Online Appendix

Parental Altruism, the Samaritan's Dilemma, and the College Decision

Agustín Díaz Casanueva*

June, 2026

This Online Appendix collects supplementary material for "Parental Altruism, the Samaritan's Dilemma, and the College Decision." It contains the PSID sample-construction detail—variable list, construction pipeline, restrictions, definitions, and merge attrition (Section OA1)—the NLSY97 and HRS data-construction details (Sections OA2–OA3), additional descriptive tables on the PSID parent wealth portfolio and consumption (Sections OA4–OA5), a within-family fixed-effects robustness check for the consumption fact (Section OA6), and the full six-stage numerical solution algorithm summarized in the paper's Solution Algorithm appendix (Section OA7). References to "the paper" or "the main text" point to the main document; OA-numbered cross-references are internal to this appendix.

---

*Central Bank of Chile, adiaz@bcentral.cl.

OA1 PSID Sample Construction: Variables, Pipeline, and Attrition

This section completes the PSID Sample Construction appendix of the paper: the full list of measures drawn from the PSID, the six-step construction pipeline, the sample restrictions, the variable definitions and key constructed variables, the consumption measure, and the merge-attrition accounting.

OA1.1 Variables Used in the Analysis

Table OA1.1 lists the measures drawn from the PSID, grouped by concept, together with the survey years over which each is available and the data product from which it is obtained. Variable names in the table are descriptive; they correspond to underlying PSID concepts rather than to any particular survey code.

Table OA1.1. Measures Drawn from the PSID

Measure	**Waves**	**Source**
Panel A: Demographics		
Age of household head	1999–2023	Family File
Age of spouse or partner	1999–2023	Family File
Number of children in the family unit	1999–2023	Family File
Age of youngest child	1999–2023	Family File
State of residence	1999–2023	Family File
Race of household head	1999–2023	Family File
Education of household head	1999–2023	Family File
Education of spouse or partner	1999–2023	Family File
Panel B: Income		
Total family income	1999–2023	Family File
Labor income of household head	1999–2023	Family File
Labor income of spouse or partner	1999–2023	Family File
Panel C: Wealth		
Total net worth (including home equity)	1999–2023	Wealth Module
Net worth excluding home equity	1999–2023	Wealth Module
Home equity	1999–2023	Wealth Module
Retirement accounts (IRAs and annuities)	1999–2023	Wealth Module
Liquid savings	2003–2023	Family File
Stocks and mutual funds	1999–2023	Wealth Module
Vehicles	1999–2023	Wealth Module
Other assets	1999–2023	Wealth Module
Other real estate	1999–2023	Family File
Student loan debt	2011–2023	Family File
Other debts	2013–2023	Family File

continued on next page

Table OA1.1 (continued)

Measure	**Waves**	**Source**
Panel D: Employment		
Employment status of head	1999–2023	Family File
Currently working (indicator)	1999–2023	Family File
Reason for leaving last job	1999–2023	Family File
Industry of employment	2003–2023	Family File
Occupation of employment	2003–2023	Family File
Panel E: Consumption		
Food expenditure	1999–2023	Family File
Housing expenditure	1999–2023	Family File
Transportation expenditure	1999–2023	Family File
Health-care expenditure	1999–2023	Family File
Clothing expenditure	2005–2023	Family File
Recreation expenditure	2003–2023	Family File
Panel F: Individual-Level		
Relationship to household head	1999–2023	Individual File
Position within the family unit	1999–2023	Individual File
Age of individual	1999–2023	Individual File
Years of completed education	2013–2023	Individual File

Notes: All monetary measures are deflated to 2016 dollars using the CPI-U. Wealth components for 1999–2007 use the PSID imputed wealth supplement; from 2009 onward they are recorded directly in the family interview file. Liquid savings is unavailable before 2003. Industry and occupation classifications change granularity in 2017; I harmonize them to a consistent classification (see below).

OA1.2 Sample Construction

The analysis sample is built in six sequential steps.

Step 1: Assemble the family-year panel. For each of the thirteen waves I extract the wave-specific family-level measures, place them on a common set of definitions, and stack the waves into a single panel with one observation per family per wave. Supplementary measures available only in the companion family release—liquid savings and detailed industry codes for 2003–2015—are merged by household and wave. I harmonize the industry and occupation classifications across the 2017 change in granularity, retaining a coarser classification that is consistent across all waves, and I define an involuntary job separation as a separation due to a plant closing or a layoff.

Step 2: Assemble the individual-year panel. I reshape the individual file into a person-by-year panel containing each member's position in the family unit, relationship to the household head, age, and (from 2013 onward) years of completed education. A time-invariant person identifier, constructed from each member's 1968 family lineage number and within-lineage person number, links individuals across waves and underlies the parent–child match.

Step 3: Construct the household panel. The family-year and individual-year panels are merged by household and wave. I retain current family-unit members who are the household head, the legal spouse, or a cohabiting partner. Including spouses ensures that a parent who is not the household head—typically the mother in a two-parent household—can still be matched to her children. Rare duplicate person-year records are resolved by giving preference to head status.

Step 4: Attach intergenerational links. Using the Family Identification Mapping System, each child is linked to the household of one parent. When more than one parent is available, the link follows a fixed priority: biological father, then adoptive father, then biological mother, then adoptive mother. Children with no matchable parent are dropped.

Step 5: Identify college attendance. I combine three approaches to maximize coverage across all thirteen waves. For 2013 onward, where years of completed education are recorded at the individual level, a child is coded as attending college if completed education lies strictly between high-school completion and a bachelor's degree at ages 18–24, or if completed education rises between consecutive waves above the high-school level before age 26. For 1999–2011, where individual education is not yet recorded, children aged 18–22 whose highest education observed in later waves exceeds high school are coded as having attended college. For the 2015 and 2017 waves, the direct enrollment indicators from the Transition into Adulthood Supplement take precedence and override the education-based codes.

Step 6: Form the parent-household sample. The child–parent records are merged with the household panel, collapsed to one observation per parent household and wave, and augmented with the number of children currently enrolled in college. I then merge state-level college costs—in-state public tuition and private tuition by state and year, both from the Integrated Postsecondary Education Data System (IPEDS)—and deflate all monetary measures to 2016 dollars using the CPI-U. Finally, I assign each parent to a permanent-income group based on average real household income observed while the head is aged 26–59 (with overall average income used as a fallback when the head is never observed in that range). The groups are: $30,000 to $60,000; $60,000 to $100,000; $100,000 to $160,000; and above $160,000. Households with average income below $30,000 or otherwise unclassifiable are dropped.

OA1.3 Sample Restrictions

I drop observations meeting any of the following conditions:

1. Missing or invalid demographic information (age of head or spouse not reported).
2. Missing or invalid education for the head or spouse.
3. Top-coded or missing wealth or income (any wealth or income component recorded with a PSID missing or refused code).
4. Negative total net worth.
5. Internally inconsistent portfolios (total net worth below home equity, or below other real estate).
6. Implausible portfolio shares (home-equity share below -40%, other-real-estate share below -40%, or vehicle share above 50% of total net worth).
7. Average real income below $30,000 or otherwise outside the classification bounds.

OA1.4 Variable Definitions

Household consumption is the sum of expenditures on food at home, food away from home, housing (rent or imputed rent for homeowners), utilities, transportation, health care, education, and child care; from 2005 onward, clothing, recreation, and vacation are also included. Income is total household income before taxes. Wealth is total net worth: the sum of home equity, other real estate, vehicles, farm and business assets, stocks, checking and savings accounts, bonds, retirement accounts, and other assets, minus all debts. Inter-vivos transfers are recorded separately for transfers from parents to children and from children to parents, and include cash gifts, help with expenses, and regular financial support. The key explanatory variable throughout the empirical analysis is child college attainment, an indicator for whether the child holds a bachelor's degree or higher. To condition on parental resources, I construct within-cohort parent income quartiles, ranking each parent against others born in the same year.

