

College as a Commitment Device: Parental Altruism and the Samaritan's Dilemma*

Agustín Díaz Casanueva[†]

June, 2026

Latest Version

PRELIMINARY DRAFT — COMMENTS WELCOME

Abstract

Why do parents invest so heavily in their children's college education? I argue that college acts as a commitment device against a Samaritan's dilemma: anticipating parental support, children under-save and over-consume, and college mitigates this by permanently raising the child's income, shrinking the set of states in which parents optimally transfer. Consistent with this mechanism, in matched parent-child data parents of college children consume more and transfer less often (though in larger amounts when they do), with the consumption difference larger than the fall in realized transfers—a pattern in line with the precautionary resources parents release when they no longer need to stand ready to support the child. I formalize the mechanism in a continuous-time dynastic model with endogenous college choice, estimated by the simulated method of moments, and use it to decompose the effect of the lack of commitment on enrollment into two opposing forces: anticipated support during college subsidizes attendance, while anticipated lifetime support cushions the child who forgoes college and depresses it. The balance of these forces—and hence the direction of the enrollment distortion—varies across the joint distribution of ability and parental wealth. The welfare consequences of removing the friction—and of policies such as free college—depend on whether parents can commit, since the same transfers that distort the college decision also insure children against income risk.

JEL: D15, D64, I22, I23

Keywords: Parental altruism, college attendance, Samaritan's dilemma, intergenerational transfers, moral hazard, dynastic models

*Previously circulated as “The Role of Parental Altruism in Parents Consumption, College Financial Support, and Outcomes in Higher Education.” I am grateful to my committee members Dirk Krueger, José Víctor Ríos Rull, and Andrew Shephard for their invaluable guidance. I also thank seminar participants at the University of Pennsylvania, Universidad Católica Alumni Workshop, the Midwest Macro Meetings, and the Chilean Economics Society for helpful comments. The views expressed herein are solely those of the author and do not necessarily reflect the views of the Central Bank of Chile. All errors are my own.

[†]Central Bank of Chile, adiaz@bcentral.cl, Address: Agustinas 1180, Santiago, Chile.

1 Introduction

Parents in the United States transfer substantial resources to their adult children. Inter-vivos transfers—cash gifts, tuition payments, housing assistance—flow overwhelmingly downward, from parents to children, and persist well into adulthood (McGarry, 1999, 2016). These transfers are infrequent—in the Health and Retirement Study, only a small minority of parent-child pairs report a positive transfer in any given wave—but when they occur, they are sizable. Parental transfers are also a central input into children’s human capital formation: families remain the primary source of college financing in the United States, and the anticipation of parental financial support is a key determinant of whether children enroll (Belley and Lochner, 2007; Brown et al., 2012; Abbott et al., 2019). From a public policy perspective, this link matters: government financial aid crowds out private parental contributions, and this crowding out is sizable precisely because family transfers are responsive to the child’s enrollment decision (Abbott et al., 2019). Understanding why parents transfer, and how these transfers—both realized and anticipated—shape children’s choices, is therefore fundamental both for explaining college attendance and for designing effective education policy.

This paper studies these questions through the lens of the Samaritan’s dilemma (Lindbeck and Weibull, 1988). Altruistic parents who cannot credibly commit to withholding future support create an incentive problem: children who anticipate parental transfers under-save and over-consume, knowing that their parents will step in when resources are low. This strategic response induces parents to make larger and more frequent transfers than they would under commitment: parents support their children *because* they care, but the anticipation of that support distorts the decisions that determine how much support will be needed. Because transfers are closely tied to the college decision, the distortion extends to human capital investment.

I argue that college education plays a distinctive role in this dynamic. From the parent’s perspective, a dollar spent on tuition is fundamentally different from a dollar transferred directly. Tuition generates a permanent shift in the child’s income level, moving the child away from the states of the world where parental support would be triggered. A direct transfer, by contrast, provides temporary relief and may itself exacerbate moral hazard by rewarding under-saving. College thus serves as a commitment device: by investing in the child’s human

capital, parents can shrink the *transfer region*—the set of states where they would optimally provide support—thereby reducing the lifetime costs of the Samaritan’s dilemma. The net effect on college attendance, however, is theoretically ambiguous. Anticipated transfers during the college years subsidize enrollment, raising attendance, but anticipated lifetime transfers also provide implicit insurance to children who forgo college, reducing their incentive to invest in human capital. The structural model quantifies these opposing forces and shows that the sign of the moral hazard effect on college varies across the ability–wealth distribution.

I formalize this mechanism in a continuous-time, non-cooperative dynastic model, building on [Barczyk and Kredler \(2014a,b\)](#). Parents choose consumption and a non-negative transfer flow to their adult children; children choose consumption and, at the start of the dynasty, whether to attend college. The interaction between parent and child is non-cooperative: each agent optimizes taking the other’s policy as given, while the parent internalizes the child’s welfare through an altruism parameter. In equilibrium, the state space is endogenously partitioned into a transfer region, where parents actively support their children, and a no-transfer region, where the child is self-sufficient. The boundary between these regions—and hence the insurance value of parental support—is shaped by the child’s education, income, and accumulated wealth. College contracts the transfer region by permanently raising the child’s expected earnings, thereby reducing the insurance value of future parental transfers. The child’s college decision therefore depends not only on the human capital return, but also on how education reshapes the future financial relationship with the parent.

I document two stylized facts from the Panel Study of Income Dynamics (PSID) and the Health and Retirement Study (HRS), comparing parents of college and non-college children while conditioning on parent income. *First*, parent consumption is markedly higher when the child attends college. *Second*, college lowers the probability of any within-family transfer, while conditional amounts are weakly larger. Both line up with the model’s central mechanism, that college shrinks the transfer region: the consumption gain is far larger than the implied fall in realized transfers, a pattern consistent with parents releasing resources because they no longer need to *stand ready* to support the child rather than because they pay fewer transfers. The estimated model quantifies this option-value channel directly.

I estimate the model by the simulated method of moments, targeting college attainment

by ability and parent wealth, the college premium, transfer patterns, and intergenerational wealth dynamics. I then use the estimated model to decompose the effect of the lack of commitment on enrollment relative to an efficient benchmark in which the parent commits to education-contingent transfers at the start of the dynasty. The decomposition separates the enrollment subsidy that anticipated in-college support provides from the cushion that anticipated lifetime support extends to children who forgo college, and shows that the balance of the two forces—and hence the direction and incidence of the enrollment distortion—varies across the joint distribution of ability and parental wealth. Removing the friction is not unambiguously beneficial, because the transfers that distort the college decision also insure children against income risk; dynastic welfare rises only when commitment preserves that insurance rather than severing the transfer relationship altogether. The welfare effects of a free-college policy likewise depend on whether parents can commit.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 presents the model. Section 4 documents the stylized facts. Section 5 describes the estimation strategy. Section 6 discusses the model fit and the role of parental altruism. Section 7 introduces the full commitment benchmark, formalizes the two opposing forces of moral hazard on college, and presents the quantitative decomposition. Section 8 conducts the welfare analysis. Section 9 concludes.

2 Related Literature

This paper draws on several literatures: intergenerational transfers and family insurance, the role of parental resources in major life-cycle investments, saving and wealth decumulation in retirement, and the determinants of college attainment.

A large body of work studies how families insure members against income risk. Since [Becker and Tomes \(1979, 1986\)](#), two broad motives for intergenerational transfers have been proposed: altruism, supported by evidence that transfers flow disproportionately to less well-off children ([McGarry, 1999, 2016](#)), and strategic exchange, as in [Bernheim et al. \(1985\)](#). [Fu et al. \(2026\)](#) estimate a dynamic model of two-way time transfers—grandparental childcare reciprocated by elderly care—and find that the exchange motive operates mainly through these time transfers, while altruism runs asymmetrically from parents toward children and

financial flows remain overwhelmingly downward; both findings are consistent with the one-sided altruism over financial transfers modeled here. The seminal tests of [Altonji et al. \(1992\)](#) and [Hayashi et al. \(1996\)](#) reject perfect insurance within the extended family, but [Attanasio et al. \(2018\)](#) finds substantial insurance potential between parents and adult children, and [Kaplan \(2012\)](#) shows that the option to move back home insures young workers against labor market risk. Most directly related to this paper, [Boar \(2021\)](#) identifies dynastic precautionary savings as a distinct insurance channel: parents accumulate wealth to buffer their children against permanent income uncertainty. I build on this framework: in my model college education attenuates this dynastic precautionary motive by raising the child's income and shrinking the transfer region, so parents of college children hold less wealth against the contingency of supporting them.

Intergenerational transfers, however, are not limited to insurance against income shocks. Parents also provide resources for large, lumpy investments such as college attendance and home purchases, and they transmit wealth through inter-vivos gifts and bequests. This broader view is important because parental transfers may affect children's outcomes not only by smoothing consumption, but also by relaxing borrowing constraints, changing portfolio choices, and shaping exposure to risk. [Nishiyama \(2002\)](#) studies altruistic and accidental bequests in an overlapping-generations model and shows how bequest motives affect wealth accumulation. [Barczyk et al. \(2023\)](#) emphasize the interaction between housing, family insurance, and the timing of transfers, showing that illiquid housing can help explain why many transfers are delayed until death. Closest to this broader mechanism, [Brandsaas \(2025\)](#) develops a life-cycle model with housing and parent-child transfers and finds that parental transfers account for a sizable share of young-adult homeownership by relaxing down-payment constraints, reducing housing risk, and interacting with commitment frictions. This paper builds on the insight that family transfers affect children's long-run investment, but studies this mechanism in the context of human capital accumulation.

Because parents here hold wealth in part to stand ready to support adult children, the paper also connects to the literature on saving and wealth decumulation in retirement, where older households draw down wealth far more slowly than a life-cycle model without family motives predicts, and the income-rich decumulate most slowly of all. [De Nardi et al. \(2010\)](#)

trace much of this slow drawdown to longevity and medical-expense risk, while [Love et al. \(2009\)](#) show that annualized wealth is roughly flat or rising over retirement and that intended bequests are needed to match the profiles of higher-income households in the Health and Retirement Study. Family links do much of the remaining work: [De Nardi \(2004\)](#) shows that voluntary bequests are required to reproduce both the slow late-life drawdown and the concentration of wealth at the top, [De Nardi and Yang \(2014\)](#) that they generate the dispersion of wealth observed at retirement, and [Lockwood \(2018\)](#) that a bequest motive can rationalize why retirees decline to insure late-life risks. Closest to the motive I emphasize, [Ameriks et al. \(2016\)](#) separate bequest from precautionary savings and find that parents hold wealth to assist children when the need is greatest. This literature models the family motive largely as a terminal bequest or treats old-age risk as exogenous; I instead micro-found a forward-looking inter-vivos motive—parents retain wealth because of the option value of an endogenous transfer region in which supporting a child is optimal—and show that this region, and hence the predicted pace of decumulation, depends on the child’s college. The estimation disciplines this margin directly by targeting the pace of late-life drawdown in the PSID (Section 5).

A separate literature asks whether financial constraints bind for college attendance. [Cameron and Heckman \(1998\)](#) and [Keane and Wolpin \(2001\)](#) found that, conditional on ability, family income played a small role, but [Belley and Lochner \(2007\)](#) document a growing importance of family resources in the 2000s, and [Lochner and Monge-Naranjo \(2011\)](#) provide evidence of binding constraints for some students. [Brown et al. \(2012\)](#) emphasize that the expected family contribution is neither legally guaranteed nor universally provided, implying that students with similar observed family resources may face very different access to parental support. [Hotz et al. \(2018\)](#) provide direct evidence that parental income and housing wealth affect college attendance, parental financing, and graduation, and that parental financing may affect the subsequent indebtedness of both parents and children. On the quantitative side, [Restuccia and Urrutia \(2004\)](#) study how parental investments drive intergenerational earnings persistence, [Caucutt and Lochner \(2020\)](#) analyze the interaction between borrowing constraints and the timing of human capital investments, and [Heckman and Mosso \(2014\)](#) provide a comprehensive treatment of how family environments shape skill development.

Closest to my setting, [Abbott et al. \(2019\)](#) develop a general equilibrium model with parental transfers and show that crowding out of private contributions by government financial aid is quantitatively important. My paper adds to this work by providing a mechanism—moral hazard reduction through the Samaritan’s dilemma—through which parental wealth affects college enrollment beyond direct cost subsidies and paternalism.

The model builds on the quantitative literature on family dynamics without commitment, including [Laitner \(1988\)](#) and [Lindbeck and Weibull \(1988\)](#). Within this tradition, [Nishiyama \(2002\)](#) studies how bequests and inter-vivos transfers jointly shape the wealth distribution, and [Lee and Seshadri \(2019\)](#) develop a dynastic model with parental investments in children’s human capital to study the intergenerational transmission of economic status. I adopt the continuous-time framework of [Barczyk and Kredler \(2014a,b\)](#), who show that modeling simultaneous consumption and transfer decisions in continuous time yields a well-defined equilibrium with an endogenous transfer region—the set of states where the parent optimally provides support. [Barczyk and Kredler \(2018\)](#) and [Barczyk et al. \(2023\)](#) apply this framework to long-term care and housing decisions; I extend it by embedding endogenous college decisions, which allows me to study how the transfer region—and the moral hazard it generates—interacts with education choice.

