. A P90–P10 move at age 10 raises adult income by approximately 19%; the same move at 18 by 12%, and at 23 by 8%. The narrower P75–P25 move is worth about half as much at every age, and the decline is close to linear, mirroring the exposure profile itself. Two features of these magnitudes deserve emphasis. First, they are causal: the sorting component is netted out throughout, so the figure reports what the move contributes, not what movers’ families would have transmitted anywhere. Second, they do not require long-range migration: the gaps are computed across municipalities, most of which lie within the same province. Hence, a substantial

⁵⁰Appendix P.9 discusses the caveats: units, outcome ages, and selection normalizations differ across studies, and refining the geography changes not only how precisely place quality is measured but also the type of move identifying the coefficient. We therefore read the cross-country comparison as indicative and the within-Spain contrast—one model, one sample, two resolutions—as the informative one.

⁵¹The causal share at age a is $\hat{b}_{23} + \hat{\gamma}_{\text{exp}}(23 - a) - \hat{\delta}_{\text{sel}}$, net of selection by construction; one percentile of rank is worth $dI/dp \approx \text{€}273$ on a baseline income of $I(40.3) = \text{€}10,755$. The construction is deliberately conservative on both margins: the gaps use the partially pooled quality surface with population weights, since raw municipal estimates would widen the percentiles through estimation noise alone, and gains are priced at the mature age-28 distribution rather than the compressed age-24 base (Appendix P.8).

part of the opportunity available to a low-income Spanish family sits within reach of its own labor market—in contrast to policy discussions that equate opportunity-improving moves with migration across regions.

Figure 13: Implied Causal Income Gains by Age at Move and Size of the Opportunity Gap



Notes: For each age at move $a \in [10, 23]$, points show the implied causal gain in adult income from moving to a higher-opportunity municipality, expressed as a percentage of the baseline age-28 income of a low-income non-mover ($I(40.3) = \text{€}10,755$). The causal share at age a is $\hat{b}_{23} + \hat{\gamma}_{\text{exp}}(23 - a) - \hat{\delta}_{\text{sel}}$ from the linear-hinge fit of Section 5.2, multiplied by the destination-vs-origin gap in latent municipal quality at parental rank 40.3 (8.2 percentile points for a P75–P25 move, 16.5 for a P90–P10 move; population-weighted percentiles of the partially pooled quality surface) and priced at the local own-rank gradient of the age-28 income distribution ($dI/dp \approx \text{€}273$). Bars are 95% credible intervals computed from joint posterior draws of $(\hat{b}_{23}, \hat{\gamma}_{\text{exp}}, \hat{\delta}_{\text{sel}})$; the quality gaps and the rank-to-income mapping are treated as fixed. Construction details in Appendix P.8.

6 Conclusion

This paper provides the first estimates of intergenerational income and wealth mobility in Spain based on administrative tax records linking parents and children over more than two decades. By constructing measures directly comparable to those developed in the modern literature on intergenerational mobility, we place Spain, for the first time, within the international landscape of economic opportunity. On average, the strength with which children's adult income tracks that of their parents (a rank–rank slope of 0.274) and the probability of moving from the bottom to the top quintile (10.4%) situate Spain in the middle of the international ranking, above the United States and Brazil, close to Italy and France, and well below Canada, Australia and the Scandinavian economies.

Where Spain stands apart is at the top. Persistence at the upper end is exceptionally strong in income and stronger still in wealth: a child born into the top percentile is roughly fifty times as likely to reach it in adulthood as a child from the bottom decile, to our knowledge the highest such ratio among countries with comparable administrative evidence. Our linkage of tax records to the universe of inheritance and gift filings sheds light on one channel sustaining this pattern, and the mechanism is one of access rather than amount. Conditional on receiving, a transfer raises a child's wealth rank by a similar amount at every parental origin, and a uniform shift of that kind does not steepen the relationship between children's ranks and their parents'. What does is that access is deeply unequal, and only on one margin: gift receipt rises sevenfold between the bottom half and the top percentile of the parental distribution, while inheritance receipt is nearly flat. Equalizing gift receipt across parental origins lowers the wealth rank–rank slope by 12.5% and the gradient in the probability of reaching the top percentile of wealth by 26.6%; equalizing inheritance receipt changes neither. Inter vivos giving thus operates as a vehicle for remaining at the top across generations rather than for reaching it from below.

A second contribution is methodological, and concerns the estimation of mobility when local samples are small. Spain, like most countries, has a population distributed unevenly across its territory, so that half of its municipalities contain too few children for conventional estimation and credible local measures of upward mobility have so far been out of reach. The same constraint binds more severely on the causal side, because the childhood-exposure design takes an estimate of each area's quality as an input and regresses movers' outcomes on it as if it were known, so noise in sparse cells passes into the exposure estimates and attenuates them. We address both problems with Bayesian hierarchical models that pool information across nested geographic levels in proportion to how precisely each area is measured. The former delivers the first reliable estimates of intergenerational mobility at the municipality and neighborhood level in Spain—a granularity that opens new questions about the local determinants of economic opportunity, and that gives policymakers a map at the scale on which interventions are actually designed. The second estimates place quality jointly with the exposure coefficients, inside the movers design itself, which removes the attenuation without discarding

the areas and moves the design depends on, and so allows the causal question to be asked at that same scale. Neither estimator is specific to Spain: most countries lack the population-wide records for which the standard procedures were developed, and both could extend evidence on childhood environments to a considerably wider set of settings.

The resulting map shows that economic opportunity in Spain is deeply territorial. Strong historical regional divides—absolute upward mobility varies by more than 25 percentile points along a northeast–southwest gradient—coexist with stark differences across municipalities within the same province and across neighborhoods within major cities, which, unlike in the United States, rank among the country’s engines of upward mobility. The movers design shows that most of this variation is causal: 64% of the gap in upward mobility between Spanish municipalities is attributable to childhood exposure to place, and the annual exposure effect at the labor-market scale is close to the United States benchmark. Selection behaves very differently across scales, however. Families who cross labor markets are barely selected on the quality of where they land, while those who move a short distance to a better municipality nearby are strongly selected, so much of the opportunity available to a low-income family lies within reach of its own local labor market rather than at the end of a long-distance move. Among local characteristics, cross-class economic connectedness is the strongest correlate of upward mobility, ahead of local labor-market conditions and educational attainment. Together, this evidence strengthens the case for place-based policies in low-opportunity regions and, within cities, for policies that reduce residential segregation and build the social capital of children from low-income families ([Chetty et al., 2026b](#)).

Future research will be needed, above all, to study the determinants of the evolution of intergenerational mobility. Because our design measures child income at a fixed age, birth cohort and labor-market entry conditions cannot yet be separated. Fully disentangling them will require longer panels—in which successive cohorts are observed repeatedly over the life cycle—as well as complementary sources of panel data and information on educational attainment and occupation, which would make it possible to distinguish changes in economic opportunity from changes in the composition and career trajectories of successive cohorts. As these data accumulate, it will become possible to establish how economic opportunity is changing over time, and why. Understanding the intensive margin of family transfers, and the mechanisms through which places shape children’s outcomes, are natural companions to that agenda.

References

- Acciari, Paolo, Alberto Polo, and Giovanni L. Violante**, “And Yet It Moves: Intergenerational Mobility in Italy,” *American Economic Journal: Applied Economics*, July 2022, 14 (3), 118–163.
- Adermon, Adrian, Mikael Lindahl, and Daniel Waldenström**, “Intergenerational Wealth Mobility and the Role of Inheritance: Evidence from Multiple Generations,” *The Economic Journal*, July 2018, 128 (612), F482–F513.
- Agencia Estatal de Administración Tributaria**, “Estadística de los Declarantes del IRPF, Ejercicio 1998,” https://www.agenciatributaria.es/AEAT.internet/datosabiertos/catalogo/hacienda/Estadistica_de_los_declarantes_del_IRPF.shtml 1999. Published statistics on personal income tax filers by income bracket, deductions, and type of income. Accessed May 2026.
- Aghion, Philippe, Ufuk Akcigit, Ari Hyytinen, and Otto Toivanen**, “The social origins of inventors,” Technical Report, National Bureau of Economic Research 2017.
- Arellano-Bover, Jaime**, “The Effect of Labor Market Conditions at Entry on Workers’ Long-Term Skills,” *The Review of Economics and Statistics*, September 2022, 104 (5), 1028–1045.
- , “The Effect of Labor Market Conditions at Entry on Workers’ Long-Term Skills,” *The Review of Economics and Statistics*, 2022, 104 (5), 1028–1045.
- Audoly, Richard, Rory McGee, Sergio Ocampo, and Gonzalo Paz-Pardo**, “The Life-Cycle Dynamics of Wealth Mobility,” Staff Report 1097, Federal Reserve Bank of New York 2025. First circulated in April 2024; revised August 2025.
- Badenes-Plá, Nuria and Borja Gambau-Suelves**, “The New “Minimum Vital Income” in Spain: Distributional and Poverty Effects in the Presence and Absence of Regional Minimum Income Schemes,” EUROMOD Working Paper EM22/20, EUROMOD at the Institute for Social and Economic Research 2020.
- Bavel, Jan Van, Sarah Moreels, Bart Van de Putte, and Koen Matthijs**, “Family Size and Intergenerational Social Mobility during the Fertility Transition: Evidence of Resource Dilution from the City of Antwerp in Nineteenth Century Belgium,” *Demographic Research*, 2011, 24, 313–344.
- Beekers, Lieke**, “Sources of Regional Variation in Intergenerational Mobility: Evidence from the Netherlands,” CentER Discussion Paper 2024-015, CentER, Center for Economic Research, Tilburg University July 2024.
- Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen**, “Who becomes an inventor in America? The importance of exposure to innovation,” *The Quarterly Journal of Economics*, 2019, 134 (2), 647–713.

- Bernardi, Fabrizio and Carlos J Gil-Hernández**, “The Social-Origins Gap in Labour Market Outcomes: Compensatory and Boosting Advantages Using a Micro-Class Approach,” *European Sociological Review*, January 2021, 37 (1), 32–48.
- Black, Sandra E and Paul J Devereux**, “Recent developments in intergenerational mobility,” Technical Report, National Bureau of Economic Research 2010.
- Black, Sandra E., Paul J. Devereux, Fanny Landaud, and Kjell G. Salvanes**, “The (Un)Importance of Inheritance,” *Journal of the European Economic Association*, 2025, 23 (3), 1060–1094.
- Blanden, Jo**, “CROSS-COUNTRY RANKINGS IN INTERGENERATIONAL MOBILITY: A COMPARISON OF APPROACHES FROM ECONOMICS AND SOCIOLOGY: CROSS-COUNTRY RANKINGS IN INTERGENERATIONAL MOBILITY,” *Journal of Economic Surveys*, February 2013, 27 (1), 38–73.
- Boserup, Simon Halphen, Wojciech Kopczuk, and Claus Thustrup Kreiner**, “Intergenerational wealth mobility: Evidence from danish wealth records of three generations,” *Univ. of Copenhagen mimeo*, 2013.
- Bourdieu, Pierre**, “The Forms of Capital,” in John G. Richardson, ed., *Handbook of Theory and Research for the Sociology of Education*, Westport, CT: Greenwood Press, 1986, pp. 241–258.
- , *Distinction: A social critique of the judgement of taste*, Harvard university press, 1987.
- Boustan, Leah, Mathias Fjællegaard Jensen, Ran Abramitzky, Elisa Jácome, Alan Manning, Santiago Pérez, Analysia Watley, Adrian Adermon, Jaime Arellano-Bover, Olof Åslund et al.**, “Intergenerational Mobility of Immigrants in 15 Destination Countries,” Technical Report, National Bureau of Economic Research 2025.
- Britto, Diogo G. C., Alexandre Fonseca, Paolo Pinotti, Breno Sampaio, and Lucas Warwar**, “Intergenerational Mobility in the Land of Inequality,” *The Review of Economics and Statistics*, 2026. Forthcoming.
- Budí-Ors, Tomás and Josep Pijoan-Mas**, “Macroeconomic development, rural exodus, and uneven industrialization,” 2022.
- Böhlmark, Anders, Matthew Lindquist, and Anders Bohlmark**, “Life-Cycle Variations in the Association between Current and Lifetime Income: Replication and Extension for Sweden,” *Journal of Labor Economics*, February 2006, 24, 879–900.
- Cervini-Plá, María**, “Exploring the Sources of Earnings Transmission in Spain,” *Hacienda Pública Española / Review of Public Economics*, 2013, 204 (1), 45–66.
- , “Intergenerational Earnings and Income Mobility in Spain,” *Review of Income and Wealth*, 2015, 61 (4), 812–828.

- Chuard, Patrick**, “The Effect of Immigration on Intergenerational Income Mobility of Natives,” *SSRN Electronic Journal*, 2021.
- Chyn, Eric and Lawrence F. Katz**, “Neighborhoods Matter: Assessing the Evidence for Place Effects,” *Journal of Economic Perspectives*, November 2021, 35 (4), 197–222.
- Collado, M. Dolores, Ignacio Ortuño-Ortín, and Andrés Romeu**, “Long-Run Intergenerational Social Mobility and the Distribution of Surnames,” UMUFAE Economics Working Paper 36768, Universidad de Murcia 2013.
- Connolly, Marie, Miles Corak, and Catherine Haeck**, “Intergenerational Mobility between and within Canada and the United States,” *Journal of Labor Economics*, 2019, 37 (S2), S595–S641.
- , —, and —, “Intergenerational Mobility Between and Within Canada and the United States,” *Journal of Labor Economics*, 2019, 37 (S2), S595–S641.
- , —, and —, “Intergenerational Mobility Between and Within Canada and the United States,” p. 47.
- Corak, Miles**, “Income Inequality, Equality of Opportunity, and Intergenerational Mobility,” *Journal of Economic Perspectives*, August 2013, 27 (3), 79–102.
- , “The Canadian Geography of Intergenerational Income Mobility,” *The Economic Journal*, October 2020, 130 (631), 2134–2174.
- and **Andrew Heisz**, “The Intergenerational Earnings and Income Mobility of Canadian Men: Evidence from Longitudinal Income Tax Data,” *Journal of Human Resources*, 1999, 34 (3), 504–533.
- Cuesta, Maite Blázquez**, “Low-Wage Employment and Mobility in Spain,” Working Paper, ZEW 2005.
- Dahl, Molly W and Thomas DeLeire**, “The association between children’s earnings and fathers’ lifetime earnings: estimates using administrative data,” 2008.
- Deutscher, Nathan**, “Place, Peers, and the Teenage Years: Long-Run Neighborhood Effects in Australia,” *American Economic Journal: Applied Economics*, April 2020, 12 (2), 220–249.
- and **Bhashkar Mazumder**, “Intergenerational mobility across Australia and the stability of regional estimates,” *Labour Economics*, October 2020, 66, 101861.
- Eriksen, Jesper and Martin D. Munk**, “The geography of intergenerational mobility — Danish evidence,” *Economics Letters*, April 2020, 189, 109024.
- Escobar, Laura De Pablos and María Gil Izquierdo**, “Intergenerational educational and occupational mobility in Spain: does gender matter?,” *British Journal of Sociology of Education*, July 2016, 37 (5), 721–742.

- Eurostat**, “Living conditions in Europe - income distribution and income inequality,” 2024. Data extracted in November 2024. Accessed on May 11, 2025.
- Fagereng, Andreas, Magne Mogstad, and Marte Rønning**, “Why Do Wealthy Parents Have Wealthy Children?,” *Journal of Political Economy*, March 2021, 129 (3), 703–756.
- Fernández-Kranz, Daniel and Núria Rodríguez-Planas**, “The Perfect Storm: Graduating during a Recession in a Segmented Labor Market,” *ILR Review*, 2018, 71 (2), 492–524.
- Gaubert, Cecile, Patrick Kline, Damian Vergara, and Danny Yagan**, “Place-Based Redistribution,” 2025.
- Gelman, Andrew, Aki Vehtari, Daniel Simpson, Charles C. Margossian, Bob Carpenter, Yuling Yao, Lauren Kennedy, Jonah Gabry, Paul-Christian Bürkner, and Martin Modrák**, “Bayesian Workflow,” November 2020. arXiv:2011.01808.
- Geys, Benny, Friedrich Heinemann, and Alexander Kalb**, “Local Governments in the Wake of Demographic Change: Evidence from German Municipalities,” *FinanzArchiv: Public Finance Analysis*, 2008, 64 (4), 434–457.
- González-Val, Rafael, Pau Insa-Sánchez, Julio Martinez-Galarraga, and Daniel A. Tirado-Fabregat**, “Market Access, the Skill Premium and Human Capital in Spain (1860–1930),” *Journal of Regional Science*, 2026, 66 (1), 201–228.
- Güell, Maia, José V Rodríguez Mora, and Christopher I Telmer**, “The informational content of surnames, the evolution of intergenerational mobility, and assortative mating,” *The Review of Economic Studies*, 2015, 82 (2), 693–735.
- Gurkova, Ekaterina**, “Human Capital Accumulation and the Long-Term Effects of Temporary Sectoral Shocks,” 2025. Job market paper, University of California, Los Angeles.
- Haider, Steven and Gary Solon**, “Life-cycle variation in the association between current and lifetime earnings,” *American Economic Review*, 2006, 96 (4), 1308–1320.
- Heidrich, Stefanie**, “Intergenerational mobility in Sweden: a regional perspective,” *Journal of Population Economics*, October 2017, 30 (4), 1241–1280.
- Hilger, Nathaniel G.**, “Parental Job Loss and Children’s Long-Term Outcomes: Evidence from 7 Million Fathers’ Layoffs,” *American Economic Journal: Applied Economics*, 2016, 8 (3), 247–283.
- Hugalde, Adriana Sánchez**, “Movilidad intergeneracional de ingresos y educativa en España (1980–90),” Working Paper 2004/1, Institut d’Economia de Barcelona 2004.
- Insa-Sánchez, Pau**, “Spatial Inequality of Opportunity in Access to Secondary Education in Nineteenth-Century Spain,” *European Review of Economic History*, 2024, 28 (1), 67–90.

- Instituto Nacional de Estadística**, “Padrón Municipal Continuo: Principales Series de Población desde 1998,” https://www.ine.es/dyngs/INEbase/es/operacion.htm?c=Estadistica_C&cid=1254736177012 1998. Population by municipality, age, sex, and nationality. Accessed May 2026.
- , “Censo de Población y Viviendas 2001,” <https://www.ine.es/censo2001/> 2001. Population and Housing Census. Tables on household structure and family composition. Accessed May 2026.
- Instituto Nacional de Estadística (INE)**, “Cálculo de variaciones del Índice de Precios de Consumo (sistema IPC base 2021),” <https://www.ine.es/varipc/index.do> 2025. Accessed: May 11, 2025.
- Johnston, Drew, Theresa Kuchler, Ayush Kumar, Johannes Stroebel, Manas Kulkarni, and Mike Bailey**, “Social Capital Around the World: Measurement, Determinants, and Effects,” 2026. Forthcoming.
- Kahn, Lisa B.**, “The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy,” *Labour Economics*, 2010, 17 (2), 303–316.
- Kenedi, Gustave and Louis Sirugue**, “Intergenerational Income Mobility in France: A Comparative and Geographic Analysis,” *SSRN Electronic Journal*, 2022.
- Laliberté, Jean-William**, “Long-Term Contextual Effects in Education: Schools and Neighborhoods,” *American Economic Journal: Economic Policy*, May 2021, 13 (2), 336–377.
- McElreath, Richard**, *Statistical rethinking: a Bayesian course with examples in R and Stan* CRC texts in statistical science, 2 ed., Boca Raton: Taylor and Francis, CRC Press, 2020.
- Ministerio de Transportes, Movilidad y Agenda Urbana**, “Metodología de delimitación de las Áreas Urbanas en España 2022,” Technical Report, Ministerio de Transportes, Movilidad y Agenda Urbana, Gobierno de España, Madrid 2022. Accessed: May 11, 2025.
- Moretti, Enrico**, “Human Capital Externalities in Cities,” in J. Vernon Henderson and Jacques-François Thisse, eds., *Handbook of Regional and Urban Economics*, Vol. 4, Elsevier, 2004, pp. 2243–2291.
- , “Workers’ Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions,” *American Economic Review*, 2004, 94 (3), 656–690.
- Nybom, Martin and Jan Stuhler**, “Interpreting Trends in Intergenerational Mobility,” *Journal of Political Economy*, August 2024, 132 (8), 2531–2570.
- , **Toru Kitagawa, and Jan Stuhler**, “Measurement error and rank correlations,” Technical Report, The IFS April 2018.
- OECD**, *A Broken Social Elevator? How to Promote Social Mobility*, Paris: OECD Publishing, 2018.

—, *OECD Economic Surveys: Spain 2018*, Paris: OECD Publishing, 2018.

Oreopoulos, Philip, Marianne Page, and Ann Huff Stevens, “The Intergenerational Effects of Worker Displacement,” *Journal of Labor Economics*, 2008, 26 (3), 455–483.

—, **Till von Wachter, and Andrew Heisz**, “The Short- and Long-Term Career Effects of Graduating in a Recession,” *American Economic Journal: Applied Economics*, 2012, 4 (1), 1–29.

Österberg, Torun, “Intergenerational Income Mobility in Sweden: What Do Tax-Data Show?,” *Review of Income and Wealth*, 2000, 46 (4), 421–436.

Oto-Peralías, Daniel, “Frontiers, warfare and economic geography: The case of Spain,” *Journal of Development Economics*, September 2020, 146, 102511.

— and **Diego Romero-Ávila**, “The economic consequences of the Spanish Reconquest: the long-term effects of Medieval conquest and colonization,” *Journal of Economic Growth*, December 2016, 21 (4), 409–464.

Piketty, Thomas, “Capital & Ideology,” *Harvard University Press (HUP)*, 2020.

Pozo, Diego Del and Lorena Moreno, “Intergenerational Mobility in Ecuador: A First Approach,” 2025.

Putnam, Robert D., *Bowling Alone: The Collapse and Revival of American Community*, New York: Simon & Schuster, 2000.

Sacerdote, Bruce, “Peer Effects with Random Assignment: Results for Dartmouth Roommates,” *The Quarterly Journal of Economics*, 2001, 116 (2), 681–704.

Salvà-Mut, Francesca, Caterina Thomás-Vanrell, and Elena Quintana-Murci, “School-to-work transitions in times of crisis: the case of Spanish youth without qualifications,” *Journal of Youth Studies*, May 2016, 19 (5), 593–611.

Santiago-Caballero, Carlos, “Intergenerational Occupational Mobility in Nineteenth-Century Spain (Valencia), 1841–1870,” *Revista de Historia Económica / Journal of Iberian and Latin American Economic History*, 2021, 39 (2), 219–264.

Schulenberg, René, Celine van Essen, Esther Hamelink, and Arjan Lejour, “A Silver Spoon or a Helping Hand? Intergenerational Wealth Mobility in the Netherlands,” CPB Publication, CPB Netherlands Bureau for Economic Policy Analysis July 2025. Translated version of the Dutch publication *Een steuntje in de rug: vermogensmobiliteit van ouder op kind*.

Solon, Gary, “Intergenerational mobility in the labor market,” in “Handbook of labor economics,” Vol. 3, Elsevier, 1999, pp. 1761–1800.

—, “Cross-country differences in intergenerational earnings mobility,” *Journal of Economic Perspectives*, 2002, 16 (3), 59–66.

Soria-Espin, Javier, “Intergenerational Mobility, Gender Differences and the Role of Out-Migration: New Evidence from Spain,” *SSRN Electronic Journal*, 2022.

Tapia, Francisco Beltran and Santiago de Miguel Salanova, “Class, literacy and social mobility: Madrid, 1880–1905,” *The History of the Family*, January 2021, 26 (1), 149–172.

Tapia, Francisco J. Beltrán, Alfonso Díez Minguela, Julio Martinez-Galarraga, and Daniel A. Tirado-Fabregat, “The Roots of Land Inequality in Spain,” IFCS - Working Papers in Economic History WH 31728, Universidad Carlos III de Madrid, Instituto Figuerola 2021.

Tirado, Daniel A., Alfonso Díez-Minguela, and Julio Martinez-Galarraga, “Regional inequality and economic development in Spain, 1860–2010,” *Journal of Historical Geography*, October 2016, 54, 87–98.

van der Weide, Roy, Christoph Lakner, Daniel Gerszon Mahler, Ambar Narayan, and Rakesh Ramasubbaiah, “Intergenerational Mobility Around the World,” *SSRN Electronic Journal*, 2021.