OA1.5 Key Constructed Variables

Portfolio shares. Each asset class—retirement accounts, home equity, other real estate, and vehicles—is expressed as a share of total net worth and used to screen out internally inconsistent portfolios (see the sample restrictions above).

College-cost instrument. The instrument used in the empirical analysis interacts the (log) in-state tuition in the parent's state of residence with the number of the parent's children currently enrolled in college, so that exposure to tuition costs is increasing in the number of children of college age.

Involuntary job separation. An indicator equal to one when the head left the previous job because of a plant closing or a layoff, and zero otherwise.

Real values. All monetary values are deflated to 2016 dollars using the CPI-U for all urban consumers. Log transformations are applied where indicated in the regression specifications.

OA1.6 Consumption Data

PSID consumption expenditure data are available for every wave from 1999 onward. The consumption measure aggregates spending across six categories: food (at home and away), housing (rent, utilities, property taxes, and insurance), transportation (vehicle costs, fuel, parking, and public transit), health care (insurance premiums and out-of-pocket spending), clothing (available from 2005), and recreation (available from 2003). Clothing and recreation are added to the aggregate only in the years in which they are collected.

OA1.7 Merge Attrition

Table OA1.2 documents the sample size at each stage of construction: from the full PSID individual file, through the intergenerational link to parents, to the merge with household-level consumption, income, and wealth. Panel B breaks the parent–child links down by relationship type, and Panel C quantifies the observations lost at each merge stage. The main source of attrition is the requirement that the matched parent appears as a household head or spouse in the family interview file during the sample period.

Table OA1.2. PSID Sample Construction and Merge Attrition

Stage	Observations	Unique IDs	% Retained
Panel A: Building the Analysis Sample			
Individual file (all persons × years)	441 158	45 386	—
After college identification	441 152	45 386	100.0
Matched to FIMS (child–parent link)	358 874	34 916	81.3
Matched to parent household (hh_data)	148 541	10 492	41.4
Collapsed to parent-HH × year	71 896	10 492	—
After sample restrictions	62 320	10 145	86.7
Panel B: Parent Identification Source			
Linked via biological father	297 284		
Linked via adoptive father	4114		
Linked via biological mother	56 856		
No identifiable parent (dropped)	0		
Panel C: Merge Losses			
Individuals not in FIMS	82 278		
Child–years without parent in hh_data	210 333		
Parent–years without state tuition data	16 488		
Final analysis sample	62 320	10 145	
of which: HH–years with child in college	6566		

Notes: Children are linked to parents via the Family Identification Mapping System. The parent–child link follows a hierarchical priority: biological father, adoptive father, biological mother, adoptive mother. “Unique IDs” refers to unique persons (first two rows) or unique parent households (remaining rows). “% Retained” is relative to the previous stage. All monetary measures are winsorized at the 1st and 99th percentiles and deflated to 2016 dollars.

OA2 NLSY97 Sample Construction

This appendix describes the construction of the NLSY97 analysis sample. The National Longitudinal Survey of Youth 1997 (NLSY97) tracks a nationally representative cohort of 8,984 Americans born between 1980 and 1984, interviewed annually from 1997 through 2011 and biennially thereafter. I use the NLSY97 for two purposes: constructing college graduation rates by ability and parent-wealth quartile (the sixteen cell-level moments used in estimation), and estimating the education-specific income process. The exact survey variable codes underlying each measure are documented in the replication package.

OA2.1 Data Sources

Baseline and demographics. The baseline survey provides each respondent's gender, race and ethnicity, birth year, and census region, together with parental education for the biological mother and father and family income during the late 1990s and early 2000s.

Cognitive ability and academic performance. Cognitive ability is measured by the math-verbal percentile from the computer-adaptive aptitude battery administered in 1999. High-school transcripts provide SAT and ACT scores and overall high-school grade point average.

Schooling module. The schooling module records detailed college enrollment and financing information at the college-by-term-by-year level: tuition charged, room and board, grants and scholarships, total loans, out-of-pocket expenses, credit hours attempted, and withdrawal from college.

Income and transfer module. The income module records annual labor income and annual hours worked (used to estimate the income process) and, for the college-age years, total parental transfers and parental allowances.

OA2.2 Variables Used in the Analysis

Table OA2.1 lists the measures drawn from the NLSY97, grouped by concept, together with their coverage and source. Variable names are descriptive and correspond to underlying survey concepts rather than to any particular survey code.

Table OA2.1. Measures Drawn from the NLSY97

Measure	**Years**	**Source**
Panel A: Demographics and Ability		
Gender	1997	Baseline
Race and ethnicity	1997	Baseline
Birth year	1997	Baseline
Math-verbal aptitude percentile	1999	Aptitude battery
SAT verbal and math scores	Cumulative	Transcript
ACT composite score	Cumulative	Transcript
Overall high-school GPA	Cumulative	Transcript
Panel B: Family Background		
Biological mother's highest grade	1997	Baseline
Biological father's highest grade	1997	Baseline
Gross family income	1997–2003	Annual
Census region at baseline	1997	Baseline
Panel C: College Enrollment and Costs		
Tuition charged per term	1997–2017	Schooling
Grants and scholarships	1997–2017	Schooling
Total student loans	1997–2011	Schooling
Out-of-pocket college expenses	1998–2001	Schooling
Room and board	1998–2001	Schooling
Credit hours attempted	1997–2017	Schooling
Panel D: Education Attainment		
Highest grade attended	1998–2003	Annual
Highest degree ever received	Cumulative	Cumulative
Withdrawal from college	Annual	Schooling
Panel E: Income Process		
Annual labor income	1997–2011	Income
Annual hours worked (all jobs)	1997–2007	Employment
Panel F: Parental Transfers		
Total parental transfers	1997–2003	Income
Parental allowance	1997–2003	Income

Notes: Cumulative measures are computed by NLS staff from the respondent's full survey history. The aptitude percentile is the math-verbal score from the computer-adaptive battery administered in 1999. Schooling-module measures are recorded by college enrollment, term, and survey year. Family income refers to the gross household income of the respondent's family of origin.

OA2.3 Sample Construction

Step 1: Assemble the respondent panel. I retain core demographics (gender, race, birth year, and census region), parental education for the biological mother and father, cognitive ability (the 1999 math-verbal aptitude percentile), academic performance (SAT and ACT scores and high-school GPA from transcripts), and family income from 1997 to 2003.

Step 2: Build the college-spell panel. The detailed schooling measures are organized into panels indexed by individual, college enrollment, term, and year—tuition, grants and scholarships, total loans, out-of-pocket expenses, room and board, credit hours, and withdrawal—and merged into a unified college-spell panel.

Step 3: Identify education attainment. College attendance is identified from the highest grade attended (available 1998–2003): a respondent is coded as having attended college if the maximum grade attended exceeds high school. College graduation is defined as attaining a bachelor's degree by age 25, using the cumulative highest-degree measure.

Step 4: Estimate the income process. Using annual labor income and annual hours worked, I construct hourly wages by education group. Following Abbott et al. (2019), I estimate the high-school-to-college income ratio and the education-specific income standard deviations, which provide three of the moments used in estimation.

Step 5: Assemble parental transfers. Parental financial transfers and allowances are organized into a person-year panel for 1997–2003. These measure financial support from parents during the college-age years and discipline the transfer structure in the model.

OA2.4 Key Constructed Variables

Ability quartiles. I rank respondents by their 1999 math-verbal aptitude percentile and partition them into quartiles, restricting the sample to respondents with a non-missing aptitude score.

Parent-wealth quartiles. Total household net worth at age 17 is used to construct within-cohort parent-wealth quartiles. Together with the ability quartiles, these define the 4×4 grid of college-graduation rates used as the core estimation moments.

College-graduation rates. The sixteen cell-level graduation rates (ability quartile by parent-wealth quartile) are the primary target moments in estimation; they trace how college attainment varies across the joint distribution of ability and parental resources.