The Samaritan’s dilemma that drives this moral hazard was formalized and analyzed in dynamic settings by [Lindbeck and Weibull \(1988\)](#) and [Bruce and Waldman \(1990\)](#). In my model, the distortion takes the form of strategic under-saving by the child, which raises the parent’s lifetime transfer costs. College mitigates the dilemma by permanently shifting the child’s income, reducing the frequency with which the child enters the transfer region. The net effect on college attendance, however, is ambiguous, because anticipated transfers both subsidize enrollment and cushion the non-college path.

3 Model

This section presents a continuous-time, non-cooperative dynastic model without commitment that captures the interactions between altruistic parents and their adult children. The model extends the framework of [Barczyk and Kredler \(2014a,b\)](#) by incorporating endogenous college decisions, allowing me to study how family transfers affect college financial support,

graduation, and parent retirement consumption. I first describe the demographic structure and income processes, then present the coupled parent-child optimization problem, the parent’s transfer decision, and the college decision.

3.1 Demographics and Timing

Time is continuous. Each dynasty consists of one parent and one child who overlap for a finite period, as illustrated in Figure 1. The child enters the model at age 18, at which point the parent is 42. Parent and child are coupled for the parent’s entire remaining life: the overlap runs from child age 18 / parent age 42 to the parent’s death at age 72, when the child is 48. The parent works until age 66 ($t = T^{\text{ret}}$), earning labor income $y_p(t)$, and then receives social security income $\text{SS}(e_p)$ until death at age 72 ($t = T$, where $T = T^{\text{ret}} + 6$). The retired parent remains coupled with the child, so transfers $\tau(t) \geq 0$ are available throughout the overlap $[0, T]$, including during the parent’s retirement; the coupled problem is solved over the parent’s full remaining life. Death is deterministic at age 72 (the model has no mortality risk). At the parent’s death, a fraction α_{beq} of remaining parent wealth transfers to the child as a bequest; the remaining $(1 - \alpha_{\text{beq}})$ share is not transmitted.¹ The child, now 48, then continues alone—working until its own retirement at 66 and consuming social security until death at 72.

At the beginning of the dynasty, the child makes a one-time, irreversible college decision. Children who attend college spend four years (ages 18–22) earning a fraction ℓ_C of full-time income while paying annual tuition ϕ . After college, the child enters the labor force with a permanently higher income process. Children who do not attend college enter the labor market immediately at age 18.

Both parent and child can save in a risk-free bond paying interest rate r . The only debt instrument is the student loan: during and after college the child pays $r + \iota_s$ on any negative

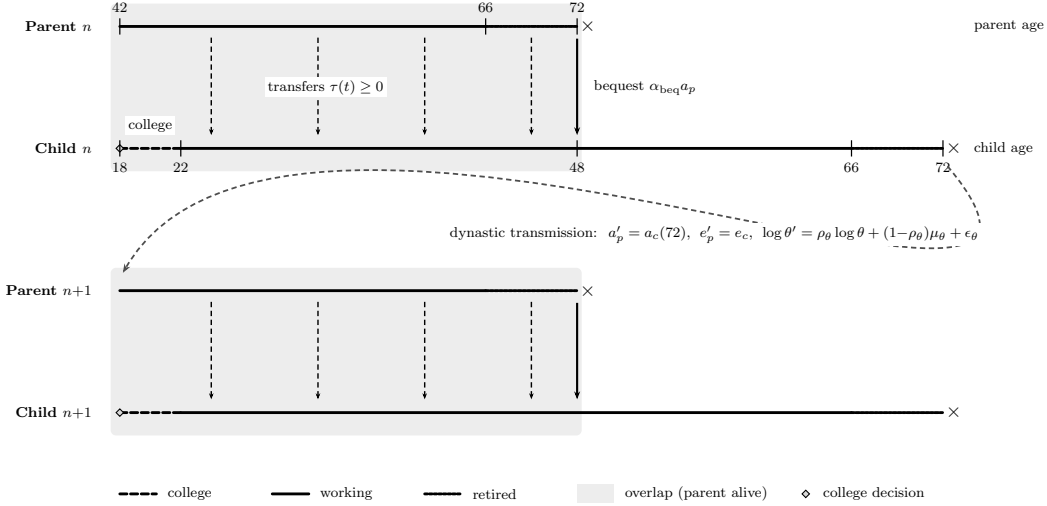
¹The bequest fraction α_{beq} separates the two roles that terminal wealth plays in the model. The warm-glow motive ϕ_b requires parents to hold substantial wealth at death—this is what matches the absence of decumulation at ages 60–70—while the intergenerational wealth moments (the rank–rank slope and child wealth conditional on parent wealth) discipline how much of that wealth children actually receive. The untransmitted share $(1 - \alpha_{\text{beq}})$ is a reduced-form stand-in for end-of-life wealth absorption that the model’s compressed lifespan omits: death is deterministic at 72, so the later years in which out-of-pocket medical and long-term-care spending absorb much of retirement wealth (De Nardi et al., 2010; Lockwood, 2018) are not modeled, and the wealth that would be consumed in them is treated as leaving the dynasty rather than reaching the child.

balance ($a_c < 0$), and this college debt must be amortized to zero by the time the child becomes a parent (age 42). The child cannot otherwise borrow as an adult, so outside the student-loan window $a_c \geq 0$. The parent faces a non-negativity constraint $a_p \geq 0$ and always earns r . Markets are incomplete: no state-contingent insurance is available against the child's income risk. The parent and child are connected through three financial links: a one-time college-conditional gift paid at enrollment (Section 3.6), the transfer flow during the overlap period, and a bequest at the parent's death, when a fraction α_{beq} of remaining parent wealth transfers to the child.

Two features deserve comment. First, altruism is one-sided: the parent cares about the child's welfare, but the child does not internalize the parent's utility. Second, transfers flow only downward—from parent to child—and the overlap is finite, ending at the parent's death (age 72). A natural alternative motive for financing college is that parents invest in their children's human capital expecting support in return—an exchange motive in the tradition of [Bernheim et al. \(1985\)](#). I abstract from this motive because intergenerational financial transfers in the United States flow overwhelmingly from parents to children.²

²[McGarry \(1999\)](#) documents that inter-vivos transfers from parents to children are an order of magnitude larger than transfers in the reverse direction, and [McGarry \(2016\)](#) confirms the asymmetry persists over the life cycle; old-age consumption is financed primarily through Social Security, pensions, and private savings. [Fu et al. \(2026\)](#) document two-way transfers of *time*—grandparental childcare and elderly care—but estimate a pronounced asymmetry in altruism, running from parents toward children.

Figure 1. Dynasty Timeline



Notes: Two consecutive generations of the dynasty. Within each generation, the parent (upper line) and child (lower line) are coupled for the parent’s entire remaining life (parent ages 42–72, child ages 18–48; shaded region), during which transfers $\tau(t) \geq 0$ flow from parent to child (dashed vertical arrows); the retired parent (ages 66–72, dotted) can still transfer. The dashed segment on each child line is the college phase (ages 18–22), preceded by the college decision (\diamond). At the parent’s death (child age 48), a fraction α_{beq} of remaining wealth transfers to the child as a bequest (solid vertical arrow); the rest is not transmitted. The child then lives alone (ages 48–72), and its terminal state seeds the next generation’s parent (curved dashed arrow), closing the dynasty as a stationary recursion. Event labels and ages are shown once, on generation n .

3.2 State Variables and Income Processes

The state of a dynasty at any instant t during the overlap phase is $(a_p(t), a_c(t), z(t); \theta, z_p, e_p, e_c)$, where a_p and a_c are parent and child wealth, z is the child’s stochastic productivity shock, θ is the child’s cognitive ability, z_p is the parent’s permanent productivity component (the productivity state the parent carried over as the previous generation’s child), and $e_p, e_c \in \{HS, C\}$ denote the education levels of parent and child.

Child income. The child’s labor income at time t is:

$$y_c(t) = w_{e_c} \cdot \left(\frac{\theta}{\bar{\theta}} \right)^{\lambda_{e_c}} \cdot \exp(\gamma(t, e_c) + z(t))$$

where $h(\theta, e_c) = (\theta/\bar{\theta})^{\lambda_{e_c}}$ maps ability into human capital with an education-specific ability gradient λ_{e_c} , following [Abbott et al. \(2019\)](#), and $\bar{\theta} = \exp(\mu_\theta)$ is the population mean ability.³ The gradient $\lambda_C > \lambda_{HS}$ implies that cognitive ability and college education are complements: higher-ability individuals benefit more from college. The education-specific wage w_{e_c} captures the relative price of education-specific labor: $w_{HS} = w$ and $w_C = w \cdot \omega_C$. The deterministic age-income profile is quadratic in age, $\gamma(t, e_c) = \gamma_{1,e} t - \gamma_{2,e} t^2$, and $z(t)$ follows an Ornstein-Uhlenbeck process:

$$dz = -\kappa_{e_c} z dt + \sigma_{e_c} dW$$

with education-specific mean-reversion rate κ_{e_c} and diffusion σ_{e_c} ([Appendix D](#)). The channel through which college reshapes the child’s income distribution, and hence the transfer region, is the income *level*: a higher ability gradient ($\lambda_C > \lambda_{HS}$), a steeper age-income profile ($\gamma_{1,C} > \gamma_{1,HS}$), and the relative wage ω_C jointly raise expected income and keep college-educated children far from the states where parental transfers would be triggered. During college ($t \in [0, 4]$ for $e_c = C$), the child works part-time, earning a fraction ℓ_C of full-time earnings. Because the degree has not yet been completed, these earnings are paid at the high-school level—the wage w_{HS} , ability gradient λ_{HS} , and age profile $\gamma(t, HS)$ —and the college premium applies only after graduation.⁴

Parent income. Parent income depends on education, age, and a permanent productivity component z_p , fixed over the parent’s life: before retirement, $y_p(t) = w_{e_p} \cdot (\theta/\bar{\theta})^{\lambda_{e_p}} \cdot \exp(\gamma(t, e_p) + z_p \nu(t))$, where $\nu(t) = \sqrt{1 + g_{z_p} \max(0, \text{age}_p(t) - 42)}$ is a deterministic age loading with $\nu = 1$ at age 42, and $y_p(t) = \text{SS}(e_p)$ after retirement at age 66. The parent draws no new idiosyncratic income shocks, so the child’s stochastic income z is the only source of within-dynasty uncertainty that drives transfers.⁵

³The normalization by $\bar{\theta}$ ensures that the human capital component equals one at mean ability for both education groups, so that the average college premium comes entirely from the life-cycle profile and wage price ω_C , not from the level effect of $\lambda_C \neq \lambda_{HS}$ evaluated at a non-zero mean $\log \theta$.

⁴In-college earnings therefore rise only weakly with ability, so high-ability students cannot self-finance net tuition out of part-time work and parental resources remain relevant throughout the ability distribution.

⁵The component z_p is the realization of the persistent productivity state the parent held at the dynastic hand-off (age 42), when it was the previous generation’s child; from that point on the state is no longer updated.

Ability transmission. Cognitive ability is transmitted across generations through a log-normal AR(1) process: $\log \theta' = \rho_\theta \log \theta + \mu_\theta(1 - \rho_\theta) + \epsilon_\theta$, with $\epsilon_\theta \sim N(0, \sigma_\theta^2)$ and μ_θ the unconditional mean of log ability. The drift term $\mu_\theta(1 - \rho_\theta)$ ensures that the stationary distribution of ability has mean μ_θ , generating realistic intergenerational persistence while allowing for regression to the mean.

Intergenerational earnings transmission. In addition to ability and bequests, the model includes a third intergenerational link operating through initial earning capacity. A child whose parent attended college enters the labor market with a permanently higher initial productivity draw: $z_0 \sim N(\zeta_{pe} \mathbf{1}\{e_p = C\}, \sigma_{z_0, e_c}^2)$, where $\zeta_{pe} \geq 0$ is a parent-education earnings premium and σ_{z_0, e_c} is an education-specific entry dispersion. The model begins at age 18, so the childhood through which parental background shapes earning capacity—investments in skills, schooling quality, and environments (Restuccia and Urrutia, 2004; Lee and Seshadri, 2019)—is pre-determined; ζ_{pe} is its reduced form, the counterpart of the parental-education dependence of children’s skill endowments in Abbott et al. (2019). Because the early-career income advantage accumulates permanently into the wealth stock even as the mean-reverting shock z decays, this channel routes part of the parent–child wealth correlation through inherited earning capacity rather than liquid transfers, leaving the child’s age-18 college-financing constraint untouched. Absent this channel, the wealth-transmission moments would load entirely on ability transmission and on the financial channels—transfers and bequests—overstating the role of family resources in college attendance.