APPENDIX


A Descriptive Statistics

Table A1: Summary Statistics: Baseline Sample (Cohorts 1980–86)

	Mean (1)	Std. Dev. (2)	Median (3)
<i>Parents:</i>			
Household Income (1998–2000 avg.)	29,247	37,212	21,435
Parent 1 Individual Income	22,990	29,789	18,027
Fraction Married	90.4%	29.4%	
Parent 1 Year of Birth	1,953.3	5.9	1,954.0
Household Gross Wealth (2022)	216,480	2,411,743	98,333
<i>Children:</i>			
Individual Income (2021–2022 avg.)	29,231	44,751	23,286
Fraction Single	85.6%	35.1%	
Year of Birth	1,983.2	2.0	1,983.0
Individual Gross Wealth (2022)	71,971	800,873	45,488
Number of Children	1,580,994		


Notes: Summary statistics for the baseline sample (1980–86 birth cohorts). Parental household income is the mean of total pre-tax household income over 1998–2000; children’s individual income is the mean of gross taxable income over 2021–2022. All monetary values are nominal € and not comparable across generations without CPI adjustment (Spain CPI 2000→2022 ≈ 1.45). Gross wealth is total declared assets in 2022: children at ages 36–42, parents at ages ≈ 68 –72. Sample excludes the Basque Country and Navarre (foral regimes). See Section 3 for details on sample construction and income definitions. See Appendix H for robustness checks with varying sample definitions

Figure A1: Income Profiles Across the Parental and Child Distributions



Notes: This figure plots average income in euros at each percentile of the parent and child income distributions. Income is measured at the household level for parents (averaged over 1998–2000) and at the individual level for children in 2022. Percentiles are computed within cohort-year and gender cells. The steep upward slope at the top of the distribution in both panels illustrates strong income concentration and motivates the use of rank-based approaches to mitigate the influence of extreme values.


Figure A2: Child and Parent Income Distribution



Notes: This figure shows kernel density plots of adjusted household income for parents (left panel) and adjusted individual income for children (right panel), both expressed in euros and trimmed at the 99th percentile. Excluding the top 1% reduces the influence of extreme values and allows for a clearer view of the mass of the income distribution.

B National Results

Figure A3: National Quintile Transition Matrix




Notes: Each cell shows the probability that a child born into a given parental income quintile (column) reaches a given child income quintile (row) in adulthood. Darker shading indicates higher probability. A strong diagonal indicates low mobility. Income definitions follow Section 2.3.

Figure A4: National Decile Transition Matrix




Notes: This figure presents the national-level decile transition matrix in the form of a heatmap, illustrating the probability of children transitioning across income quantiles based on their parents' household income decile. Each row represents the children's individual income decile in adulthood, while each column corresponds to the parental household income decile. The percentages inside each cell indicate the probability that a child born into a specific parental income decile reaches a given decile in adulthood.

Figure A5: Minimum Basic Income (IMV) Receipt by Parental Income



Notes: The figure plots the share of children receiving the Minimum Basic Income (*Ingreso Mínimo Vital*, IMV) in 2022, by parental household income ventile and child gender. The IMV is a non-contributory benefit introduced in 2020 that targets households in situations of poverty and social exclusion. Income definitions follow Section 2.3.

Figure A6: Mean Child Income Rank vs. Parental Income Rank at the National Level by Gender



Notes: This figure shows a nonparametric binned scatter plots depict the mean child income percentile by parental household income percentile separately for male (orange) and female (purple) children in the baseline sample. Lines show OLS fits: the estimated rank–rank slope is 0.260 (SE: 0.001) for males and 0.285 (SE: 0.001) for females.

Table A2: Intergenerational Mobility at the National Level in Spain by Subgroups


	Dependent variable:											
	Child Income Percentile											
	Baseline (1)	Baseline (Fixed Parental Age) (2)	Male (3)	Female (4)	2nd Gen Migrant (5)	1st Gen Migrant (6)	Child Income Percentile First Born (7)	Not First Born (8)	Married Parents (9)	Unmarried Parents (10)	Divorced Parents (11)	Extended Sample (1980-1990) (12)
Parental Income Rank	0.274*** (0.001)	0.273*** (0.001)	0.260*** (0.001)	0.285*** (0.001)	0.276*** (0.008)	0.251*** (0.008)	0.277*** (0.004)	0.279*** (0.004)	0.280*** (0.001)	0.242*** (0.008)	0.247*** (0.006)	0.277*** (0.001)
Constant	36.650*** (0.044)	33.492*** (0.124)	41.336*** (0.062)	32.152*** (0.062)	36.049*** (0.551)	36.743*** (0.411)	36.552*** (0.208)	35.549*** (0.255)	36.261*** (0.048)	37.433*** (0.366)	38.664*** (0.272)	36.527*** (0.035)
AUM	43.505	40.327	47.83	39.289	42.959	43.017	43.471	42.524	43.258	43.472	44.828	43.443
PQ5Q1	10.394	10.394	8.22	12.71	9.755	12.356	10.178	10.078	9.627	12.18	12.208	10.204
Observations	1,580,994	1,544,194	788,463	792,519	14,085	13,219	68,801	45,620	1,429,504	17,322	33,169	2,495,136
R ²	0.075	0.076	0.070	0.082	0.072	0.064	0.078	0.081	0.078	0.052	0.050	0.077
Adjusted R ²	0.075	0.076	0.070	0.082	0.072	0.064	0.078	0.081	0.078	0.052	0.050	0.077

Note:

*p<0.1; **p<0.05; ***p<0.01


Each cell in this table reports the coefficient from a univariate OLS regression of child income percentile on parental income percentile, with standard errors in parentheses. Column (1) reports baseline results using the national sample of children born between 1980 and 1986, where parental income is averaged over 1998–2000 and child income is averaged over 2020–2022 at ages 36–42. Column (2) adds a control for parental age at the time of the child's birth. Columns (3) and (4) restrict the sample to male and female children, respectively. Columns (5) and (6) examine second-generation migrants (children born in Spain with foreign-born fathers) and first-generation migrants (children born abroad and raised in Spain). Columns (7) and (8) divide the sample by birth order, identifying first-born and not-first-born children. First-born children are identified as the eldest child within households where multiple children are present, and the age gap between siblings is small (indicating closely spaced births). In households with only one child or where the age gap between siblings is unusually large, the first-born status is not assigned to ensure accuracy. Columns (9), (10), and (11) separate children by parental marital status: married, unmarried, or divorced. Column (12) extends the baseline sample to include children born between 1980 and 1990. AUM (Absolute Upward Mobility) measures the expected rank in the income distribution of children whose parents are at the 25th percentile. PQ5Q1 is the probability of moving from the bottom income quintile (Q1) to the top income quintile (Q5).

Figure A7: Gender Composition Across Child Income Percentiles




Notes: The figure plots, for each child income percentile, the share of that percentile comprised of male (orange) and female (purple) children, expressed as a percentage of the national average gender share (dashed line at 100%). Values above 100% indicate over-representation relative to the overall gender distribution; values below 100% indicate under-representation. For example, at the lowest income percentiles (0–20), females exceed their national average share by up to 25pp, whereas males are underrepresented in the bottom quintile. Above the 50th percentile, male representation gradually increases—peaking at over 140% in the top percentile—while female representation falls below 80%, illustrating a pronounced gender gap in access to top-income positions.

Figure A8: Quintile Transition Matrix by Gender




Notes: This figure shows the quintile transition matrix by gender, illustrating the probability of children transitioning across income quintiles based on their parents' household income quintiles. The left panel represents outcomes for men, while the right panel depicts outcomes for women. Each bar on the x-axis corresponds to a parental quintile, with the stacked segments representing the distribution of children's individual income across destination quintiles. The percentages shown within each segment indicate the likelihood of a child moving from the corresponding parental quintile to a specific destination quintile. Child individual income and parental household income are defined as in the baseline national-level figures. Both metrics follow the methodology and definitions outlined in the baseline national-level analyses.

Figure A9: Intergenerational Mobility in International Perspective



(a) Rags-to-Riches Probability (Bottom Quintile to Top Quintile)



(b) Rank-Rank Slope


Notes: Panel (a) shows the probability that a child born into the bottom income quintile reaches the top quintile. Panel (b) displays the rank-rank slope from a regression of child income rank on parental income rank. All estimates come from studies using administrative tax data and comparable birth cohorts. See Appendix Table A3 for sources and definitions.

Table A3: Rank-Rank Slope in International Perspective

Country	RRS	# obs.	Data	Child Income Definition ¹	Child Cohort	Child Age or Year at Income Measurement	Parent Age or Year at Income Measurement	Source
Australia	0.215	1,025,800	Federal income tax returns	Average total pretax family income	1978-1982	2011-2015	1991-2001	Deutscher and Mazumder (2020, Table 2)
Brazil	0.546	1,304,586	Federal Revenue of Brazil	Pretax and informal income	1988-1990	2012-2014	2010-2016	Britto et al. (2020, Figure 3)
Canada	0.242	N/A	Canadian census data	Average total pretax individual income	1980-1982	2011-2012	1996-2000	Connolly and Corak (2020, Table 7)
China	0.443	22,313	China Family Panel Studies (CFPS)	Average total pretax individual income	1987-1988	2010-2016	2010-2016	Yi Fan et al (2019, Table 2)
Denmark	0.253	≈410,000	Danish register data	Average total pretax individual income	1973-1979	2010-2012	when child between 7-15	Landersø and Heckman (2017, Table A6)
Ecuador	0.272	3,819,846	Statistical Registry of Social Security (REESS)	Labor (formal and informal) income	1987-1998	2018-2023	2006-2011	Del Pozo and Moreno (2025, p.3)
France	0.303	64,571	Permanent Demographic Sample	Parents: (predicted) household wage; Children: average total pretax household income	1972-1981	2010-2016 (between 35-45)	1978-1980	Kenedi and Sirugue (2023, p.3)
Germany	0.245	1,128	German Socio-Economic Panel	Average total pretax individual income	1957-1976	between 8-29	1984-1986	Bratberg et al. (2017, Table 3)
Italy	0.22	1,719,483	Electronic Database of Personal Income Returns	Total pretax individual income	1979-1983	2016-2017	1998-2000	Acciari et al (2022, Figure 1)
Spain	0.274	1,492,107	Spanish Opportunity Atlas	Total pretax individual income	1980-1986	2021-2022	1998-2000	Soria and Medina (2025, Figure 1)
Sweden	0.197	927,008	Swedish Social Insurance Agency	Sum of taxable income	1968-1976	32-34	16-40	Heidrich (2017, Table 2)
Sweden	0.215	252,745	Swedish registry administrative data	Average pretax household income	1957-1964	1920-1950	1978-1990	Bratberg et al. (2017, p.75)
Switzerland	0.14	923,262	Social Security Earnings Records	Average total pretax individual income	1967-1984	30-33	when child between 15-20	Chuard-Keller and Grassi (2021, Fig 1)
US	0.341	9,867,736	Federal income tax records	Average total pretax family income	1980-1982	2011-2012	1996-2000	Chetty et al (2014, Table 1)

¹ The parent income definition is always at the family level.


Figure A10: Conditional Probability of Reaching the Top 1% by Parental Percentile Across Countries



Notes: Each point shows the probability that a child reaches the top 1% of adult income, conditional on parental income percentile. The dashed line at 1% denotes perfect equality. Country estimates from Britto et al. (2026) (Brazil), Connolly et al. (n.d.) (Canada), Acciari et al. (2022) (Italy), and Chetty et al. (2014b) (United States).


C Intergenerational Wealth Mobility

Figure A11: Parental Wealth vs. Parental Income as Predictors of Child Wealth




Notes: Mean child gross wealth percentile plotted against the parental *income* percentile (red) and the parental *gross wealth* percentile (blue), with OLS fits. The non-monotonicity at the bottom of the parent-wealth distribution—children of the seemingly “poorest-in-wealth” parents hold above-median wealth—reflects transfer-depleted donors: parents who have already given assets away appear artificially poor in the 2021–2022 snapshot. This mechanical endogeneity, together with the absence of any parental wealth measure for the late 1990s, is why parents are ranked by their pre-determined 1998–2000 income throughout the analysis. See Tables A5 and A6.

Figure A12: Mean Child Income and Gross Wealth Rank vs. Parental Income Rank




Notes: Panel (A) plots the mean child income percentile and Panel (B) the mean child gross wealth percentile against the parental household income percentile, together with the OLS fit. The rank–rank slope is 0.274 for income and 0.226 for gross wealth (standard errors below 0.001); each panel also reports the slope restricted to percentiles 10–90 and absolute upward mobility, the predicted child rank at the 25th parental percentile. Child outcomes are 2021–2022 averages (ages 35–42); parental income is the 1998–2000 household average. Ranks are recomputed within the baseline estimation sample ($N = 1,501,948$): within birth cohort for children, pooled at the household level for parents.

Figure A13: Quintile Transition Matrices: Child Income and Gross Wealth




Notes: Each cell reports the probability that a child born to parents in a given quintile of the parental household income distribution reaches each quintile of the child income distribution (Panel A) or the child gross wealth distribution (Panel B); rows of each column sum to one within parental quintile. Top-quintile persistence is 36.6% for income and 35.4% for wealth; the bottom-to-top probability is 10.5% and 11.7%, respectively. The apparent similarity partly reflects that parental income is an imperfect proxy for the parental resources relevant to wealth (see Figure A14). Sample as in Figure A12.

Figure A14: Joint Distribution of Income and Gross Wealth Ranks: Children and Parents




Notes: Each cell reports the share of individuals in a given income decile (columns) belonging to each gross wealth decile (rows); columns sum to one, and both panels share a common colour scale. Panel (A) refers to children (rank correlation 0.53), Panel (B) to parents (0.55). Concentration along the diagonal indicates positive but far-from-perfect association: the income and wealth distributions identify overlapping yet distinct groups, at both the top and the bottom. Parental gross wealth is measured in 2021–2022 and is used here for descriptive purposes only.

Figure A15: Probability of Reaching the Top 1%: Income vs. Gross Wealth




Notes: Probability that a child reaches the top 1% of the child income (red) and child gross wealth (blue) distribution, by parental household income percentile; the dashed line at 1% denotes perfect equality of opportunity. The implied TTRPR—the probability at the top parental percentile divided by the mean probability in the bottom parental decile—is 49.2 for income and 62.8 for gross wealth on this sample. Both series are computed on the baseline estimation sample and are internally comparable; the income figure therefore differs slightly from that reported in the national income analysis.

Figure A16: Transfer-Associated Rank Gaps: Wealth vs. Income




Notes: For each five-percentile bin of parental income, the difference in mean child rank between each transfer category and no-transfer children, for child gross wealth (left) and child income (right). Averaged over the parental distribution, the wealth gaps are +19.2 (only gift), +8.1 (only inheritance) and +23.5 (both) percentiles; the corresponding income gaps are +5.3, -0.2 and +5.6. The wealth gaps are close to constant in parental origin over most of the distribution—the gift gap stands at 17.4 ranks at the tenth parental percentile and 21.3 at the ninetieth—and widen within the top decile, the intensive margin showing through in data that record only the extensive margin of receipt. Transfers are therefore associated with a level shift in asset position rather than with differential earnings capacity. Descriptive associations; see Figure A17 for regression versions.

Figure A17: Transfer Receipt and Child Rank: Regression Estimates




Notes: Coefficients on transfer-receipt indicators with 95% confidence intervals, for child gross wealth rank (left) and child income rank (right). Inheritance estimates are identified off children with no observed gift; gift estimates off children with no observed inheritance; the omitted group in both cases is children with no observed transfer. Specifications control for parental income rank (1), add cohort fixed effects and gender (2), and add childhood-region fixed effects (3). Inheritance receipt shifts the wealth rank (about +8.4 with full controls) but not the income rank (+0.2), the pattern expected when receipt is driven by parental mortality rather than by selection on the child's earnings capacity; gift receipt shifts both (+20.0 and +6.0), indicating positively selected recipients. Parental mortality is itself correlated with socioeconomic status, so this comparison bounds rather than eliminates selection. See Table ??.

Figure A18: Gift Timing and Child Outcomes




Notes: Rank differences between gift-only recipients and no-transfer children, by age at first gift, controlling for parental income rank, cohort fixed effects and gender. The analysis is restricted to the age window observable for every baseline cohort given the 2002–2022 transfer registry; outside that window, age at gift partly identifies birth cohort. Wealth gaps are essentially flat in the age at receipt, as are income gaps: the transfer-associated advantage is a level shift rather than an accumulating return, consistent with the transferred asset itself constituting the gap rather than with differential accumulation on top of it.

Figure A19: Age at Observed Transfers, by Type




Notes: Distribution of age at the first observed transfer for gift-only and inheritance-only recipients, over the 2002–2022 registry window. Most observed first transfers predate the 2021–2022 wealth measurement window, so the transfer-associated wealth differences of Figure A16 reflect assets received well before wealth is measured rather than contemporaneous receipt. Because the registry begins in 2002 and cohorts differ in age at that date, the observable age range varies by cohort; see the notes to Figure A18.

Figure A20: Access to Family Transfers on the Binary Margins



Notes: Share of children receiving at least one gift and at least one inheritance, by parental household income percentile. This is the counterpart of Figure ?? on the two binary margins rather than the four mutually exclusive categories. The distinction matters for statements about access: because the categories of Figure ?? are exclusive, high-origin children who receive both transfer types are counted in the combined category rather than the inheritance-only one, which makes inheritance-only receipt appear to decline with parental income. On the binary margins, gift receipt rises from 6.2% of children of bottom-half parents to 41.8% of children of top-percentile parents, a nearly sevenfold gradient, while inheritance receipt rises only from 17.6% to 22.5%. Gifts operate on a choice margin available to parents with assets to spare; inheritances on a mortality margin that reaches all families. All decompositions in Table ?? use these margins. Sample and definitions as in Figure 5.

Figure A21: The Parental Income Gradient in Child Wealth Under Equal Access to Transfers



Notes: The observed profile of mean child gross wealth rank against parental income rank, together with three counterfactual profiles. Writing that profile as a weighted average of the category-specific profiles, $E[y | p] = \sum_k \pi_k(p) \mu_k(p)$, where $\pi_k(p)$ is the share of children of parental rank p in transfer category k and $\mu_k(p)$ their mean child rank, each counterfactual replaces $\pi_k(p)$ by its population value $\bar{\pi}_k$ while leaving the conditional means $\mu_k(p)$ as observed. Each line therefore shows the gradient that would obtain if every parental origin faced the same probability of receiving the stated type of transfer, holding fixed what recipients do with what they receive. Equalizing gift access lowers the rank–rank slope from 0.226 to 0.198; equalizing inheritance access leaves it unchanged; equalizing both delivers essentially the gift result. The in-figure box reports each slope and its change relative to the observed one. Because only the extensive margin of receipt is observed, transfer amounts are unavailable, and most bequests are not yet realized, these are lower bounds on the role of transfers in the gradient. Sample and definitions as in Figure 5.

Table A4: Sample and Measurement: Wealth Analysis

	Value
Analysis sample	Baseline estimation sample (1980–86)
Number of children	1,501,948
Gross wealth observed	100.0%
Wealth measurement window	2021–2022 (ages 35–42)
Transfer observation window	2005–2022
Median age at first gift	34 (IQR 29–37)
Median age at first inheritance	33 (IQR 28–37)
First transfers occurring before 2021	81.3%

Notes: The wealth analysis uses the same baseline estimation sample as the income analysis: gross wealth is observed from administrative property and financial registries rather than wealth-tax filings, so its coverage does not select on wealth-tax liability and no children are lost. Transfers are the stock of gifts and inheritances recorded in the transfer-tax registry over the observation window. Wealth is a snapshot near the end of that window; because most transfers predate it and most parents remain alive, realized bequests are only partially observed and all transfer-related results are a lower bound.

Table A5: Rank–Rank Slopes: Parental Income vs. Wealth as Predictors

Parental ranking variable	Child outcome	
	Income	Gross wealth
Parent household income (1998–2000)	0.274	0.226
Parent household gross wealth (2021–2022)	0.261	0.313


Notes: Each cell reports the rank–rank slope of the column child outcome on the row parental ranking variable (baseline estimation sample, $N = 1,501,948$; all standard errors < 0.001). Each outcome is most strongly predicted by the same parental dimension, and wealth is the most persistently transmitted of the two. Parental wealth is measured after many *inter vivos* transfers, which mechanically depresses the measured wealth of donor parents; the parent-wealth slopes are therefore attenuated and constitute a lower bound. This endogeneity—visible as non-monotonicity at the bottom of the parent-wealth distribution (Appendix Figure A11)—together with the absence of any parental wealth measure for 1998–2000, is why parents are ranked by income throughout the analysis.

Table A6: Parental Ranking Variables and Child Outcomes: Robustness to the Timing of Parental Wealth

Parental ranking variable	Child outcome	
	Income	Gross wealth
Parent household income (1998–2000)	0.274	0.226
Parent household gross wealth (2021–2022)	0.261	0.313
Parent household gross wealth (2016)	0.263	0.338


Notes: Rank–rank slopes on the baseline estimation sample; all standard errors are below 0.001. Parental wealth measured in 2016 precedes gifts made between 2016 and 2022 and is therefore less depleted by transfers than the later measure, though gifts made between 2005 and 2015 still reduce it and a single year carries more transitory variation than a two-year mean. That the parent-wealth→child-wealth slope *rises* with the earlier measure, while the parent-wealth→child-income slope is unchanged, is the pattern predicted by transfer-driven attenuation and the opposite of what classical measurement error would produce.

Figure A22: Transfer Composition Across the Child Gross Wealth Distribution




Notes: Each vertical bar decomposes the children at a given percentile of the child gross wealth distribution into the four mutually exclusive transfer categories of Figure 5. The share with any observed transfer rises from about 11% at the bottom of the distribution to 68% in the top percentile. Inheritance-only receipt is broad-based, fluctuating between roughly 10% and 17% through most of the distribution; gift-only receipt is rare below the median and expands sharply in the top decile; and the combination of both is concentrated almost exclusively at the very top. This figure conditions on the child's own wealth position, which is itself raised by receipt, so the gradient it displays is in part definitional; the access evidence is Figures ?? and A20, which condition on parental origin instead. The counterpart across the child income distribution is Figure A23. Sample and definitions as in Figure 5.

Figure A23: Transfer Composition Across the Child Income Distribution




Notes: Counterpart of Figure A22 for the child *income* distribution: each vertical bar decomposes the children at a given income percentile into the four mutually exclusive transfer categories. The gradient is far flatter than across the wealth distribution, and the sharp expansion of gift receipt at the top is largely absent. Transfer receipt sorts children by asset position much more strongly than by earnings, the same channel specificity visible in Figure 5. Sample and definitions as in Figure 5.

Figure A24: Transfer Composition Within the Top 1% of Child Gross Wealth




Notes: Transfer composition at each position within the top 1% of the child gross wealth distribution, from the bottom of the top percentile (left) to the top 0.01% (right). The share with any observed transfer continues to rise within the top percentile, approaching 80% (dashed line) toward its upper end, and the combination of gift and inheritance accounts for a growing share of that total. Sample and definitions as in Figure 5.

Figure A25: Origin of the Top 1% in Child Gross Wealth, by Parental Income Rank and Transfer Status




Notes: The top panel plots, for children in the top 1% of the child gross wealth distribution, the share originating from each parental income percentile, overall and separately by transfer category; the dashed line at 1% is the benchmark under independence. Origin selection is severe in every category—the profiles are flat and near zero through most of the parental distribution and rise steeply in the top percentiles—but the categories are ordered systematically throughout the upper tail. The bottom panel reports, for each category, the share coming from the top parental income percentile: 28.1% of gift-only recipients, 25.6% of gift-and-inheritance recipients, 16.9% of all top-1% children, 10.8% of inheritance-only recipients and 7.6% of no-transfer children. Children with no observed transfer are the least concentrated of the four, though even they remain overrepresented at the top parental percentile roughly eightfold. Because the remaining 99 parental percentiles are 99 times as populous as the top one, composition shares of this kind should be read alongside the conditional probabilities of Figure ???. Transfer status is post-determined; shares describe composition, not effects. Sample and definitions as in Figure 5.

Figure A26: Transfer Composition of Top-1% Income Children, by Parental Origin



Notes: Counterpart of Figure 8 for children in the top 1% of child *income*; origin groups and labelling as there. Transfer receipt is markedly less prevalent among top earners than among top wealth holders at every origin: 30% of top-1%-income children of bottom-half parents have received a transfer against 53% of their top-wealth counterparts, and 67% of those of top-percentile parents against 85%. Bottom-half origins also contribute a larger share of the income elite (16.4%) than of the wealth elite (15.6%) despite the lower receipt rate. Both patterns are consistent with transfers operating on asset position rather than on earnings capacity, as in Figure 5. Transfer status is post-determined; shares describe composition, not effects.


Figure A27: Reaching the Top 1% of Gross Wealth Under Equal Access to Transfers



Notes: The probability of reaching the top 1% of the child gross wealth distribution by parental income percentile, observed and under the equal-access counterfactuals of Figure A21, which hold receipt rates fixed at their population level across all parental origins while leaving conditional outcomes as observed. Panel (A) covers the full parental distribution; Panel (B) restricts to its top decile, where essentially all of the variation lies, with larger markers. The dashed line at 1% is the benchmark under independence. Equalizing access to gifts lowers the gradient in this probability by roughly a quarter, against an eighth for the average wealth rank in Figure A21; equalizing access to inheritances leaves it close to unchanged. Transfers therefore bear on mobility disproportionately at the top of the wealth distribution, and do so through the gift margin. As throughout, receipt is post-determined and only the extensive margin is observed, so these are accounting counterfactuals rather than estimates of the effect of transfers.


D Robustness of National Results

Figure A29: Lifecycle Bias: Intergenerational Income Correlation by Number of Cohorts Included and Age Range




Notes: Each line shows the estimated parent-child rank-rank slope for successive birth-cohort groups (1980–1990), where child income is measured over varying age windows ending in 2022. Shades of blue correspond to the number of years of income used (1 to 4 years). The horizontal dashed line marks our baseline cohort sample (1980–1986 measured at ages 36–42) with slope 0.274. The tight clustering of slope estimates around the baseline—regardless of cohort range or income-measurement span—confirms the robustness of our core sample choice and suggests minimal bias from alternative age or cohort selections.