Average tuition. The gross annual cost of college (\$12,200) is computed from average tuition reported in the NLSY97 schooling module, following Abbott et al. (2019). The model uses a net annual tuition (after institutional discounts) of $\phi = \$6{,}700$, with the remaining gap covered by the endogenous financial-aid schedule $g(a_p)$ described in the main text.

OA3 HRS Sample Construction

This appendix describes the construction of the HRS analysis sample. The Health and Retirement Study (HRS) is a biennial longitudinal survey of Americans over age 50, conducted by the University of Michigan since 1992. I use the RAND HRS Longitudinal File, which harmonizes variables across all waves (1992–2022), together with four supplements: the RAND Family-Kids file, the RAND Family-Resident file, the Consumption and Activities Mail Survey (CAMS), and the supplemental survey on college support. Bequest information is taken from the HRS exit interview. The exact survey variable codes underlying each measure are documented in the replication package.

OA3.1 Data Sources

RAND HRS Longitudinal File. The core source is the RAND HRS longitudinal file, in which RAND staff harmonize variable definitions, handle skip patterns, and construct consistent wealth, income, and health measures across the sixteen waves spanning 1992–2022. From this file I draw respondent and spouse demographics, household wealth and its components, household income and its sources, and respondents' subjective probabilities of leaving a bequest.

RAND Family-Kids file. This supplement records information about each respondent's children—demographics (age, gender, education, marital status), geographic proximity, financial transfers in both directions, care provision, and bequest intentions—for waves spanning 1992–2014. I organize it into a child-by-wave panel.

RAND Family-Resident file. This supplement provides household composition by wave—the number of living children, sons, and daughters—which I organize into a respondent-by-wave panel.

Consumption and Activities Mail Survey (CAMS). CAMS is a supplemental mail survey that collects detailed consumption and spending data for a subset of HRS households, biennially from 2001 to 2017. It provides total household spending, durable and non-durable spending, housing expenditure, and mortgage payments.

Supplemental survey on college support. This supplement collects retrospective information on parents' financial contributions to their children's college education. For each child who attended college it records tuition, room and board, years of college, and the share of each cost paid by the parent, which I use to construct the total parental contribution to college in 2016 dollars.

Exit interview. The HRS exit interview, conducted with a proxy informant after a respondent's death, records bequest and inheritance information. I use it to construct child-level bequests (dollar transfers and shares of the estate) and housing inheritance (whether each child inherited the main residence and its imputed value).

College tuition data. Regional in-state college tuition by census division and year is merged with HRS respondents using their census division of residence, providing the regional college-cost variation used in the analysis.

OA3.2 Variables Used in the Analysis

Table OA3.1 lists the measures drawn from the HRS, grouped by concept, together with their coverage and source. Variable names are descriptive and correspond to underlying survey concepts rather than to any particular survey code.

Table OA3.1. Measures Drawn from the HRS

Measure	**Coverage**	**Source**
Panel A: Respondent Demographics		
Age of respondent	1992–2022	RAND HRS
Age of spouse	1992–2022	RAND HRS
Gender	Time-invariant	RAND HRS
Race	Time-invariant	RAND HRS
Years of education	Time-invariant	RAND HRS
Highest degree	Time-invariant	RAND HRS
Birth year	Time-invariant	RAND HRS
HRS sample cohort	Time-invariant	RAND HRS
Census division of residence	1992–2022	RAND HRS
Panel B: Household Wealth		
Total net worth	1992–2022	RAND HRS
Net home value	1992–2022	RAND HRS
Retirement accounts (IRA/Keogh)	1992–2022	RAND HRS
Checking and savings	1992–2022	RAND HRS

continued on next page

Table OA3.1 (continued)

Measure	**Coverage**	**Source**
Vehicles	1992–2022	RAND HRS
Business and farm assets	1992–2022	RAND HRS
Other real estate	1992–2022	RAND HRS
Panel C: Income		
Total household income	1992–2022	RAND HRS
Social Security retirement income	1992–2022	RAND HRS
Social Security disability / SSI	1992–2022	RAND HRS
Pension income	1992–2022	RAND HRS
Panel D: Bequest Expectations		
Probability of bequest $\geq$ \$10K	1992–2022	RAND HRS
Probability of bequest $\geq$ \$100K	1992–2022	RAND HRS
Probability of bequest $\geq$ \$500K	1992–2022	RAND HRS
Probability of any bequest	1992–2022	RAND HRS
Panel E: Child Information		
Child's years of education	1992–2014	Family-Kids
Child's age	1992–2014	Family-Kids
Child's gender	Time-invariant	Family-Kids
Child's marital status	1992–2014	Family-Kids
Child's income bracket	1992–2014	Family-Kids
Child lives within ten miles	1992–2014	Family-Kids
Panel F: Inter-Vivos Transfers		
Parent gave a transfer (indicator)	1992–2014	Family-Kids
Amount the parent gave	1992–2014	Family-Kids
Child gave a transfer (indicator)	1992–2014	Family-Kids
Amount the child gave	1992–2014	Family-Kids
Panel G: College Support		
Tuition and board paid by the parent	Retrospective	College support
Number of children who attended college	Retrospective	College support
Parent contributed a positive amount	Retrospective	College support
Panel H: Consumption		
Total household consumption	2001–2017	CAMS
Non-durable spending	2001–2017	CAMS
Housing expenditure	2001–2017	CAMS
Panel I: Bequests		
Dollar bequest to child	Exit interview	Exit interview
Share of estate to child	Exit interview	Exit interview
Housing inheritance value	Exit interview	Exit interview

Notes: The Family-Kids and Family-Resident supplements cover waves spanning 1992–2014. Monetary measures from the RAND HRS file are nominal at the source and are deflated to 2016 dollars in the analysis. The college-support measures are retrospective, covering each child's entire college period. Child income brackets are reported by the parent respondent; I use the bracket midpoint.

OA3.3 Sample Construction

Step 1: Assemble the respondent panel. I reshape the RAND HRS longitudinal file into a person-by-wave panel containing time-invariant identifiers and demographics (birth year, gender, race, education, religion) together with wave-varying measures of wealth, income, health, bequest expectations, family structure, employment, and spouse characteristics.

Step 2: Organize the child file. The Family-Kids supplement is reshaped from one record per child to a child-by-wave panel, with a child identifier that links each child to the parent respondent. Measures are grouped into child demographics, child-to-parent transfers, parent-to-child transfers, and bequest intentions.

Step 3: Add household composition. The Family-Resident supplement, which records the number of living children, sons, and daughters by wave, is merged into the respondent panel to provide the child-count breakdown.

Step 4: Add consumption. CAMS is organized into a respondent-by-wave panel of total household spending, durable and non-durable spending, housing expenditure, mortgage payments, and total household consumption. CAMS covers a subsample of HRS households biennially from 2001 to 2017.

Step 5: Construct college support. For each child who attended college, the total parental contribution is the sum of the real tuition paid by the parent (annual tuition, deflated to 2016 dollars, multiplied by years of college and the share of tuition the parent paid) and the real room-and-board paid by the parent (annual room and board, deflated, multiplied by years and the share the parent paid). Records with missing cost or share information are dropped; an explicit “did not contribute” response is recoded to zero. The total family contribution sums across all children in the household.

Step 6: Identify widowhood. I extract each respondent’s year of death and attach it to the spouse’s record, which yields, for each respondent, the year of the partner’s death and supports the analysis of widowhood.

Step 7: Construct bequests from the exit interview. The exit interview provides two kinds of child-level bequest information. Financial bequests are the dollar amounts and shares of the estate transferred to each child; survey missing codes are removed, and when the estate share is missing but dollar amounts are available, the share is imputed as each child's portion of total transfers. Housing inheritance is constructed for respondents whose main residence passed to children rather than to a surviving spouse: I identify which children received the home and, when the home was left to all children jointly, divide its imputed value equally among them.