Asset-return risk. Following Barczyk and Kredler (2014a), each agent’s wealth carries a small multiplicative shock,

$$da_p = \dot{a}_p dt + \sigma_a a_p dB_p, \quad da_c = \dot{a}_c dt + \sigma_a a_c dB_c,$$

where B_p and B_c are independent standard Brownian motions, the drifts \dot{a}_p, \dot{a}_c are given in (8)–(9), and the volatility vanishes at the borrowing constraint ($\sigma_a a = 0$ at $a = 0$). The shock admits an interpretation as modest idiosyncratic asset-return risk, but its essential role is to regularize the equilibrium. Barczyk and Kredler (2014a) show that a diffusion in

the *wealth* state renders the value functions C^2 almost everywhere and the transfer boundary smooth, restoring existence of a pure-strategy Markov equilibrium and eliminating the non-smoothness—the “blast from the past” of [Barczyk and Kredler \(2021\)](#)—that arises in deterministic and discrete-time formulations. The income shock z alone does not deliver this: it does not smooth the transfer boundary, which lives in the wealth state. I set σ_a to the minimal magnitude (about three percent per year) that suffices for the regularization, so the wealth distribution is shaped by saving and bequests rather than by return risk.

3.3 The Coupled Parent-Child Problem

During the overlap period, parent and child simultaneously choose consumption and savings, while the parent additionally chooses a non-negative transfer flow $\tau(t) \geq 0$. The interaction is non-cooperative: each agent optimizes taking the other’s policy as given, but the parent internalizes the child’s welfare through the altruism parameter $\eta > 0$.

The continuous-time formulation follows [Barczyk and Kredler \(2014a\)](#), who show that simultaneity of moves in continuous time yields a well-defined equilibrium—an advantage over discrete-time formulations where the timing protocol (who moves first within a period) can affect the equilibrium outcome.

3.3.1 Parent’s Problem

The parent chooses consumption c_p and transfer flow $\tau \geq 0$ to maximize lifetime utility, weighting the child’s instantaneous utility by η . The parent’s value function $V^p(a_p, a_c, z; t)$ satisfies:

$$\begin{aligned} \rho V^p = \max_{c_p, \tau \geq 0} & \left\{ u(c_p) + \eta u(c_c^*) + V_{a_p}^p [r a_p + y_p - c_p - \tau] + V_{a_c}^p [\tilde{r}_c a_c + y_c - c_c^* + \tau] \right. \\ & + V_z^p (-\kappa z) + \frac{1}{2} \sigma_z^2 V_{zz}^p + \frac{1}{2} \sigma_a^2 a_p^2 V_{a_p a_p}^p + \frac{1}{2} \sigma_a^2 a_c^2 V_{a_c a_c}^p + V_t^p \\ & \left. + \lambda_g [V^p(a_p - \tau_L^*, a_c + \tau_L^*, z; t) - V^p(a_p, a_c, z; t)] \right\} \end{aligned} \quad (1)$$

where ρ is the continuous-time discount rate, $u(c) = c^{1-\gamma}/(1-\gamma)$ is CRRA utility with $\gamma = 1.5$, c_c^* is the child’s equilibrium consumption policy, taken as given by the parent, and \tilde{r}_c is the child’s effective interest rate, which accounts for student debt (defined in [Section 3.3.5](#)).

Consumption is bounded below by a floor $\underline{c} = \$3,900$ (in 2016 dollars), representing the effective consumption guaranteed by means-tested government transfers (SSI, SNAP, Medicaid), following De Nardi (2004). The terms involving V_z^p and V_{zz}^p capture the effect of income uncertainty on the parent's value function. The final term is the lump-sum transfer channel: at Poisson arrivals with intensity λ_g , the parent may execute a one-time wealth rebalancing τ_L^* toward the child (Section 3.3.4).

3.3.2 Child's Problem

The child chooses consumption c_c to maximize own lifetime utility:

$$\begin{aligned} \rho V^c = \max_{c_c} & \left\{ u(c_c) + V_{a_c}^c [\tilde{r}_c a_c + y_c - c_c + \tau^*] + V_{a_p}^c [r a_p + y_p - c_p^* - \tau^*] \right. \\ & + V_z^c (-\kappa z) + \frac{1}{2} \sigma^2 V_{zz}^c + \frac{1}{2} \sigma_a^2 a_p^2 V_{a_p a_p}^c + \frac{1}{2} \sigma_a^2 a_c^2 V_{a_c a_c}^c + V_t^c \\ & \left. + \lambda_g [V^c(a_p - \tau_L^*, a_c + \tau_L^*, z; t) - V^c(a_p, a_c, z; t)] \right\} \end{aligned} \quad (2)$$

where τ^* and c_p^* are the parent's equilibrium policies, taken as given, and the jump term reflects that the child's state moves with the *parent's* lump-sum policy τ_L^* at transfer opportunities—altruism is one-sided, so the child receives but does not choose. The child tracks parent wealth a_p because it affects expected future transfers: wealthier parents are more likely to provide support.

3.3.3 Transfer Decision

The parent's optimal transfer is characterized by a complementarity condition:

$$\tau^* \geq 0, \quad V_{a_p}^p \geq V_{a_c}^p, \quad \tau^* (V_{a_p}^p - V_{a_c}^p) = 0 \quad (3)$$

This partitions the state space into two regions. In the *no-transfer region* (\mathcal{N}), the parent's marginal value of own wealth exceeds the marginal value of child wealth ($V_{a_p}^p > V_{a_c}^p$), and $\tau^* = 0$. In the *transfer region* (\mathcal{T}), these marginal values are equalized ($V_{a_p}^p = V_{a_c}^p$) and $\tau^* > 0$. The boundary between \mathcal{N} and \mathcal{T} is a free boundary, endogenously determined as part of the equilibrium. The parent transfers when the child is relatively poor (low a_c), faces

negative productivity shocks (low z), or when the parent is relatively wealthy (high a_p).

This complementary-slackness condition is the continuous-time analog of the transfer decision in discrete-time dynastic models, but with two key advantages. First, the flow transfer carries no within-period commitment: a discrete-time transfer is implicitly a promise held fixed for the length of the period, whereas the flow can be revised at every instant. Second, the boundary of the transfer region provides a clean characterization of the states where moral hazard is operative. Flow transfers are not the only instrument: at stochastically arriving opportunities the parent can also move wealth in discrete amounts (Section 3.3.4), and the same free boundary governs when it chooses to do so.

Transfer size. Within \mathcal{T} , the transfer is determined at the consumption margin. The static optimality condition $u'(c_p) = \eta u'(c_c)$ implies an altruistic consumption target $c_c = \eta^{1/\gamma} c_p$ for the child, and the parent funds the gap between this target and the child's own flow resources $R_c = \tilde{r}_c a_c + y_c$, capped at the consumption level c_c^{opt} the child would choose for itself from its first-order condition:

$$\tau^* = \max\left\{0, \min\left(c_c^{\text{opt}} - R_c, \eta^{1/\gamma} c_p - R_c\right)\right\}. \quad (4)$$

The child consumes $R_c + \tau^*$ inside \mathcal{T} : the transfer is consumption support, not a wealth transfer. The cap at c_c^{opt} reflects this—any gift beyond the child's desired consumption would be saved by the child, converting parent wealth into child wealth.⁶

The flow transfer never moves the parent's wealth stock, for two reasons. Mechanically, τ is a flow: over an interval dt the child receives τdt , and moving a discrete amount of wealth in an instant would require an unbounded flow, which lies outside the flow strategy space. Economically, the absence of commitment makes pre-funding unattractive: wealth placed in the child's account is controlled by the child, who would consume it too quickly and under-save against future shocks. By retaining the wealth, the parent preserves the option to support the child if and when the child is actually constrained, so the flow provides liquidity

⁶Transfers are therefore not confined to children exactly at the borrowing constraint, as in the no-labor-income setting of [Barczyk and Kredler \(2014b\)](#) where gifts flow only to a recipient with zero wealth. With flow income, the gift (4) is positive whenever the child's flow resources fall short of both its desired consumption and the altruistic target, $R_c < \min(c_c^{\text{opt}}, \eta^{1/\gamma} c_p)$, so the transfer region concentrates at low child wealth and adverse income draws but its boundary lies at strictly positive a_c (Figure 2).

period by period rather than capital, and the parent continues to optimize its own saving inside \mathcal{T} ($\dot{a}_p = ra_p + y_p - c_p - \tau^*$). Appendix I derives these properties formally.

3.3.4 Lump-Sum Transfer Opportunities

The flow transfer is consumption support by construction—capped at the child’s desired consumption, it props up the child’s spending in bad states but never adds a dollar to the child’s balance sheet. Observed inter-vivos giving is not so confined: parents make discrete, sizeable transfers—help with a house down payment, a wedding, support around the birth of a grandchild—that move *wealth* between generations well before any bequest. In the HRS, transfers of \$500 or more in a two-year window occur for 5 to 25 percent of parent-child pairs depending on parent wealth (Section 4.3), and their lumpiness is intrinsic to how they arrive: they are occasioned by life events, not paid as annuities. A model whose only wealth-moving instruments are a tuition gift at age 18 and a bequest at the parent’s death is asked to match intergenerational wealth transmission with an instrument set that, by construction, contributes nothing to it in between.

There is also a methodological reason why this channel is absent from continuous-time altruism models, and a discipline on how to restore it. [Barczyk and Kredler \(2014a\)](#) observe that discrete-time dynastic models implicitly assume the donor can commit to a lump-sum transfer for the length of a period—commitment smuggled in by the time aggregation itself—and the continuous-time reformulation removes it; but in removing the artificial commitment it also removes the lump, leaving only the flow. I restore the lump without restoring the commitment: at exogenous Poisson arrivals with intensity λ_g , the parent may execute a one-time transfer $\tau_L \in [0, \max(a_p, 0)]$, chosen at the moment of the opportunity to maximize its own value,

$$\tau_L^*(a_p, a_c, z; t) = \arg \max_{\tau_L \in [0, \max(a_p, 0)]} V^p(a_p - \tau_L, a_c + \tau_L, z; t). \quad (5)$$

Because the arrival is exogenous, the parent chooses neither the timing nor a schedule: the exercise is an instantaneous action, not a promise, so no commitment is violated; and because opportunities arrive smoothly in time, the channel introduces none of the anticipation discontinuities that deterministic transfer dates would create ([Barczyk and Kredler, 2021](#)). The

arrival intensity has a natural interpretation as the frequency of life events at which discrete transfers occur.

The optimality condition for (5) ties the lump to the same free boundary that governs the flow. An interior $\tau_L^* > 0$ requires $V_{a_c}^p > V_{a_p}^p$ at the pre-transfer state—strictly inside \mathcal{T} —and the optimal lump moves the state onto the boundary where the two marginal values are equalized. Lump-sum transfers are therefore infrequent, sizeable, and state-contingent: they fire when the child has fallen deep into the transfer region, and they rebalance the dynasty’s wealth in one shot rather than drip-feeding consumption. This is precisely the empirical object the extensive-margin transfer moments measure, and it gives altruism a channel through which it can *raise* child wealth—and hence the intergenerational wealth correlation—rather than only eroding the child’s own saving through the Samaritan’s dilemma. The remaining lump-sum transfers in the model are the college-conditional gift at enrollment (Section 3.6) and the bequest at death; lump opportunities begin at the child’s college-exit age, so the college gift remains the only lump instrument during the schooling years.

3.3.5 Consumption Policies and Wealth Dynamics

The first-order conditions for consumption are:

$$u'(c_p^*) = V_{a_p}^p, \quad (6)$$

$$u'(c_c^*) = V_{a_c}^c. \quad (7)$$

Under CRRA utility, these invert to $c_p^* = (V_{a_p}^p)^{-1/\gamma}$ and $c_c^* = (V_{a_c}^c)^{-1/\gamma}$. The drifts of the two wealth processes (the deterministic part of the laws of motion $da_p = \dot{a}_p dt + \sigma_a a_p dB_p$ and $da_c = \dot{a}_c dt + \sigma_a a_c dB_c$) are:

$$\dot{a}_p = r a_p + y_p(t) - c_p - \tau, \quad (8)$$

$$\dot{a}_c = \tilde{r}_c a_c + y_c(t, z) - c_c + \tau, \quad (9)$$

where $a_p \geq 0$. The child’s wealth, by contrast, can be negative: a child who attends college may finance tuition and consumption with student loans, accumulating a debt balance ($a_c < 0$) that it carries into working life (Section 3.7). The child’s effective interest rate therefore

depends on the sign of its wealth,

$$\tilde{r}_c = r + \iota_s \mathbf{1}\{a_c < 0\},$$

so the child earns r on savings and pays the student-loan rate $r + \iota_s$ on any outstanding debt. The student loan is the only debt instrument: during college the constraint is relaxed to $a_c \geq -\bar{b}$, and the accumulated balance is repaid so that $a_c \geq 0$ by the time the child becomes a parent (age 42); the child cannot otherwise borrow as an adult.