Figure A28: Lifecycle Bias: Intergenerational Income Correlation by Age at Which Child's Income Is Measured



Notes: Each line plots the estimated parent-child rank-rank slope for a given birth cohort as the age at which the child's income is measured varies from 21 to 42. Child income is averaged over a two-year window centered on the age shown, while parental income is fixed as the three-year average from 1998–2000. The dashed horizontal line marks the baseline slope of 0.274 from our main sample (children aged 36–42 for cohorts 1980–1986). The convergence of all cohort-specific slopes toward the baseline line after age 36 demonstrates that measuring child income at these older ages minimizes life-cycle bias in mobility estimates.


Figure A30: Attenuation Bias: Rank-Rank Slopes by Number of Years Used to Measure Parent Income



Notes: This figure plots the estimated parent-child rank-rank slope for our baseline cohort (1980–1986, children aged 36–42) as we vary the number of years of parental income used to compute the mean (from 1 to 6 years). Using only one year of parental income yields a slope of 0.261, whereas averaging over six years increases the slope to 0.280.


E Evolution of Intergenerational Mobility Across Cohorts

Figure A31: Rank–Rank Slope by Birth Cohort, Child Income Measured at Age 31




Notes: The figure plots the estimated rank–rank slope for each birth cohort from 1980 to 1990, with child income averaged over a two-year window centered on age 31 and parental income fixed at the 1998–2000 household average. The upper axis reports the corresponding calendar year of income measurement, which increases mechanically with birth cohort: age, birth cohort, and measurement year are collinear by construction. The vertical dotted line marks the 1983 cohort, which reached age 25—our approximation of labor-market entry—in 2008. Confidence intervals are omitted for legibility; with more than 200,000 observations per cohort, the slopes are precisely estimated and their sampling variability is small relative to the cross-cohort differences shown.

Figure A32: Cohort Profile of the Rank–Rank Slope at Different Measurement Ages



Notes: Each line plots the cohort-specific rank–rank slope with child income measured at a different fixed age (28 through 31). Levels shift downward at younger ages because rank–rank slopes are attenuated earlier in the life cycle (Figure A28); the object of interest here is the *shape* of the cohort profile, not its level. The U-shape is present at every measurement age, indicating that it is not an artifact of fixing measurement at age 31. Fixing the age holds life-cycle position constant across cohorts but does not break the collinearity between cohort and measurement year, and therefore does not identify the channel generating the pattern.

Figure A33: Mean Child Income Rank of Children from p25 and p75 Families by Birth Cohort



Notes: The figure decomposes the movement in the cohort-specific rank–rank slope into the mean adult income rank of children from families at the 25th and 75th percentiles of the parental income distribution, computed from the fitted cohort-specific regression lines. Child income is measured at age 31. The rise in persistence for the later cohorts reflects both a decline in the mean rank of children from p25 families and an increase for children from p75 families.


F The Geography of Intergenerational Mobility in Spain

Figure A34: Map of Provinces in Spain, Coloured by Region




Notes: This figure presents a map of provinces in Spain colored by regions (Comunidad Autónoma). Regions in Spain have their own parliaments and a high degree of independence in important areas of public policy as health, education or transports.

Figure A35: Map of Municipalities in Spain, Colored by Province




Notes: This figure presents a map of municipalities in Spain colored by provinces.

Figure A36: Bayesian Hierarchical vs. Non-Parametric Estimates of Rags-to-Riches Probability Across Provinces




Notes: This figure compares two estimation approaches for the probability that a child from the bottom parental income quintile reaches the top child income quintile ($P(Q_5|Q_1)$) across Spanish provinces. Blue points are estimates from the Bayesian hierarchical ordered logistic model described in Section 4.1 and Appendix A.3; black points are traditional non-parametric estimates. Horizontal lines denote 95% confidence intervals. The hierarchical model yields substantially more precise estimates, particularly for provinces with smaller sample sizes, by borrowing strength across units. Geographic patterns closely mirror those observed for absolute upward mobility (μ_{25}).

Figure A37: Geographic Variation in Absolute Upward Mobility (μ_{25}) by Urban Area




Notes: This figure displays estimates of absolute upward mobility, defined as the mean income percentile in adulthood for children whose parents were at the 25th percentile of the national parental income distribution. Estimates are shown for each Urban Area. Darker shades of blue indicate higher upward mobility (children reach higher income percentiles on average), while darker shades of red indicate lower upward mobility. Income definitions and sample construction follow the baseline methodology outlined in Section 3.

Figure A38: Geographic Variation in Absolute Upward Mobility (μ_{25}) by Province



Notes: This figure displays estimates of absolute upward mobility, defined as the mean income percentile in adulthood for children whose parents were at the 25th percentile of the national parental income distribution. Estimates are shown for each province. Darker shades of blue indicate higher upward mobility (children reach higher income percentiles on average), while darker shades of red indicate lower upward mobility. Income definitions and sample construction follow the baseline methodology outlined in Section 3.


Figure A39: Geographic Variation in Absolute Upward Mobility (μ_{25}) by Autonomous Community



Notes: Bayesian hierarchical estimates of absolute upward mobility (μ_{25}) by autonomous community. Blue shades indicate higher upward mobility; red shades indicate lower mobility. Income definitions and sample construction follow Section 3.

Figure A40: Geographic Variation in Absolute Upward Mobility (μ_{25}) at the ZIP Code Level

Absolute Upward Mobility in Spain at the ZIP code Level




Notes: This map displays Bayesian hierarchical estimates of absolute upward mobility (μ_{25}) at the ZIP code level across Spain. μ_{25} is the expected mean child income percentile for children from families at the 25th percentile of the parental income distribution. The color gradient ranges from blue (higher upward mobility) to red (lower upward mobility). Outside urban areas, most municipalities correspond to a single ZIP code, so estimates at the two levels are often identical. Within cities, ZIP-code-level variation reveals finer neighborhood-level differences in mobility prospects.

Table A7: Intergenerational Mobility in the Largest 50 Urban Areas of Spain


Urban Area Name	AUM	Population	P(Q3 Q1)	P(Q5 Q1)	RRS
Lleida	51.4	169620	0.22	0.19	0.20
Barcelona	50.21	5156625	0.22	0.18	0.22
Soria	49.89	39398	0.21	0.13	0.21
Ávila	49.28	57744	0.21	0.13	0.21
Manresa	48.96	104907	0.22	0.16	0.27
Madrid	48.85	6196700	0.26	0.17	0.23
Tarragona - Reus	48.68	386600	0.24	0.18	0.25
Guadalajara	48.50	161683	0.23	0.11	0.25
Girona	48.45	161582	0.26	0.13	0.26
Burgos	47.94	177935	0.23	0.15	0.21
Cuenca	47.29	54690	0.22	0.14	0.20
Zaragoza	47.28	754084	0.24	0.12	0.24
Segovia	47.19	69273	0.28	0.13	0.22
Teruel	47.15	35890	0.25	0.13	0.21
Sant Feliú de Guixols	46.49	61787	0.23	0.13	0.20
Albacete	46.31	173329	0.24	0.11	0.23
Logroño	46.25	176883	0.24	0.11	0.23
Melilla	46.23	86487	0.27	0.14	0.18
León	46.02	192082	0.26	0.13	0.21
Puertollano	45.84	47035	0.26	0.15	0.17
Aranjuez	45.74	59607	0.26	0.14	0.21
Zamora	45.47	66174	0.25	0.13	0.18
Palma de Mallorca	45.46	561538	0.26	0.12	0.24
Ourense	45.28	131000	0.21	0.11	0.18
Palencia	45.27	93853	0.28	0.11	0.24
Blanes - Lloret de Mar	45.17	83084	0.27	0.15	0.15
Salamanca	45.00	183965	0.30	0.13	0.18
Huesca	44.99	53132	0.27	0.10	0.20
Lugo	44.90	98276	0.23	0.12	0.19
València	44.86	1564253	0.25	0.11	0.22
Talavera de la Reina	44.80	94028	0.25	0.12	0.22
Valladolid	44.61	405071	0.29	0.14	0.21
Asturias	44.17	801390	0.23	0.13	0.22
Castellò de la Plana	44.10	306215	0.26	0.11	0.21
A Coruña	43.91	416345	0.23	0.11	0.22
Toledo	43.90	123509	0.25	0.10	0.22
Alcoi	43.87	79829	0.24	0.10	0.21
Sagunt	43.72	72837	0.24	0.11	0.21
Ponferrada	43.72	80485	0.18	0.09	0.21
Santiago de Compostela	43.53	147632	0.23	0.09	0.23
Ceuta	43.28	83179	0.26	0.13	0.14
Santander - Torrelavega	43.25	321695	0.26	0.11	0.22
Gandía	43.20	123545	0.27	0.10	0.14
Cartagena	42.83	215826	0.27	0.11	0.21
Lorca	42.71	94404	0.22	0.11	0.20
Ciudad Real	42.02	90114	0.28	0.08	0.23
Murcia	41.91	664058	0.25	0.10	0.22
Elda - Petrer	41.75	86894	0.21	0.10	0.19
Ferrol	41.58	128413	0.23	0.09	0.23

Figure A41: Economic Connectedness Coverage by Municipality Population Decile




Notes: Bars show the share of municipalities in each population decile for which economic connectedness data is available. The blue line and points show the share of total 2001 Census population covered. Numbers above bars indicate the average municipality population in each decile. Coverage is nearly complete for the top three population deciles but limited for smaller municipalities.

Figure A42: Economic Connectedness and Upward Mobility at the ZIP-Code Level




Notes: Each point represents a Spanish ZIP code. The horizontal axis measures upward economic connectedness from [Johnston et al. \(2026\)](#); the vertical axis measures upward income mobility (AUM). Points are colored by whether their AUM is above or below the national median. The fitted line is from a linear regression. The regression coefficient, standard error, and number of ZIP codes are reported in the annotation box. Economic connectedness data is available for ZIP codes covering predominantly medium and large municipalities (see Appendix Figure [A41](#) for coverage by population decile).

Figure A43: Top Three Correlates of Upward Mobility by Category




Notes: Each point represents the coefficient from a univariate OLS regression of AUM on one standardized covariate. Variables are grouped by thematic category; the three strongest predictors within each category are displayed. Both OLS and Bayesian estimates are shown. Horizontal lines denote 95% confidence intervals.

Figure A44: The Great Gatsby Curve in Large Cities in Spain



Notes: Standardized correlation between the Gini index and the OLS rank-rank slope across municipalities, computed at increasing population thresholds. The analysis compares the original Gini and a bottom-99% Gini that excludes extreme top incomes. The association between inequality and persistence emerges only in municipalities with populations above approximately 20,000, highlighting the urban specificity of the Great Gatsby Curve in Spain.

Figure A45: Correlation Between Inequality and Intergenerational Persistence by Municipality Size



Notes: This figure plots the standardized correlation between the Gini index of income inequality and the OLS rank-rank slope (a measure of intergenerational persistence) across municipalities, conditional on different population thresholds. The x-axis shows increasing population cutoffs, and the y-axis reports the corresponding correlation. The analysis compares two Gini measures: the original Gini and a bottom 99% Gini that excludes extreme top incomes. While both measures yield similar patterns, the association between inequality and persistence only becomes salient in municipalities with populations exceeding 20,000, highlighting the urban specificity of the Great Gatsby Curve in Spain.

Table A8: Economic Connectedness and Upward Mobility: Controlling for Income, Inequality, and Geography

	EC (1)	+ Income (2)	+ Income + Gini (3)	+ Income + Gini p99 (4)	+ CCAA FE (5)	+ Province FE (6)
Panel A. Unweighted						
Upward economic connectedness	0.634*** (0.018)	0.641*** (0.029)	0.643*** (0.030)	0.643*** (0.029)	0.298*** (0.042)	0.234*** (0.046)
Municipal income level		-0.010 (0.030)	-0.013 (0.031)	-0.011 (0.030)	-0.046 (0.030)	-0.027 (0.030)
Municipal Gini			0.005 (0.021)			
Municipal Gini, winsorized at p99				0.008 (0.019)	0.078*** (0.019)	0.075*** (0.020)
Municipalities	1,807	1,807	1,807	1,807	1,807	1,807
R^2	0.402	0.402	0.402	0.402	0.566	0.603
Adjusted R^2	0.401	0.401	0.401	0.401	0.562	0.592
Panel B. Population-weighted						
Upward economic connectedness	0.777*** (0.053)	0.693*** (0.072)	0.652*** (0.065)	0.639*** (0.069)	0.226*** (0.042)	0.163*** (0.047)
Municipal income level		0.099 (0.104)	0.182** (0.083)	0.162 (0.102)	0.034 (0.028)	0.048* (0.027)
Municipal Gini			-0.094* (0.055)			
Municipal Gini, winsorized at p99				-0.172*** (0.040)	0.052** (0.021)	0.056** (0.022)
Municipalities	1,807	1,807	1,807	1,807	1,807	1,807
R^2	0.506	0.511	0.516	0.527	0.795	0.828
Adjusted R^2	0.506	0.511	0.515	0.526	0.793	0.823


Notes: The dependent variable is upward mobility at the municipality level, based on OLS estimates. All variables are standardized to have mean zero and unit standard deviation in the estimation sample. Municipal income level is log-transformed before standardization. Columns (1)–(4) sequentially add municipal income and inequality controls. Columns (5) and (6) add autonomous-community and province fixed effects, respectively, to the specification in column (4). Panel A reports unweighted specifications. Panel B reports specifications weighted by 2001 municipal population. Robust HC1 standard errors are reported in parentheses. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A9: Economic Connectedness and Conditional Correlates of Upward Mobility

	EC (1)	+ Unemp. (2)	+ Education (3)	+ Youth dep. (4)	+ Agric. emp. (5)	+ CCAA FE (6)	+ Province FE (7)
Panel A. Unweighted							
Upward economic connectedness	0.634*** (0.018)	0.485*** (0.021)	0.454*** (0.028)	0.426*** (0.028)	0.394*** (0.029)	0.009 (0.048)	-0.043 (0.054)
Unemployment rate		-0.268*** (0.024)	-0.266*** (0.024)	-0.252*** (0.024)	-0.251*** (0.024)	-0.150*** (0.025)	-0.125*** (0.027)
Average education, age 30–39			0.045* (0.025)	0.034 (0.025)	0.010 (0.026)	0.116*** (0.027)	0.113*** (0.028)
Youth dependency				-0.073*** (0.021)	-0.059*** (0.021)	-0.111*** (0.028)	-0.095*** (0.029)
Agricultural employment					-0.092*** (0.022)	-0.070*** (0.023)	-0.081*** (0.024)
Partial correlation	0.634	0.478	0.367	0.334	0.301	0.005	-0.020
Municipalities	1,807	1,807	1,807	1,807	1,807	1,807	1,807
R^2	0.402	0.451	0.452	0.456	0.460	0.584	0.616
Adjusted R^2	0.401	0.451	0.451	0.454	0.459	0.580	0.605
Panel B. Population-weighted							
Upward economic connectedness	0.777*** (0.053)	0.609*** (0.060)	0.479*** (0.052)	0.413*** (0.052)	0.431*** (0.052)	-0.026 (0.048)	-0.101** (0.051)
Unemployment rate		-0.432*** (0.047)	-0.449*** (0.049)	-0.428*** (0.049)	-0.424*** (0.048)	-0.200*** (0.026)	-0.156*** (0.025)
Average education, age 30–39			0.144** (0.060)	0.089* (0.052)	0.098* (0.054)	0.124*** (0.024)	0.113*** (0.022)
Youth dependency				-0.157*** (0.040)	-0.160*** (0.041)	-0.127*** (0.030)	-0.139*** (0.027)
Agricultural employment					0.054 (0.042)	-0.027 (0.026)	-0.071*** (0.023)
Partial correlation	0.712	0.588	0.413	0.351	0.345	-0.018	-0.066
Municipalities	1,807	1,807	1,807	1,807	1,807	1,807	1,807
R^2	0.506	0.571	0.586	0.600	0.601	0.815	0.844
Adjusted R^2	0.506	0.570	0.586	0.599	0.600	0.813	0.840


Notes: The dependent variable is upward income mobility at the municipality level, based on OLS estimates. All variables are standardized to have mean zero and unit standard deviation in the estimation sample. The table sequentially adds the leading municipal correlates of upward mobility, excluding the broad socioeconomic status index. Columns (6) and (7) add autonomous-community and province fixed effects, respectively, to the full top-five specification. Panel A reports unweighted specifications. Panel B reports specifications weighted by 2001 municipal population. Robust HC1 standard errors are reported in parentheses. The partial correlation row reports the association between upward economic connectedness and upward mobility after residualizing both variables on the controls and fixed effects included in each column. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure A46: Evolution of Absolute Upward Mobility (μ_{25}) in the 50 Largest Municipalities: 1980 vs. 1990 Cohorts



Notes: Orange markers show μ_{25} for the 1980 birth cohort; teal markers show μ_{25} for the 1990 birth cohort. Municipalities are ranked by population according to the 2001 Census. μ_{25} is the expected mean child income percentile for children whose parents were at the 25th percentile of the national parental income distribution. Income definitions and sample construction follow Section 3.

Figure A47: Evolution of Absolute Upward Mobility (μ_{25}) at the Municipality Level: 1980 vs. 1990 Birth Cohorts



Notes: This figure maps absolute upward mobility (μ_{25}) at the municipality level for children born in 1980 (Panel a) and 1990 (Panel b). μ_{25} is the expected mean child income percentile for children whose parents were at the 25th percentile of the national parental income distribution. Both panels use a common color scale: blue shades indicate higher mobility and red shades lower mobility. Estimates are derived from the Bayesian hierarchical model described in Section 4.1. The comparison reveals a widespread deterioration in mobility prospects for the later cohort, with the expansion of low-mobility areas (red) particularly visible in southern and central Spain. Several traditionally high-mobility zones in the North-East retain comparatively better outcomes but also exhibit visible reductions in the intensity of upward mobility.

G Intergenerational Linkage, Coverage, and Representativeness

This appendix describes how the Spanish Opportunity Atlas links parents to children in administrative tax records, documents the population coverage of the resulting dataset, and evaluates its representativeness. We draw on multiple external data sources—the 1998 *Padrón Municipal Continuo* (Spain’s continuous population register) (Instituto Nacional de Estadística, 1998), the 2001 Census, *Censo de Población y Viviendas* (Instituto Nacional de Estadística, 2001), and the Spanish Tax Agency’s 1998 published statistics of the personal income tax (*Estadística de los Declarantes del IRPF*) (Agencia Estatal de Administración Tributaria, 1999)—to benchmark our sample against the target population of children born between 1980 and 1990 in the *Territorio de Régimen Fiscal Común* (hereafter “common tax territory”), which excludes the Basque Country and Navarra due to their independent fiscal administrations.

G.1 Intergenerational Linkage Procedure and Panel Structure

The Spanish Opportunity Atlas is built from anonymized administrative records constructed by the Spanish Tax Agency (AEAT) specifically for this research project. We do not have direct access to the universe of tax records, rather, AEAT selected and linked the relevant populations following a protocol designed to maximize the representativeness and identification of parent-child pairs for the 1980–1990 birth cohorts, and delivered the resulting anonymized dataset to the research team. We describe the construction process and examine coverage along two dimensions: the share of the population captured (subsection A.2) and the representativeness of the income distribution (subsection A.3).

Sample construction by the Spanish Tax Agency. The construction of the intergenerational linkage proceeded in two stages, working backwards from children to parents. In the first stage, AEAT identified a sample of target children from the fiscal year 2003. The year 2003 was chosen as the starting point rather than 1998 because the 1998 tax records lack complete identifying information (name, surname, and date of birth) for all dependents. Moreover, the Spanish tax identification number (*NIF*) is mandatory only from age 14 onward; for younger children without a NIF, identification requires matching on name, surname, and date of birth using probabilistic algorithms. By 2003, all children in the 1980–1990 cohorts were aged 13–23, ensuring that the vast majority had been assigned a NIF.⁵² This process yielded 2,869,611 children born between 1980 and 1990 who appeared as dependents (*descendientes*) in the *Modelo 100* (personal income tax return) of fiscal year 2003 together with their parents tax records.

In the second stage, Spanish Tax Agency traced these parents backwards to the fiscal year 1998 by searching

⁵²Children born in 1990 were 13 in 2003, just below the mandatory NIF threshold of 14. AEAT reports using matching algorithms for this age group, though some identification loss is possible for the youngest cohort.

for their tax returns in that year, managing to find 2,712,065 of the initial 2,869,611 (94.5%) children, forming the core sample of target children. The remaining 157,546 children (5.5%) could not be linked because their parents did not file a tax return in 1998, likely cases of parents who began filing after 1998 (recent immigrants, young parents entering the labor market, or parents below the filing threshold in 1998). From the matched parent-child pairs, AEAT constructed a household file containing 2,022,118 unique fiscal households, each identified by a household number. The number of households is smaller than the number of children because siblings share the same household.

Household composition. Each fiscal household in 1998 contains at least one target child and the associated parent(s). AEAT classified household members into three categories: P (primary tax filer), H (target child), and D (dependent). The delivered dataset contains 9,568,605 individuals: 2,712,065 target children (H), 2,549,042 primary filers (P), and 4,307,498 other members of the original household (D). The category D is a residual classification encompassing any household member who is neither the primary filer nor a target child, including spouses, siblings of the target child, grandparents, and grandchildren. Using a reclassification procedure based on age differences between household members, we identify 1,278,360 spouses of the primary filer among the dependents, bringing the total number of identified parental figures to 3,827,402 (2,549,042 primary filers plus 1,278,360 reclassified spouses). We also identify 1,287,073 siblings of target children. The remaining dependents include 1,264,339 children of the target children (i.e., grandchildren of the parents), 273,565 grandparents, and smaller numbers of other relatives. Table [A10](#) summarizes these categories.

Panel structure. The Spanish Tax Agency then followed all 9.57 million individuals across the tax records from 1998 to 2022, creating a 25-year longitudinal panel. The nature of the available records changes over time:

- 1998–2007 (IRPF period): Only individuals who filed a personal income tax return (*Modelo 100*) are observed. Since filing was required only above certain income thresholds (see below), individuals with very low or no declared income are not observed in these years. During this period, especially in the first years, the vast majority of individuals for whom we have tax records are parents, since children were not old enough to work (only 0.5% appear to have worked in 1998). An important institutional feature is that 72% of primary filers (P) are male (1,831,968 men versus 714,396 women), reflecting the cultural practice at time in which the husband was typically recorded as the primary declarant. In households with a single P, the mother usually appears as a dependent (D) rather than as a co-filer, a pattern that motivates the spouse reclassification procedure described above. Crucially, approximately 92% of parental tax returns in our sample are filed jointly (*declaración conjunta*), meaning that total household income—including the earnings of both spouses—is captured in the return even when only

Table A10: Household Composition of the Atlas Population

Category	Role	N	%
Panel A: Original AEAT classification			
Target children	Children born 1980–1990 (H)	2,712,065	28.3
Primary tax filers	Tax filer parent (P)	2,549,042	26.6
Dependents	Other household members (D)	4,307,498	45.0
Panel B: Reclassified dependents^a			
Spouses	Partner of the primary filer ^b	1,278,360	13.4
Grandchildren	Children of target children ^c	1,264,339	13.2
Close siblings	Sibling, ± 5 years of H ^d	560,314	5.9
Mid-range siblings	Sibling, 6–10 years of H	539,664	5.6
Distant siblings	Sibling, 11–15 years of H	187,095	2.0
Older siblings	Sibling, >15 years older than H	29,687	0.3
Grandparents	Ascendants of the parent ^e	273,565	2.9
Parent's relatives	Sibling or in-law of parent ^f	85,628	0.9
Other ^g	Unclassified or ambiguous	89,846	0.9
Total		9,568,605	100.0

^a Dependents (D) are reclassified using the difference in birth year between each dependent and the parents (P) and target children (H) in the same fiscal household. Let Δ_P = dependent's birth year – mean birth year of P, and Δ_H = dependent's birth year – mean birth year of H.

^b $|\Delta_P| \leq 10$, household has exactly one P, and dependent is of opposite sex to P. 94.8% are female, consistent with fathers being the primary filer in 72% of households.

^c $\Delta_H > 15$ (born more than 15 years after the target child). Mean birth year: 2014.

^d Siblings satisfy $15 \leq \Delta_P \leq 55$ (plausible child of P) and $|\Delta_H| \leq 15$ (born within 15 years of H). Subcategories are defined by $|\Delta_H|$: close (≤ 5), mid-range (6–10), distant (11–15).

^e $\Delta_P \leq -20$ (born 20 or more years before the parent). Mean birth year: 1929.

^f $|\Delta_P| \leq 10$ but household already has two P, or dependent is of the same sex as P.

^g Includes 42,164 ambiguous cases ($10 < \Delta_P < 15$), 13,320 residual cases, and 8,006 individuals lacking birth year information. Together these represent less than 1% of the sample. An additional 139 individuals belong to the 401 households with no identified parent.

one parent appears as the primary filer. This substantially mitigates concerns about incomplete income measurement during the IRPF period.