Step 8: Add regional college costs. Average in-state college tuition by census division and year is organized into a region-by-year panel and merged with HRS respondents using their census division of residence.

OA3.4 Sample Restrictions

I impose the following restrictions, analogous to those applied to the PSID sample:

1. **Age restrictions:** parents over age 50 and children over age 26.
2. **Interview status:** only respondents with completed interviews in a given wave.
3. **Couple households:** for analyses requiring spouse information, I restrict to coupled households.
4. **Missing data:** survey "don't know," "refused," and "missing" responses are carried through the construction stage and handled within each specific analysis, so that wealth and income are not discarded prematurely.

The final HRS analysis sample contains 19,179 parent–child pairs and 98,861 observations.

OA4 Parent Wealth Portfolio

Table OA4.1 provides a detailed breakdown of parent income and wealth by income group. The table complements the summary variables reported in the parent demographics table in the main text.

Table OA4.1. Parent Income and Wealth by Income Group (PSID, 2016 Dollars)

	By Average HH Income Group				
	$30–60	$60–100	$100–160	$160+	All
Household income	44,223	76,829	120,169	242,277	102,569
	(21,535)	(37,964)	(51,447)	(253,591)	(121,904)
Head labor income	23,523	40,976	64,728	123,504	53,930
	(19,369)	(28,162)	(42,902)	(123,970)	(64,202)
Total wealth	102,555	200,522	367,076	965,941	326,288
	(230,260)	(377,059)	(536,534)	(1,052,497)	(612,283)
Home equity	48,321	76,931	127,798	223,253	102,056
	(86,276)	(105,763)	(136,410)	(195,989)	(138,176)
IRA/annuity	6,437	20,026	49,226	107,888	35,668
	(38,692)	(68,967)	(108,878)	(163,964)	(98,914)
Non-housing fin. wealth	53,469	121,496	234,059	698,124	214,727
	(182,888)	(318,143)	(453,318)	(895,431)	(505,467)
Savings	7,950	16,013	36,321	68,743	25,676
	(31,182)	(39,546)	(61,142)	(87,368)	(56,070)
Vehicles	12,549	19,741	24,701	32,096	20,737
	(16,601)	(19,952)	(22,565)	(28,691)	(22,256)
Other real estate	7,938	13,934	22,799	55,426	20,513
	(42,810)	(57,200)	(75,043)	(122,965)	(73,634)
Student loans	1,736	4,134	7,932	7,755	4,881
	(7,821)	(12,628)	(18,488)	(20,246)	(14,759)
Observations	15 096	17 356	13 011	7863	53 326
Unique parent households	2718	2770	1833	1041	8362

Notes: Means with standard deviations in parentheses. All monetary variables deflated to 2016 dollars using the CPI-U and winsorized at the 1st and 99th percentiles. Income groups defined by parent average real household income observed between ages 25 and 60.

OA5 Parent Consumption Detail

Table OA5.1 decomposes total household consumption into major expenditure categories. Housing constitutes the largest category across all income groups, followed by transportation and food. Discretionary categories—clothing, recreation, and vacation—are only available from the 2005 wave onward and increase sharply with income.

Table OA5.1. Parent Household Consumption by Income Group (PSID, 2016 Dollars)

	By Average HH Income Group				
	$30–60	$60–100	$100–160	$160+	All
Total consumption	36,426	46,131	57,243	73,130	50,076
	(18,374)	(20,389)	(24,489)	(33,277)	(26,252)
Food	7,915	9,211	10,582	12,872	9,718
	(4,854)	(4,944)	(5,296)	(6,269)	(5,470)
Housing	15,030	18,677	24,487	33,215	21,206
	(9,071)	(10,757)	(13,588)	(18,356)	(13,873)
Transportation	9,310	11,880	13,811	15,562	12,166
	(8,213)	(8,709)	(9,844)	(11,999)	(9,661)
Health care	2,108	3,402	4,195	5,027	3,469
	(3,027)	(3,649)	(4,157)	(4,832)	(3,950)
Clothing (2005+)	1,356	1,577	1,912	2,808	1,776
	(1,704)	(1,762)	(1,907)	(2,626)	(1,988)
Recreation (2005+)	503	807	1,178	1,805	957
	(920)	(1,211)	(1,464)	(1,897)	(1,396)
Vacation (2005+)	835	1,482	2,376	3,886	1,870
	(1,505)	(2,020)	(2,583)	(3,604)	(2,554)
Observations	15 096	17 356	13 011	7863	53 326
Unique parent households	2718	2770	1833	1041	8362

Notes: Means with standard deviations in parentheses. Consumption data available from the PSID beginning in 1999. All expenditure categories deflated to 2016 dollars and winsorized at the 1st and 99th percentiles. Clothing, recreation, and vacation are available from 2005 onward only. Total consumption is the sum of all listed categories.

OA6 Additional Robustness: Parent Fixed Effects

As additional robustness for the consumption stylized fact in the main text, I exploit parent fixed effects to identify the college effect from within-family (between-sibling) variation. For consumption, the specification uses child-pair-year observations: within a given family and year, siblings who differ in college attainment generate variation in the college indicator while holding all parent-level characteristics constant.

Table OA6.1 reports parent fixed effects regressions of log parent consumption on the child's college indicator, at the child-pair-year level. The positive and significant coefficient confirms that the consumption difference documented in the main text is not driven by unobserved parental characteristics: even within the same family, parent consumption is higher in years when the college-educated child's characteristics are more salient.

Table OA6.1. Parent Consumption and Child College Status: Parent Fixed Effects (PSID)

	(1) OLS + Income Q FE	(2) Parent FE
College (BA+)	0.068^{***}	0.027^{**}
	(0.018)	(0.013)
Controls	Yes	Yes
Parent FE	No	Yes
Parent income Q FE	Yes	–
Observations	13,841	13,841
R^2 / Within R^2	0.469	0.840

Notes: Regressions of log parent household consumption on child college indicator (BA+), at the child-pair-year level. Column 1: OLS with parent income quartile fixed effects and comprehensive controls. Column 2: parent fixed effects with child age polynomial and education controls. Standard errors clustered at the parent household level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

OA7 Solution Algorithm (Full Detail)

This appendix gives the full numerical solution algorithm summarized in the Solution Algorithm appendix of the main text. The model is solved using an upwind finite difference scheme with implicit time-stepping, following Achdou et al. (2022) and the transfer-region methods of Barczyk and Kredler (2014a,b). The algorithm proceeds in six stages: grid construction, solving the child-alone HJB, solving the coupled parent-child HJB system, computing college choice probabilities, forward simulation, and moment computation for SMM estimation.

OA7.1 State Space and Grid Construction

The model features three continuous state variables during the overlap period: parent wealth $a_p \in [0, \bar{a}_p]$, child wealth $a_c \in [\underline{a}_c, \bar{a}_c]$, and the persistent income shock $z \in [-\bar{z}, \bar{z}]$. In addition, child ability θ and education $e \in \{\mathrm{HS}, \mathrm{C}\}$ are discrete states that index separate HJB problems.

Asset grids. Both wealth grids use exponential spacing to concentrate points near the origin, where policy functions exhibit greater curvature. For N grid points from $a_{\min}$ to $a_{\max}$ with growth factor $g > 0$:

$$a_i = a_{\min} + a_{\max} \cdot \frac{(1+g)^{i-1} - 1}{(1+g)^{N-1} - 1}, \qquad i = 1, \ldots, N. \tag{1}$$

The parent grid has $N_{a_p} = 30$ points over $[0, 4{,}000]$ (thousands of dollars). The child grid has $N_{a_c} = 34$ points: 10 points in the negative (student-debt) region $[\underline{a}_c, 0)$ and 24 exponentially spaced points over $[0, \bar{a}_c]$. The negative segment mirrors the exponential spacing of the positive segment, so that nodes are dense near zero—where amortizing balances spend most of their time—and sparse toward the debt floor. This spacing matters for convergence: with a uniform debt segment, college-attendance cells fail to converge as the segment is refined, whereas the mirrored-exponential segment is converged at 10 points. We set the child-grid floor to $\underline{a}_c = -25$, below the cumulative student borrowing limit $\bar{b}_0 = 18$ (the college financial system in the main text) so the constraint is bracketed by interior grid points, and $\bar{a}_c = 2{,}500$.