3.4 Child-Alone Problem and the Dynastic Link

When the parent dies (parent age 72, child age 48), the child solves a consumption-savings problem with their own accumulated wealth—augmented by the bequest—for the remainder of their life:

$$\begin{aligned} \rho V^{c,\text{alone}} = \max_{c_c} & \left\{ u(c_c) + V_{a_c}^{c,\text{alone}} [\tilde{r}_c a_c + y_c(t, z) - c_c] \right. \\ & \left. + V_z^{c,\text{alone}}(-\kappa z) + \frac{1}{2} \sigma^2 V_{zz}^{c,\text{alone}} + \frac{1}{2} \sigma_a^2 a_c^2 V_{a_c a_c}^{c,\text{alone}} + V_t^{c,\text{alone}} \right\} \end{aligned} \quad (10)$$

The child lives alone until its own death at age 72, working until 66 and then drawing social security. The dynasty is then closed as a stationary recursion: a new dynasty begins with a parent whose initial wealth equals the child's terminal wealth a_c at age 72, whose education is $e'_p = e_c$, and whose child has ability θ' drawn from the intergenerational transmission process.

3.5 Terminal Conditions

At the parent's death ($t = T$, parent age 72, child age 48), the coupled system terminates and the child transitions to the child-alone problem. The boundary conditions at $t = T$ are:

$$V^p(a_p, a_c, z; T) = \phi_b \frac{(a_p + \underline{a}_{\text{beq}})^{1-\gamma}}{1-\gamma} + \eta \cdot V^{c,\text{alone}}(a_c + \alpha_{\text{beq}} a_p, z; 0), \quad (11)$$

$$V^c(a_p, a_c, z; T) = V^{c,\text{alone}}(a_c + \alpha_{\text{beq}} a_p, z; 0). \quad (12)$$

The child's value function transitions to the child-alone problem (10), evaluated at $t = 0$ of the child-alone clock—the moment of the parent's death. Equation (12) states this continuity:

the child carries forward its own wealth a_c plus the bequest $\alpha_{\text{beq}} a_p$.

The parent's terminal value (11) has two components. The first, $\phi_b (a_p + \underline{a}_{\text{beq}})^{1-\gamma}/(1-\gamma)$ with $\phi_b \geq 0$, is a warm-glow term following De Nardi (2004) that captures the parent's direct utility from holding wealth at the end of the overlap period.⁷ The second component, $\eta \cdot V^{c,\text{alone}}(a_c + \alpha_{\text{beq}} a_p, z; 0)$, reflects the parent's altruistic concern for the child's continuation utility after the parent's death. The child's continuation value is evaluated at effective wealth $a_c + \alpha_{\text{beq}} a_p$, where $\alpha_{\text{beq}} \in [0, 1]$ is the fraction of parent terminal wealth that transfers to the child as a bequest at death. This bequest channel creates two motives for parental wealth accumulation: the warm-glow motive ϕ_b and the altruistic motive through α_{beq} , which disciplines the intergenerational wealth correlation.

3.6 College Decision

At the beginning of the dynasty ($t = 0$), two decisions are taken in sequence: the parent makes a one-time inter-vivos gift, and the child then decides whether to attend college. Let $V_e^c(a_p, a_c)$ and $V_e^p(a_p, a_c)$ denote the child's and parent's equilibrium continuation values under education $e \in \{C, HS\}$, evaluated at $z = 0$.

Child's decision. Given the parent's college gift (described below), the child weighs its own value gain from attending—the college branch evaluated at the post-gift state, the high-school branch at the no-gift state—

$$G \equiv V_C^c - V_{HS}^c,$$

against the psychic cost of college, which follows the Abbott et al. (2019) linear-in-log specification with an idiosyncratic shock:

$$\kappa(\theta, \varepsilon_\kappa) = (\kappa_0 + \kappa_\theta \log \theta + \sigma_\kappa \varepsilon_\kappa) \cdot \bar{V}, \quad \varepsilon_\kappa \sim N(0, 1),$$

⁷The shifter $\underline{a}_{\text{beq}} \geq 0$ makes bequests a luxury good in the sense of De Nardi (2004): with $\underline{a}_{\text{beq}} > 0$ the marginal propensity to bequeath rises with wealth, so poor parents leave little while the rich bequeath a rising share—the mechanism that generates the fat right tail of the wealth distribution.

where ε_κ is drawn once per individual at $t = 0$ and $\bar{V} = \text{mean}_{a_p, \theta} |V_C^c - V_{HS}^c|$ is a normalization constant that makes the dimensionless parameters κ_0 , κ_θ , and σ_κ comparable in scale to the college value gap. The gradient $\kappa_\theta < 0$ implies that higher-ability individuals face lower psychic costs (following Cunha et al. 2005 and Heckman et al. 2006). The child attends whenever the value gain exceeds the realized cost, $G \geq \kappa(\theta, \varepsilon_\kappa)$. Integrating over the normally distributed cost shock yields a probit college-attendance probability:

$$\Pr(e_c = C \mid a_p, a_c, \theta) = \begin{cases} \Phi\left(\frac{G/\bar{V} - \kappa_0 - \kappa_\theta \log \theta}{\sigma_\kappa}\right) & \text{if } G > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

where Φ is the standard normal CDF and σ_κ governs the dispersion of college choices: a larger σ_κ flattens the mapping from the value gap to attendance and, through favorable cost draws, raises enrollment among low-ability types. The decision is the child's alone—consistent with one-sided altruism, the child does not weigh the parent's value in choosing its education.

Parent's college gift and paternalism. Before the child's cost shock is realized, the parent chooses a one-time college-conditional gift τ_C , paid at enrollment, anticipating how it shifts the child's choice:

$$\tau_C^* = \arg \max_{\tau_C \in [0, \min(a_p, \bar{\tau}_C)]} \Pr(C \mid \tau_C) [V_C^p(a_p - \tau_C, \tau_C) + \xi] + (1 - \Pr(C \mid \tau_C)) V_{HS}^p(a_p, 0), \quad (14)$$

where $\Pr(C \mid \tau_C)$ is the probit (13) with the college branch evaluated at the post-gift state $(a_p - \tau_C, a_c = \tau_C)$ and the high-school branch at $(a_p, 0)$. The conditional gift requires no commitment: payment and enrollment are a single transaction—the parent finances enrollment at the moment it occurs—so there is no promise to renege on, and a child who chooses high school receives no lump sum at $t = 0$. The gift is capped at the dynasty's four-year net tuition bill, $\bar{\tau}_C = 4 \phi_{\text{net}}(y_p^0, \theta)$: it is *enrollment financing*, the earmarked support the NLSY97 aid data measure, not a general wealth transfer. Without the cap the enrollment transaction would double as the parent's only opportunity to move wealth before the lump-sum channel opens at college exit, and the gift would scale with altruism far beyond any schooling

cost; with it, general altruistic rebalancing is routed through the Poisson opportunities of Section 3.3.4, where the transfer-probability moments discipline it. The parameter $\xi \geq 0$ captures a paternalistic motive following [Abbott et al. \(2019\)](#): the parent enjoys additional utility ξ if the child attends college, over and above the altruistic valuation already embedded in V_C^p .⁸ Paternalistic parents—particularly wealthy ones, whose marginal utility of consumption is low—use the gift to raise the probability that the child enrolls, the force that generates college attendance among low-ability children of wealthy parents. During the schooling years the college gift is the only lump-sum instrument; after college exit, support moves through the state-contingent flow $\tau(t)$, the Poisson lump-sum opportunities of Section 3.3.4, and ultimately the bequest.

Parental altruism enters the college decision through the continuation values V_C and V_{HS} . Parents who anticipate supporting their child during and after college raise the value of attending, and because the continuation values incorporate the entire future trajectory of transfers—which depends on the child’s income path and hence on education—parent wealth affects enrollment beyond the direct cost channel. Anticipated transfers, however, raise both V_C and V_{HS} : they subsidize college enrollment but also cushion non-college children against income risk. The net effect of the Samaritan’s dilemma on the college decision is therefore theoretically ambiguous; Section 7.1 formalizes this result.

3.7 College Financial System

I model the college financial system following [Abbott et al. \(2019\)](#). The annual cost of attending college is $\phi = \$6,700$, reflecting average net tuition (after institutional discounts) from the NLSY97. During college, students work a fraction $\ell_C = 0.56$ of full-time hours.

Financial aid is modeled as a continuous expected-family-contribution (EFC) schedule keyed on parental *income* at college entry, following [Abbott et al. \(2019\)](#). Let $y_p^0(\theta, e_p)$ denote the parent’s income when the child enters college (parent age 42). The annual grant

⁸In [Abbott et al. \(2019\)](#) the paternalistic motive operates through the parents’ inter-vivos transfer, which can induce enrollment; the same logic applies here. Because the parent cannot commit to future transfers, the utility term alone would not affect any decision; the conditional gift is the instrument through which the paternalistic motive acts, and it relaxes the child’s college-financing constraint at entry. An enrolling child enters the overlap with $a_c = \tau_C^*$ and the parent with $a_p - \tau_C^*$.

is:

$$g(y_p^0) = \bar{g} \cdot \max\left(0, 1 - \frac{y_p^0}{\bar{y}_{\text{efc}}}\right)$$

where \bar{g} is the maximum grant (for zero-income families) and \bar{y}_{efc} is the parental-income level at which aid phases out completely. This continuous specification captures the essential feature of the federal financial aid system: grants decline smoothly with family income, generating an income-dependent effective price of college.

During college, students can borrow up to a cumulative limit $\bar{b}(a_p) = \bar{b}_0 + b_{\text{slope}} \cdot a_p$ at the student loan rate $r + \iota_s$, following [Abbott et al. \(2019\)](#). The base limit $\bar{b}_0 = \$23,000$ reflects the average federal student loan cap; the collateral channel $b_{\text{slope}} \geq 0$ allows children of wealthier parents to borrow more (e.g., through parental co-signing). This is implemented by allowing $a_c \geq -\bar{b}(a_p)$ during $t \in [0, 4]$; the accumulated balance then amortizes so that $a_c \geq 0$ by the time the child becomes a parent (age 42). The child's wealth dynamics during college are:

$$\dot{a}_c = (r + \iota_s) a_c + \ell_C \cdot y_c(t, z) - c_c - \phi_{\text{net}}(y_p^0) + \tau, \quad a_c \geq -\bar{b}(a_p)$$

where $\phi_{\text{net}}(y_p^0) = \phi - g(y_p^0)$ is net tuition. This system generates heterogeneity in the effective cost of college across the parental income distribution. Parental transfers supplement institutional aid, and as [Abbott et al. \(2019\)](#) document, government grants can crowd out private contributions.

3.8 Equilibrium

A Markov-Perfect Equilibrium (MPE) consists of value functions V^p, V^c , consumption policies c_p^*, c_c^* , a transfer-flow policy τ^* , a lump-sum transfer policy τ_L^* , and a college enrollment rule e_c^* , for each (θ, e_p, e_c) pair, such that: (i) given c_c^*, τ^* , and τ_L^* , the parent's HJB (1), the complementary-slackness condition (3), and the lump-sum optimality condition (5) are satisfied; (ii) given c_p^*, τ^* , and τ_L^* , the child's HJB (2) is satisfied; (iii) after the parent's death ends the overlap ($t > T$), $V^{c, \text{alone}}$ satisfies (10); (iv) the college rule follows from (13); and (v) boundary conditions (11)–(12) and borrowing constraints hold at all times.

The equilibrium is solved backward: first the child-alone problem, then the coupled system

during the overlap period, and finally the college choice at $t = 0$. Within the coupled system, the transfer is obtained in closed form at each state—the bounded gift (4)—so the transfer region is determined pointwise rather than by iterating on the complementary-slackness condition, and the value functions are updated by implicit time-stepping. Appendix I derives the equilibrium properties, including the Samaritan’s dilemma distortion in continuous time, and Appendix J describes the numerical solution algorithm.

3.9 Empirical Implications

College reshapes the transfer region \mathcal{T} —the set of states in which the parent optimally supports the child—and the boundary shift maps into family outcomes observable in matched parent-child data. I first derive why college contracts \mathcal{T} , then separate the implications that the descriptive evidence of Section 4 can document from the one that only the estimated model recovers.

Why college shrinks the transfer region. Recall that the parent transfers when the marginal value of own wealth equals the marginal value of child wealth ($V_{a_p}^p = V_{a_c}^p$). Because the parent weighs the child’s utility by η in the objective (1), $V_{a_c}^p$ is increasing in η : more altruistic parents have a larger transfer region. This condition is more likely to bind when the child is relatively poor or faces negative income shocks. College raises the child’s expected income through two channels: a higher ability gradient ($\lambda_C > \lambda_{HS}$) and a steeper life-cycle earnings profile. The resulting increase in child wealth pushes states away from the transfer boundary, so college-educated children spend less time in \mathcal{T} . Higher expected income therefore contracts \mathcal{T} for college-educated children relative to high-school-educated children. This is a property of the solved equilibrium rather than a closed-form result: the boundary of \mathcal{T} is a free boundary determined numerically (Appendix I), and the contraction holds throughout the estimated model.