- 2008–2022 (ERM period): Beginning in 2008, the data source expands to include the *Estadística de Rendimientos y Retenciones del Trabajo* (ERM), which captures all individuals with any income subject to withholding, including those who do not file a tax return. This dramatically increases coverage: approximately 8 to 9 million individuals from our sample are observed annually. Since we measure children’s income over 2021–2022, well within the ERM period, the income of target children is observed for approximately 98% of the sample regardless of their filing status. This ensures that our estimates of intergenerational mobility are not affected by selective non-filing among the younger generation.

G.2 Population Coverage

Target population. According to the 1998 Population Census ([Instituto Nacional de Estadística, 1998](#)), there were 5,342,429 Spanish nationals aged 8 to 18 residing in Spain, corresponding to the 1980–1990 birth cohorts. Subtracting the Basque Country (240,947) and Navarra (63,041), whose tax records are administered separately (see below), the target population in the common tax territory is 5,038,441 children.

From the target population to the estimation sample. Table [A11](#) traces the path from this target population to the sample underlying our estimates, distinguishing three layers. The first is the raw linkage: the Atlas contains 2,712,065 target children identified as dependents in the 2003 tax records and traced back to their parents’ 1998 returns, plus 284,230 siblings born in the same cohorts, for a total of 2,996,295 individuals, or 59.5% of the target population.⁵³ The second layer is the analysis dataset, which retains children whose parental income in 1998–2000 is available and passes the quality filters described in Section 2; it contains 2,495,136 children, or 49.5% of the target population. The third layer is the estimation sample after the bottom-5% child-income trim: 2,370,385 children (47.0%) in the extended 1980–1990 sample, and 1,501,948 children (43.4%) in the baseline 1980–1986 cohorts. Every coverage statement in this paper refers explicitly to one of these layers.

Coverage by birth cohort. Figure [A48](#) shows that coverage varies substantially across cohorts, following an inverted-U shape: at the Atlas layer it rises from 40.3% for the 1980 cohort to a peak of 73.2% for 1987, then declines to 65.0% for 1990. The pattern is mechanical and follows directly from the construction procedure. Children born in 1980 were 23 years old in 2003 and many had already left the parental household, so they no longer appear as dependents in the records from which the linkage starts; children born after 1987 were

⁵³Of these, 11,036 target children (0.4%) lack a valid birth year and are excluded from all cohort-level figures; the mass of invalid values is concentrated immediately adjacent to the 1980–1990 window, consistent with coding errors in the recorded date of birth.

Table A11: From Target Population to Estimation Sample

	All cohorts (1980–1990)		Baseline (1980–1986)	
	N	% of target	N	% of target
Target population (Padrón 1998) ^a	5,038,441	100.0	3,457,992	100.0
Atlas linkage (target children + siblings) ^b	2,985,259	59.2	1,894,292	54.8
Analysis dataset (income-quality filters) ^c	2,495,136	49.5	1,580,994	45.7
Estimation sample (bottom-5% child-income trim)	2,370,385	47.0	1,501,948	43.4

^a Spanish nationals born 1980–1990 residing in the common tax territory (excluding the Basque Country and Navarra). Source: Padrón Municipal Continuo 1998, INE.

^b Children identified as dependents in the 2003 tax records and traced back to their parents' 1998 returns, plus reclassified siblings born in the same cohorts. Counts refer to cohort-assignable individuals; an additional 11,036 target children (0.4% of the published Atlas total of 2,996,295) lack a valid birth year and are excluded from cohort-level figures.

^c Excludes children whose parental income in 1998–2000 is unavailable and parental households dropped by the negative-income and imputed-wealth (“false poor”) filters described in Section 2. The attrition rate of this step is stable across cohorts (15.9 to 16.9% of the Atlas layer), indicating that income-quality filtering does not interact with the mechanical linkage margin.


increasingly below the age at which the tax identification number (*NIF*) becomes reliable, complicating identification. Two features of the figure deserve emphasis. First, the three layers move in parallel: the attrition rate of the income-quality filters is essentially constant across cohorts, between 15.9 and 16.9% of the Atlas layer, so the filters do not interact with the mechanical linkage margin. Second, precisely because the intensity of this margin varies so strongly across cohorts, the cohort dimension provides a natural laboratory for testing whether non-coverage distorts the mobility estimates, an exercise we develop below.

Family structure. A second source of non-coverage is that some children cannot be linked to parental tax records even in principle. According to the 2001 Census ([Instituto Nacional de Estadística, 2001](#)), 4.3% of children aged 10–19 did not live with any parent, and 15.6% lived with a single parent who may not have filed a tax return. Among the 80.2% of children living with both parents (77.9% in married couples plus 2.3% in cohabiting partnerships), Atlas coverage rises to 74.2%. The remaining gap at this margin corresponds to households in which no parent filed in 1998, which we examine next.

Parental filing and the income margin. Non-filing is the one source of non-coverage that is tied to income by construction, so it deserves particular scrutiny. Under Article 96 of Ley 18/1991, in force for fiscal year 1998, individuals were exempt from filing if their total gross income was below €7,212 (1,200,000 pesetas; approximately €12,500 in 2022 prices) and consisted exclusively of labor earnings plus capital income and capital gains jointly below €1,502; the limit was €7,513 for pensions and joint returns, and any self-employment income triggered a filing obligation regardless of amount. The threshold corresponds to roughly 1.3 times the annual statutory minimum wage.⁵⁴ Only households whose adults all earned below this level

⁵⁴The 1998 *Salario Mínimo Interprofesional* was approximately €5,570 per year including supplementary payments, or roughly €9,700 in 2022 prices.

Figure A48: Coverage of the target population by birth cohort



Notes: The figure shows, for each birth cohort, the share of the target population covered at three successive layers of the data: the raw Atlas linkage (children identified as dependents in the 2003 tax records and traced to their parents' 1998 returns, plus reclassified siblings), the analysis dataset after the income-quality filters of Section 2, and the estimation sample after the bottom-5% child-income trim. The denominator is the number of Spanish nationals of each cohort residing in the common tax territory according to the 1998 Population Census (Table A11). Source: Spanish Opportunity Atlas and 1998 Population Census, INE.


could legally abstain from filing, so any missing parents are concentrated at the very bottom of the income distribution. The relevant question for our estimates is therefore not how many parents are observed in total, but whether many poor parents are absent. Two independent pieces of evidence indicate that this group is small.

The first is the income composition of parent-filers in 1998. The AEAT's published tax statistics ([Agencia Estatal de Administración Tributaria, 1999](#)) report that 6,487,633 of the 15,424,100 filers claimed the deduction for cohabiting dependent children under age 25 (*Partida 069*). Figure A49 and Table A12 break down filers by income bracket and parental status. Among filers with taxable income below €3,005, only 1.9% claimed the child deduction, and such filers account for a negligible share of all declarants; the share of filers with children rises monotonically with income, reaching 51.6% in the €12,020–18,030 bracket and over 60% above €30,050. The population at risk of non-filing is confined by construction to households whose adults all earned below 1.3 times the minimum wage, and within the observed filer distribution dependent children are nearly absent from precisely those income levels. A large reservoir of low-income parents hidden just below the threshold is therefore difficult to reconcile with the data, although households operating entirely in the informal economy cannot be ruled out.

The second piece of evidence comes from the 1999 tax reform. Ley 40/1998, effective for fiscal year 1999, tripled the filing threshold from €7,212 to €21,035, releasing a large share of low-income taxpayers from the obligation to file. If low-income parents had appeared in the 1998 records only because filing was compulsory,

they should vanish once it became voluntary. Instead, attrition among the linked parents in the Atlas was materially lower than in the general population: between 1998 and 1999 the number of filers nationwide fell by roughly 11%, against only 9.1% among Atlas parents. The gap reflects the *mínimo por descendientes* introduced by the same reform, which made filing financially attractive for parents with dependent children even below the new threshold: the implied voluntary filing rate among below-threshold parents is approximately 82%, compared with 48% in the general below-threshold population. Low-income parents remained observable in the tax records at rates far above what the filing rules alone would predict.

Figure A49: Distribution of tax filers by income bracket and parental status



Notes: The figure shows the distribution of the 15.4 million IRPF filers in fiscal year 1998 across taxable income brackets (grey bars) and, within each bracket, the filers claiming the deduction for cohabiting dependent children under age 25, *Partida 069* (blue bars); both series are expressed as shares of all filers. The dashed vertical line marks the 1998 filing threshold of €7,212 under Article 96 of Ley 18/1991 (approximately €12,500 in 2022 prices). Filers below the threshold appear in the data because they filed voluntarily or were obliged to file for other reasons, such as self-employment income. Source: Estadística de los Declarantes del IRPF 1998, AEAT.

Reconciling child and parental coverage. Child coverage and parental coverage measure different margins and should not be expected to coincide. Child coverage (59.5%) is an individual-level rate over the entire census population. Parental coverage is a household-level rate: a household enters the Atlas whenever at least one of its children was still a fiscal dependent in 2003, even if older siblings had already left and are lost to the linkage. A back-of-the-envelope calculation illustrates the difference. Combining the census count of target children with the Atlas averages of 1.99 children and 1.26 primary filers per household yields on the order of 3.2 million parent-filers with at least one child aged 8 to 18; the Atlas contains 2,549,042, implying that roughly four in five filing households appear in the data. The gap between this household-level figure and the 59.5% child-level rate is precisely the within-household loss of older children documented in Figure A48: the missing individuals are concentrated among children who had left the parental household by 2003, not among poor families. The remaining, structurally unreachable sources are children not residing with any

Table A12: Distribution of Tax Filers by Income Bracket and Parental Status, 1998

Taxable income bracket (€)	All filers		Filers with children		Rate (%)
	N	%	N	%	
< 3,005	954,424	6.2	17,679	0.1	1.9
3,005–6,010	2,144,920	13.9	388,487	2.5	18.1
6,010–9,015	3,164,788	20.5	1,202,118	7.8	38.0
9,015–12,020	2,788,025	18.1	1,291,376	8.4	46.3
12,020–18,030	3,098,047	20.1	1,597,182	10.4	51.6
18,030–30,050	2,322,315	15.1	1,364,115	8.8	58.7
30,050–60,101	790,412	5.1	521,708	3.4	66.0
60,101–150,253	143,478	0.9	94,720	0.6	66.0
> 150,253	17,691	0.1	10,248	0.1	57.9
Total	15,424,100	100.0	6,487,633	42.1	42.1


Notes: Brackets are in taxable income (*base imponible*), converted from millions of pesetas at 1M ptas = €6,010.12. “Filers with children” are those claiming the deduction for cohabiting unmarried dependents under age 25 (*Partida 069*). Percentages in both columns are computed as shares of the total number of filers (15,424,100), so that the two distributions are directly comparable. “Rate” is the share of all filers in each bracket who claim the child deduction. Source: Estadística de los Declarantes del IRPF 1998, AEAT ([Agencia Estatal de Administración Tributaria, 1999](#)).

parent (4.3%) and children of non-filing parents, a group that the filing evidence above indicates is small.

Coverage and cohort-level intergenerational mobility estimates. If non-coverage selected children on unobservables related to mobility, cohorts covered at 40% and cohorts covered at 73% should deliver systematically different estimates. Figure A50 shows no trace of such a divergence. Across the eleven cohorts, absolute upward mobility varies by less than one half of a percentile rank in total, and the rank-rank slope evolves smoothly and monotonically. The sharpest test comes from the reversal of coverage after 1987: coverage falls by 8.2 percentage points between the 1987 and 1990 cohorts, yet the rank-rank slope continues rising (+0.007) and μ_{25} continues falling (−0.18), on the same trajectory as when coverage was increasing by 18.1 points between 1983 and 1987. Estimates track the labor-market entry conditions of each cohort documented in Section 4.4, not its coverage. Formally, the raw cross-cohort association between the slope and coverage disappears once a linear cohort trend is included; the coefficients are reported in the notes to Figure A50.

Population reweighting. Unlike other contexts, as the Canadian administrative data for which Statistics Canada provides income-based weights that raise effective coverage to 88–99.5% ([Connolly et al., 2019b](#)), the Atlas does not come with sampling weights. We therefore construct post-stratification weights from the census, assigning each observation the ratio of its cell’s population share to its sample share, first by birth cohort and then by cohort and gender. Table A13 shows that reweighting leaves the national estimates essentially unchanged: the baseline rank-rank slope moves from 0.2755 to 0.2748 and μ_{25} from 43.48 to 43.56.

Figure A50: Cohort-specific intergenerational mobility estimates



Notes: Panel (a) shows the cohort-specific rank-rank slope and Panel (b) absolute upward mobility (μ_{25}), estimated on the estimation sample with income ranks computed within cohort; shaded areas are 95% confidence intervals. Child income is measured over the common 2021–2022 window, so cohort-level differences reflect both age-at-measurement and genuine cohort effects, and levels are not directly comparable to the fixed-age estimates of Figure A31. The raw cross-cohort association between the rank-rank slope and cohort coverage is +0.0004 (SE 0.0001) and becomes –0.0001 (SE 0.0001) once a linear cohort trend is included; for μ_{25} the corresponding figures are –0.011 (SE 0.003) and +0.002 (SE 0.003). These associations are computed over eleven cohort-level observations and are reported for completeness only.

Because weights are constant within cohort, within-cohort income ranks are unaffected; the exercise therefore isolates the effect of demographic composition and cannot, by construction, correct for income-selective non-coverage within cohorts. That margin is bounded by the filing evidence above and stress-tested by the simulations below.

Table A13: Robustness of National Mobility Estimates to Population Reweighting


	Rank–rank slope	(SE)	μ_{25}	(SE)
<i>Panel A: Baseline sample (cohorts 1980–1986)</i>				
Unweighted	0.2755	(0.0008)	43.48	(0.03)
Cohort weights	0.2749	(0.0008)	43.49	(0.03)
Cohort $\times$ gender weights	0.2748	(0.0008)	43.56	(0.03)
<i>Panel B: Extended sample (cohorts 1980–1990)</i>				
Unweighted	0.2792	(0.0006)	43.38	(0.02)
Cohort weights	0.2783	(0.0006)	43.40	(0.02)
Cohort $\times$ gender weights	0.2784	(0.0006)	43.43	(0.02)

Notes: Each observation is weighted by the ratio of its cell’s population share to its sample share. Cohort weights use the Padrón 1998 counts of Spanish nationals in the common tax territory by birth cohort; cohort $\times$ gender weights additionally use INE female population shares by cohort. Because weights are constant within cohort, within-cohort income ranks are unchanged: the exercise isolates the effect of cohort and gender composition and cannot, by construction, correct for income-selective non-coverage within cohorts, which is examined separately in this appendix. Standard errors are conventional OLS standard errors. Unweighted estimates differ marginally from the headline estimates in the main text (0.274) due to small differences in rank construction; the comparison across weighting schemes is internally consistent.

Demographic composition of the sample. Gender is the only demographic characteristic for which an exact population benchmark is available at the cohort level. Figure A51 compares the female share of each cohort in the analysis dataset with the corresponding INE population share. Women are over-represented by 1.5 to 2.3 percentage points in the low-coverage cohorts born 1980–1983, and the sample composition converges to the population benchmark (deviations below 0.6 points) for cohorts born after 1985, precisely as the mechanical linkage margin closes. This is the signature of earlier male emancipation from the parental household: young men who had entered the labor market by 2003 drop out of fiscal dependency and hence out of the linkage. The deviation runs in the direction predicted by the mechanical margin rather than by income-based selection, and the cohort-by-gender weights of Table A13 confirm that it has no material effect on the estimates.

Simulated selective attrition. To quantify how strong selective non-coverage would have to be before it mattered, Table A14 reports simulations in which the best-covered cohort (1987) is artificially thinned to the coverage rate of the worst-covered cohort (1980, retaining 55% of the original sample), with the probability of remaining in the sample tilted against low percentiles of parental income, child income, or both. Three results

Figure A51: Female share by birth cohort: Atlas sample and population



Notes: The figure compares the female share of each birth cohort in the analysis dataset (solid blue) with the female share in the population of Spanish nationals of the same cohort (dashed grey; 1998 Population Census, INE). Population shares are national and include the Basque Country and Navarra (5.7% of the total); gender ratios are essentially constant across regions.

stand out. First, selection on parental income alone barely moves the estimates: even when bottom-decile families are retained at 12.5% against 92.5% for top-decile families, the rank-rank slope changes by $+0.004$. This reflects a general property of rank-based estimators: truncating the regressor distribution, the worst case implied by the filing threshold, is largely absorbed by the re-computation of ranks. Second, selection on the child's own income, the channel through which early emancipation of low-earning children could operate, shifts the slope by at most $+0.02$, with the effect saturating even under extreme selection. Third, under the harshest joint scenario the measured slope falls by 0.05 : removing families in which both generations have low income compresses both margins of the joint distribution and attenuates the estimated gradient. The selection actually observed in the data, namely the cross-cohort stability of Figure A50 and gender deviations of at most 2.3 percentage points, lies far below the magnitudes required in any of these scenarios.

Implications for the intergenerational mobility estimates. Two margins of non-coverage should be kept distinct. The dominant margin, children identified as dependents in 2003 rather than 1998, is mechanical: it varies sharply across cohorts for reasons unrelated to income, its missing individuals are older children who had left the parental household rather than poor families, and the evidence above shows that the estimates are invariant to it, whether tested through cohort variation, population reweighting, or observable composition. The second margin, households in which no parent filed in 1998, is income-selective by construction but quantitatively small, as indicated by the near-absence of dependent children among below-threshold filers and by the high voluntary filing rate of parents around the 1999 reform. To the extent that some severely disadvantaged households remain unobserved at this margin, the simulations in Table A14 indicate that their omission would compress the parental distribution from below and, if anything, attenuate our estimates: under this interpretation, the intergenerational persistence reported in this paper is a conservative estimate for Spain in this period.

Table A14: Simulated Selective Attrition: How Strong Would Selection Have to Be?

	Implied keep rate (%)		Δ Rank–rank slope		$\Delta \mu_{25}$	
	Bottom decile	Top decile	Mean	(SD)	Mean	(SD)
<i>Panel A: Selection on parental income (regressor)</i>						
$\gamma = 0.00$	55.0	55.0	-0.0001	(0.0019)	0.00	(0.05)
$\gamma = 0.02$	33.6	75.2	0.0122	(0.0016)	-0.31	(0.04)
$\gamma = 0.05$	12.5	92.5	0.0040	(0.0016)	-0.10	(0.04)
<i>Panel B: Selection on child income (outcome)</i>						
$\gamma = 0.00$	55.0	55.0	-0.0009	(0.0015)	0.02	(0.04)
$\gamma = 0.02$	33.6	75.2	0.0223	(0.0015)	-0.57	(0.04)
$\gamma = 0.05$	12.5	92.5	0.0216	(0.0018)	-0.55	(0.05)
<i>Panel C: Joint selection on parental and child income</i>						
$\gamma = 0.00$	55.0	55.0	-0.0003	(0.0020)	0.01	(0.05)
$\gamma = 0.02$	34.1	74.8	0.0049	(0.0012)	-0.13	(0.03)
$\gamma = 0.05$	12.9	92.1	-0.0534	(0.0015)	1.36	(0.04)

Notes: Each scenario takes the best-covered cohort (1987) and randomly deletes observations until 55% remain, matching the lowest observed cohort coverage (1980). The probability of remaining in the sample is logistic in the percentile indicated by each panel, with slope γ ; $\gamma = 0$ corresponds to fully random attrition. Ranks are recomputed among survivors before re-estimation, as in the actual pipeline. Cells report means and standard deviations over 20 replications, expressed as deviations from the full-cohort estimates. Selection on the regressor (Panel A) leaves rank-based estimates essentially unchanged even under extreme selection, illustrating why filing-threshold-type truncation of the parental distribution is largely innocuous. Selection on the outcome (Panel B) is the channel that can bias estimates, and its effect saturates at roughly +0.02 in the rank–rank slope. Under the most severe joint scenario (Panel C), the measured slope falls, implying that plausible selective attrition would lead to understating intergenerational persistence.

Exclusion of the Basque Country and Navarra. The Basque Country and Navarra operate independent tax administrations under their historical fiscal agreements (*concierto* and *convenio económico*, respectively). Tax records from these territories are not part of AEAT’s national database, making intergenerational linkage impossible for their residents. According to the 1998 Population Census, the two communities account for 303,988 children aged 8 to 18, or 5.7% of the national total.

H Sample Construction

This appendix examines the sample filters applied to the bottom of the child income distribution. Our baseline sample trims the bottom 5% of the within-cohort child income rank distribution (Section 2.2). Here we document the measurement anomaly that motivates this filter, use administrative wealth records to characterize precisely which individuals drive it, construct an alternative, narrower filter that excludes only individuals whose reported income is directly contradicted by their own asset holdings, and show that the paper’s substantive conclusions are unchanged across all filtering choices.

H.1 The anomaly at the bottom of the child rank distribution

Figure A52 plots the mean parental income rank at each percentile of the child income rank distribution, for the bottom 30 child percentiles and under four alternative samples defined below. In an unfiltered sample, the relationship between child and parental ranks should be monotonically increasing: children at lower income ranks should, on average, come from poorer families. This is what we observe between the 8th and 30th child percentiles. Below the 8th percentile, however, the pattern reverses sharply. In the sample with no bottom trim, children at the very lowest income ranks have parents at the 60th percentile of the parental income distribution on average, 18 rank points *above* the level prevailing among children at ranks 10 to 30. Children in the bottom two percentiles report near-zero annual income (below €1,000 in virtually all cases), yet their measured family background is more affluent than that of children at the 10th or 20th percentile.

This inverted pattern is the signature of a composition problem rather than of poverty: the very bottom of the child rank distribution mixes genuinely poor individuals with a distinct group whose Spanish taxable income is near zero for reasons unrelated to economic deprivation, such as unreported self-employment earnings, income earned abroad, or resources held as assets rather than received as taxable income. Because these individuals enter the rank-rank regression as extreme low-outcome observations attached to high-rank parents, their presence attenuates measured persistence, while removing too many bottom-ranked children risks discarding genuine poverty. The remainder of this appendix uses wealth records to separate the two groups.




Figure A52: Parental income rank across the bottom of the child rank distribution

Notes: Mean parental household income rank by child income rank (bottom 30 percentiles), extended sample (cohorts 1980–1990), under the four samples described in Appendix H.2. Ranks are within-cohort percentiles. In the samples without a bottom trim, mean parental rank rises steeply below the 8th child percentile, reaching the 60th percentile among children with the lowest reported incomes. The child false-poor filter removes roughly two thirds of this spike; the baseline 5% trim removes it entirely, along with all other bottom-ranked children.

H.2 False-poor observations: definition and filters

We compare four samples. The first three are nested and cumulative; the fourth is the paper’s baseline, which branches from the second.

Sample (1) applies only the negative-income filter, excluding parental households that report negative income in any year of the 1998–2000 window. Sample (2) additionally excludes *parental false-poor* households: those reporting 1998 income in the bottom 10% of the national distribution while their imputed housing wealth lies in the top 40%, a combination indicating substantial unreported resources rather than genuine low income. Sample (3) additionally excludes *child false-poor* individuals, defined symmetrically on the child side: individuals whose income rank lies in the bottom 5 percentiles of the within-cohort distribution while their *own* gross wealth, observed in the administrative wealth module and averaged over 2016–2022, lies in the top 25% of the sample. Percentile ranks are recomputed within cohort after each exclusion. The baseline sample instead applies a uniform trim of the bottom 5% of the child rank distribution to sample (2), without conditioning on wealth.

The two false-poor filters embody a single principle, applied to each generation: an observation is excluded when its reported income is directly contradicted by its own asset holdings. Importantly, the child false-poor filter conditions exclusively on the child’s own income and own wealth. It does not use any information about the child’s parents, so it cannot mechanically remove downwardly mobile children of affluent families: a child of rich parents with genuinely low income *and* no assets of their own remains in the sample.

Figure A53 illustrates the income-wealth dissonance that defines both groups. Excluded child false-poor individuals sit at the 3rd percentile of income but the 87th percentile of own gross wealth on average, with a median income of €480 against a median wealth of €84,658; bottom-ranked children who are retained hold a mean wealth rank of 28. The parental panel shows the same qualitative contrast (55th versus 31st wealth percentile), attenuated because the parental filter relies on a 1998 housing-wealth proxy while wealth is observed only from 2016 onward, two decades during which mortality and inter vivos transfers erode the correspondence.




Figure A53: Income-wealth dissonance among excluded and retained low-income observations

Notes: Mean within-cohort income rank (grey) and own gross-wealth rank (red) for excluded false-poor observations and for low-income observations retained in the sample, by generation. Child income is measured in 2021–2022, parental income in 1998, and gross wealth (real estate, deposits, financial and other assets, before liabilities) is averaged over observed years 2016–2022. The dotted line marks the median. Excluded groups combine bottom-ranked income with above-median own wealth; retained low-income groups are poor in both dimensions.

H.3 Who is excluded by each filter

Table A15 characterizes the observations removed by each filter. The two false-poor groups display the dissonance described above in their respective generations. The child false-poor group is small (16,752 individuals, 0.6% of the no-trim sample) but extreme: a median own wealth of €84,658 against a median income of €480, and, although the filter never conditions on family background, these individuals turn out to come from affluent families, with parents at the 61st percentile of 1998 income and the 65th percentile of 2016–2022 wealth on average. The connection between the bottom-rank anomaly and family affluence therefore emerges from the data rather than being imposed by the filter, and it accounts for most of the spike in Figure A52.