Income shock grid. The Ornstein-Uhlenbeck process $dz = -\kappa_e z\,dt + \sigma_e\,dW$ is discretized on a uniform grid of $N_z = 9$ points spanning ±3 stationary standard deviations:

$$z_k \in \left\{-3\,\frac{\sigma_e}{\sqrt{2\kappa_e}}, \dots, +3\,\frac{\sigma_e}{\sqrt{2\kappa_e}}\right\}, \qquad k = 1, \dots, N_z.$$

The grid width is set to the maximum stationary standard deviation across education types to ensure both HS and college processes are well-covered.

Ability grid. Child ability follows an AR(1) in logs across generations: $\log\theta' = \rho_\theta \log\theta + \mu_\theta(1-\rho_\theta) + \sigma_\theta\varepsilon$. We discretize this using the Tauchen (1986) method with $N_\theta = 5$ grid points spanning ±3 stationary standard deviations of $\log\theta$.

Time grid. The backward induction uses $N_t = 60$ equally spaced time steps over the parent's full remaining life ($T = 30$ years, parent ages 42–72, child ages 18–48), yielding $\Delta t = 0.5$ years per step. The implicit scheme is unconditionally stable, so this step size is accurate at half the cost of a finer grid. The coupled parent-child system is solved over the entire horizon $[0, T]$: the parent retires at $T^{\text{ret}} = 24$ (age 66) and switches to social security income, but remains coupled to the child and can continue to transfer until death at T (age 72).

OA7.2 Stage 1: Child-Alone HJB

After the parent dies (child age 48), the child solves a standard consumption-savings problem independently for the remainder of life ($T_{\text{child alone}} + T_{\text{child retire}} = 24$ years, child ages 48–72, working until 66 and then on social security). The HJB equation, with the same multiplicative wealth diffusion as the coupled problem, is:

$$\begin{aligned}\rho V^{c,\text{alone}}(a,z,t) = \max_{c\geq 0}\Big\{ & u(c) + V_a^{c,\text{alone}}\left[r\,a + y(t,z) - c\right] + V_z^{c,\text{alone}}(-\kappa z) \\ & + \tfrac{1}{2}\sigma^2 V_{zz}^{c,\text{alone}} + \tfrac{1}{2}\sigma_a^2 a^2 V_{aa}^{c,\text{alone}} + V_t^{c,\text{alone}}\Big\}.\end{aligned} \tag{2}$$

We solve this problem separately for each ability type θ and education level $e_c \in \{\text{HS}, \text{C}\}$, giving $2 \times N_\theta = 10$ independent HJB problems that are solved in parallel.

Terminal condition. At the child's death (age 72, i.e. $t = T_{\text{child}}$ on the child-alone clock), the terminal value is $V^c(a, z, T_{\text{child}}) = u(c_T)/\rho$ where $c_T = \max\{r\,a + \text{SS}_{e_c} + 0.1\,a,\ \epsilon\}$ represents consuming all remaining resources.

Optimal consumption via the first-order condition. At each grid point and time step, the FOC $u'(c^*) = V_a^c$ yields the candidate consumption $c^* = (V_a^c)^{-1/\gamma}$, where γ is the CRRA coefficient. We compute V_a^c using upwind finite differences: the *forward* difference $\partial_a^+ V$ and *backward* difference $\partial_a^- V$ yield two candidate consumption levels, c^{fwd} and c^{bwd}, with associated savings drifts $s^{\text{fwd}} = ra + y - c^{\text{fwd}}$ and $s^{\text{bwd}} = ra + y - c^{\text{bwd}}$. The upwind scheme selects:

$$c^* = \begin{cases} c^{\text{fwd}} & \text{if } s^{\text{fwd}} > 0 \quad \text{(saving: use forward difference)}, \\ c^{\text{bwd}} & \text{if } s^{\text{bwd}} < 0 \quad \text{(dissaving: use backward difference)}, \\ ra + y & \text{otherwise (drift} \approx 0\text{: consume resources).} \end{cases}$$

This upwind selection ensures numerical stability by using the derivative in the direction from which information flows.

Implicit time stepping. Given optimal policies at time step n, we construct the sparse infinitesimal generator $\mathbb{A}^n$ that encodes the upwind finite-difference operator in a and the discretized OU operator in z (combining upwind drift and central-difference diffusion terms). The implicit update solves:

$$\left[\left(\frac{1}{\Delta t} + \rho\right)\mathbf{I} - \mathbb{A}^n\right]\mathbf{V}^n = \mathbf{u}^n + \frac{1}{\Delta t}\mathbf{V}^{n+1}, \tag{3}$$

where $\mathbf{V}^n$ and $\mathbf{u}^n$ are the vectorized value function and flow utilities at step n, $\mathbf{I}$ is the identity matrix of dimension $N_{a_c} \times N_z$, and $\mathbf{V}^{n+1}$ is the value at the next (future) time step. This system is solved via sparse LU factorization.

OU generator. The infinitesimal generator for the OU process on the uniform z-grid combines a central-difference approximation for diffusion with upwind differencing for the drift:

$$
\begin{aligned}
(\mathbb{A}_z)_{k,k+1} &= \frac{\sigma^2/2}{\Delta z^2} + \frac{\max(-\kappa z_k, 0)}{\Delta z}, \\
(\mathbb{A}_z)_{k,k-1} &= \frac{\sigma^2/2}{\Delta z^2} + \frac{\max(\kappa z_k, 0)}{\Delta z}, \\
(\mathbb{A}_z)_{k,k} &= -(\mathbb{A}_z)_{k,k+1} - (\mathbb{A}_z)_{k,k-1},
\end{aligned}
\tag{4}
$$

with reflecting boundary conditions at $k = 1$ and $k = N_z$.

Wealth-diffusion generator. The multiplicative wealth shock $\sigma_a a\, dB$ contributes a second-order term $\frac{1}{2}\sigma_a^2 a^2 V_{aa}$, discretized by a central second difference on the (non-uniform) wealth grid. With diffusion coefficient $D_i = \frac{1}{2}\sigma_a^2 a_i^2$, forward spacing $\Delta_i = a_{i+1} - a_i$ and backward spacing $\Delta_{i-1} = a_i - a_{i-1}$, the generator entries are

$$
\begin{aligned}
(\mathbb{A}_a)_{i,i+1} &= D_i \frac{2}{\Delta_i(\Delta_i + \Delta_{i-1})}, \\
(\mathbb{A}_a)_{i,i-1} &= D_i \frac{2}{\Delta_{i-1}(\Delta_i + \Delta_{i-1})}, \\
(\mathbb{A}_a)_{i,i} &= -(\mathbb{A}_a)_{i,i+1} - (\mathbb{A}_a)_{i,i-1}.
\end{aligned}
\tag{5}
$$

The coefficient D_i vanishes at $a_i \leq 0$ (no volatility at the borrowing constraint, following Barczyk and Kredler, 2014a) and at the grid boundaries, so no extra boundary condition is required. Because D_i depends only on the grid, this block is assembled once and reused across all time steps; it preserves the generator's M-matrix structure, so the implicit step remains stable.

OA7.3 Stage 2: Coupled Parent-Child HJB

During the overlap period, the parent and child interact strategically. The parent's HJB equation in the main text governs the parent's problem, with the transfer decision characterized by the complementary-slackness condition described there. That is, the parent transfers whenever the marginal value of an additional dollar in the child's account exceeds the marginal value of retaining it.