Figure 2 illustrates this mechanism in the (a_c, a_p) state space. The boundary $\partial\mathcal{T}$ separates states where transfers are positive (to the left) from states where the parent saves independently (to the right). College shifts this boundary to the left: at any given level of parent wealth, the child needs to be poorer before transfers become optimal. The region $\mathcal{T}_{HS} \setminus \mathcal{T}_C$ —states where a high-school child would receive transfers but a college child would

not—measures how far college retracts the transfer region. Together with the overconsumption zone bordering it—where the child under-saves in anticipation of entering \mathcal{T} (Appendix I)—this contraction captures the scope for moral hazard that college eliminates.




Figure 2. Transfer Region Boundary: College vs. High School

Notes: Schematic illustration of the transfer region \mathcal{T} in the (a_c, a_p) state space. In the transfer region (shaded), the parent optimally transfers to the child ($\tau > 0$). The solid boundary corresponds to a college-educated child; the dashed boundary corresponds to a high-school-educated child. College shrinks \mathcal{T} because higher expected income keeps child wealth further from the transfer boundary.

The contraction of \mathcal{T} yields two predictions, both conditional on parent income, that the descriptive evidence can document. The first is a consumption release: a smaller transfer region means parents spend less of their lifetime making transfers, so resources that would have flowed to the child under the high-school path are instead available for own consumption, and parent consumption is higher when the child attends college. The second concerns the two margins of transfers. Because a larger negative shock is required to push a college child

into \mathcal{T} , the probability of transferring falls; conditional on transferring, the states reached lie further inside \mathcal{T} , so the optimal transfer τ^* is weakly larger. The model thus predicts a negative extensive-margin effect and a zero-or-slightly-positive intensive-margin effect. Both are sign predictions that Section 4 documents.

The same logic carries a further implication, now about magnitudes rather than signs. Parents hold wealth against the option of supporting a child who falls into \mathcal{T} , so college releases resources even in states where no transfer is ever paid. The consumption release therefore reflects both fewer realized transfers and a smaller precautionary burden, and the model predicts that the first exceeds the second: the consumption gain from college is larger than the fall in realized transfers. This is the continuous-time option-value channel of [Barczyk and Kredler \(2014b\)](#), in which the family inefficiency operates before transfers flow rather than through the transfers themselves. The estimated model quantifies this channel directly (Section 6).

4 Empirical Evidence

This section documents how parent consumption and inter-vivos transfers vary with the child’s college attainment in matched parent-child data (PSID, HRS, and NLSY97). Each fact compares outcomes for parents whose children attended college with those whose children did not, conditioning on parent income to absorb the parent’s own resources. The patterns are descriptive: college attainment is jointly determined with parental resources, preferences, and expectations, and the option value of parental support is unobserved, so these correlations cannot by themselves identify the mechanism. Their role is twofold. They establish that family finances reorganize around the child’s education in the direction the transfer-region mechanism predicts; and, because none of them is targeted in estimation, they serve as out-of-sample checks on the model. Section 6.4 confronts the estimated model with each of these untargeted moments.

4.1 Data

The empirical analysis draws on three data sources. The Panel Study of Income Dynamics (PSID) provides matched parent-child panels with consumption, income, wealth, and inter-

vivos transfer data from 1999 onward. I link parents to adult children using the Family Identification Mapping System (FIMS). The summary-statistics sample (Tables 1–2) comprises 53,326 parent-year observations from 8,362 households with heads aged 25–60; the parent-consumption analysis of Section 4.2 further restricts to parents aged 50+ with adult children aged 26+, collapsed to the parent-year level, yielding 15,830 observations. The Health and Retirement Study (HRS, 1992–2014 Family-Kids file) offers a larger sample of older parent-child pairs with detailed transfer information—548,239 parent-to-child wave observations, of which 78,541 involve a positive transfer—which I use to study the extensive and intensive margins of inter-vivos transfers. The National Longitudinal Survey of Youth 1997 (NLSY97) provides cognitive ability scores (ASVAB) and parental wealth for 5,400 individuals.

4.1.1 Summary Statistics

Tables 1–2 report summary statistics for the PSID sample, stratified by average household income group. Table 1 presents parent demographics and key financial variables: parents in higher income groups are older on average, more likely to hold a college degree, more likely to own their home, and hold substantially more wealth. Total wealth ranges from \$103,000 for the lowest income group to \$966,000 for the highest, with home equity accounting for roughly half of wealth for low-income families but a smaller share for wealthy families. Total consumption expenditures increase monotonically with income, from \$36,000 to \$73,000. A detailed breakdown of the wealth portfolio is provided in Appendix Table E.1, and a consumption decomposition by expenditure category in Appendix Table F.1. Table 2 presents child-level statistics, including the college attendance rate among 18–24 year olds by parental income group, which rises steeply with parental resources.

Table 1. Parent Demographics and Financial Characteristics by Income Group (PSID)

| | By Average HH Income Group | | | | All |
|---|----------------------------|-----------------------|-----------------------|-------------------------|-----------------------|
| | \$30–60 | \$60–100 | \$100–160 | \$160+ | |
| <i>Panel A: Demographics</i> | | | | | |
| Age of head | 45.6 (16.2) | 45.2 (14.4) | 46.8 (13.4) | 49.1 (13.2) | 46.3 (14.6) |
| Children in household | 1.3 (1.4) | 1.2 (1.3) | 1.1 (1.2) | 1.0 (1.1) | 1.1 (1.3) |
| Total children | 2.1 (1.3) | 2.0 (1.0) | 2.1 (1.0) | 2.2 (1.0) | 2.1 (1.1) |
| Head has college degree | 0.07 (0.26) | 0.19 (0.39) | 0.42 (0.49) | 0.69 (0.46) | 0.29 (0.45) |
| White | 0.44 (0.50) | 0.63 (0.48) | 0.77 (0.42) | 0.85 (0.36) | 0.64 (0.48) |
| Black | 0.45 (0.50) | 0.29 (0.45) | 0.18 (0.38) | 0.09 (0.28) | 0.28 (0.45) |
| Homeowner | 0.50 (0.50) | 0.71 (0.45) | 0.84 (0.36) | 0.85 (0.36) | 0.70 (0.46) |
| <i>Panel B: Income, Wealth, and Consumption</i> | | | | | |
| Household income | 44,223 (21,535) | 76,829 (37,964) | 120,169 (51,447) | 242,277 (253,591) | 102,569 (121,904) |
| Head labor income | 23,523 (19,369) | 40,976 (28,162) | 64,728 (42,902) | 123,504 (123,970) | 53,930 (64,202) |
| Total wealth | 102,555 (230,260) | 200,522 (377,059) | 367,076 (536,534) | 965,941 (1,052,497) | 326,288 (612,283) |
| Total consumption | 36,426 (18,374) | 46,131 (20,389) | 57,243 (24,489) | 73,130 (33,277) | 50,076 (26,252) |
| 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. Income groups defined by parent average real household income (2016 dollars) observed between ages 25 and 60. Panel A reports demographics; Panel B reports income, wealth, and consumption. “Children in household” is the number of children under age 18 in the parent’s PSID family unit; “Total children” is the total number of children linked to the parent household via the Family Identification Mapping System (FIMS). All monetary variables in 2016 dollars, winsorized at the 1st and 99th percentiles. A detailed wealth portfolio breakdown is in Appendix Table E.1.

Table 2. Children Descriptive Statistics by Parent Income Group (PSID)

| | By Parent's Average HH Income Group | | | | All |
|--|-------------------------------------|-----------------|-----------------|-----------------|-----------------|
| | \$30–60 | \$60–100 | \$100–160 | \$160+ | |
| <i>Panel A: Child Demographics</i> | | | | | |
| Age | 17.9 (19.0) | 15.9 (15.7) | 16.7 (13.6) | 17.4 (13.4) | 16.9 (16.0) |
| Number of siblings in sample | 1.9 (1.7) | 1.6 (1.2) | 1.5 (1.2) | 1.6 (1.1) | 1.6 (1.4) |
| College-age (18–24) | 0.14 (0.34) | 0.14 (0.35) | 0.16 (0.36) | 0.17 (0.38) | 0.15 (0.36) |
| <i>Panel B: Education</i> | | | | | |
| Years of education (age 18+) | 13.7 (11.8) | 13.7 (9.9) | 13.8 (7.8) | 14.4 (8.1) | 13.9 (9.7) |
| Max years of education (age 25+) | 16.1 (14.8) | 15.7 (11.5) | 15.9 (9.9) | 16.5 (8.9) | 16.0 (11.9) |
| Currently in college | 0.05 (0.21) | 0.06 (0.23) | 0.08 (0.27) | 0.10 (0.29) | 0.07 (0.25) |
| Ever attains post-secondary ed. | 0.33 (0.47) | 0.37 (0.48) | 0.47 (0.50) | 0.56 (0.50) | 0.41 (0.49) |
| <i>Panel C: College Rate Among 18–24 Year Olds</i> | | | | | |
| College attendance rate | 0.32 (0.47) | 0.40 (0.49) | 0.47 (0.50) | 0.54 (0.50) | 0.42 (0.49) |
| College-age observations | 4308 | 4984 | 4270 | 2907 | 16 469 |
| Total child–year observations | 31 731 | 35 484 | 27 026 | 16 962 | 111 203 |
| Unique children | 5680 | 5560 | 3784 | 2186 | 17 210 |

Notes: Means with standard deviations in parentheses. Each observation is a child-year. Children linked to parents via FIMS. Income groups refer to parent average real household income (2016 dollars). “Years of education” reported conditional on age ≥ 18 ; “Max years of education” conditional on age ≥ 25 . “College attendance rate” computed conditional on the child being aged 18–24. “Ever attains post-secondary ed.” indicates maximum observed years of education exceeds 12. Years of education are directly observed only from 2013 onward; for earlier waves, college attendance is inferred from a proxy that combines the typical college-age window with the individual’s eventual attainment (maximum years of education observed in 2013 or later exceeding 12). Individuals not observed in any post-2012 wave cannot be assigned by this proxy and are excluded from the pre-2013 attendance measure.

4.2 Parent Consumption and Child College Status

The model implies that, conditional on their own income and wealth, parents whose children attend college should consume more (the consumption-release implication of Section 3.9). College shifts children out of the transfer region \mathcal{T} by permanently raising their income, freeing parents from contingent obligations and raising their consumption even when no transfers are actually flowing.

I examine this implication by estimating the following regression on the PSID sample at the parent-year level:

$$C_{p,t} = \beta_0 + \beta_1 \text{ShareCollege}_{p,t} + \beta_2 N_{p,t} + \delta_{Q_p^{inc}} + \beta_X \mathbf{X}_{\mathbf{p},t} + \varepsilon_t + \epsilon_{p,t}$$

where $C_{p,t}$ is household consumption (in 2016 dollars) of parent p at time t , $\text{ShareCollege}_{p,t}$ is the fraction of parent p 's children with a bachelor's degree or higher, $N_{p,t}$ is the total number of children, $\delta_{Q_p^{inc}}$ are parent income quartile fixed effects (within-cohort ranking), and $\mathbf{X}_{\mathbf{p},t}$ is a comprehensive set of controls: parent total wealth, non-financial wealth, household income, wealth quartile, labor force participation, household size, head's birth year, education, U.S. state, a cubic in head's age, housing tenure, and race. The specification includes year fixed effects ε_t .

Table 3 presents the results. Column 1 shows that a higher share of college-educated children is associated with higher parent consumption. Columns 2 and 3 use alternative definitions—any child with a BA, or all children with a BA—yielding consistent results. Column 4 uses consumption levels rather than logs: a move from no college-educated children to all corresponds to roughly \$2,500 more in annual parent consumption. The specification conditions on parent income, wealth, and education, so the association is not mechanically the level effect of parental resources; it remains a within-group correlation rather than a causal estimate, and the magnitude is the object the estimated model is later asked to reproduce out of sample (Section 6.4).

Table 3. Parent Consumption and Child College Status (PSID)

| | (1) | (2) | (3) | (4) |
|----------------------------|---------------------|---------------------|---------------------|---------------------|
| | Log C | Log C | Log C | Level C (000s) |
| Share of Children with BA+ | 0.081*** (0.019) | | | 2.470*** (0.866) |
| Number of Children | 0.014* (0.008) | 0.010 (0.008) | 0.016* (0.008) | 0.170 (0.287) |
| Any Child with BA+ | | 0.080*** (0.017) | | |
| All Children with BA+ | | | 0.056*** (0.017) | |
| Controls | Yes | Yes | Yes | Yes |
| Year FE | Yes | Yes | Yes | Yes |
| Parent income Q FE | Yes | Yes | Yes | Yes |
| Observations | 15,830 | 15,830 | 15,830 | 15,831 |
| R^2 | 0.558 | 0.558 | 0.557 | 0.357 |

Notes: OLS regressions of parent household consumption on child college indicators. “Share of Children with BA+” is the fraction of children with a bachelor’s degree or higher. The dependent variable in Column 4 is total consumption in thousands of 2016 dollars. Controls include a cubic in parent age, household income, non-housing wealth, total wealth, wealth quartile, and fixed effects for year, education, parent income quartile, gender, housing tenure, state, labor force status, race, household size, and birth year. Standard errors clustered at the parent household level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

4.3 Inter-Vivos Transfers

If college frees parents from contingent obligations, do realized transfers fall correspondingly? Answering requires decomposing transfers into their extensive and intensive margins, which map directly onto the model’s transfer-region structure (Section 3.9).