The contrast with the blind 5% trim is instructive. The trimmed group is eight times larger (131,181 individuals) and exhibits no income-wealth dissonance: mean parental income rank of 49, own wealth rank of

35, and a median own wealth of €8,641. The uniform trim therefore removes a small mismeasured minority together with a majority of genuinely poor individuals drawn from across the entire parental distribution.

Table A15: Characteristics of observations excluded by each sample filter

	Parental false-poor	Child false-poor	Blind 5% child-rank trim
Individuals excluded	71,037	16,752	131,181
<i>Parental characteristics</i>			
Mean income rank (1998)	24.2	61.5	49.3
Mean gross-wealth rank (2016–22)	54.6	64.9	48.0
Median household income (1998)	€8,267	€29,017	€19,959
Median gross wealth	€101,096	€180,038	€77,834
<i>Child characteristics</i>			
Mean income rank (2021–22)	48.3	2.8	3.0
Mean gross-wealth rank (2016–22)	53.0	87.1	35.3
Median income	€20,312	€480	€798
Median gross wealth	€28,328	€84,658	€8,641
Share with income < €1,000	0.03	0.65	0.55

Notes: Each column describes the observations removed by one filter, relative to the sample to which the filter is applied. *Parental false-poor:* households in the bottom 10% of reported 1998 income whose imputed housing wealth lies in the top 40% nationally. *Child false-poor:* individuals in the bottom 5 percentiles of the within-cohort child income rank whose own gross wealth lies in the top 25% of the sample. *Blind 5% child-rank trim:* all individuals in the bottom 5 percentiles of the child rank, regardless of wealth (the paper’s baseline treatment). Income ranks are within-cohort percentiles; wealth ranks are computed over the full sample from gross wealth (real estate, deposits, financial and other assets, before liabilities) averaged over observed years 2016–2022. The two false-poor groups exhibit the same income–wealth dissonance in their respective generations: near-bottom reported income combined with wealth far above the median. The blindly trimmed group instead shows no such dissonance (median own wealth €8,641), indicating that the blind trim removes mostly genuinely poor individuals alongside the mismeasured minority.

H.4 Intergenerational mobility estimates across samples

Table A16 reports national mobility estimates for the baseline cohorts (1980–1986) under the four samples, and Figure A54 overlays the corresponding rank-rank relationships; Figures A55 and A56 zoom into the bottom and top of the parental distribution, and Figure A57 summarizes top-percentile persistence.

Three results stand out. First, absolute mobility is essentially invariant to the filtering choice. Predicted child rank at the 25th parental percentile ranges from 43.5 to 44.2 across all four samples, the predicted rank at the median is identical to one decimal (50.4), and the probability of moving from the bottom to the top quintile varies between 10.4% and 10.8%. The paper’s conclusions about absolute upward mobility do not depend on how the bottom of the child distribution is treated.

Second, the rank-rank slope is more sensitive, rising from 0.247 in sample (1) to 0.250 in sample (2), 0.259 in sample (3), and 0.274 in the baseline. Part of the gap between the baseline and the no-trim samples is mechanical rather than substantive: because ranks are recomputed after trimming, dropping the bottom 5%

stretches the surviving ranks from the interval $[5, 100]$ back to $[1, 100]$, rescaling the slope by a factor of approximately $100/95$. In the extended sample, this mechanical component accounts for 0.013 of the 0.024 gap between the no-trim and baseline slopes, with the remaining 0.011 attributable to the excluded observations themselves. The targeted filter in sample (3), which removes only 0.6% of observations, involves a negligible mechanical component.

Third, measures of persistence at the extremes move in the expected directions. Relative to sample (2), excluding child false-poor individuals raises the mean rank of children of top-percentile parents from 71.4 to 72.7, reflecting the removal of affluent-family children whose near-zero reported income is contradicted by their assets; the baseline trim raises it further to 74.1, and raises bottom-decile persistence from 0.136 to 0.155, because it also removes genuinely poor children who would otherwise populate the diagonal of the transition matrix.

Table A16: Intergenerational mobility estimates under alternative sample filters

Sample	N	Households	RRS	μ_{25}	μ_{50}	Quantile & Decile Transitions			Top persistence	
						Q1→Q1	Q1→Q5	P10→P10	$\bar{r} p100$	$\Pr(\geq p99)$
(1) Negative parental income excluded	1,712,274	1,335,643	0.247	44.2	50.4	0.269	0.108	0.135	71.3	0.186
(2) = (1) + parental false-poor excluded	1,664,156	1,297,446	0.250	44.1	50.4	0.270	0.105	0.136	71.4	0.186
(3) = (2) + child false-poor excluded	1,651,058	1,289,397	0.259	43.9	50.4	0.275	0.104	0.142	72.7	0.189
Baseline: (2) + bottom-5% child-rank trim	1,580,994	1,246,817	0.274	43.5	50.4	0.283	0.104	0.155	74.1	0.192

Notes: Unweighted pooled OLS estimates on within-cohort percentile ranks, baseline cohorts 1980–1986, child income averaged over 2021–2022 and parental household income over 1998–2000. Rows are cumulative: the first excludes only parental households with negative income in 1998–2000; the second additionally excludes *parental false-poor* households, whose reported 1998 income lies in the bottom 10% while their imputed housing wealth lies in the top 40% nationally; the third additionally excludes *child false-poor* individuals, whose income rank lies in the bottom 5 percentiles of the within-cohort distribution while their own gross wealth (real estate, deposits, and financial assets averaged over 2016–2022) lies in the top 25% of the sample—i.e., individuals whose near-zero reported income is contradicted by their own asset holdings. Percentile ranks are recomputed within cohort after each exclusion. The final row is the paper’s baseline, which replaces the child false-poor filter with a blind trim of the bottom 5% of the child rank distribution. RRS is the rank-rank slope; μ_{25} and μ_{50} are predicted child ranks at parental percentiles 25 and 50; Q1→Q1, Q1→Q5, and P10→P10 are quintile/decile transition probabilities; $\bar{r} | p100$ is the mean child rank among children of top-percentile parents, and $\Pr(\geq p99)$ their probability of reaching the top two percentiles. Standard errors (omitted) are below 0.001 for the RRS and 0.03 for μ_{25} in all rows.




Figure A54: Rank-rank relationship under alternative sample filters

Notes: Mean child income rank by parental income rank (100 bins), extended sample (1980–1990), unweighted, under the four samples of Appendix H.2. Lines are OLS fits; the annotation box reports the rank-rank slope and predicted rank at the 25th parental percentile for each sample. The dotted vertical line marks the 25th percentile.




Figure A55: Rank-rank relationship at the bottom of the parental distribution




Figure A56: Rank-rank relationship at the top of the parental distribution

Notes: Filters that remove low-income children raise measured top persistence, because a disproportionate share of near-zero-income individuals have parents in the upper part of the distribution.




Figure A57: Top-percentile persistence across sample filters

H.5 Additional robustness: trim depth and parental income windows

The comparison above varies the *type* of bottom filter. As a complementary exercise, Table A17 holds the baseline uniform-trim approach fixed and varies its *depth*, together with the window over which parental income is measured. Replacing the 5% trim with 3% or 2% trims moves the rank-rank slope by at most 0.007 and absolute upward mobility by at most 0.2 percentile points; replacing the three-year parental income average with a single year lowers the slope by 0.003 to 0.008, consistent with classical attenuation from measuring permanent income with one year of data. Together with Table A16, these results bracket the baseline estimate from both directions: intermediate trim depths yield estimates within rounding of the baseline, while the limiting no-trim case and the targeted wealth-based filter yield modestly lower slopes for the reasons documented above, with absolute mobility essentially unchanged throughout.

Table A17: National rank-rank slope and absolute upward mobility under alternative parent-year configurations

Configuration	N	RRS	SE	AUM	SE
Standard (parents 1998–2000, drop $\leq$ 5% kids)	1,580,993	0.274	0.0008	43.5	0.029
Parents 1998 only	1,580,654	0.271	0.0008	43.6	0.029
Parents 1999 only	1,523,134	0.266	0.0008	43.7	0.030
Parents 2000 only	1,545,595	0.266	0.0008	43.7	0.030
Standard but drop $\leq$ 3% kids	1,614,149	0.271	0.0008	43.6	0.029
Standard but drop $\leq$ 2% kids	1,630,809	0.267	0.0008	43.7	0.029

Notes: National-level rank-rank slope (RRS) and absolute upward mobility (AUM, predicted child rank at parent percentile 25) under alternative parent-year configurations and bottom-trim thresholds. The standard configuration is the canonical sample used throughout Section 3.6: parental income averaged across IRPF 1998–2000, child income averaged across 2021–2022, restricted to children born 1980–1986, with the bottom 5% of the within-cohort child rank distribution dropped, and excluding parents flagged as “false poor” by the housing-wealth proxy ($\text{VAR6} > 312.48$ within the bottom 10% of 1998 income) or with negative IRPF income in any of 1998–2000. Single-year parent configurations replace the 3-year mean with parental income from a single IRPF year, holding the kid set, child-side measurement, and exclusions fixed. The bottom two rows replace the 5% kid-rank trim with 3% and 2% trims respectively, holding the parent configuration at the 3-year mean. Standard errors are conventional OLS standard errors; the AUM standard error uses the linear combination $\text{Var}(\alpha + 25\beta)$. The single-year RRS estimates are 0.003–0.008 below the 3-year baseline, consistent with classical attenuation from measuring permanent income with one year of data. The looser-trim rows are within 0.007 of the standard; the limiting case of no trim, together with a narrower wealth-based alternative to the trim, is analyzed in Appendix H See Section 3.6 for the corresponding figure on attenuation across averaging windows from one to six years.

Finally, we verify that the paper’s headline claim about elite persistence is not an artifact of the bottom trim. Figure A58 recomputes the relative probability of reaching the top 1% of the child income distribution—for children of top-1% parents versus children of bottom-decile parents—replacing the baseline sample with sample (3). The two Spanish estimates are indistinguishable: excluding child false-poor individuals lowers both the numerator and the denominator proportionally, leaving the ratio unchanged at 51.4. Because sample (3) applies no uniform trim, its top percentile corresponds to the true top 1% of the linked sample, so the cross-country comparison is not affected by the rank-rescaling concerns discussed above. Spain’s position relative to the United States, Italy, and Canada is identical under either sample.




Figure A58: Relative probability of reaching the top 1%: robustness to the child false-poor filter

Notes: Each bar reports the ratio between the probability that a child of top-1% parents reaches the top 1% of their cohort's income distribution and the average probability that a child of bottom-decile parents does so. For Spain, both the baseline sample (bottom-5% child-rank trim) and sample (3) of Appendix H.2 (child false-poor filter) are shown; both use cohorts 1980–1986 and individual child income. Underlying values: in the baseline sample, 12.47% of children of top-1% parents reach the top 1%, against an average of 0.24% among children of bottom-decile parents (ratio 51.4); in sample (3), the corresponding figures are 12.36% and 0.24% (ratio 51.4). For the United States, 9.65% versus 0.20% (ratio 48.3); for Italy, 13.65% versus 0.36% (ratio 38.4); for Canada, 9.09% versus 0.44% (ratio 20.6). International estimates are drawn from published transition matrices: [Chetty et al. \(2014c\)](#) for the United States, [Acciari et al. \(2022\)](#) for Italy, and [Corak \(2020\)](#) for Canada; income concepts and child ages at measurement differ across studies, so cross-country magnitudes should be interpreted as indicative.

H.6 Interpretation and choice of baseline

The evidence in this appendix supports two conclusions. The bottom-rank anomaly is real and driven to a large extent by income mismeasurement among individuals from affluent families, which justifies some treatment of the bottom of the child distribution. At the same time, the child false-poor filter of sample (3) is, in our view, the conceptually preferable treatment: it removes the smallest number of observations and families of any filter considered (16,752 individuals in 8,027 households, against 131,181 individuals in 127,013 households under the uniform trim), it targets exclusively the individuals whose reported income is demonstrably contradicted by their own assets, it retains the genuinely poor children that the uniform trim discards, and it involves essentially no mechanical rank-rescaling.

An important residual ambiguity nevertheless remains. Even after excluding child false-poor individuals, a smaller spike in mean parental rank persists at the very first child percentile (Figure A52). The individuals behind it report near-zero income and hold no significant assets in Spain, yet come from comparatively affluent families. With the information available we cannot determine whether they represent genuine downward mobility from the upper part of the parental distribution, in which case they belong in the sample, or individuals whose income and assets are simply not observed in Spanish records, most plausibly emigrants of the post-2008 wave, in which case their measured rank is not informative. Because this group cannot be cleanly classified, we retain the 5% trim as the baseline throughout the paper and present sample (3) as the preferred robustness benchmark. As Table A16 shows, the choice is not consequential for the paper's conclusions: absolute mobility estimates are essentially identical across samples, and the baseline rank-rank slope (0.274) is modestly above the targeted filter estimate (0.259), with roughly half of the gap between the baseline and the unfiltered slope reflecting mechanical rank rescaling rather than substantive differences in measured persistence.

I Representativeness of the Targeted Income Distribution

This final part assesses whether the income records underlying our analysis are representative of the income distribution of the relevant Spanish population. The difficulty of the exercise is not measurement but population definition: the Atlas child records cover individuals who could be linked to a parental fiscal household in the common fiscal territory, and no external source reproduces exactly that population. We therefore use two benchmarks that depart from it in opposite directions, one a household survey and one a tabulation of the full administrative population, and read the resulting comparisons as bounds. We focus throughout on income year 2022, the most recent year for which child income is available in the Atlas.⁵⁵

⁵⁵Our main estimates measure child income as the average of 2021 and 2022. We validate the 2022 cross-section alone because the survey benchmark is a repeated cross-section rather than a panel, so a two-year average cannot be constructed for the same individuals. Averaging two adjacent years reduces idiosyncratic year-to-year variation but leaves the shape and central location of

The first benchmark is the Spanish Survey on Living Conditions (*Encuesta de Condiciones de Vida*, ECV), the standard nationally representative source for the individual income distribution in Spain and the Spanish component of EU-SILC, which reports the detailed individual labor-income components required to harmonize its income concepts with administrative tax records.⁵⁶ We compare it with the Atlas child records born between 1980 and 1986, the cohorts used in the main intergenerational mobility analysis. We first describe the two data sources and the population that each one covers, then construct a survey benchmark whose population is as close as possible to the population that the administrative records can contain, then define a harmonized labor income concept measurable in both sources, and finally compare the resulting distributions for the full set of linked records and for the estimation samples used in the paper. Section I.1 then repeats the comparison against the administrative population register.

Data sources and their populations. The two sources differ both in how they measure income and in whom they cover, and both differences matter for the comparison. The ECV is a household survey representative of the population residing in private dwellings in Spain. Its strengths are a well-defined target population and nationally calibrated person weights. Its limitations are those of any survey: a modest sample size for narrow birth cohorts, item and unit non-response, recall error, rounding, and imputation, all of which are most severe in the tails of the income distribution.⁵⁷ The Atlas records, by contrast, are administrative tax data in which the main labor income components are third-party reported, essentially eliminating recall and rounding error, and the number of observations is three orders of magnitude larger. Their limitation is the population frame: the records capture fiscal income, so that income earned informally or abroad appears as zero, and the child records used here are those of the parents-children linked Atlas panel rather than of the full Spanish birth cohort.

We compare the ECV with three administrative samples that correspond to successive stages of the sample construction used in the paper. The first, which we label all linked children, contains every individual born between 1980 and 1986 in the linked Atlas panel, that is, every child who could be identified in a parental fiscal household, with no further restriction; it includes filers and non-filers and comprises 1,867,822 individuals. This is the only sample intended to be representative of the linked population, and it is the natural benchmark for the representativeness exercise: any divergence at this stage reflects the coverage and measurement properties of the administrative source itself, before any analytical choice.⁵⁸

the cross-sectional distribution essentially unchanged, so the 2022 distribution is an accurate proxy for the object measured in the main analysis.

⁵⁶The 2023 survey wave collects income for the previous calendar year; throughout this appendix we therefore label the benchmark ECV 2022, referring to the income reference year.

⁵⁷For the 1980–1986 cohorts, the ECV contains 6,115 adult observations before any restriction and 4,602 in the comparable benchmark defined below, so that statistics at the 99th percentile rest on fewer than 50 observations and should be read with corresponding caution.

⁵⁸This sample conditions on linkage to a parental household but not on any income-based filter. The extent to which the linked panel covers the full birth cohort is documented separately in the coverage appendix; the present exercise evaluates the income distribution conditional on linkage. The two exercises are consistent: the 1.87 million linked children correspond to roughly 59 percent of the 3.18 million individuals represented by the comparable survey benchmark defined below.

The second and third samples add the filters that define the estimation sample of the paper, and they differ from each other only in how they treat the very bottom of the income distribution. Both start from the linked sample and first remove parental households with negative fiscal income in 1998–2000, together with parental *false-poor* households: households whose reported 1998 income places them near the bottom of the distribution but whose imputed housing wealth places them near the top, so that their low reported income is contradicted by their observed assets and most plausibly reflects mismeasurement rather than genuine poverty. The two samples then diverge in how they treat low-income children. The second sample, which excludes child false-poor observations, applies to children the same logic used for parents: it removes individuals whose within-cohort income rank lies at or below 5 but whose own gross wealth, computed from real estate, deposit, and financial asset holdings, places them in the top quarter of the wealth distribution, that is, children whose near-zero reported income is contradicted by substantial personal assets. This leaves 1,651,058 individuals. The third sample is the baseline of the paper, which instead removes the bottom 5 percent of the within-cohort child income-rank distribution outright, regardless of wealth, leaving 1,580,994 individuals. The exact thresholds and the order in which ranks are recomputed after each exclusion are described in Appendix A16.⁵⁹

The distinction between the two estimation samples is what the figures below make visible at the bottom of the distribution. The wealth-based filter removes only the small group of children whose reported poverty is inconsistent with their assets, so it prunes spurious zeros while retaining genuine non-earners; the blind trim removes the entire bottom of the rank distribution, and therefore discards genuinely low-income children alongside the mismeasured minority.

Construction of the comparable survey benchmark. The raw ECV represents residents of Spain regardless of origin or region, while the Atlas cannot contain two groups of residents by construction. First, the Basque Country and Navarre operate foral tax regimes whose records are not part of the state tax administration data, so their residents are absent from the Atlas. Second, individuals born abroad who arrived in Spain as adults cannot be linked to a parental fiscal household in 1998–2003 and therefore cannot enter the Atlas child sample. We construct a comparable ECV benchmark by removing both groups from the survey: households resident in the Basque Country or Navarre, identified through the regional variable of the household file, and individuals not born in Spain, identified through the country-of-birth variable of the person register. Table A18 reports the resulting population accounting. The foral restriction removes 5.2 percent of the weighted population of these cohorts, and the country-of-birth restriction removes a further 24.4 percentage points, leaving 70.4 percent of the original weighted population, or 3.18 million represented individuals. The size of the second restriction reflects the large weight of adult immigration in these cohorts and anticipates its quantitative importance for the comparison. It also makes this benchmark a conservative one: the

⁵⁹The false-poor exclusions are motivated and analyzed in detail in the filtering appendix. Here they serve only to trace how each successive filter moves the income distribution of the estimation sample.

country-of-birth restriction removes more of the immigrant population than the Atlas itself does, since the Atlas retains foreign-born individuals who arrived as children and appear as dependents in a Spanish fiscal household, a group representing less than 3 percent of these cohorts. The comparable ECV therefore describes a population slightly more native, and slightly better paid, than the population the Atlas can contain.

Table A18: Construction of the comparable ECV benchmark

Step	N	Weighted N	Share of row 1 (%)
ECV 2022, cohorts 1980–1986	6,115	4,518,576	100.0
Excluding Basque Country and Navarre	5,689	4,285,326	94.8
Excluding foreign-born (comparable benchmark)	4,602	3,180,754	70.4

Notes: Each row reports the unweighted number of observations and the population represented by the ECV cross-sectional person weights (PB040) for individuals born 1980–1986 in the ECV 2023 wave, whose income variables refer to calendar year 2022, after each successive restriction. The second row drops households resident in the Basque Country and Navarre (household file variable DB040 equal to ES21 or ES22), territories under foral tax regimes that are not covered by the Atlas administrative records. The third row additionally drops individuals not born in Spain (person register variable RB280 different from 1, including 11 observations with missing RB280), who cannot appear in the Atlas child sample because linkage requires presence in a parental fiscal household in 1998–2003; foreign-born individuals who arrived as children and appear as dependents in a Spanish fiscal household are retained in the Atlas and represent less than 3 percent of these cohorts. The resulting sample in the third row is the comparable ECV benchmark used throughout this appendix.

Income harmonization. We compare a single harmonized concept, individual gross labor income including self-employment, defined as the sum of employee earnings, retirement pensions, unemployment benefits, and self-employment income. This is the broadest labor-market concept that can be constructed consistently in both sources, and it is the closest, among the definitions available in the survey, to the total gross income measure used in the main analysis.⁶⁰ We do not attempt to validate capital income. In the ECV, capital income components are measured at the household rather than the individual level, are strongly affected by non-response and imputation, and are concentrated precisely in the sparse upper tail of the survey; no individual-level comparison of capital income between the two sources would be informative. Labor income, by contrast, accounts for the large majority of total gross income at the ages observed here, so the harmonized concept covers most of the income measure used in the paper.

Results. Table A19 reports summary statistics for the five samples, and Figure A59 displays the corresponding percentiles. Three findings emerge. First, the linked administrative records track the comparable survey distribution closely from the median upward. The median is 20,395 euros in the linked sample and 20,312 euros in the comparable ECV, a difference of 0.4 percent. The gap remains small across the upper half: 33,537 against 31,994 euros at the 75th percentile, a difference of 4.8 percent, 46,740 against 45,035 at the 90th, or 3.8 percent, and 105,273 against 103,940 at the 99th, or 1.3 percent. Panel A of Figure A59 shows the two distributions tracking each other across the reported percentiles.

⁶⁰In the Atlas ERM 2022 records, the concept sums salaries (M11), retirement pensions (M12), taxable and exempt unemployment benefits (M13 and M152), and self-employment income from economic activities (M4). In the ECV, it sums gross monetary and non-monetary employee income (PY010G, PY020G, PY021G), gross unemployment benefits (PY090G), gross old-age benefits (PY100G), and gross monetary self-employment income (PY050G). In both sources, component sums treat missing values as zero, and negative values, which arise from self-employment losses, are not censored.

Second, the comparison is sensitive to the population definition, not to the data source. Against the unrestricted ECV, whose median is 18,622 euros, the same linked sample would appear shifted upward by roughly 10 percent. The entire shift is accounted for by the two population restrictions, with the country-of-birth restriction doing most of the work: excluding the foreign-born raises the survey median from 18,622 to 20,312 euros, while excluding the foral territories operates in the opposite direction and is smaller. The apparent upward shift in the administrative data therefore reflected the presence in the survey of a large population that the administrative child sample cannot contain.

Third, the remaining divergence is confined to the bottom of the distribution. At the 25th percentile the linked sample records 10,698 euros against 11,897 in the comparable benchmark, 10.1 percent below it, and its share of exact zeros is higher, 9.4 against 7.4 percent. Only 0.4 percentage points of the administrative zeros correspond to individuals with no geographic record in 2022, the group that pools emigrants, deceased individuals, and individuals without any administrative trace in that year, so most of the excess is attributable to residents whose fiscal labor income is genuinely zero.⁶¹ The two sources measure genuinely different objects at this margin: part of this income exists but is informal and hence invisible to the tax records, while non-employed individuals are also the group most affected by survey non-response.


Estimation samples. Panel B of Figure A59 and the last two rows of Table A19 document how the estimation filters move the samples relative to the benchmark. Both operate on the bottom of the distribution and leave the top nearly unchanged. At the 25th percentile the sample excluding false-poor observations stands 7.8 percent above the benchmark and the baseline 18.3 percent above it, against 10.1 percent below for the linked sample; at the median the corresponding figures are 6.4 and 10.2 percent, and at the 75th percentile 8.5 and 10.6 percent. At the 99th percentile the three samples differ from the benchmark by 1.3, 4.1, and 5.7 percent respectively. The extensive margin shows the same ordering: the sample excluding false-poor observations retains a zero-income share of 4.8 percent, removing individuals whose reported zero income is contradicted by their own asset holdings while preserving a substantial share of genuine non-earners, whereas the baseline trim mechanically reduces it to 2.4 percent. These figures quantify the trade-off between the two treatments of the bottom documented in Appendix A16 and should be read jointly with it: the level differences reported here do not mechanically affect rank-based mobility statistics, which are computed on within-sample ranks, but they delimit the population for which absolute statements about the bottom of the distribution apply.⁶²

Against this benchmark the linked Atlas child records reproduce the survey distribution of labor income

⁶¹ A geographic record in 2022 is a non-missing municipality, province, or autonomous community code in the annual geographic panel of the administrative records. The indicator is informative but imperfect in both directions: individuals outside Spain may retain an address carried in the administrative data, while 0.5 percent of individuals with positive income also lack a geographic record.

⁶² A level shift common to all incomes leaves within-sample percentile ranks unchanged. Sample selection affects rank-based estimates only insofar as it alters the joint distribution of parental and child ranks; the sensitivity of the rank-rank slope and of absolute upward mobility to these same filters is reported in Appendix A16.