Terminal condition. At $t = T$ (parent age 72, the parent's death), the coupled system terminates. The parent's terminal value combines a warm-glow term over remaining wealth and altruistic concern for the child's continuation value (the parent's terminal condition in the main text). The child's terminal value is $V^c(a_p, a_c, z, T) = V^{c,\text{alone}}(a_c + \alpha_{\text{beq}}\, a_p, z)$, the child-alone value function from Stage 1 evaluated at $t = 0$ with effective wealth including the bequest.

Transfer allocation in closed form. We solve the coupled system backward in time. At each time step n and grid point (i, j, k) the transfer is obtained in *closed form*: because altruism is one-sided (only the parent gives) and the recipient's marginal value enters through a static first-order condition, the transfer region $\mathcal{T}^n$ and the gift size are determined pointwise, without iterating on the region. The allocation is the continuous-time analog of the rule in Barczyk and Kredler (2014a) and the `GetBKalloc` routine of Barczyk et al. (2023), specialized to one-sided altruism.

1. **Own-consumption first-order conditions.** Each agent's unconstrained optimal consumption is recovered from its own marginal value of wealth using the upwind selection above:

   $$c_p = (V^p_{a_p})^{-1/\gamma}, \qquad c_c^{\text{opt}} = (V^c_{a_c})^{-1/\gamma},$$

   with implied parental saving $s_p = ra_p + y_p - c_p$.

2. **Bounded gift.** The static transfer FOC $u'(c_p) = \eta\, u'(c_c)$ implies a target child consumption $c_c = \eta^{1/\gamma} c_p$, and hence a first-best gift $g_{\text{fb}} = \eta^{1/\gamma} c_p - R_c$, where $R_c = \tilde{r}_c\, a_c + y_c$ is the child's own resources. The gift is capped at the level that lifts the child to its *own* optimum (the second-best ceiling) and floored at zero:

   $$\tau = \max\Big\{\, 0,\ \min\big(\, c_c^{\text{opt}} - R_c,\ \ \eta^{1/\gamma} c_p - R_c \,\big) \Big\}. \tag{6}$$

   The point lies in the transfer region, $(i, j, k) \in \mathcal{T}^n$, if and only if $\tau > 0$. The lower bound is the one-sided no-negative-transfer constraint; the upper bound prevents a wealthy parent from transferring more than brings the child to its own optimum, which is what disciplines the gift to a finite, liquidity-insurance role.

3. **Wealth drifts.** The parent funds the gift out of its own saving and continues to optimize its wealth—it is *not* frozen in the transfer region:

$$\dot{a}_p = s_p - \tau, \qquad c_c = \min\left(c_c^{\text{opt}},\ R_c + \tau\right), \qquad \dot{a}_c = R_c + \tau - c_c.$$

Since $\tau \leq c_c^{\text{opt}} - R_c$, the child consumes its resources plus the entire gift, so $\dot{a}_c = 0$ inside the transfer region: the transfer is pure liquidity insurance to a constrained child, consistent with the mechanism in the main text.

4. **Construct the generator and solve.** Given the policies and drifts, assemble the sparse infinitesimal generator $\mathbb{A}^n$ over the full state space (a_p, a_c, z) of dimension $N_{a_p} \times N_{a_c} \times N_z$. The generator combines upwind drift operators in both wealth dimensions, the multiplicative wealth-diffusion blocks in a_p and a_c, and the OU operator in z.

A key computational insight is that the parent and child value functions share the *same* state transitions—both wealth drifts and the OU process are identical for both agents given the equilibrium policies. Therefore, the system matrix

$$\mathbf{B} = \left(\frac{1}{\Delta t} + \rho\right)\mathbf{I} - \mathbb{A}^n$$

is factorized *once* (via sparse LU), and used for two triangular solves with different right-hand sides:

$$\mathbf{B}\,\mathbf{V}^{p,n} = \mathbf{u}^p + \frac{1}{\Delta t}\,\mathbf{V}^{p,n+1} + \mathbf{J}^{p,n}, \tag{7}$$

$$\mathbf{B}\,\mathbf{V}^{c,n} = \mathbf{u}^c + \frac{1}{\Delta t}\,\mathbf{V}^{c,n+1} + \mathbf{J}^{c,n}, \tag{8}$$

where $\mathbf{u}^p = u(c_p) + \eta\, u(c_c)$ and $\mathbf{u}^c = u(c_c)$ are the flow utilities and $\mathbf{J}^{p,n}$, $\mathbf{J}^{c,n}$ carry the lump-sum jump terms defined in the next step.

5. **Lump-sum stage (operator splitting).** The Poisson lump-sum channel (described in the main text) is handled *explicitly*, so the nonlocal jump operator never enters the matrix $\mathbf{B}$ and the symbolic factorization is reused unchanged. At each time step with the channel active (child age ≥ 22), the lump policy is computed on the already-known

continuation values: at each node,

$$\tau_L^*(i,j,k) = \arg \max_{\tau_L \in \mathcal{M}(a_{p,i})} V^{p,n+1}(a_{p,i} - \tau_L,\ a_{c,j} + \tau_L,\ z_k)\,,$$

where the candidate set $\mathcal{M}(a_p)$ combines a fixed menu of small lumps with fractions of the parent's available wealth, all bounded by $\max(a_p, 0)$, and off-node candidates are valued by bilinear interpolation in the two wealth dimensions. Because the value function is concave, interpolation *under*states off-node values while $\tau_L = 0$ is evaluated exactly, so discretization error biases against spurious lumps. The jump contributions to (7)–(8) are then

$$\mathbf{J}^{p,n} = \lambda_g\big[\mathcal{M}\mathbf{V}^{p,n+1} - \mathbf{V}^{p,n+1}\big]\,, \qquad \mathbf{J}^{c,n} = \lambda_g\big[\mathcal{M}\mathbf{V}^{c,n+1} - \mathbf{V}^{c,n+1}\big]\,,$$

where $\mathcal{M}V$ denotes evaluation at the rebalanced state $(a_p - \tau_L^*, a_c + \tau_L^*, z)$ under the *parent's* policy in both lines (one-sided altruism: the child's value jumps with the parent's choice). The explicit treatment is unconditionally well-behaved here: the per-step jump weight is $\lambda_g \Delta t = 0.125$ at the baseline calibration, far inside the stability region, and the splitting follows the standard treatment of nonlocal terms in Achdou et al. (2022).

Regularity and the role of stochastic timing. The exogenous, exponentially-timed arrival of lump opportunities keeps the lump channel from disturbing the smoothness that the wealth diffusion delivers. The regularity proposition in the paper's Solution Algorithm appendix states this regularity result precisely and gives its argument: under the Poisson channel the value functions remain C^2 away from the free boundary $\partial\mathcal{T}$, so the Markov-perfect equilibrium is computed without the "blast from the past" non-smoothness of deterministic-date formulations (Barczyk and Kredler, 2021).

This procedure is repeated for each combination of child ability θ, child education $e_c \in \{\mathrm{HS}, \mathrm{C}\}$, parent education $e_p \in \{\mathrm{HS}, \mathrm{C}\}$, and the parent's permanent productivity component z_p (on a grid of $N_{z_p} = 5$ nodes), yielding $N_\theta \times 2 \times 2 \times N_{z_p} = 100$ coupled HJB problems that are solved in parallel across threads.

College-period modifications. When the child is in college (age < 22), two modifications apply: (i) the child earns a fraction ℓ_C of full-time income and pays net tuition $\phi_{\text{net}}(y_p^0) = \phi - \bar{g} \cdot \max(0,\, 1 - y_p^0/\bar{y}_{\text{efc}})$, which varies continuously with parent *income* at college entry through the EFC system; and (ii) the child may borrow up to $\bar{b}(a_p)$ at the student loan rate $r + \iota_s$.