The key feature of inter-vivos transfers in the HRS is that most parent–child pairs report zero transfers in any given wave: only 14.4% of observations involve a positive transfer from parent to child (Table 4). Conditional on a positive transfer, however, amounts are substantial—\$8,161 on average, with a median of \$2,897 and a 90th percentile of \$18,394.

This pattern—a large mass at zero with a heavy right tail—is what the model predicts. The transfer region \mathcal{T} is entered infrequently, but when entered, transfers are non-trivial insurance events.

Table 4. Distribution of Inter-Vivos Transfers from Parents to Children (HRS)

| | |
|---------------------------------|----------|
| Pr(Transfer > 0) | 0.144 |
| E(Transfer Transfer > 0) | \$8,161 |
| Median(Transfer Transfer > 0) | \$2,897 |
| P75(Transfer Transfer > 0) | \$7,968 |
| P90(Transfer Transfer > 0) | \$18,394 |
| E(Transfer, unconditional) | \$1,172 |
| Observations | 548,239 |
| Observations with Transfer > 0 | 78,541 |

Notes: Wave-level statistics for parent-to-child transfers from HRS 1992–2014 (Family-Kids file, waves 1–12). Sample restricted to parent age ≥ 50 , child age ≥ 26 . Transfer amounts in real 2016 US dollars. Unconditional mean includes zeros.

I decompose the within-family college effect along these two margins. Table 5 reports linear probability models for $\Pr(\text{Transfer} > 0)$ with parent fixed effects:

$$\mathbf{1}[IVT_{ijt} > 0] = \beta_0 + \beta_1 \text{College}_j + \gamma_t + \beta_X \mathbf{X}_{jt} + \delta_{Q_i^p} + \varepsilon_i + \epsilon_{ijt}$$

College attainment is associated with a 3.1 percentage point lower probability of any parent-to-child transfer—a 21% reduction relative to the mean transfer probability of 14.4% (Table 5, Column 1). Column 2 interacts college with parent income quartile and shows a steep gradient: the college coefficient is -1.0 pp for Q1 parents, -2.8 pp for Q2 ($-0.010 - 0.018$), -3.4 pp for Q3, and -4.2 pp for Q4. This gradient is consistent with the mechanism: wealthier parents have larger transfer regions ex ante—they can afford to support children in more states of the world—so college contracts their transfer region by more.

Table 5. Extensive Margin: Effect of College on Transfer Probability (HRS)

| | Extensive Margin (Linear Probability) | |
|----------------|---------------------------------------|----------------------|
| | (1) | (2) |
| | Pr(Parent → Child) | Pr(Parent → Child) |
| College (BA+) | -0.031*** (0.002) | -0.010*** (0.003) |
| College × Q2 | | -0.018*** (0.004) |
| College × Q3 | | -0.024*** (0.004) |
| College × Q4 | | -0.032*** (0.005) |
| Parent FE | Yes | Yes |
| Wave FE | Yes | Yes |
| Dep. Var. Mean | 0.144 | 0.144 |
| Observations | 548,164 | 548,164 |
| Within R^2 | 0.012 | 0.012 |

Notes: Linear probability models. Dependent variable is an indicator for any positive transfer in the HRS wave. All specifications include parent fixed effects, wave dummies, child cohort controls, and parent income quartile dummies. Standard errors clustered at parent level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

Table 6 conditions on positive transfers and is consistent with the model’s intensive-margin implication. Conditional on being in the transfer region, college children receive \$557 more per transfer event (Column 1, $p < 0.10$). Column 2 reveals heterogeneity: for Q1 parents, college children receive \$906 *less* conditional on entry, but for wealthier parents the effect is positive and large—\$958 more for Q2 ($-906 + 1,864$), and \$826 more for Q4. This pattern is consistent with the boundary-shift mechanism: because college moves the transfer boundary $\partial\mathcal{T}$ to the left in the (a_c, a_p) space, the states that push a college-educated child into \mathcal{T} involve worse realizations of child income. Conditional on entry, the child is further from the boundary, and the optimal transfer is therefore larger—particularly for wealthier parents

whose transfer regions are larger and whose boundary shift from college is more pronounced.

Table 6. Intensive Margin: Transfer Amount Conditional on Positive Transfer (HRS)

| | Intensive Margin (Amount Transfer > 0) | |
|----------------|--|------------------|
| | (1) | (2) |
| | Parent → Child | Parent → Child |
| College (BA+) | 557* (288) | -906* (486) |
| College × Q2 | | 1864*** (607) |
| College × Q3 | | 965* (542) |
| College × Q4 | | 1732** (684) |
| Parent FE | Yes | Yes |
| Wave FE | Yes | Yes |
| Dep. Var. Mean | 8161 | 8161 |
| Observations | 78,541 | 78,541 |
| Within R^2 | 0.002 | 0.002 |

Notes: OLS regressions on the subsample with positive transfers. Dependent variable is the real transfer amount (2016 US\$). All specifications include parent fixed effects, wave dummies, child cohort controls, and parent income quartile dummies. Standard errors clustered at parent level in parentheses. * $p < .10$, ** $p < .05$, *** $p < .01$.

The 3.1 percentage point fall in transfer probability associated with college corresponds to roughly \$253 per year in expected transfers ($0.031 \times \$8,161$), an order of magnitude smaller than the \$2,500 consumption difference of Section 4.2. The comparison is suggestive rather than decisive—the two figures come from different samples and neither is causal—but it is the pattern the option-value channel generates: college lowers the burden of standing ready to support the child, releasing resources beyond what flows through realized transfers. Section 6.4 reports the model’s counterpart of both numbers.

Altogether, college is associated with higher parent consumption and with less frequent but conditionally larger transfers. The reduced form cannot quantify the mechanism: college attainment is jointly determined with parental resources and transfers, and the option value of future support is unobserved. The model is therefore estimated on a distinct set of moments—transfer amounts, the college-attainment gradients by parent wealth and income, the pace of parent decumulation, and the intergenerational wealth distribution—while the patterns documented here are held out as untargeted checks (Section 6.4).

5 Estimation

I estimate the model in three stages. In the first stage, I set parameters that are well-identified from external data or the existing literature, including the income-process parameters, which I take from [Abbott et al. \(2019\)](#). In the second stage, I map the discrete-time income-process estimates to their continuous-time counterparts. In the third stage, I estimate the remaining fourteen structural parameters by the simulated method of moments (SMM), matching 38 targeted moments.

5.1 Externally Set Parameters

Table 7 reports the parameters set in the first stage. The coefficient of relative risk aversion is $\gamma = 1.5$, following [Abbott et al. \(2019\)](#). The annual real interest rate is $r = 0.03$, following [Daruich and Kozłowski \(2020\)](#). The continuous-time discount rate ρ is estimated internally by SMM to match the wealth-to-income ratio and wealth distribution moments. The college costs, borrowing limits, and financial-aid schedule parameters (\bar{g} , \bar{y}_{efc}) are set externally to the [Abbott et al. \(2019\)](#) values, as described in Section 3.7. The income-process parameters—the education-specific ability gradients λ_e , age profiles, OU mean-reversion rates κ_e and diffusions σ_e , and entry dispersions—and the intergenerational ability parameters (ρ_θ , σ_θ , μ_θ) are likewise taken from [Abbott et al. \(2019\)](#). The relative price of college labor ω_C and the unsecured borrowing wedge ι_s are *estimated* by SMM (Section 5), identified by the high-school-to-college income ratio and the child negative-net-worth share, respectively; they are not externally set. Following [Barczyk and Kredler \(2014a\)](#), the multiplicative wealth-diffusion volatility is fixed at $\sigma_a = 0.03$ per year—the minimal randomization that delivers

C^2 value functions and a smooth transfer boundary (Section 3)—rather than estimated, so the wealth distribution is disciplined by saving and bequests. The lump-sum transfer arrival intensity is likewise set externally at $\lambda_g = 0.25$ per year—one opportunity every four years on average—rather than estimated: because arrivals are independent of the dynasty’s state, λ_g sets only a wealth-blind ceiling of $1 - e^{-2\lambda_g} \approx 0.39$ on the probability of a transfer in any two-year observation window, while the wealth *gradient* of transfer incidence must come from the size of the region where the parent chooses to exercise the opportunity—the free boundary that altruism η governs. Fixing λ_g therefore keeps the identification of η from the transfer-probability moments clean. Retirement income is estimated directly from PSID data on households where the respondent is retired and over age 67, computing the average sum of Social Security, SSI, disability, and employer pension income by education group.

Table 7. Externally Set Parameters

| Parameter | Description | Value | Source |
|--------------------------------------|---------------------------------------|---------|------------------------------|
| <i>Preferences</i> | | | |
| r | Real interest rate | 0.03 | Daruich and Kozlowski (2020) |
| γ | CRRA risk aversion | 1.5 | Abbott et al. (2019) |
| \underline{c} | Consumption floor (\$000s/yr) | 3.9 | De Nardi (2004) |
| σ_a | Wealth-diffusion volatility (/yr) | 0.03 | Barczyk and Kredler (2014a) |
| <i>Demographics (years)</i> | | | |
| T^{ret} | Parent working, overlap (child 18–42) | 24 | — |
| $T - T^{\text{ret}}$ | Parent retired, overlap (child 42–48) | 6 | — |
| T_{alone} | Child alone, working (48–66) | 18 | — |
| $T_{\text{retire},c}$ | Child retirement (66–72) | 6 | — |
| <i>Ability gradients</i> | | | |
| λ_C | Returns to ability, college | 0.797 | Abbott et al. (2019) |
| λ_{HS} | Returns to ability, HS | 0.517 | Abbott et al. (2019) |
| <i>Earnings process (annual)</i> | | | |
| ρ_C^{ann} | Persistence, college | 0.966 | Abbott et al. (2019) |
| ρ_{HS}^{ann} | Persistence, HS | 0.952 | Abbott et al. (2019) |
| $(\sigma_\eta^{\text{ann}})^2$ | Innovation variance (both groups) | 0.017 | Abbott et al. (2019) |
| σ_C | OU diffusion, college | 0.133 | Abbott et al. (2019) |
| σ_{HS} | OU diffusion, HS | 0.134 | Abbott et al. (2019) |
| \bar{w} | Avg. annual earnings | \$70k | Census |
| <i>College costs</i> | | | |
| ϕ_C | Annual tuition (before aid) | \$6.7k | Abbott et al. (2019) |
| τ_C | Time worked in college | 0.56 | Census |
| <i>Intergenerational ability</i> | | | |
| ρ_θ | Persistence of log ability | 0.36 | Abbott et al. (2019) |
| σ_θ | Std. dev. of ability innovation | 0.40 | Abbott et al. (2019) |
| μ_θ | Mean of log ability | 1.30 | Abbott et al. (2019) |
| <i>Financial aid & borrowing</i> | | | |
| \bar{g} | Max. annual grant (\$000s) | 2.82 | Abbott et al. (2019) |
| \bar{y}_{efc} | EFC income phase-out (\$000s) | 72.0 | Abbott et al. (2019) |
| \bar{b}_0 | Student loan limit (base) | \$23.0k | Federal |
| ι_s | Student loan premium | 0.031 | Abbott et al. (2019) |
| <i>Retirement income (annual)</i> | | | |
| SS_C | Soc. Security + pension, college | \$14.3k | PSID |
| SS_{HS} | Soc. Security + pension, HS | \$13.0k | PSID |

Notes: Dollar amounts in 2016 dollars. The continuous-time income process is set externally from Abbott et al. (2019) (Table 3.2, males): persistence maps to the mean-reversion rate via $\kappa_e = -\ln \rho_e^{\text{ann}}$, and the OU diffusion σ_e is chosen so the stationary variance $\sigma_e^2/(2\kappa_e)$ equals the AGMV persistent-shock variance $(\sigma_\eta^{\text{ann}})^2/(1 - (\rho_e^{\text{ann}})^2)$, giving $\sigma_{HS} = 0.134$, $\sigma_C = 0.133$ (Appendix D). The diffusions are *not* estimated. The model carries no transitory income shock. The financial aid schedule $g(y_p^0) = \bar{g} \cdot \max(0, 1 - y_p^0/\bar{y}_{\text{efc}})$ is a function of parent *income* at college entry; \bar{g} and \bar{y}_{efc} are set externally to Abbott et al. (2019), not estimated. Retirement income from PSID (respondent retired, age > 67). The student loan is the only debt instrument; its base limit \bar{b}_0 is cumulative over four college years, and the effective limit $\bar{b}(a_p) = \bar{b}_0 + b_{\text{slope}} \cdot a_p$ has slope b_{slope} estimated by SMM (Table 9). The effective student rate is $r + \iota_s = 0.061$.