Figure A59: Labor income including self-employment: Atlas samples and the comparable ECV benchmark



Notes: Selected percentiles of gross annual labor income including self-employment for individuals born 1980–1986. Panel A compares the comparable ECV benchmark, which excludes the Basque Country, Navarre, and the foreign-born and is weighted with the survey person weights, with the full sample of linked Atlas children, which is unweighted. Panel B adds the two estimation samples described in Table A19. Percentiles are computed separately within each sample.

Table A19: Labor income including self-employment in Atlas 2022 and ECV 2022

Source	N	Weighted N	Mean	p25	p50	p75	p90	p95	p99	Zero (%)	>50k (%)
ECV 2022, unrestricted	6,115	4,518,576	22,259	10,400	18,622	29,228	43,192	57,013	95,654	9.6	6.8
ECV 2022, comparable population	4,602	3,180,754	24,046	11,897	20,312	31,994	45,035	59,333	103,940	7.4	7.5
Atlas, all linked children	1,867,822	–	24,593	10,698	20,395	33,537	46,740	60,749	105,273	9.4	8.4
Atlas, estimation sample excl. false poor	1,651,058	–	26,207	12,824	21,603	34,719	48,083	62,555	108,195	4.8	9.0
Atlas, baseline estimation sample (5% trim)	1,580,994	–	27,347	14,073	22,393	35,384	48,888	63,533	109,888	2.4	9.4

Notes: Income statistics for individuals born 1980–1986 under the labor income including self-employment concept of Table ??, defined as gross employee earnings plus retirement pensions, unemployment benefits, and self-employment income. ECV statistics refer to income year 2022 (survey wave 2023) and are weighted with the cross-sectional person weights; the weighted N reports the represented population. The comparable ECV population excludes households in the Basque Country and Navarre and individuals not born in Spain (Table A18). Atlas statistics are unweighted and computed from ERM 2022 administrative records. All linked children covers every record born 1980–1986 in the linked Atlas panel, i.e., children identified in parental fiscal households; it is not the full Spanish birth cohort. The estimation sample excluding false poor starts from the linked sample after excluding parental households with negative 1998–2000 income and parental false-poor households, and additionally excludes children whose within-cohort income rank is at or below 5 while their own gross-wealth rank, computed from real estate, deposit, and financial asset holdings averaged over 2016–2022, exceeds 75. The baseline estimation sample instead excludes the bottom 5 percent of the within-cohort child income-rank distribution. Percentiles are computed separately within each source and concept; zero and negative values are included.

closely from the median upward, with the divergence concentrated among zero and near-zero fiscal incomes at the bottom, where the two sources measure genuinely different objects. Because the benchmark removes more of the immigrant population than the Atlas itself does, this comparison should be read as a lower bound on the divergence between the linked panel and its target population. We now turn to a second benchmark that errs in the opposite direction.

I.1 Comparison with the Full Administrative Population

At our request, the tax agency produced a percentile tabulation of individual gross income in 2022 for the full population born between 1980 and 1990 resident in the common fiscal territory, covering filers and non-filers identified through information returns and excluding holders of foreign identification numbers. Because the tabulation comes from the same administrative source and the same underlying definitions as our own records, measurement is held fixed and any difference between the two is attributable to population composition alone. This is the cleanest available test of whether linkage to a parental fiscal household selects on income.

The benchmark is informative precisely where the survey benchmark is weak, and weak where the survey benchmark is informative. It contains no sampling error and covers 4,727,839 individuals, against which our linked child records for the same cohorts represent 62.3%. It excludes foreigners currently holding a foreign identification number, but it retains naturalized immigrants who arrived in Spain as adults and who cannot be linked to a parental fiscal household in our data. It therefore includes a population that the Atlas cannot contain, and unlike the survey benchmark it under-restricts rather than over-restricts.⁶³ Its treatment of individuals without any income record is not fully known: the first percentile threshold lies at 0.51 euros and

⁶³We match the concept by computing total gross income net of realized capital gains (M01 minus M5) in our own records, and we restrict all series to the 1980–1990 cohorts covered by the tabulation. Restricting our records to the 1980–1986 cohorts used in the main analysis raises the median by 4.2%, so the wider cohort range accounts for only a small part of the differences reported below and does not affect the ordering of the samples.

the second at 25.56 euros, implying that fewer than 2% of its records lie at or below zero, against 5.4% of exact zeros among linked children. We therefore report percentiles from the 25th upward and read the extensive margin separately.

Figure A60 and Table A20 report the comparison. The linked child records lie between 4 and 9% above the administrative population across the reported percentiles: 4.1% at the 25th, 6.9% at the median, 9.2% at the 75th, and 4.3% at the 99th. The two estimation samples lie further above, by 12.5 and 16.6% at the median respectively, and the gap narrows toward the top in every sample. The extensive margin displays the same ordering documented against the survey benchmark: the share with zero income falls from 5.4% among linked children to 1.3% once the false-poor exclusions are applied and to 0.5% under the baseline trim, against fewer than 2% in the population tabulation.

Table A20: Gross income in the Atlas child samples and in the administrative population, 1980–1990


Source	p10	p25	p50	p75	p90	p95	p99
Administrative population (AEAT)	4,140	11,501	19,571	31,019	44,964	58,612	102,558
Atlas, all linked children	2,201	11,977	20,920	33,861	47,727	62,384	106,964
Atlas, estimation sample excl. false poor	5,898	13,840	22,026	35,000	49,130	64,195	109,894
Atlas, baseline estimation sample (5% trim)	8,364	14,899	22,823	35,661	50,008	65,231	111,535

Notes: Individual gross income excluding capital gains in 2022 for individuals born 1980–1990 resident in the common fiscal territory. The administrative population row is a percentile tabulation of the full population provided by the tax agency, covering filers and non-filers identified through information returns and excluding holders of foreign identification numbers; reported values are the upper bounds of each percentile group. The Atlas rows are computed from the same administrative records for the linked child panel and for the two estimation samples, using the same income definition, defined as total gross income M01 minus net capital gains M5, and the same cohort range. All linked children covers every record born 1980–1990 in the linked Atlas panel, that is, children identified in parental fiscal households. The estimation sample excluding false poor removes parental households with negative 1998–2000 income, parental false-poor households, and children whose within-cohort income rank is at or below 5 while their own gross-wealth rank exceeds 75. The baseline estimation sample instead removes the bottom 5 percent of the within-cohort child income-rank distribution. Atlas statistics are unweighted. Zero and negative values are included.

Bracketing the divergence. The two benchmarks depart from the Atlas target population in opposite directions, and this is what makes them informative jointly. The comparable survey benchmark removes every foreign-born individual, including those who arrived as children and do appear in our records, so it describes a population slightly more native and better paid than the one the Atlas can contain; against it, the linked child records stand 0.4% above the median. The administrative population register removes only current holders of foreign identification numbers, so it retains naturalized immigrants who arrived as adults, earn less on average, and cannot be linked to a parental fiscal household; against it, the same records stand 6.9% above the median. The divergence between the linked panel and its true target population is therefore bounded between roughly 0–7% at the median, and by construction closer to the lower bound to the extent that naturalized adult immigrants are numerous in these cohorts.

Beyond the bound itself, the two comparisons agree on the shape of the divergence and on the ranking of the samples, which is the more relevant result. In both, the gap is largest below the median and narrows monotonically toward the top, a pattern consistent with a coverage margin operating on early emancipation

Figure A60: Gross income excluding capital gains: Atlas samples and the administrative tax population



Notes: Selected percentiles of individual gross income excluding capital gains in 2022, for individuals born 1980–1990 resident in the common fiscal territory. The population series is a percentile tabulation of the full administrative population provided by the tax agency; reported values are the upper bounds of each percentile group. The Atlas series are computed from the same administrative records, using total gross income net of realized capital gains, for the linked child panel and for the two estimation samples described above. All series include zero and negative values. Atlas statistics are unweighted.

rather than with any shortfall of high-income records, and inconsistent with an upward selection concentrated among high earners. In both, the estimation sample that excludes false-poor observations remains closer to the population at the extensive margin than the baseline trim, which removes genuine non-earners along with the mismeasured minority. Taken together, the two exercises support the use of the Atlas for the income distribution of the linked population, delimit the population to which absolute statements about the bottom of the distribution apply, and identify the treatment of the bottom as the sample-construction choice most worth revisiting.

J Calculation of the Top-Tail Relative Persistence Ratio (TTRPR)

We compute the TTRPR in three steps:

1. **Compute conditional probabilities for each parental percentile.** Let p_i denote the event that parents lie in the i th income percentile, and let $y \in \text{Top1}$ denote the event that the child reaches the top 1% of the adult income distribution. For each $i = 1, \dots, 100$ we calculate

$$q_i = P(y \in \text{Top1} \mid p_i) = \frac{\#\{\text{children with } p_i \text{ and } y \in \text{Top1}\}}{\#\{\text{children with } p_i\}} \times 100\%.$$

Numerically, for $i = 100$ (the top 1% of parents), we obtain

$$q_{100} = P(y \in \text{Top1} \mid p_{100}) \approx 12.47\%.$$

2. **Aggregate probability for the bottom decile.** To measure the probability of reaching the top 1% from the bottom 10% of parents, we average the ten conditional probabilities $q_1, \dots, q_{10}$:

$$\bar{q}_{1:10} = \frac{1}{10} \sum_{i=1}^{10} q_i \approx 0.24\%.$$

3. **Form the TTRPR as a ratio of conditional probabilities.** The Top-Tail Relative Persistence Ratio is defined as

$$\text{TTRPR} = \frac{P(y \in \text{Top1} \mid p_{100})}{P(y \in \text{Top1} \mid p \in \text{Bottom10})} = \frac{q_{100}}{\bar{q}_{1:10}} \approx \frac{12.47\%}{0.24\%} \approx 51.4.$$

This indicates that a child born into the top 1% is about 51 times more likely to attain the top 1% than a child born into the bottom 10%.

K Methodology of the Bayesian hierarchical model

K.1 Bayesian hierarchical model for local-level AUM and RRS

We estimate absolute upward mobility ($\mu_{25,g}$) and the rank-rank slope (β_g) at fine geographic resolution by fitting a Bayesian hierarchical linear model. Independent OLS regressions per locality produce noisy and often implausible estimates in small cells; the hierarchical structure allows information to be shared across geographies, stabilizing local estimates while leaving data-rich cells essentially unaffected.

We model the adult income rank of child i , y_i , as a linear function of their parental rank, p_i , with intercept and slope allowed to vary across geographies. For the municipality-level model, municipalities (indexed by m) are nested within provinces (indexed by r); the ZIP-level model substitutes ZIP codes (z) for municipalities under the same hierarchy. The full specification is:

$$y_i \sim \text{Normal}(\mu_i, \sigma), \quad (8)$$

$$\mu_i = (\alpha + \alpha_r + \alpha_m) + (\beta + \beta_r + \beta_m) p_i, \quad (9)$$

$$\alpha_r \sim \text{Normal}(0, \tau_{\alpha r}), \quad \alpha_m \sim \text{Normal}(0, \tau_{\alpha m}), \quad (10)$$

$$\beta_r \sim \text{Normal}(0, \tau_{\beta r}), \quad \beta_m \sim \text{Normal}(0, \tau_{\beta m}). \quad (11)$$

Here α and β are population-level parameters anchoring the regression line nationally; α_r, β_r and α_m, β_m are province- and municipality-level deviations, drawn from common Normal distributions whose standard deviations $\tau_{\alpha r}, \tau_{\alpha m}, \tau_{\beta r}, \tau_{\beta m}$ are themselves estimated from the data. The variance components and the residual standard deviation σ receive weakly informative Exponential(1) priors. We do not introduce an autonomous-community level: several autonomous communities in Spain coincide with a single province, which would prevent identification of separate variance components at the two levels.

The model is fit via Hamiltonian Monte Carlo using `Stan` via the `brms` interface, with four chains and standard convergence diagnostics.

After fitting, we recover local mobility measures from the posterior. For a municipality m within province r , the local rank-rank slope is

$$\beta_g = \hat{\beta} + \hat{\beta}_r + \hat{\beta}_m,$$

and absolute upward mobility (the predicted child rank when $p_i = 25$) is

$$\mu_{25,g} = (\hat{\alpha} + \hat{\alpha}_r + \hat{\alpha}_m) + 25 (\hat{\beta} + \hat{\beta}_r + \hat{\beta}_m),$$

where hatted quantities denote posterior means. The ZIP-level model produces analogous estimates with z

in place of m .

K.2 Hierarchical ordered logistic model for quintile transitions

To estimate the full 5×5 matrix of intergenerational quintile-to-quintile transition probabilities at fine geographic resolution — including the “rags-to-riches” probability $P(Y = Q_5 \mid X = Q_1)$, where Y and X denote the child’s and parent’s income quintiles — we fit a Bayesian hierarchical ordered logistic model. This specification respects the ordinal nature of the outcome and pools information across geographic units and parental quintiles, addressing the cell-sparsity problem that makes raw transition counts unstable.

The observed outcome for child i is $y_i \in \{1, \dots, K\}$ with $K = 5$, modelled as the discretization of a latent index $y_i^* = \phi_i$:

$$P(y_i = k \mid x_i, g_i) = P(c_{g_i, k-1} < y_i^* \leq c_{g_i, k}),$$

where g_i indexes the geographic group (municipality or province), x_i is the parent’s income quintile, and $c_{g_i, k}$ are group-specific cutpoints with $c_{g_i, 0} = -\infty$ and $c_{g_i, K} = +\infty$.

The latent index ϕ_i has three structural features:

Group-specific monotonic effect of parental income. The effect of the parent’s quintile x_i on the latent index has a flexible but monotonic shape, parameterized so that higher parental quintiles weakly raise the predicted child outcome. Letting $h_{g, x}$ denote the shape function at group g and parental quintile x , we set $h_{g, 1} = 0$, $h_{g, J} = 1$, and

$$h_{g, j} = \sum_{l=1}^{j-1} q_{g, l} \quad \text{for } j = 2, \dots, J-1,$$

where the increments $\mathbf{q}_g = (q_{g, 1}, \dots, q_{g, J-1})$ are derived from underlying parameters $\boldsymbol{\eta}_{q, g}$ via a softmax transformation. The contribution of the monotonic effect to the latent index is $\beta_g \cdot h_{g, x_i}$, where β_g is the group-specific magnitude of the gradient.

Predictor-level scaling (heteroskedasticity). The variance of the latent index — equivalently, the spacing between cutpoints relative to the predictor effect — may differ across parental quintiles. We allow this through a scaling factor δ_{x_i} that multiplies both the linear predictor and the cutpoints:

$$\phi_i' = \delta_{x_i} \cdot \beta_{g_i} \cdot h_{g_i, x_i}, \quad \tilde{c}_{g_i, k} = \delta_{x_i} \cdot c_{g_i, k}.$$

The scaling factor itself is monotonic across x_i , parameterized as $\delta_{x_i} = \exp(-\sigma_{\text{mono}, x_i})$ with the $\sigma_{\text{mono}, j}$ derived analogously from a softmax transformation of underlying parameters $\boldsymbol{\eta}_\sigma$.

Hierarchical priors on group-level parameters. The group-level parameters — shape $\eta_{q,g}$, magnitude β_g , and cutpoints $c_{g,k}$ — are assigned hierarchical priors that pool toward national-level means:

$$\begin{aligned}\eta_{q,g} &= \mu_q + \text{diag}(\tau_q) \zeta_{q,g}, \\ \beta_g &= \beta + \tau_\beta \zeta_{\beta,g}, \\ c_{g,1} &= \mu_{c1} + \tau_{c1} \zeta_{c1,g},\end{aligned}$$

with similarly structured non-centered parameterizations for the cutpoint log-gaps. Here β , μ_q , μ_{c1} are national-level means and τ_q , τ_β , τ_{c1} are between-group standard deviations; ζ are unit-Normal auxiliaries.

We fit the model in `Stan` via Hamiltonian Monte Carlo. Once the posterior is obtained, the predicted probability that a child from parental quintile x in group g ends up in quintile k is

$$P(y = k \mid x, g) = \Psi(\tilde{c}_{g,k} - \phi'(x, g)) - \Psi(\tilde{c}_{g,k-1} - \phi'(x, g)),$$

where Ψ is the logistic cumulative distribution function. Point estimates are posterior means; uncertainty is summarized by posterior credible intervals. The full $K \times J$ transition matrix is recovered by evaluating this expression for every (k, x) pair within each group.

L Validation of the Bayesian Hierarchical Estimator

We assess the gain in estimator accuracy from Bayesian partial pooling over cell-by-cell OLS using a Monte Carlo exercise on the IGM panel itself. Let g index ZIP codes and i index children. For each ZIP code with at least $N_{\min} = 5,000$ children in our cohort window (25 such ZIPs), we treat the OLS rank-rank fit

$$y_i = \tilde{\alpha}_g + \tilde{\beta}_g p_i + e_i$$

on the *full* within-ZIP sample as the local ground truth, and define the ground-truth absolute upward mobility as $\tilde{\mu}_{25,g} = \tilde{\alpha}_g + 25 \tilde{\beta}_g$. For computational tractability and to keep the joint Bayesian fit at a manageable size, we restrict the simulation to the $G_{\text{truth}} = 10$ most-populous ZIP codes among those 25; varying this restriction (results not shown) leaves the qualitative pattern unchanged.

To simulate the data scarcity that small ZIPs and municipalities face in practice, we then draw $R = 100$ independent random subsamples of size n from each ground-truth ZIP, where $n \in \{25, 50, 100, 250, 500, 1,000\}$. For each subsample we refit three estimators:

1. **OLS without pooling.** A separate OLS rank-rank regression $y_i = \alpha_g + \beta_g p_i + \varepsilon_i$ fit within each ZIP.

This is the standard cell-by-cell approach used in much of the geographic-mobility literature.

2. **Bayesian partial pooling, ZIP-level only.** A hierarchical model with random intercepts and slopes by ZIP:

$$y_i \sim \text{Normal}(\mu_i, \sigma), \quad \mu_i = (\alpha + \alpha_g) + (\beta + \beta_g) p_i,$$

with $\alpha_g \sim \text{Normal}(0, \tau_\alpha)$, $\beta_g \sim \text{Normal}(0, \tau_\beta)$, and weakly informative priors on the population-level parameters and variance components.

3. **Full Bayesian hierarchy (province + ZIP).** As (2), with an additional province-level random intercept and slope α_r, β_r nested above the ZIP-level effects.

For each replication we record point estimates of α_g , β_g , and $\mu_{25,g}$ (we use posterior medians for the Bayesian estimators), and compute squared errors against the local ground truth. The reported RMSE in Figure A61 is

$$\text{RMSE}(n, m) = \sqrt{\frac{1}{R G_{\text{truth}}} \sum_{r=1}^R \sum_{g=1}^{G_{\text{truth}}} (\hat{\mu}_{25,g}^{(r)}(n, m) - \tilde{\mu}_{25,g})^2},$$


where $m \in \{\text{OLS}, \text{Bayes-ZIP}, \text{Bayes-Province}\}$.

Implementation. Bayesian models are estimated using `brms` and Stan via Hamiltonian Monte Carlo (4 chains, 1,000 iterations each). Subsamples are drawn without replacement within each ZIP code, and the same set of ground-truth ZIP codes is used across all replications.

Per-parameter results. Figure A62 reports the corresponding RMSE for each of the five quantities the rank-rank specification produces: the intercept α_g , the slope β_g , and three points along the predicted child-rank curve, $\mu_{25,g}$, $\mu_{50,g}$, and $\mu_{100,g}$. Two patterns generalize across all five parameters. First, the Bayesian estimators dominate OLS at every sample size below $n \approx 500$, and the gap is largest for the slope — the parameter where small-sample noise produces the most implausible OLS values (negative or near-zero slopes in cells where the population slope is around 0.3). Second, the gain from adding the province level over a flat ZIP hierarchy is concentrated in the intercept and in $\mu_{25,g}$, where regional baselines vary most and pooling toward a province-specific mean is more informative than pooling toward the national mean.


Bias-variance decomposition. Figure A63 reports the standard $\text{MSE} = \text{Bias}^2 + \text{Variance}$ decomposition for $\hat{\mu}_{25,g}$ at each sample size and method. The decomposition makes the trade-off explicit: OLS has near-zero bias at every sample size but very large variance at small n (> 80 percentile-points-squared at $n = 25$); the Bayesian estimators accept a small but non-zero bias from shrinkage in exchange for a 5- to 10-fold variance reduction at small n . At large n the Bayesian bias becomes the binding constraint, and OLS variance shrinks

Figure A61: Monte Carlo RMSE of $\hat{\mu}_{25,g}$ as a function of within-ZIP sample size



Notes: For each of $G_{\text{truth}} = 10$ ZIP codes with at least 5,000 children, we treat the OLS rank-rank fit on the full ZIP sample as the local ground truth and compute $\tilde{\mu}_{25,g} = \tilde{\alpha}_g + 25 \tilde{\beta}_g$. We then draw $R = 100$ random subsamples of size n from each ZIP, refit each of the three estimators on the subsample, and compute the squared error of $\hat{\mu}_{25,g}$ against the local truth. The vertical axis reports $\sqrt{\text{mean squared error}}$ averaged across replications and ZIPs. See Table A21 for the underlying numerical values and Figure A62 for the corresponding per-parameter breakdown.


Figure A62: Per-parameter RMSE: OLS versus Bayesian estimators



Notes: Each panel reports $\sqrt{\text{mean squared error}}$ of the labeled parameter relative to the within-ZIP OLS-on-full-sample ground truth, averaged across replications and the G_{truth} ground-truth ZIP codes. Same data, same estimators, and same draw scheme as Figure A61; differs only in the parameter being evaluated.

until it modestly outperforms the Bayesian estimators in net MSE — although the practical relevance of this regime is limited: as we document in Section 4.2, the vast majority of Spanish municipalities and ZIP codes never reach the sample sizes where OLS catches up.

Figure A63: Bias-variance decomposition of $\text{MSE}(\hat{\mu}_{25,g})$ by method and sample size



Notes: For each (method, sample size) cell we compute the per-ZIP bias of $\hat{\mu}_{25,g}$ relative to the local ground truth across the R replications, square it, and average across the G_{truth} ground-truth ZIP codes; the variance term is the corresponding average within-ZIP variance of $\hat{\mu}_{25,g}$. The two components sum to the mean squared error reported (as $\sqrt{\text{MSE}}$) in Figure A61.

Table A21: Monte Carlo RMSE of $\hat{\mu}_{25,g}$ by sample size and method

Sample size n	25	50	100	250	500	1,000
OLS — no pooling	9.14	6.59	4.26	2.97	1.93	1.32
Bayes — ZIP only	4.26	3.77	2.94	2.21	1.68	1.24
Bayes — Province + ZIP	4.10	3.57	2.63	1.99	1.58	1.21

Notes: Entries are $\sqrt{\text{mean squared error}}$ of $\hat{\mu}_{25,g}$ (in percentile points), averaged across $R = 100$ replications and $G_{\text{truth}} = 10$ ground-truth ZIP codes with at least 5,000 children each. Posterior medians are used as point estimates for the Bayesian estimators.

L.1 Aggregate-OLS validation of municipal AUM estimates

The Monte Carlo exercise above (Figure A61) establishes that Bayesian partial pooling delivers substantial RMSE gains over cell-by-cell OLS at the sample sizes faced by the vast majority of Spanish municipalities. A natural further question concerns the geographic pattern itself: by partially pooling municipalities toward their province mean, the hierarchical model could in principle either manufacture cross-province differences in μ_{25} that a well-powered OLS would not see, or, conversely, erase local differences that are genuinely in the data. We address this concern directly by comparing the Bayesian AUM estimates against an aggregate-OLS estimator at the unit on which the question is sharpest: the cross-province border.

Aggregate-OLS border gaps. For every adjacent province pair (A, B) , we define a *border municipality* as one whose centroid lies within a threshold distance — we report 20, 30, and 40 km — of at least one municipality in the other province. We then pool the children of all side- A border municipalities and all side- B border municipalities into a single OLS rank-rank regression with a side- A interaction,

$$y_i = \alpha + \beta p_i + \gamma \mathbf{1}[i \in A] + \delta p_i \cdot \mathbf{1}[i \in A] + \varepsilon_i, \quad (12)$$


so that the aggregate-OLS border gap $\hat{\mu}_{25,A}^{\text{OLS}} - \hat{\mu}_{25,B}^{\text{OLS}} = \gamma + 25\delta$ is identified by a single regression on thousands of children per side — median 2,600 per side at the 30 km threshold — with a proper joint standard error from the regression variance-covariance matrix. The corresponding Bayes border gap is the n -weighted average of the municipal Bayes AUMs over the same side- A town set minus the same average over side- B . Both estimators see the same children and aggregate over the same town sets; they differ only in how the within-town estimation handles partial pooling.