Borrowing constraints. The child's wealth drift is constrained to be non-negative at the borrowing limit: if $a_c \leq \underline{a}_c + \epsilon$ and $\dot{a}_c < 0$, we set $\dot{a}_c = 0$. During college, $\underline{a}_c = -\bar{b}$; after college, $\underline{a}_c = 0$. The parent faces a natural borrowing constraint $a_p \geq 0$.

OA7.4 Stage 3: College Choice

At $t = 0$, the decision is solved in two stages on each (a_p, θ, e_p, z_p) node. First, the parent's college gift: candidate gifts τ_C combine a fixed menu of small gifts, the non-negative child-grid nodes, and their midpoints, all bounded by the parent's wealth a_p (the college-decision section of the main text); off-node candidates are valued by bilinear interpolation in (a_p, a_c), while the high-school branch is fixed at $(a_p, 0)$; for each candidate, the child's college-choice probit and the parent's college-gift objective from the main text are evaluated at $z = 0$, $t = 0$, and the maximizing τ_C^* is selected. No HJB re-solve is required: the coupled solutions already cover the (a_p, a_c) plane. Second, the child's choice: the value gain $G = V_C^c(a_p - \tau_C^*, \tau_C^*) - V_{HS}^c(a_p, 0)$ is compared to the realized psychic cost $\kappa(\theta, \varepsilon_\kappa) = (\kappa_0 + \kappa_\theta \log\theta + \sigma_\kappa \varepsilon_\kappa) \cdot \bar{V}$ with $\varepsilon_\kappa \sim N(0, 1)$, where $\bar{V} = \text{mean}_{a_p,\theta}|V_C^c - V_{HS}^c|$ (evaluated at $a_c = 0$) normalizes the cost to the scale of the value gap. Because the cost shock is Gaussian, the attendance probability $\Phi\big((G/\bar{V} - \kappa_0 - \kappa_\theta \log\theta)/\sigma_\kappa\big)$ is available in closed form, so no numerical integration is required. In the simulation, each dynasty re-solves the gift stage at its own (e_p, z_p) node with the ex-ante probit, and the child's choice is then deterministic given its drawn ε_κ; an enrolling child enters the overlap with $a_c = \tau_C^*$ and the parent with $a_p - \tau_C^*$, while a high-school child starts at $a_c = 0$ with the parent's wealth unchanged.

Boundary monotonicity correction. At very low parent wealth ($a_p \approx 0$), the coupled HJB can produce spurious college probabilities due to convergence difficulties at the corner of the state space. We enforce monotonicity by scanning from the bottom of the a_p grid upward: if the college probability at point i exceeds that at point $i+1$ for the first few grid points, we cap it at the first interior-monotone value. This correction affects at most the bottom third of the grid.

OA7.5 Stage 4: Dynastic Simulation

We simulate $M = 4{,}000$ dynasties across $N_{\text{coh}} = 11$ overlapping generations by forward-simulating the model's discretized law of motion at the same time step used to solve the HJB system.

Initialization. The first cohort begins at median ability $\theta = \theta_{(N_\theta+1)/2}$. Subsequent generations draw ability from the AR(1) process: $\log\theta' = \rho_\theta \log\theta + \mu_\theta(1-\rho_\theta) + \sigma_\theta\varepsilon$. Each new generation's parent wealth equals the previous generation's terminal child wealth (clamped to $[0, \bar{a}_p]$). The initial income shock z_0 is drawn from an education-specific entry distribution $z_0 \sim \mathcal{N}\big(\zeta_{pe}\,\mathbf{1}\{e_p = C\},\ \sigma_e^2/2\kappa_e + \sigma^2_{z_0,e_c}\big)$: the variance combines the OU stationary variance with the Abbott et al. (2019) entry-dispersion component σ_{z_0,e_c}, and the mean carries the parent-college earnings premium ζ_{pe} (the model section of the main text).

College decision. For each dynasty m in cohort g, we interpolate the college probability $\Pr(\mathrm{C} \mid a_p^{m,g}, \theta^{m,g})$ from the solved grid and draw education status from a Bernoulli distribution.

Forward simulation. Given education choices and initial conditions, we simulate the state dynamics forward in discrete time steps. At each time step $n = 1, \ldots, N_{\text{steps}}$:

1. **Lump-sum arrival**: with probability $1 - e^{-\lambda_g \Delta t}$ per step (active from child age 22), a transfer opportunity arrives; the lump policy τ_L^* is interpolated at the current state and executed as a discrete rebalancing, $a_p \mathrel{-}= \tau_L^*$, $a_c \mathrel{+}= \tau_L^*$, *before* the flow policies are applied—the discrete analog of jump-then-drift, consistent with the HJB timing in which the flow policies already price arrivals through the jump term.

2. **Interpolate policies**: trilinear interpolation of (c_p, c_c, τ) from the solved HJB policy grids at the (post-jump) state (a_p^n, a_c^n, z^n).

3. **Compute income**: parent income y_p (deterministic or Social Security if retired) and child income y_c (net of tuition during college), including the education-specific transitory shock with standard deviation $\sigma_{\varepsilon,e}$ (estimated by SMM; it enters realized income but not the state, with the Jensen correction $\mathbb{E}[e^{\varepsilon}] = 1$).

4. **Update wealth** (drift plus the multiplicative wealth-diffusion shock):

$$a_p^{n+1} = a_p^n + \big(r_{\text{eff}}(a_p^n)\, a_p^n + y_p - c_p - \tau\big)\,\Delta t + \sigma_a\, a_p^n \sqrt{\Delta t}\, \xi_p^n,$$
$$a_c^{n+1} = a_c^n + (\tilde{r}_c\, a_c^n + y_c - c_c + \tau)\,\Delta t + \sigma_a\, a_c^n \sqrt{\Delta t}\, \xi_c^n, \tag{9}$$

   with independent $\xi_p^n, \xi_c^n \sim \mathcal{N}(0,1)$; the diffusion term is set to zero whenever $a \leq 0$, matching the HJB.

5. **Update income shock**: $z^{n+1} = z^n - \kappa_e\, z^n\, \Delta t + \sigma_e \sqrt{\Delta t}\, \xi^n$, where $\xi^n \sim \mathcal{N}(0,1)$.

6. **Enforce constraints**: clamp a_p to $[0, \bar{a}_p]$ and a_c to $[\underline{a}_c, \bar{a}_c]$ (with $\underline{a}_c = -\bar{b}$ during college, $\underline{a}_c = 0$ after). Clamp z to the grid bounds.

Reproducibility. All random draws (initial z_0, college choice, OU innovations, transitory shocks, the parent and child wealth-diffusion innovations ξ_p, ξ_c, and the lump-sum arrival uniforms) are pre-generated sequentially before the threaded simulation loop, ensuring bit-identical results regardless of the number of parallel threads; the arrival uniforms are drawn last and only when $\lambda_g > 0$, so the $\lambda_g = 0$ economy reproduces the no-lump model draw-for-draw.

OA7.6 Stage 5: Moment Computation and SMM Estimation

The model is estimated by the simulated method of moments. It produces 43 moments—30 targeted—in the eight blocks defined in the estimation section of the main text (college attendance, 16; income, 3; transfers, 5; wealth, 4 targeted and 8 untargeted; education transmission, 1; decumulation, 1; college by parent income, 4 untargeted; negative net worth, 1 untargeted). All are computed from a settled cohort (the second-to-last generation) to avoid initial-condition transients. Two computations are specific to the simulation: the transfer *premium* differences the in-college transfer (the enrollment gift plus consumption-support flows) against an equal-length post-college window, and the transfer-*probability* moments count two-year windows from child age 26 in which cumulated flow plus lump-sum transfers reach the $500 HRS threshold.