5.2 Income Process Estimation

The income process in continuous time is an Ornstein-Uhlenbeck process $dz = -\kappa_e z dt + \sigma_e dW$ with education-specific parameters. The entire income process is set externally from the male estimates in [Abbott et al. \(2019\)](#): I map their persistence to the continuous-time mean-reversion rate κ_e and choose the diffusion σ_e so that the stationary variance $\sigma_e^2/(2\kappa_e)$ reproduces their persistent-shock variance. The diffusion coefficients are therefore *not* estimated within the structural routine; the income process is calibrated first-stage and held fixed.⁹

In discrete time, log labor income is decomposed as in [Abbott et al. \(2019\)](#): $\log \epsilon_j = \lambda_e \log \theta + \gamma_{e,j} + z_j$, where λ_e is the education-specific ability gradient. The ability gradient and age-income profile are taken from [Abbott et al. \(2019\)](#), who estimate λ_e from the NLSY79 by regressing log hourly wages on log AFQT scores, separately by education group, and estimate the deterministic age-income profile $\gamma_{e,j}$ from PSID data using a quadratic in age. The income shock persistence is also taken from [Abbott et al. \(2019\)](#), who estimate it by minimum distance on the autocovariance structure of wage residuals by age and education:

$$\begin{aligned} z_{iat}^e &= \log y_{it} - \widehat{f^e(a_{it})} - \hat{\lambda}_e \log(\text{AFQT}_i) \\ z_{iat}^e &= \rho_e^{\text{ann}} z_{i,a-1,t-1}^e + \eta_{iat}^e \\ \eta_{iat}^e &\sim N(0, (\sigma_\eta^e)^2), \quad z_{i0t}^e \sim N(0, (\sigma_{z_0}^e)^2) \end{aligned}$$

I map the discrete-time persistence to its continuous-time mean-reversion counterpart κ_e and the innovation variance to the OU diffusion σ_e by matching the autocovariance structure; [Appendix D](#) derives the mapping formulas. [Table 8](#) reports the discrete-time estimates and their continuous-time images. The mean-reversion rates are similar across education groups ($\kappa_{HS} = 0.049$, $\kappa_C = 0.035$), and the diffusions are nearly identical ($\sigma_{HS} = 0.134$, $\sigma_C = 0.133$). Because college earnings are more persistent, the implied stationary standard deviation of the persistent shock is in fact slightly *higher* for college graduates (0.51 versus 0.43 for high school), so the model's transfer-region contraction is driven by the income *level*, not by a reduction in income risk ([Section 3](#)).

⁹[Abbott et al. \(2019\)](#) estimate the income process separately by gender. I use their male estimates throughout, as the model abstracts from gender differences.

Table 8. Income Process and Age Profile

| | High-School | College | Source |
|--|-------------|---------|--|
| <i>Ability gradient</i> | | | |
| λ_e | 0.517 | 0.797 | AGMV, Tbl. 3.1 |
| <i>Deterministic age profile</i> | | | |
| $\gamma_{1,e}$ | 0.0673 | 0.1197 | AGMV, Tbl. E.1 |
| $\gamma_{2,e} \times 1000$ | -0.670 | -1.191 | AGMV, Tbl. E.1 |
| <i>Persistent shock — discrete-time AR(1)</i> | | | |
| ρ_z^{ann} | 0.952 | 0.966 | AGMV, Tbl. 3.2 |
| $(\sigma_\eta^{\text{ann}})^2$ | 0.017 | 0.017 | AGMV, Tbl. 3.2 |
| $\sigma_{z_0}^2$ (entry var.) | 0.059 | 0.094 | AGMV, Tbl. 3.2 |
| <i>Persistent shock — continuous-time OU (mapped, fixed)</i> | | | |
| κ_e | 0.049 | 0.035 | $-\ln \rho_z^{\text{ann}}$ |
| σ_e | 0.134 | 0.133 | $\sqrt{2\kappa_e(\sigma_\eta^{\text{ann}})^2/(1 - (\rho_z^{\text{ann}})^2)}$ |
| σ_{z_0} (entry s.d.) | 0.243 | 0.307 | $\sqrt{\sigma_{z_0}^2}$ |
| stationary s.d. $\sigma_e/\sqrt{2\kappa_e}$ | 0.43 | 0.51 | implied |
| <i>Relative price of college labor (SMM)</i> | | | |
| ω_C | — | 0.642 | SMM (Tbl. 9) |

Notes: The table collects all parameters of the income process and reports the source of each. “AGMV” denotes [Abbott et al. \(2019\)](#) (male estimates); the cited table number refers to that paper. The ability gradient λ_e , the deterministic age profile $(\gamma_{1,e}, \gamma_{2,e})$, and the discrete-time persistent-shock parameters $(\rho_z^{\text{ann}}, \sigma_\eta^{\text{ann}}, \sigma_{z_0})$ are taken directly from AGMV. The continuous-time mean-reversion rates κ_e and OU diffusions σ_e are mapped from the AGMV persistence ρ_z^{ann} and innovation variance σ_η^{ann} as described in the text and Appendix D; the model carries no transitory shock ($\sigma_{\varepsilon,e} = 0$, as in AGMV). The only income-process parameter estimated by SMM is the relative price of college labor ω_C (ω is normalized to one for high school), reproduced here from Table 9. The implied stationary standard deviation of the persistent shock, $\sigma_e/\sqrt{2\kappa_e}$, is 0.43 for high school and 0.51 for college.

Figure D.1 in Appendix D illustrates the implied life-cycle income profiles by education, including the deterministic age-income component, expected income for a median-ability worker, the role of OU shock volatility, and the complementarity between ability and education.

5.3 Ability Gradient

The ability gradient λ_e in the human capital function $h(\theta, e) = (\theta/\bar{\theta})^{\lambda_e}$ governs the returns to cognitive ability by education level. I take the estimates of λ_e directly from [Abbott et al. \(2019\)](#), who estimate them from the NLSY79 by regressing log hourly wages (purged of the PSID age profile) on log AFQT scores, separately by education group. The key estimates are $\hat{\lambda}_{HS} = 0.517$ and $\hat{\lambda}_C = 0.797$, implying that a one-log-point increase in ability raises college earnings by 80% but high-school earnings by only 52%. This complementarity between ability and education is central to the mechanism: the college premium is strongly increasing in ability, so removing parental transfers has heterogeneous effects across the ability distribution. These parameters are reported in [Table 7](#).

5.4 SMM Estimation

The remaining fourteen parameters—parent altruism η , psychic cost parameters κ_0, κ_θ , and σ_κ , the discount rate ρ , the warm-glow weight ϕ_b , the luxury-bequest shifter $\underline{a}_{\text{beq}}$, the bequest fraction α_{beq} , the parent income-fanning rate g_{zp} , the paternalistic motive ξ , the relative price of college labor ω_C , the collateral channel slope b_{slope} , the unsecured borrowing wedge ι_s , and the intergenerational earnings premium ζ_{pe} —are estimated by the simulated method of moments (SMM; [McFadden, 1989](#); [Pakes and Pollard, 1989](#)). SMM is well suited to this setting because the model does not admit closed-form expressions for the moments of interest: the coupled HJB system and the endogenous transfer region must be solved numerically. Following [Duffie and Singleton \(1993\)](#) and [Lee and Ingram \(1991\)](#), I simulate the model at each candidate parameter vector and match the simulated moments to their data counterparts.

Let $\mathbf{x} \in \mathbb{R}^{14}$ denote the parameter vector. For each \mathbf{x} , I solve the model and simulate $M = 4,000$ dynasties, computing the model counterparts of 38 targeted moments $\mathbf{m}^{\text{sim}}(\mathbf{x})$ (four further moments—the two education-specific income standard deviations, the variance of the inverse hyperbolic sine of parent wealth, and the p90/p50 ratio of parent wealth—are computed but left untargeted as validation checks). The estimator minimizes a block-diagonal percentage-deviation loss:

$$\hat{\mathbf{x}} = \arg \min_{\mathbf{x}} \sum_g \frac{w_g}{n_g} \sum_{m \in g} \left[\frac{m_m^{\text{sim}}(\mathbf{x}) - m_m^{\text{data}}}{m_m^{\text{data}}} \right]^2, \quad (15)$$

where the targeted moments are partitioned into blocks: college attendance rates by ability \times parent-wealth quartile ($n_g = 16$, $w_g = 3$), income ($n_g = 3$: the high-school-to-college income ratio, with the two education-specific income standard deviations carrying zero weight as untargeted validation moments), transfers ($n_g = 5$), intergenerational wealth ($n_g = 11$, with the p90/p50 parent-wealth ratio and the variance of the inverse hyperbolic sine of parent wealth carrying zero weight as untargeted validation moments), decumulation ($n_g = 1$), college graduation by parent-income quartile ($n_g = 4$), and the fraction of children with negative net worth ($n_g = 1$); all blocks other than the college-rate block carry weight $w_g = 1$. The group-size normalization $1/n_g$ ensures that larger groups do not mechanically dominate. College moments receive a weight of three because they are the core target of the model; down-weighting them leads to well-fitting wealth distributions but implausibly flat college-ability gradients. While an optimal weighting matrix based on the inverse of the variance-covariance of the moments would be asymptotically efficient, it requires a first-stage consistent estimate that is computationally expensive in this setting. The block-diagonal weighting provides a transparent alternative and is standard in the structural life-cycle literature ([Abbott et al., 2019](#); [Boar, 2021](#)).

5.4.1 Target Moments

The 38 targeted moments are organized into the following blocks, each chosen to provide leverage on specific parameters. Several blocks are the estimation counterparts of the facts in Section 4: the transfer block disciplines the level of parental support whose margins Section 4.3 documents, and the decumulation moment targets the slow pace of late-life wealth drawdown documented in the retirement-saving literature ([Love et al., 2009](#); [De Nardi et al., 2010](#)).

College graduation rates (16 moments). College graduation rates by ability quartile \times parent wealth quartile, from the NLSY97. These moments form the core of the estimation: they trace out how college attainment varies across the joint distribution of ability and parental resources, which is precisely the variation the model is designed to explain.

Income ratio (1 moment). The high-school-to-college mean income ratio, computed from individual gross earnings in the NLSY97. The cross-sectional income standard deviations by education are computed as well but left *untargeted* (the income process is fixed externally to [Abbott et al., 2019](#), so they serve as a validation check rather than a target).¹⁰

Transfer moments (5 moments). The mean college support received *conditional on attendance* and the probability of a post-college inter-vivos transfer by parent wealth quartile. The support moment comes from the NLSY97 college financial-aid questions, which record parental aid earmarked for schooling and not expected to be repaid; because the question is asked only of enrollees, the data object is conditional on attendance. The model counterpart is the matching earmarked object: the gift paid at enrollment, averaged across college attendees. Two definitional choices matter for identification. First, conditioning on attendance: an unconditional mean would conflate the attendance *rate*, already targeted by the twenty college-rate moments, with the transfer *size*, taxing the transfer parameters twice for any attendance misfit. Second, matching the earmarked concept: the model’s in-college flow transfers are general consumption support—the analog of allowances and in-kind room and board, which the NLSY97 does not measure with usable coverage—so including them in the model moment would compare a total-support object against an aid-only target. The flows are disciplined instead by the moments they do map into: the parent consumption and wealth paths, the college block, and the extensive margin below. The four transfer-probability moments are the HRS extensive margin of Section 4.3: the probability that a parent aged 50 or older gives \$500 or more to a non-coresident child aged 26 or older within a two-year survey wave, by parent-wealth quartile. The model counterpart splits the post-college overlap into two-year windows and counts a window as positive if cumulated flow transfers plus any lump-sum transfer reach the same threshold; these moments discipline the dynamic-game transfers—and through them the lump-sum channel and the free boundary—separately from

¹⁰Because each agent in the model represents a single individual—not a household—the income targets use individual gross earnings (wages and salary) from the NLSY97, not household income. This is consistent with the approach in [Boar \(2021\)](#) and [Lee and Seshadri \(2019\)](#), who also target individual labor income in single-agent dynastic models. The income ratio (0.69) is computed from the cross-section of individuals aged 25–37 to maximize sample size, while the (untargeted) standard deviations ($\sigma_{HS} = 27.7\text{K}$ and $\sigma_C = 41.7\text{K}$) are computed from ages 30–37. All figures are in thousands of 2016 dollars, deflated by the CPI-U. The sample conditions on positive earnings above \$5,000 and defines college graduates as those with a bachelor’s degree or higher.

the college-support level.¹¹

Wealth (10 targeted moments). From the PSID linked sample: the aggregate wealth-to-income ratio, which disciplines the overall level of wealth accumulation; the parent-child wealth rank-rank slope ($\hat{\beta}^{rr} = 0.372$); mean child wealth conditional on parent wealth quartile (4 moments); median wealth for parents and children; the variance of the inverse hyperbolic sine of child wealth; and the fraction of parents with net worth below \$25,000 (0.28), which disciplines the left mass of the parent wealth distribution. Two further dispersion measures of the parent wealth distribution are computed but left *untargeted* as validation checks: the p90/p50 ratio (7.36) and the variance of the inverse hyperbolic sine of parent wealth (roughly 40% of whose data value is generated by the negative-wealth tail, which the model rules out by the parent constraint $a_p \geq 0$).¹²

Decumulation (1 moment). The parent wealth decumulation rate over ages 60–70.