Figure A64 reports the result across all adjacent province pairs that survive a minimum-sample filter (≥ 500 children and ≥ 3 municipalities on each side), at the three border-town thresholds. The pattern is robust: at all three thresholds the Bayes and OLS border gaps correlate at $r \in [0.84, 0.88]$, with a regression slope of 0.66–0.71. Both estimators identify the same provinces as having sharp borders — the largest well-powered case, the Madrid–Toledo border, gives $\widehat{\text{Gap}}^{\text{OLS}} = 3.70$ and $\widehat{\text{Gap}}^{\text{Bayes}} = 3.66$ pp on 54,000 Madrid-side children — and the same provinces as essentially indistinguishable across the line (Madrid–Ávila: -0.19 vs. 0.84). Where the two estimators differ, the Bayesian estimator is consistently the conservative version (slope < 1): partial pooling slightly attenuates aggregate border gaps relative to a well-powered OLS, never amplifies them. The mean absolute Bayes–OLS discrepancy is 0.81 pp at the 30 km threshold, and the paired t-test rejects equality ($p < 0.01$ at every threshold), but the direction of disagreement is the safe one for any inferential claim about geographic mobility: the map in Figure 9(b) is an *understatement* of the spatial heterogeneity that aggregate OLS reveals, not an overstatement.

Shrinkage is driven by sample size. The small aggregate attenuation in Figure A64 has a transparent source: the deviation between the Bayesian and OLS municipal estimates is essentially governed by municipal sample size. Figure A65 makes this explicit. Panel (a) plots the Bayesian AUM against the raw unpooled OLS estimate for all 2,991 municipalities with at least 25 children, with each point coloured by $\log_{10} n$: well-sampled municipalities sit on the identity line, while sparsely-sampled municipalities are pulled toward the higher-level mean (the dotted horizontal at 44.7 pp, the n -weighted national mean). Panel (b) plots the signed shrinkage $\hat{\mu}_{25,g}^{\text{Bayes}} - \hat{\mu}_{25,g}^{\text{OLS}}$ against n on a log scale; the cloud collapses toward zero as n grows, with the empirical 90th-percentile envelope (orange long-dashed) tracing the characteristic $1/n$ shape of Bayesian shrinkage. Quantitatively, the median $|\hat{\mu}_{25,g}^{\text{Bayes}} - \hat{\mu}_{25,g}^{\text{OLS}}|$ decreases monotonically from 3.66 pp in the smallest-


sample decile ($n = 25\text{--}31$) to 0.38 pp in the largest ($n \geq 680$), and the Bayes–OLS correlation rises from 0.64 to 0.99 over the same range. The slight slope-below-unity in Figure A64 is the cumulative consequence: averaging the Bayesian municipal estimates across border municipalities of mixed sample sizes preserves the cross-pair pattern (high correlation) but inherits a small amount of shrinkage from the sparsely-sampled towns within the aggregate.

Figure A64: Aggregate-OLS vs. Bayes border gap, robustness across border-town distance thresholds



Notes: Each point is an adjacent province pair. The horizontal axis is the aggregate-OLS border gap from a single rank-rank regression on the pooled children of side-A and side-B border municipalities with a side-A interaction (equation (12); gap = $\gamma + 25\delta$ with joint SE from the regression vcov). The vertical axis is the n -weighted mean of the municipal Bayes AUMs over the same side-A town set minus the corresponding mean over side-B. Border municipalities are defined as those with a centroid in the other province within the indicated threshold distance. Grey error bars are ± 1.96 SE on each axis; dashed line is the identity; the blue line is the OLS fit (slope and r reported in the facet label) with a 95% confidence band. Sample-size filter: at least 500 children and at least 3 municipalities on each side.

Figure A65: Bayesian shrinkage at the municipal level is driven by sample size



Notes: Both panels include all 2,991 municipalities with at least 25 children in the cohort sample. Panel (a): hierarchical Bayes AUM (μ_{25}) against the raw unpooled OLS estimate from a within-municipality rank-rank regression, each point coloured by the log of the municipal child count. Dashed line is the identity; dotted horizontal is the n -weighted national mean AUM (44.7 pp). Panel (b): signed shrinkage $\hat{\mu}_{25,g}^{\text{Bayes}} - \hat{\mu}_{25,g}^{\text{OLS}}$ against the municipal child count on a log scale, with orange long-dashed lines tracing the empirical 90th-percentile envelope across n -quantile bins.

M Definition of Urban Areas

Definition and Coverage of Urban Areas

Our analysis at the "urban area" level utilizes the definition of Functional Urban Areas (Áreas Urbanas Funcionales, AUF) established by the Spanish Ministry of Transport, Mobility and Urban Agenda (MITMA) in its 2022 methodology ([Ministerio de Transportes, Movilidad y Agenda Urbana, 2022](#)). This approach, harmonized with Eurostat guidelines, defines an AUF as a densely populated core city (or cities) along with its surrounding commuting zone, reflecting functional economic integration.

Delineation Criteria for Functional Urban Areas

The MITMA ([Ministerio de Transportes, Movilidad y Agenda Urbana, 2022](#)) methodology employs a multi-layered set of criteria combining demographic, economic, and spatial information:

- Large Urban Areas (Grandes Áreas Urbanas): These generally require a principal municipality with at least 50,000 inhabitants. Exceptions are made for provincial capitals like Soria and Teruel, which are included despite smaller populations due to their central role in their respective provinces.
- Constituent Municipalities in Plurimunicipal Urban Areas: Each municipality within such an AUF must typically have a minimum of 1,000 inhabitants, though exceptions are made based on regional or functional significance.
- Data Sources for Delimitation: The classification primarily draws on recent municipal-level data, including the 2011 Population and Housing Census, the 2021 Active Population Survey (EPA), and the 2021 Population Nomenclator (*Padrón Continuo*).
- Smaller Urban Areas: For urban areas not classified as Large Urban Areas and centered around municipalities with populations between 5,000 and 20,000, a four-step filtering process is applied:
 1. Assessment of the "nuclear" population (core population minus dispersed inhabitants).
 2. Exclusion of municipalities with negative demographic evolution between 1960 and 2021.
 3. Application of an economic filter based on the share of employment in the service sector, an indicator of urbanization.
 4. Use of a "potential accommodation" indicator, which combines secondary and vacant dwellings (adjusted by a national occupancy index) with the core population.

This approach aims to ensure that the delineated urban areas reflect both physical density and functional economic integration. According to [Ministerio de Transportes, Movilidad y Agenda Urbana \(2022\)](#), this method-

ology classifies 1,084 municipalities as urban, representing 13.3% of all Spanish municipalities. These urban municipalities are home to 82.5% of the national population but occupy only 20.7% of Spain's territory.

Population Coverage of Urban Areas in Our Analysis

We exclude the Basque Country and Navarre from our geographic mobility analyses since we do not have tax records from those regions. Consequently, for calculations involving urban area population coverage, we adjust both the numerator (population in urban areas) and the denominator (relevant national population). The official [Ministerio de Transportes, Movilidad y Agenda Urbana \(2022\)](#) definition includes four main urban areas located entirely within the Basque Country or Navarre: Bilbao, Donostia–San Sebastián, Vitoria-Gasteiz (all in the Basque Country), and Pamplona (Navarre). (The Irún–Hondarribia area, also in the Basque Country, is typically considered part of the San Sebastián AUF or as a minor urban area and is also excluded by extension). After excluding these four AUFs, our analysis considers the remaining 78 official urban areas. The total population of these included urban areas is approximately 37.1 million inhabitants (based on 2022 population data). The corresponding national population for Spain, excluding residents of the Basque Country and Navarre, is approximately 44.5 million (2022 data). Therefore, in the context of our study, approximately 83.3% of Spain's population (excluding the Basque Country and Navarre) resides within these 78 defined urban areas.

N Classification of Municipality-Level Characteristics

This appendix presents the classification of the municipal-level correlates used in our geographical analysis of intergenerational mobility. The variables span a wide range of dimensions that capture the socioeconomic environment in which children grow up. They are grouped into nine broad thematic categories based on their conceptual relevance: Demographics and Household Structure, Education, Housing, Income & Wealth, Labor Market Conditions, Migration, Sectoral Composition, Social Capital, Segregation and Commuting, and Urban Amenities.

The correlates are drawn from three types of sources. The first and most comprehensive is the 2001 Population and Housing Census (*Censo de Población y Viviendas*, INE), which provides detailed municipality-level information on educational attainment, occupational structure, household composition, housing tenure, and demographic characteristics. The second is the *Atlas Digital de las Áreas Urbanas*, compiled by the Ministry of Transport, Mobility and Urban Agenda (MITMA), which contains indicators on labor market conditions, commuting patterns, land use, and urban infrastructure measured between 2001 and 2007. The third source is municipal income and inequality data from the Spanish Tax Agency (AEAT). The vast majority of variables are measured between 2000 and 2007, coinciding with the childhood and adolescent years of the cohorts in our sample.

In addition to these administrative sources, we incorporate a novel measure of social capital derived from anonymized Facebook friendship data: economic connectedness (EC), defined as the share of friendships that cross socioeconomic lines within a community (Chetty et al., 2022a,b). We use the cross-country data release from Johnston et al. (2026), which provides EC estimates at the ZIP-code level for Spain. We aggregate these to the municipality level by averaging across ZIP codes within each municipality (see Appendix ?? for details on aggregation and coverage). This measure is classified under the Social Capital category because it captures a dimension of local social structure—bridging ties across socioeconomic groups—that is conceptually distinct from the labor market, human capital, and income variables available in administrative records. While the Facebook-derived EC data are measured in a more recent period than the Census and MITMA variables, Chetty et al. (2022a) show that the cross-sectional distribution of economic connectedness is highly persistent over time, reflecting stable features of community social structure rather than transitory fluctuations. Below, we describe each thematic category and the types of indicators included.

Demographics and Household Structure. This category captures the age composition and dependency burden of each municipality. Variables include the child dependency index, youth share of the working-age population, and youth dependency ratios. These measures characterize the demographic structure that children grow up in, indicating the potential demand on educational and social services as well as the density of peer networks. This category also includes measures of household composition and the prevalence of multi-generational or large families, such as the average number of children per household and the share of large family households. These indicators provide insight into caregiving structures, economic strain, and the intra-household distribution of resources that may influence children’s development (Van Bavel et al., 2011).

Education. Educational indicators reflect the prevailing level of human capital and access to learning opportunities within a municipality. Variables include the average education level among young adults (ages 30–39), post-compulsory education attainment, and the illiteracy rate. These metrics describe the skill profile of local populations and shape the educational environment that children are exposed to during their formative years, including peer effects and aspiration channels (Moretti, 2004a).

Housing. Housing variables describe both the quality and accessibility of the residential environment. This includes the total appraised value of housing, renter shares by age group, the proportion of households owning a second home, and the incidence of housing quality deficiencies. These indicators proxy for material living conditions, housing market segmentation, and the degree of residential stability, which may affect neighborhood cohesion and child outcomes.

Income & Wealth. This category summarizes the economic resources available to families and the broader local economy. Variables include average household income, the number of vehicles per household, and municipal-level Gini coefficients. These indicators capture material living standards, financial vulnerability, and the local distribution of economic resources.

Labor Market Conditions. These variables describe the strength, segmentation, and inclusivity of local labor markets. Indicators include the unemployment rate (general and youth-specific), the share of temporary contracts, and the share of households with at least one unemployed member. These metrics capture the economic opportunities available in the community, exposure to job insecurity, and the likelihood that families face persistent or cyclical income shocks.

Migration. Migration indicators reflect the composition and diversity of the municipality's population. This includes the total foreign-born population and the foreign-born share among young adults. These measures speak to patterns of geographic mobility and the social environments children may encounter in terms of linguistic, ethnic, and socioeconomic diversity.

Sectoral Composition. Sectoral variables provide information on the occupational and industrial landscape of a municipality. Indicators include the share of employment in agriculture, industry, construction, and services, as well as average occupational prestige scores. These factors influence local labor-market signaling, aspirations, and access to professional networks that may matter for children's long-term prospects.

Social Capital. This category contains economic connectedness (EC), the share of friendships between low- and high-SES individuals within a community. EC is derived from anonymized Facebook friendship networks and measures a specific form of bridging social capital: the extent to which individuals from different socioeconomic backgrounds interact socially ([Chetty et al., 2022a,b](#)). Unlike the other categories, which rely on administrative records from the early 2000s, EC is constructed from more recent social network data. We classify it separately from the income and labor market categories because it captures social interaction patterns that are conceptually orthogonal to the level or distribution of local economic resources.

Segregation and Commuting. This category captures how residents interact spatially with their environment, especially in relation to employment and services. Variables include commuting modes such as reliance on private vehicles, motorcycles, or public transit. These indicators proxy for geographic accessibility, urban infrastructure, and spatial isolation, all of which shape exposure to opportunity and the cost of access to resources.

Urban Amenities. Urban amenity variables describe environmental and public space characteristics within municipalities. These include the availability of artificial green spaces, urban tree coverage, and indicators of urban quality. These measures reflect municipal investment in the built environment, which may affect both children's physical well-being and quality of life.

Table A22: Municipality-Level Covariates Used in the Correlational Analysis

Variable	Year	Description	Source
Social Capital			
Economic connectedness	2022	Average SES rank of friends for individuals at the 25th percentile of the national SES distribution, derived from anonymized Facebook networks	Facebook ^a
Labor Market Conditions			
Households with unemployment	2001	Share of households with at least one unemployed member	MITMA
Unemployment rate (20–59)	2001	Unemployment rate among population aged 20–59	INE
Youth unemployment (<25)	2006	Share of unemployed under 25 among all unemployed	SEPE
Temporary contracts	2006	Share of temporary contracts among all contracts	SEPE
Education			
Avg. education, age 30–39	2001	Average educational attainment index among adults aged 30–39	INE
Illiteracy rate	2001	Share of illiterate population aged 16+	INE
Post-compulsory education	2001	Share of population with secondary or tertiary education	INE
Demographics and Household Structure			
Youth dependency	2001	Number of children (0–14) per 100 adults of working age	MITMA
Average household size	2001	Mean number of members per household	INE
Avg. children per household	2001	Mean number of children per family nucleus	INE
Single-parent households	2001	Share of single-parent family nuclei	INE
Sectoral Composition			
Agricultural employment	2007	Share of Social Security affiliates in agriculture	TGSS
Industrial employment	2007	Share of Social Security affiliates in industry	TGSS
Construction employment	2007	Share of Social Security affiliates in construction	TGSS
Service employment	2007	Share of Social Security affiliates in services	TGSS
Segregation and Commuting			
Car passenger commuting	2001	Share of workers commuting as car passengers	INE
Train commuting	2001	Share of workers commuting by train	INE
Work in same municipality	2001	Share of workers employed in their municipality of residence	INE
Income & Wealth			
Mean income (2000)	2000	Average household income from tax records	AEAT
Avg. cadastral value	2007	Mean appraised parcel value (thousands €)	Catastro
Vehicles per household	2001	Average number of vehicles per household	INE
Municipal Gini	2000	Gini coefficient of household income	AEAT
Migration			
Foreign-born population	2001	Share of foreign-born residents in total population	INE
Foreign-born, age 25–34	2001	Share of foreign-born among population aged 25–34	INE
Urban Amenities			
Green urban areas	2005	Share of artificial green zones and urban tree cover in total land area	SIOSE
Parks and urban green spaces	2005	Share of parks and urban green areas in total land area	SIOSE
Population density	2001	Residents per km ²	INE

Notes: This table lists the main municipality-level covariates used in the correlational analysis reported in Section 4.3. All variables are standardized to mean zero and unit standard deviation prior to estimation. Sources: INE = National Statistics Institute (2001 Census); MITMA = Ministry of Transport, Mobility and Urban Agenda (*Atlas Digital de las Áreas Urbanas*); SEPE = Public Employment Service; TGSS = General Treasury of Social Security; AEAT = Spanish Tax Agency; Catastro = General Directorate of the Land Registry; SIOSE = Land Cover Information System of Spain.

^a Economic connectedness is constructed from Johnston et al. (2026), who derive the measure from anonymized Facebook friendship networks. Available at the ZIP-code level; aggregated to municipalities by averaging across ZIP codes.

O Methodological Details of the Correlational Analysis

This appendix describes the methodology underlying the correlational analysis reported in Section 4.3. We compute two types of correlations—unweighted and population-weighted—for each combination of geographic correlate and mobility outcome. We consider two mobility outcomes: (i) the direct OLS estimate of absolute upward mobility (AUM), which are used for the baseline analysis, defined as the mean child income percentile for children whose parents are at the 25th percentile, following Chetty et al. (2014b); and (ii) its Bayesian counterpart.

Unweighted correlations. These are estimated via univariate OLS regressions of standardized AUM on each standardized covariate. This is algebraically equivalent to computing the Pearson correlation coefficient. Standard errors are obtained using robust (HC1) inference.

Population-weighted correlations. To account for varying population sizes across municipalities, we also compute population-weighted Pearson correlations using 2001 municipal population as weights. Because closed-form standard errors are not available for weighted correlations, we estimate uncertainty via a nonparametric bootstrap with 500 replications.

Sample construction and outlier filtering. For OLS-based AUM, we restrict the sample to municipalities with at least 20 observed children and apply a multi-step filtering procedure to remove extreme outliers. First, we exclude municipalities where the rank-rank slope lies outside the theoretical $[0, 1]$ range, removing 304 of 3,557 municipalities (8.55%). Second, we apply interquartile-range (IQR) filtering to the rank-rank slope, excluding 139 additional outliers (4.27%). Third, we apply IQR filtering to AUM itself, removing 30 further municipalities (0.96%). Finally, we drop 2 municipalities where the estimated rags-to-riches probability equals 1. After filtering, the OLS sample comprises 3,075 municipalities (86.45% of the original set). For Bayesian estimates, we include all municipalities with at least one observed child without filtering (5,876 municipalities).

Economic connectedness sample. The economic connectedness (EC) measure from Johnston et al. (2026) is available at the ZIP-code level for Spain. We aggregate to the municipality level by averaging EC across ZIP codes within each municipality. After merging with the OLS mobility sample and applying the outlier filters described above, the joint EC–mobility sample comprises 1,807 municipalities. This subsample covers the vast majority of medium and large population centers but excludes small rural municipalities for which EC data are unavailable or where mobility estimates are too noisy to be informative. The EC regressions reported in Section 4.3 are estimated on this subsample. Coverage details are reported in Appendix Figure A41.

Standardization. All correlates and mobility outcomes are standardized to have mean zero and unit standard deviation prior to estimation, so that regression coefficients can be interpreted as standardized associations. Municipal income is log-transformed before standardization.

P The Bayesian movers design specification: data construction, estimation, and validation

This appendix documents the Bayesian movers design of Section 5. Section P.1 describes the construction and auditing of the residential histories; Section P.2 the joint estimation model; Section P.3 the estimation sample; Section P.4 the measurement problem that motivates joint estimation and its design-based validation; Section P.5 simulation evidence; Section P.6 robustness checks; and Section P.7 the consistency of the model’s place-quality estimates with the geographic estimates of Section 4.2. Section P.8 details the construction of the income-gain estimates of Section 5.3, and Section P.9 collects the international comparison.

P.1 Residential histories and the unified movers panel

Sources. Three sources jointly cover parents’ residential locations over 1998–2022. (i) The 1998 income-tax file records each filer’s ZIP code, which we map to municipalities; this anchors every family’s *origin*. (ii) The fiscal data record each individual’s ZIP code annually from 2005 onward — the same location variable used for the descriptive geography of Section 4.2. (iii) For the intervening years 1999–2006, we recover locations from a municipal code embedded in each year’s tax record itself. Although nominally an administrative routing field, this code is in practice a municipality identifier: each year carries roughly 7,500 distinct values, 98.6% of which map to exactly one municipality, with the residual ambiguity confined to large cities that carry several sub-municipal codes. Movers over 1999–2005 (otherwise invisible) are therefore recoverable at municipal resolution.

Classification rules. We classify each parent’s full 1998–2022 history under a single set of rules, applied identically across the three sources: the origin is the first observed municipality; a move is a change of municipality followed by at least three consecutive years of stability at the new location; a mover is dated by her first valid move, and a permanent resident records no valid move. Administrative renumberings of the municipal codes (which would otherwise masquerade as moves) are detected and filtered with two tests: flows in which a code’s population transfers essentially en bloc to a successor code are merged as relabelings, and detected changes whose co-movers predominantly did *not* change municipality between the 1998 and 2005 endpoint measurements are classified as administrative and discarded for dating. Households whose records cannot

support dating (primarily families whose tax records remained stale after a mid-window move) are excluded rather than guessed at.

Cross-source audit. Because the tax-record municipal codes (1998–2006) and the annual ZIP codes (2005–2022) overlap in 2005–2006, every component of the construction can be audited against an independent source. Four findings summarize the audit. First, the two sources measure the same signal: among permanent residents in municipalities where the code-to-municipality mapping is unambiguous, the tax-record code agrees with the ZIP-code municipality for 99.8% of 1.1 million households, and 95.5% of the residual disagreements are sub-municipal code variation within large cities rather than location errors. Second, the two sources run on clocks offset by about one year: tax-record codes reflect the address at filing time (the spring after the tax year), so they date a given move roughly one year earlier than the ZIP codes do; a sensitivity that shifts all code-dated moves by one year changes no conclusion. Third, false positives are negligible: among 1.34 million stayers verified in the ZIP-code data, spurious code changes occur for 0.3%, dominated by known administrative renumberings. Fourth, stale records are confined to the excluded: among households excluded as undatable, 19% still carry their 1998 origin code in 2005 — precisely the stale records the exclusion rule was designed to remove — while recovered movers are at their destination code by 2005 in 93–97% of cases. One dating caveat survives the audit: recorded moves dated 1999 include a backlog of earlier moves re-based when the 1998–99 filing reform refreshed address records; excluding all 1999-dated moves leaves the exposure estimates unchanged (Section P.6).

P.2 The joint estimation model

Latent place quality. Place quality is a latent surface defined over municipalities, cohorts, and parental ranks:

$$\theta_m(c, x) = \alpha_{p(m),c} + \beta_{p(m),c} x + \text{dev}_m,$$

where $p(m)$ is the province of municipality m , x is the parental rank (on the unit interval), $(\alpha_{pc}, \beta_{pc})$ are province-by-cohort rank-rank intercepts and slopes, and dev_m is a cohort-invariant municipal deviation with partial pooling $\text{dev}_m \sim \text{Normal}(0, \tau_{\text{dev}})$. The restriction of dev_m to a cohort-invariant intercept is data-driven: a signal-noise decomposition of the unpooled cell estimates shows that cohort-within-municipality variation and municipality-specific slopes carry no identifiable signal in the Spanish panel, whereas the municipal level deviation does.

Permanent residents enter as sufficient statistics. For the permanents of each (municipality $\times$ cohort) group g , the Gaussian rank-rank likelihood depends on the group’s data only through $(n_g, \sum_i x_i, \sum_i y_i, \sum_i x_i^2, \sum_i x_i y_i, \sum_i y_i^2)$. We collapse the 1.32 million permanent-resident children to these exact sufficient

statistics — 46,400 group rows — and evaluate the likelihood

$$y_{ig} \sim \text{Normal}(\alpha_{p(m_g),c_g} + \text{dev}_{m_g} + \beta_{p(m_g),c_g} x_{ig}, \sigma_{\text{perm}})$$

in its summed form. This is compression of the *data* (lossless, by sufficiency) rather than compression of a first-stage *estimate* (lossy, and the source of the attenuation documented in Section P.4): no cell is filtered, no estimate is plugged in, and every municipality — including those with a handful of permanent residents — contributes exactly its information content.

Movers. Each mover’s outcome loads on the latent quality gap between destination and origin evaluated at her own parental rank:

$$y_i = \alpha_{\text{cell}(i)} + \lambda_{c_i} + h \text{dev}_{o(i)} + b_{a_i} \Delta_i + \varepsilon_i, \quad \Delta_i = \theta_{d(i)}(c_i, x_i) - \theta_{o(i)}(c_i, x_i), \quad (13)$$

with $\varepsilon_i \sim \text{Normal}(0, \sigma_{\text{mover}})$. The cells are origin-province $\times$ parental-quintile $\times$ age-at-move intercepts α_{cell} estimated as free parameters under weakly informative priors — hard fixed effects inside the Bayesian model, mirroring the identification strategy of Chetty and Hendren (2018a) and Britto et al. (2026) — and λ_c are mean-zero cohort intercepts. The term $h \text{dev}_{o(i)}$ is required by the resolution mismatch between the selection cells (province level) and the quality gap (municipal level): origin-municipality quality enters Δ_i with a negative sign, so omitting it biases the exposure coefficients downward, most severely for within-province movers, for whom origin variation is roughly half of the variance of Δ_i . The fitted loading is $\hat{h} = 1.12$ (90% CI [1.07, 1.16]). Because Δ_i is composed directly from the latent parameters — shared with the permanent residents’ likelihood — rather than plugged in, province-level uncertainty cancels automatically in within-province gaps, and no separate errors-in-variables correction is required.