Loss function. The objective is a weighted percentage-deviation loss with fixed per-moment diagonal weights,

$$\mathcal{L}(\boldsymbol{\vartheta}) = \sum_m w_m \left(\frac{m_m(\boldsymbol{\vartheta}) - \hat{m}_m}{\hat{m}_m} \right)^2, \tag{10}$$

with the weights $\{w_m\}$ fixed prior to estimation. Following Lee and Seshadri (2019), I do not weight all moments equally but combine a group-size normalization with deliberate adjustments toward the moments most informative about the model's mechanisms. The transfer and income blocks enter at their group-size-normalized weight. The sixteen college-attendance cells, the core target, carry elevated mass—downweighting them produces well-fitting wealth distributions but implausibly flat attendance gradients—and within that block the four high-wealth, low-ability cells, where the altruistic college gift shows up most clearly, are doubled. The two moments that discipline the parent wealth tail (the p90/p50 ratio and the share of parents with net worth below $25,000) and the intergenerational education gap are upweighted, while the intergenerational-wealth-transmission moments (rank–rank slope, conditional child wealth, child-wealth dispersion) and the share of children with negative net worth are left untargeted and reported only as validation. Table OA7.1 lists the exact weight on every moment. Deliberately weighting some moment groups more heavily than others is standard in the structural life-cycle literature (De Nardi et al., 2010; Abbott et al., 2019; Lee and Seshadri, 2019; Boar, 2021).

Table OA7.1. Per-Moment Weights in the SMM Objective

Moment block	#	Weight (each)
College attendance (parent-wealth Q × ability Q)	16	0.075
of which high-wealth, low-ability cells	(4)	0.150
Income: HS/C income ratio, var. log income (HS, C)	3	0.333
Transfers: mean transfer, $\Pr(\tau > 0)$ by parent-wealth Q	5	0.200
Aggregate wealth-to-income ratio	1	0.091
Median parent wealth	1	0.091
Parent decumulation rate	1	1.000
p90/p50 parent wealth	1	0.150
Share of parents with net worth < $25k	1	0.500
Intergenerational education gap	1	0.500

Notes: Per-moment diagonal weights w_m in the SMM objective (10), fixed prior to estimation. The table lists the 30 targeted moments; the remaining 13 moments are computed but left untargeted (weight 0). Within the college-attendance block, the four high-wealth, low-ability cells (parent-wealth $Q3$–$Q4$ × ability $Q1$–$Q2$) carry twice the weight of the other twelve.

Optimization. We minimize $\mathcal{L}$ over the 13 free parameters using a parallel $(\mu/\mu_w, \lambda)$-CMA-ES optimizer restricted to the free (non-fixed) subspace. The relative price of college labor ω_C lies outside this subspace: it is held fixed during the optimization, pinned to the value that reproduces the high-school-to-college income ratio (the income-process table in the estimation section of the main text), so the income-ratio moment is matched by construction rather than by a free parameter. Each candidate vector is evaluated by solving the full model (Stages 1–4) and computing the loss, and multi-start initialization ensures coverage of the parameter space.

OA7.7 Stage 6: Counterfactual Economies

Full commitment. The parent and child commit to a transfer schedule $\{\tau_t\}_{t\geq 0}$ at $t = 0$. This is solved as a single planner's problem that jointly maximizes parent and child welfare, eliminating the Samaritan's dilemma. The resulting college rates isolate the moral hazard effect against the full-commitment benchmark: $\Delta_{\mathrm{MH}}^{\mathrm{FC}} = \Pr(\mathrm{C} \mid \text{baseline}) - \Pr(\mathrm{C} \mid \mathrm{FC})$. This full-commitment comparison is distinct from the main-text decomposition, which uses the one-time-transfer economies ($\Delta_{\mathrm{MH}}^{\mathrm{pure}}$ and Δ_{targ} in the decomposition equations of the main text).

No-moral-hazard (college gift). The parent makes a one-time college-conditional gift τ_C at $t = 0$, paid only if the child attends college; thereafter the child solves the child-alone problem for its entire remaining life. The parent's problem is:

$$\max_{\tau_C \in [0, a_p]} \Pr(\mathrm{C} \mid \tau_C, \theta) \left[V_p^{\text{alone}}(a_p - \tau_C) + \eta\, V^{c,\text{alone}}(\tau_C \mid \mathrm{C}) \right] \\ + (1 - \Pr(\mathrm{C} \mid \tau_C, \theta)) \left[V_p^{\text{alone}}(a_p) + \eta\, V^{c,\text{alone}}(0 \mid \mathrm{HS}) \right],$$

so the child receives τ_C in the college state and nothing in the high-school state. Here $V_p^{\text{alone}}(\cdot)$ is the value of a single-agent parent problem—a continuous-time consumption-savings HJB over the parent's remaining life (labor income to age 66, social security thereafter, the consumption floor, the warm-glow bequest, and the same wealth diffusion $\sigma_a a\, dB$ as the baseline). Solving it as a genuine dynamic program, rather than approximating it by a permanent-income annuity, places the parent's counterfactual value on the same footing as the baseline coupled value, so the consumption-equivalent comparisons in the welfare section of the main text difference like with like. The child-alone value $V^{c,\text{alone}}(\tau_0 \mid e_c)$ is solved over the child's entire post-transfer life and likewise carries the wealth diffusion.

No-moral-hazard (menu). Same as above, but the parent commits to an education-contingent menu: a transfer τ_C paid if the child attends college and τ_{HS} paid if the child chooses high school, both optimized in $[0, a_p]$. This is the efficient one-time-transfer benchmark; it nests the unconditional transfer ($\tau_C = \tau_{HS}$) and measures the effect of removing moral hazard *plus* optimal education targeting.

Four-way decomposition. The difference in college rates across these allocations decomposes as:

$$\underbrace{\Pr(\mathrm{C} \mid \mathrm{BL}) - \Pr(\mathrm{C} \mid \mathrm{menu})}_{\Delta_{\mathrm{MH}}^{\mathrm{total}}} = \underbrace{\Pr(\mathrm{C} \mid \mathrm{BL}) - \Pr(\mathrm{C} \mid \mathrm{uncond.})}_{\Delta_{\mathrm{MH}}^{\mathrm{pure}}} + \underbrace{\Pr(\mathrm{C} \mid \mathrm{uncond.}) - \Pr(\mathrm{C} \mid \mathrm{menu})}_{\Delta_{\mathrm{targ}}}. \quad (11)$$

The signs are consistent with the main text: a positive $\Delta_{\mathrm{MH}}^{\mathrm{pure}}$ indicates that baseline moral hazard raises college attendance relative to the no-moral-hazard economy.

References

Abbott, Brant, Giovanni Gallipoli, Costas Meghir, and Giovanni L Violante, "Education policy and intergenerational transfers in equilibrium," *Journal of Political Economy*, 2019, *127* (6), 2569–2624.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Income and wealth distribution in macroeconomics: A continuous-time approach," *The Review of Economic Studies*, 2022, *89* (1), 45–86.

Barczyk, Daniel and Matthias Kredler, "Altruistically motivated transfers under uncertainty," *Quantitative Economics*, 2014, *5* (3), 705–749.

_ and _ , "A dynamic model of altruistically-motivated transfers," *Review of Economic Dynamics*, 2014, *17* (2), 303–328.

_ and _ , "Blast from the Past: The Altruism Model is Richer Than You Think," *Journal of Economic Theory*, 2021, *198*, 105364.

_ , Sean Fahle, and Matthias Kredler, "Save, Spend, or Give? A Model of Housing, Family Insurance, and Savings in Old Age," *The Review of Economic Studies*, 2023, *90* (5), 2116–2187.

Boar, Corina, "Dynastic Precautionary Savings," *The Review of Economic Studies*, 2021, *88* (6), 2735–2765.

Lee, Sang Yoon (Tim) and Ananth Seshadri, "On the Intergenerational Transmission of Economic Status," *Journal of Political Economy*, April 2019, *127* (2), 855–921.

Nardi, Mariacristina De, Eric French, and John Bailey Jones, "Why Do the Elderly Save? The Role of Medical Expenses," *Journal of Political Economy*, 2010, *118* (1), 39–75.