Parameterizations of b_a . The linear-hinge form sets $b_a = b_{23} + \gamma_{\text{exp}}(23 - a)$ for $a \leq 23$ and $b_a = \delta_{\text{sel}} + \delta_{\text{post}}(23 - a)$ for $a \geq 26$ (ages 24–25 are excluded; Section P.3). The semiparametric form estimates each b_a freely. Priors are weakly informative throughout: $\alpha_{pc} \sim \text{N}(0.35, 0.25)$, $\beta_{pc} \sim \text{N}(0.3, 0.3)$, $\tau_{\text{dev}} \sim \text{Half-N}(0, 0.05)$ with a non-centered parameterization for dev_m , $\alpha_{\text{cell}} \sim \text{N}(0.5, 1)$, $\lambda_c \sim \text{N}(0, 0.1)$, unit-normal priors on $(b_{23}, \gamma_{\text{exp}}, \delta_{\text{sel}}, \delta_{\text{post}}, h)$, and Exponential(4) on the residual scales. We sample with Hamiltonian Monte Carlo (Stan via `cmdstanr`; 4 chains, 800 warmup + 800 sampling iterations for the municipal fits, 2,000 + 2,000 for the province fit); all reported parameters have $\hat{R} \leq 1.01$ with no divergent transitions. (The one slow mixer is the origin-quality loading h in the *province*-geography fit, where it is a null nuisance — origin-province quality is absorbed by the province intercepts, and $\hat{h} = -0.04$, 90% CI [−0.15, 0.07], $\hat{R} = 1.05$; the substantive parameters of that fit have $\hat{R} \leq 1.001$ with effective sample sizes near 4,000.) Prior-sensitivity checks (tightening and loosening the slope-scale priors) leave all reported

parameters unchanged.

P.3 The standard estimation sample

Outcome and cohorts. The outcome is the child’s income rank at age 24, computed within cohort over the comprehensive fiscal census records available from 2008 onward, which cover $\geq 98.5\%$ of the linked population including true zeros — inclusion therefore requires no employment or filing condition, as in the U.S. data of [Chetty and Hendren \(2018a\)](#). This fixes the cohorts at 1984–1994. For the 1980–83 cohorts, age-24 income falls in the pre-2008 filer-only regime: coverage is 25–44% of the cohort, *declines* with parental income (39% in the bottom parental quintile vs. 29% in the top, reflecting university enrollment at 24), and varies nearly two-fold across provinces. Ranks computed on such a selected subsample are not population ranks, and we exclude these cohorts rather than mix coverage regimes; they enter the age-28 robustness check of Section [P.6](#), for which their incomes are fully covered.

Excluded ages 24–25. Moves at age 24 occur in the calendar year the outcome is earned, so the coefficient there mixes partial-year exposure, moving disruption, and selection; and because the audit of Section [P.1](#) shows recorded move dates can trail true moves by up to a year, the age-25 bin predictably contains true age-24 movers. We therefore estimate the exposure segment on ages 10–23 and the placebo on ages 26–33, displaying the age-24–25 estimates without using them — a more conservative convention than the $a > 24$ placebo of [Chetty and Hendren \(2018a\)](#) and [Britto et al. \(2026\)](#), whose location data carry the same annual-filing lag. The displayed estimates corroborate this logic: $\hat{b}_{24} = 0.38$ and $\hat{b}_{25} = 0.30$ fall between the causal segment and the placebo level. The placebo tail extends to age 33, where support ends (4–6 cohorts per age); its terminal move-year (2020, the COVID spike) falls entirely in the placebo region, and its exclusion leaves the estimates unchanged (Section [P.6](#)).

Counts. The standard sample comprises 1,951,345 children with age-24 outcomes; 1,320,889 are children of permanent residents across 5,507 municipalities, and 247,682 are children of movers, of whom 192,436 enter estimation — the age window and the exclusion of ages 24–25 account for the difference, along with a 0.2% exclusion of moves touching the foral territories (Álava, Gipuzkoa, Bizkaia, Navarra), whose permanent residents are absent from the common-regime tax records. Recovered 1999–2005 moves contribute 38% of the mover sample and essentially all moves observed at the youngest ages.

P.4 Measured attenuation and design-based validation

The measurement problem, quantified. The canonical two-stage design estimates place quality cell by cell and plugs the estimates into the mover regression as data (we refer to this as the *plug-in* estimator), guarding against noise with a minimum-cell-size filter ($N \geq 400$ in Britto et al., 2026). At Spanish municipal resolution the filter retains 0.9% of (municipality $\times$ cohort) cells, and only 5.8% of movers have both endpoint cells passing it; without the filter, the plug-in quality gaps carry a measured within-cell reliability of $\hat{\lambda} = 0.60$ (that is, sampling noise accounts for 40% of their variance). The consequence is directly visible in the raw data (Table A23): restricting a plug-in OLS to moves between ever-larger municipalities, the pre-outcome exposure coefficient climbs from 0.48 to 0.74 as reliability rises toward one — a pattern that tracks measurement quality rather than economic differences across move types. All plug-in comparators in this subsection are computed on the model’s own rank conventions and estimation panels, so the contrasts with the joint model are not confounded by sample or scale differences; the one remaining asymmetry is that the plug-in estimator requires both endpoint municipalities to have at least 20 permanent residents, which drops 8.7% of movers — disproportionately children with poorer parents and *better* outcomes — whereas the joint model retains them.

Table A23: Plug-in OLS exposure estimates by endpoint municipality size

Both endpoint municipalities have...	all (≥ 20)	≥ 200	$\geq 1,000$	$\geq 3,000$
Pre-outcome coefficient (pooled, ages 10–23)	0.475 (0.016)	0.594 (0.021)	0.663 (0.035)	0.740 (0.056)
Placebo coefficient (pooled, ages 26–33)	0.190	0.256	0.348	0.388
Measured reliability $\hat{\lambda}$ of plug-in $\hat{\Delta}$	0.60	0.81	0.94	0.97

Notes: Cell-demeaned plug-in OLS on the standard sample under the model’s rank conventions (origin-province $\times$ quintile $\times$ age cells, origin control), restricted to movers whose origin and destination municipalities both exceed the stated permanent-resident count. Conventional standard errors in parentheses (not clustered by municipality; the point estimates are the object of interest). Both coefficients rise monotonically with the reliability of the plug-in quality gap, toward the joint model’s full-sample estimates (0.63 pre-outcome, 0.32 placebo).

Design-based corroboration. Three estimators with no Bayesian content bracket the joint model. (i) *Split-sample IV*: dividing each municipality’s permanent residents randomly in half (each half subject to a proportional $n \geq 10$ usability rule, so the IV runs on the same movers as the plug-in OLS) yields two independent gap estimates; instrumenting one with the other removes the attenuation by construction. This gives a pre-outcome coefficient of 0.676 (SE 0.028) and a placebo level of 0.212 (SE 0.050), against the joint model’s 0.63 and 0.32. (ii) *Province-scale OLS*: at province resolution reliability is $\hat{\lambda} \approx 0.99$, and plain plug-in OLS agrees with the joint model fit at that geography. (iii) *Scale accounting*: the plug-in placebo (0.19) equals the joint model’s (0.32) shrunk by almost exactly the measured reliability of the plug-in gaps, so the magnitude of the attenuation matches what the measured noise implies. One agreement is deliberately *not* claimed as validation: the plug-in OLS γ (0.021) coincides with the joint estimate (0.020) only because uniform attenuation

and the age-varying reliability of the mover composition (younger ages are richer in well-measured recovered moves) happen to offset; the simulation below shows γ is attenuated by 24–39% when composition does not rescue it.

P.5 Simulation evidence: bias and noise across the reliability range


We validate the estimator on synthetic data with known truth, over a grid spanning the reliability range from just below our data to essentially exact measurement. The data-generating process mirrors the measured structure of the Spanish panel: log-normal municipality sizes (dispersion matched to the observed distribution; median size swept from 5 to 275 permanent residents across eight configurations, plus two side experiments), province and municipal quality components at measured scales, municipal quality correlated with size (a rural penalty), 70% within-province moves with destinations drawn proportionally to size, and a hinge exposure curve with known $(b_{23}, \gamma_{\text{exp}}, \delta_{\text{sel}}) = (0.45, 0.030, 0.30)$. On each of 100 replications per configuration we fit three estimators to identical data: plug-in OLS (subject to the real $N \geq 20$ usability rule), split-sample IV (per-half $n \geq 10$, so it runs on the same movers as OLS), and the joint model. Figure A66 reports bias and sampling noise against the measured reliability of each configuration; configurations are placed on the axis by their *measured* reliability, which positions the real data ($\hat{\lambda} = 0.60$) just left of the two leftmost configurations ($\lambda = 0.62, 0.65$).

Four results emerge. First, plug-in OLS fails through *bias*, not noise: at the two configurations adjacent to the real data’s reliability it understates b_{23} by 31–32%, δ_{sel} by 25–31%, and γ_{exp} by 24–40% — and because bias does not shrink with sample size while noise does, at our actual sample size the joint model roughly halves the total error. Second, the joint model is centered across the range, with posterior uncertainty calibrated within $\pm 20\%$ of the truth; at the most extreme sparsity it carries a small shrinkage bias of its own (-0.07 on b_{23} , half the OLS bias and visible in the figure). Third, split-sample IV — the natural frequentist correction — is unbiased throughout but pays a large variance penalty exactly where correction is needed: at the real data’s reliability its sampling dispersion is 2–3 times the joint model’s on the level parameters and five times on γ_{exp} (SD 0.084 vs. 0.016 against a truth of 0.030 — the exposure slope is effectively unidentified by IV at low reliability). Fourth, two side experiments confirm the mechanism: switching off the size-quality correlation leaves the OLS bias unchanged (classical measurement error suffices), and raising the within-province share worsens it (smaller true gaps for the same noise) while leaving the joint model indifferent. At high reliability all three estimators agree — the methods separate exactly where the municipal design must operate, and nowhere else.

Two limits of the exercise should be stated. The data-generating process is the joint model’s own structure (linear latent quality, cohort-invariant municipal deviations, homoskedastic errors), so the simulation establishes internal consistency and the estimator comparison under a common truth — not the model’s structural

assumptions, which rest on the design-based checks of Section P.4 and the placebo tests. And the simulation contains no dating error, so it cannot speak to the move-dating issues of Section P.1, which are handled by design in the real fits and probed by the robustness checks of Section P.6.

Figure A66: Bias and error of three estimators across the reliability range



Notes: 100 replications per configuration on the data-generating process described in the text; truth $b_{23} = 0.45$, $\gamma_{\text{exp}} = 0.030$, $\delta_{\text{sel}} = 0.30$. Top panel: bias of each estimator against the measured reliability λ of the plug-in quality gaps (vertical dashed line: the reliability measured in the Spanish data). Bottom panel: root-mean-squared error extrapolated to the real mover sample size (sampling SDs scaled by $\sqrt{20,000/192,436}$, treating all sampling variance as mover-sample variance — a heuristic, since part of the plug-in estimators' variance scales with the permanents panel instead; bias unchanged). Replications with $\hat{R} > 1.05$ in the joint fit are excluded (12 of 1,000). Open triangle: side experiment with the size-quality correlation removed; open diamond: within-province share raised to 0.85.

P.6 Robustness and specification checks

Selection-control architecture. Replacing the hard cells of equation (13) with a control-function architecture — cell effects decomposed into a Mundlak-style projection on leave-one-out cell means of the latent gap plus pooled cell deviations — moves $\hat{\gamma}_{\text{exp}}$ from 0.0204 to 0.0211 and the placebo level from 0.32 to 0.29. A 100-replication simulation study of the two architectures confirms that the fixed-effects specification is the least biased and that δ_{sel} , not γ_{exp} , is the architecture-sensitive parameter — which is why we quote δ_{sel} from the fixed-effects fit.

Dating and period shocks. Excluding all 1999-dated recovered moves (the re-basing backlog documented in Section P.1) yields $\hat{\gamma}_{\text{exp}} = 0.021$ (90% CI [0.014, 0.028]) and a placebo level of 0.33 — indistinguishable from the baseline 0.020 and 0.32. Excluding all moves recorded in 2020 (the COVID spike, which falls entirely in the placebo region for our cohorts) leaves $\hat{\gamma}_{\text{exp}}$ unchanged at 0.020 and the placebo level within one posterior SD. Shifting all code-dated moves forward one year (the filing-clock offset documented in Section P.1) translates the young-age support without changing any conclusion.

Outcome age. Refitting the joint model (semiparametric) with income measured at age 28 — which adds cohorts 1980–83, whose age-28 incomes are fully covered, and relocates the mechanical placebo boundary to 28 — reproduces the exposure profile: the per-age slope over the causal window is 0.019 (SE 0.003), the selection level is 0.316 (identical to the age-24 fit’s 0.316), and the kink sits at the new outcome age, as the design requires. This check carries particular weight given the finding of Section P.7 that early-career and mature place quality *differ* within provinces: despite that divergence in the quality surfaces, the exposure and selection estimates are invariant to the outcome age — the same-age comparison underlying the estimator delivers the immunity that Chetty and Hendren (2018a) argue for.

The province placebo. At municipal resolution the placebo is exactly flat ($\hat{\delta}_{\text{post}} = -0.000$, 90% CI [−0.020, 0.019]). At province resolution it declines mildly with age at move ($\hat{\delta}_{\text{post}} = 0.043$, 90% CI [0.006, 0.079]): cross-province movers in their late twenties and thirties show selection that weakens with age under the age-24 outcome. This pattern, which is absent at municipal scale, does not affect the placebo *level* comparison of Table A24 (the reported level, 0.07 [−0.02, 0.17], is the fitted placebo line averaged over ages 26–30, so the slope is already accounted for), but we flag it as a caveat on the province-scale selection estimate.

P.7 Consistency with the geographic estimates of Section 4.2 and the interpretation of age-24 place quality

The latent municipal quality surface estimated by the joint model can be compared to the object mapped in Section 4.2: the two models are implemented independently, differ in likelihood detail (cohort resolution, the movers’ contribution, the sufficient-statistics formulation), and are fit to overlapping but distinct samples. The like-for-like comparison uses the variant of the joint model whose outcome matches Section 4.2’s (income measured in 2021–22). Across the 5,613 municipalities with estimates in both, expected outcomes at the 25th parental percentile correlate at 0.93 (Pearson; 0.91 Spearman), rising to 0.94 among municipalities with at least 100 permanent residents; the between-province correlation is 0.95 and the within-province correlation of municipal deviations is 0.56, the expected signature of two differently-shrunk small-area estimators sharing a common target. The causal design and the descriptive geography thus rest on one internally consistent latent

geography when they measure the same outcome.

The standard sample’s *age-24* quality surface is deliberately not that object: at 24, incomes reflect early-career position — students versus early labor-market entrants — so a municipality’s age-24 deviation is not a measure of its long-run quality (within provinces the two deviations in fact correlate negatively, -0.31 ; the geographic counterpart of the lifecycle measurement patterns discussed in Section 2 and of the college example in Chetty and Hendren, 2018a). The movers design is immune to this composition because it compares movers to permanent residents *at the same age*, and the outcome-age refit of Section P.6 — identical exposure slope and selection level at ages 24 and 28 — confirms the immunity empirically. The age-24 surface should therefore not be read as a map of long-run place quality; Section 4.2’s estimates serve that purpose.

P.8 From rank gains to income gains: construction of the estimates in Section 5.3

The income-gain estimates of Figure 13 combine three ingredients: a causal share, a quality gap, and a rank-to-income price. The *causal share* of a move at age a is $\hat{b}_{23} + \hat{\gamma}_{\text{exp}}(23 - a) - \hat{\delta}_{\text{sel}}$, computed per posterior draw from the linear-hinge fit; subtracting $\hat{\delta}_{\text{sel}}$ removes the part of movers’ outcome gains attributable to sorting, so the share reflects only what the move itself contributes. The *quality gaps* are computed from the model’s own municipal quality surface evaluated at $p = 40.3$, the rank used throughout Section 5.3 to represent a low-income child (it is the expected adult rank of children whose parents sit at the 25th parental percentile): each municipality’s expected adult rank at that parental rank is the province-level prediction plus the partially pooled municipal deviation dev_m , and the P75–P25 and P90–P10 gaps (8.2 and 16.5 percentile points) are percentile ranges of this quantity across municipalities, weighted by the number of permanent-resident children. Two features of this construction guard against mechanical inflation of the gaps: using the partially pooled surface (raw cell estimates would widen the percentiles through estimation noise alone), and weighting by population (unweighted percentiles would be dominated by small rural municipalities).

The *price* converts percentile gains into euros. Because age-24 incomes are early-career and compressed — many high-trajectory individuals are still students or recent entrants — pricing rank gains at the age-24 income distribution would misstate their long-run value. We instead price at the age-28 income distribution of the same outcome universe, the oldest age fully covered for our cohorts: around the baseline rank, one percentile of position is worth $dI/dp \approx \text{€}273$ per year, on a baseline income of $I(40.3) = \text{€}10,755$. A rank gain of g percentiles therefore translates into a proportional income gain of $273 g/10,755$. The credible intervals in Figure 13 propagate the joint posterior uncertainty of $(\hat{b}_{23}, \hat{\gamma}_{\text{exp}}, \hat{\delta}_{\text{sel}})$; the quality gaps and the rank-to-income price are treated as fixed, so the intervals reflect uncertainty in the causal parameters, not in the geography.

P.9 International comparison with the most similar literature


Table A24 compares our estimates with the closest international evidence on childhood exposure effects using a movers design. Which foreign estimate each Spanish resolution should be read against depends on what a move across the relevant boundary changes for the child, not on the nominal size of the unit. Spanish provinces are administrative divisions coarser than the functional geographies used abroad: they aggregate over commuting zones rather than coincide with them, and no Spanish metropolitan area is split across provinces, so an interprovincial move is always a long-distance relocation that changes the labor market the family faces. Its counterpart is therefore not the commuting zone as a unit but the U.S. county estimate identified from moves exceeding 100 miles across commuting zones, which isolates the same kind of move: 0.036 in Spain against 0.037 in the United States. Municipalities lie at the opposite end, since metropolitan areas are divided into many of them, and the municipal coefficient is identified largely from moves that change neighborhood and housing market but not labor market. Against the U.S. within-commuting-zone county estimate, the correspondence is again close: 0.020 against 0.022. Two caveats apply: U.S. counties are far larger than Spanish municipalities and restricted to those above 250,000 inhabitants, and roughly 30% of Spanish municipal moves cross a provincial boundary, retaining a labor-market component the U.S. within-commuting-zone sample excludes by construction.

The remaining comparisons are looser because their units sit between the two Spanish resolutions. Brazilian immediate geographic regions are functionally defined around an urban network and local hub and, like commuting zones, contain rather than divide their metropolitan areas, placing them near the provincial margin. Dutch COROP regions are commuting-based and approximately NUTS 3, but subdivide the Randstad conurbation into separate units connected by dense commuting, placing them intermediate. Australian Statistical Areas 4 partition each capital city into many units and are closest in spirit to the municipal margin, though far larger in population. Their exposure estimates do not track this ordering: Brazil and the Netherlands report 0.024 and 0.021, close to the Spanish municipal value rather than to the provincial one. Dispersion across countries holding scale roughly fixed is thus comparable to dispersion across scales within Spain, and part of it plausibly reflects the outcome age, since the Dutch estimate is measured at 28 rather than 24. The Australian results are less comparable still, as they allow exposure to differ between childhood and adolescence and find the larger effect in adolescence.

Selection does not align in levels, but its dispersion is informative once units are ordered by whether they partition metropolitan areas rather than by population. Spanish provinces, the coarsest unit and never splitting a metropolitan labor market, yield the lowest statistic; commuting zones and Brazilian immediate regions yield 0.126 and 0.138; COROP regions yield 0.228; and Australian SA4s yield 0.308, essentially the Spanish municipal value of 0.32. Selection into destination thus appears concentrated in residential sorting within labor markets, which any sufficiently fine geography inherits. This ordering should be read as suggestive: the


U.S. statistic moves from 0.126 to 0.226 between equations (5) and (6) with geography held fixed, so specification alone generates variation of the same magnitude. What the Spanish credible intervals do establish is that the two resolutions differ on this margin while they do not on exposure: the exposure intervals overlap across resolutions ($[0.023, 0.048]$ against $[0.014, 0.027]$), whereas the selection intervals exclude each other ($[-0.02, 0.17]$ against $[0.18, 0.45]$). The provincial interval contains zero but also contains the U.S. commuting-zone value, and should not be read as evidence that selection is absent in interprovincial moves.

Figure A67: Childhood Exposure profiles: Spain at Municipal Scale vs. each comparison country



Notes: As in Figure A68, with the Spanish municipal-scale series (dark blue) in every panel: posterior means and 90% credible intervals of the per-age coefficients b_a from the semiparametric joint model at municipal resolution — the specification underlying Figure 12 — with ages 24–25 displayed but excluded from estimation, as in that figure. The Spanish credible intervals are narrower here than at province scale because the municipal design draws on all 192,436 movers, whereas the province-scale profile is identified primarily by cross-province movers. Comparison series, interval caveats, and the outcome-age convention are as in Figure A68.

Figure A68: Childhood exposure profiles: Spain at province scale vs. each comparison country



Notes: Each panel repeats the Spanish province-scale series (dark blue) against one comparison country. The Spanish series reports posterior means and 90% credible intervals of the per-age coefficients b_a from the semiparametric joint model fit at province geography on the estimation sample of Section 5.1; as in Figure 12, ages 24–25 are displayed but excluded from estimation. Comparison series are the baseline specifications of the corresponding studies: United States — commuting zones (Chetty and Hendren, 2018a); Brazil — immediate geographic regions (Britto et al., 2026); Netherlands — COROP regions, coefficients and standard errors as reported, displayed as 90% intervals (Beekers, 2024); Australia — SA4 regions (Deutscher, 2020). No interval estimates are available for the U.S. and Australian series, and the Brazilian intervals are reproduced from the published figure at the confidence level displayed there; interval levels are therefore not harmonized across series. The vertical dashed line marks the boundary after which moves postdate the age-24 outcome used by Spain, the United States, Brazil, and Australia; the Dutch outcome is measured at age 28, so its selection region begins later. Age support differs by design: the Spanish series begins at age 10 because the residential histories start in 1998. Series are horizontally offset within each age for legibility; the comparisons are descriptive (see the discussion in this appendix and the notes to Table A24).

Table A24: Estimated Childhood Exposure Slopes and Selection Levels in International Perspective

Country	Geography	Specification	Annual exposure $\hat{\gamma}_{\text{exp}}$	Reported selection statistic
<i>Spain: joint Bayesian model</i>				
Spain	Province	Linear-hinge age profile	0.036 [0.023, 0.048]	0.07 [−0.02, 0.17]
Spain	Municipality	Linear-hinge age profile	0.020 [0.014, 0.027]	0.32 [0.18, 0.45]
<i>United States: Chetty and Hendren (2018a)</i>				
United States	Commuting zone	Semiparametric: unrestricted b_m (eq. 5)	0.044 (0.003)	0.126
United States	Commuting zone	Parametric controls, unrestricted b_m (eq. 6)	0.038 (0.002)	0.226
United States	Commuting zone	Parametric linear hinge (eq. 7)	0.040 (0.002)	NR
United States	County; across-CZ moves	Linear hinge, age-24 outcome	0.037 (0.003)	NR
United States	County; within-CZ moves	Linear hinge, age-24 outcome	0.022 (0.003)	NR
United States	County; within-CZ moves	Linear hinge, pooled outcomes at ages ≥ 24	0.027 (0.003)	NR
<i>Other international estimates</i>				
Brazil (Britto et al., 2026)	Immediate geographic region	Baseline income-rank specification	0.0238 (0.0010)	0.138
Netherlands (Beekers, 2024)	COROP region	Baseline income-rank specification	0.0211 (0.0054)	0.2281 (0.0605)
Australia (Deutscher, 2020)	Statistical Area 4	Piecewise baseline age profile	0.011 (0.007), ages 2–11 0.042 (0.003), ages 12–24	0.308

Notes: Spanish estimates are posterior means from the joint Bayesian linear-hinge specification, estimated at the two geographic resolutions on the sample of Section 5.1. Brackets report 90% credible intervals. The semiparametric Spanish specification instead estimates each age-at-move coefficient b_a freely; regressing these coefficients on age over the pre-outcome window gives 0.020 (SE 0.004) at municipal resolution, essentially identical to the linear-hinge estimate. Ages 24–25 are excluded and the Spanish post-outcome segment is estimated over ages 26–33. The reported Spanish selection statistic is the fitted placebo line averaged over ages 26–30, computed per posterior draw; at municipal resolution the placebo is flat ($\hat{\delta}_{\text{post}} = -0.000$), so this equals $\hat{\delta}_{\text{sel}}$, while at province resolution the placebo declines mildly with age (Appendix P.6), so the averaged level lies below the segment’s value at younger placebo ages. For the United States, the first commuting-zone row corresponds to Figure IV, Panel A, and equation (5): 0.044 is the slope through the unrestricted age coefficients and 0.126 is their mean for moves after age 23. The second row corresponds to Figure IV, Panel B, and equation (6), which replaces the saturated origin controls with parametric controls but continues to estimate the b_m coefficients freely. The third row reports the microdata linear-hinge estimate from equation (7), the closest age-profile counterpart to the Spanish linear-hinge specification. Chetty and Hendren do not report its corresponding selection intercept. County estimates come from Online Appendix Table V. All included counties have populations above 250,000. The across-CZ sample also requires moves exceeding 100 miles. Adding family fixed effects and time-varying family controls produces exposure estimates of 0.033 and 0.032, respectively, for across-CZ county moves. For within-CZ moves using pooled outcomes, the corresponding estimates are both 0.029 (SE 0.025). County-level selection coefficients are estimated within these regressions but are not reported in the paper. NR therefore means not reported. International entries report each paper’s baseline income-rank specification. Parentheses report standard errors. The comparisons are descriptive because outcome ages, geographic definitions, mover samples, and selection-control architectures differ across studies.