College graduation by parent income (4 moments). College graduation rates by parent-income quartile, from the NLSY97. These discipline the income gradient of attainment, complementing the wealth gradient in the 16 joint college moments.

Negative net worth (1 moment). The fraction of children aged 35–40 with negative net worth, from the PSID. This moment disciplines the left tail of the child wealth distribution generated by student-loan debt.

5.4.2 Identification

The model is overidentified: fourteen parameters are estimated from 38 targeted moments. While all moments jointly contribute to the estimation, the identification logic rests on distinct sources of variation for each parameter group.

¹¹The HRS quartiles are computed within survey-wave cohorts; Appendix A details the sample construction for both sources.

¹²The rank-rank slope of 0.372 is consistent with the range in the literature: Charles and Hurst (2003) estimate an intergenerational wealth elasticity of 0.37, and Pfeffer and Killewald (2018) find rank correlations of 0.40 or higher when children are measured at ages 45–64. My estimate falls in the middle because children in my sample are observed at ages 35–40.

Altruism (η). The altruism parameter is identified jointly from the extensive margin of post-college transfers and the wealth gradient in college attendance. A higher η enlarges the transfer region \mathcal{T} : flow support flows in more states, and—because lump-sum opportunities are exercised exactly when the state lies strictly inside \mathcal{T} —a larger region converts more Poisson arrivals into realized transfers. Since the arrival intensity λ_g is fixed and wealth-blind, the steep wealth gradient in the HRS transfer probabilities (5.5 percent in the bottom parent-wealth quartile to 24.8 percent in the top) can only be matched by the free boundary itself, which is what disciplines η . The college-attendance wealth gradient provides complementary leverage: a higher η raises the lifetime transfer burden parents face when children do not attend college, making college more attractive to wealthy parents. Finally, because lump-sum transfers move wealth—not just consumption— η now also bears on the intergenerational wealth moments with a positive sign, where previously altruism could only depress child self-saving through the Samaritan’s dilemma.

Psychic cost ($\kappa_0, \kappa_\theta, \sigma_\kappa$). The psychic cost parameters are identified from the ability gradient in college attendance, holding wealth constant. The constant κ_0 governs the overall college rate, the log-ability gradient κ_θ determines how steeply college attendance rises with ability, and the shock scale σ_κ controls the dispersion of college choices among low-ability types: a larger σ_κ generates more college attendance at the bottom of the ability distribution by giving some individuals favorable psychic cost draws.

Discount rate (ρ). The discount rate is pinned down by the aggregate wealth-to-income ratio and median wealth levels. A lower ρ (more patient agents) generates more wealth accumulation, raising both statistics.

Warm-glow weight (ϕ_b). The warm-glow motive is identified from the parent wealth decumulation rate: a stronger warm-glow motive (ϕ_b) slows wealth drawdown in retirement, as parents derive utility from holding wealth at T . The decumulation rate of +2.5% per year (parents ages 60–70 are still accumulating) is consistent with evidence that pre-retirees continue saving due to precautionary motives and warm-glow utility (De Nardi, 2004; Hurd, 1990). The intergenerational rank-rank slope and conditional child wealth moments provide

additional leverage on ϕ_b by disciplining how parental wealth transmits across generations.

Parent income fanning (g_{zp}). The fanning rate is identified from the left mass of the parent wealth distribution—the fraction of parents with net worth below \$25,000. Faster fanning spreads parent incomes over the working life, generating a larger mass of low-income, low-saving parents; without it, the model concentrates parent wealth too tightly around the median.

Paternalism (ξ). The paternalistic motive is identified from college attendance among low-ability children of wealthy parents. A positive ξ leads parents—particularly wealthy ones, whose marginal utility of consumption is low—to use the $t = 0$ gift to induce enrollment that the child’s own value gain would not justify, raising attendance precisely where ability is low and parental resources are high. The size of the gift itself provides additional leverage through the transfer moments.

Bequest fraction (α_{beq}). The bequest fraction is identified from the intergenerational wealth moments: the rank–rank slope, conditional child wealth by parent wealth quartile, and the variance of child wealth. A higher α_{beq} strengthens the intergenerational wealth correlation by channeling more parent wealth to the child at death.

Collateral channel (b_{slope}). The collateral channel slope is identified from the wealth gradient in college attendance, particularly at the top of the parent wealth distribution. A positive b_{slope} relaxes borrowing constraints for children of wealthy parents, raising college enrollment among high-wealth families.

Relative price of college labor (ω_C). The college wage scale is identified from the high-school-to-college mean income ratio. Because ω_C scales college earnings relative to high-school earnings, it maps almost one-for-one into that ratio, holding the externally set ability gradients and life-cycle profiles fixed.

Unsecured borrowing wedge (ι_s). The student-loan rate spread is identified from the fraction of children with negative net worth and, secondarily, the parent-wealth gradient in

college attendance. A larger wedge makes carried student debt costlier, compressing the left tail of the child wealth distribution and raising the dependence of enrollment on parental resources.

Intergenerational earnings premium (ζ_{pe}). The earnings premium is identified from the intergenerational wealth moments—the rank–rank slope and mean child wealth by parent wealth quartile. Because ζ_{pe} routes the parent–child correlation through inherited earning capacity rather than liquid wealth, it lifts the conditional child-wealth profile and the rank–rank slope without relaxing the child’s age-18 college-financing constraint.

Nonhomothetic bequest shifter. The shifter in the bequest function is identified from the upper tail of the wealth distribution: the p90/p50 ratio of parent wealth, the variance of the inverse hyperbolic sine of child wealth, and conditional child wealth in the top parent quartile. Following [De Nardi \(2004\)](#), a positive \underline{a}_{beq} makes bequests a luxury good, so poor parents bequeath little while the rich bequeath a rising share of wealth, generating the fat right tail.

Table 9. SMM-Estimated Parameters

| Parameter | Description | Value |
|--|--|-------|
| <i>Preferences</i> | | |
| β | Discount factor (annual), $\beta = e^{-\rho}$ | 0.96 |
| <i>Parental altruism & paternalism</i> | | |
| η | Altruism weight | 0.11 |
| ϕ_b | Warm-glow terminal wealth weight | 3.80 |
| κ_{beq} | Luxury-bequest shifter (De Nardi, \$000s) | 74.4 |
| ξ | Paternalistic motive (utility from child's college) | 0.07 |
| <i>College psychic cost (AGMV-style)</i> | | |
| κ_0 | Psychic-cost constant | 1.7 |
| κ_θ | Psychic-cost gradient in $\log(\theta)$ | -0.3 |
| σ_κ | Psychic-cost idiosyncratic shock S.D. | 1.14 |
| <i>College returns & financing</i> | | |
| w_C | Relative price of college labor (college wage scale) | 0.60 |
| b_1 | Collateral channel: borrowing-limit slope in a_p | 0.068 |
| ι | Unsecured borrowing wedge (rate spread on debt) | 0.111 |
| <i>Bequests & intergenerational transmission</i> | | |
| α_{beq} | Bequest fraction (share of terminal wealth inherited) | 0.37 |
| Δ_{z_0} | Parental-college premium in initial productivity z_0 | 1.31 |
| g_{zp} | Parent income-type fanning rate | 0.175 |

Notes: Parameters estimated jointly by the simulated method of moments, minimizing the percentage-deviation distance between the 38 targeted simulated moments and their data counterparts. Parameter definitions and the identification argument are given in Section 5. Externally set parameters, including the income process and the financial-aid schedule, are reported in Table 7.

6 Results

6.1 Model Fit

Table 10 reports the model's fit to the 38 targeted moments. Panel A shows college graduation rates by ability and parent wealth quartile. The model matches the overall level of attainment closely—about 43% in the model against 42% in the data—and reproduces the qualitative shape over most of the distribution: graduation rises with child ability within

the bottom three wealth quartiles and, in most ability groups, with parent wealth through the third quartile. The two systematic discrepancies are that the model’s *ability* gradient is somewhat too flat and its *wealth* gradient turns non-monotone at the top (discussed below). Panel B reports income moments and Panel C transfer moments. Panel D reports the intergenerational wealth distribution moments constructed from the PSID parent–child linked sample that discipline the warm-glow weight ϕ_b and the intergenerational wealth transmission channel.

The model’s systematic failure sits at the top of the wealth distribution. Averaged over wealth, attainment rises from 28% at the lowest ability quartile to 53% at the highest—a 25-point ability gradient against 35 points in the data—so the ability gradient is somewhat too flat. The sharper miss is the wealth profile: in the data, graduation rises monotonically from 32% in the bottom parent-wealth quartile to 54% in the top, whereas in the model it rises from 40% to 50% through the third quartile and then *falls* to 41% in the top quartile. The shortfall is concentrated among able children of wealthy parents—a top-ability child of a top-wealth parent attends at 44% in the model against 74% in the data—while attendance of children of poor parents is mildly over-predicted. Panel C shows the mirror image on the transfer side: the wealth gradient of the transfer probability matches the middle quartiles but overshoots at the top (40% versus 25%), and the mean transfer is too large (\$10.6k versus \$6.8k). The juxtaposition is informative: the same wealthy parents whose support is over-predicted have children whose attendance is under-predicted. Section 7 traces both to the Samaritan cushion—the anticipated support of a wealthy parent is strongest precisely in the non-college state, compressing the college value gain at the top of the wealth distribution; when the ongoing game is replaced by a one-time transfer, top-quartile attendance rises to 56%, in line with the data (Table 15). Among the remaining moments, the high-school-to-college income ratio in Panel B is under-predicted (0.41 versus 0.69), so the model’s college premium is currently too large, while the intergenerational wealth moments in Panel D fit well, with the rank–rank slope slightly below its data counterpart (0.30 versus 0.37).

Table 10. Targeted Moments: Data and Model

| <i>Panel A: College Attainment by Parent Wealth and Child Ability (NLSY97)</i> | | | | | |
|--|----------------|----------------|----------------|----------------|--|
| Wealth Quartile \ Ability Quartile | 1 | 2 | 3 | 4 | |
| 1 | 0.22 (0.19) | 0.40 (0.24) | 0.41 (0.33) | 0.55 (0.53) | |
| 2 | 0.30 (0.24) | 0.40 (0.30) | 0.49 (0.42) | 0.50 (0.53) | |
| 3 | 0.33 (0.26) | 0.48 (0.39) | 0.53 (0.51) | 0.64 (0.63) | |
| 4 | 0.27 (0.33) | 0.46 (0.46) | 0.45 (0.62) | 0.44 (0.74) | |

| <i>Panel B: Income Moments</i> | | | <i>Panel D: Intergenerational Wealth Moments</i> | | |
|---------------------------------------|-------|------|--|-------|-------|
| | Model | Data | | Model | Data |
| High-School/College mean income ratio | 0.41 | 0.69 | Wealth-income ratio | 2.21 | 1.87 |
| High-School Var(log income) | 0.31 | 0.45 | Rank-rank slope (IG wealth) | 0.30 | 0.37 |
| College Var(log income) | 0.49 | 0.42 | $E[a_c \text{parent wQ1}]$ (\$k) | 85.7 | 67.1 |
| <i>Panel C: Transfer Moments</i> | | | $E[a_c \text{parent wQ2}]$ (\$k) | 136.2 | 112.1 |
| | Model | Data | $E[a_c \text{parent wQ3}]$ (\$k) | 192.8 | 177.6 |
| Mean transfer (\$k/yr) | 10.6 | 6.8 | $E[a_c \text{parent wQ4}]$ (\$k) | 324.7 | 340.4 |
| Pr(transfer > 0) — Wealth Q1 | 0.03 | 0.05 | Median wealth, parents (\$k) | 90.7 | 85.4 |
| Pr(transfer > 0) — Wealth Q2 | 0.10 | 0.10 | Median wealth, children (\$k) | 71.9 | 51.2 |
| Pr(transfer > 0) — Wealth Q3 | 0.19 | 0.16 | Var(IHS wealth), parents | 6.22 | 5.99 |
| Pr(transfer > 0) — Wealth Q4 | 0.40 | 0.25 | Var(IHS wealth), children | 14.75 | 12.89 |
| | | | Parent wealth growth rate | 0.027 | 0.025 |

Notes: Model moments without parentheses; data moments in parentheses. Panel A: college graduation rates by parent wealth quartile (rows) and child ability quartile (columns), NLSY97. Panel B: income moments, NLSY97. Panels C and D: transfer and intergenerational wealth moments, PSID linked sample; wealth in thousands of 2016 dollars. Moment definitions, sources, and the targeted/untargeted distinction are given in Section 5; sample construction in Appendix A.

Figure 3 complements Panel D by comparing the full wealth distribution in the model and the data. The model reproduces the right-skewed shape of the empirical wealth distribution and the concentration of wealth among higher-income households. The fit matters because parent wealth governs both the college decision and the transfer policy.

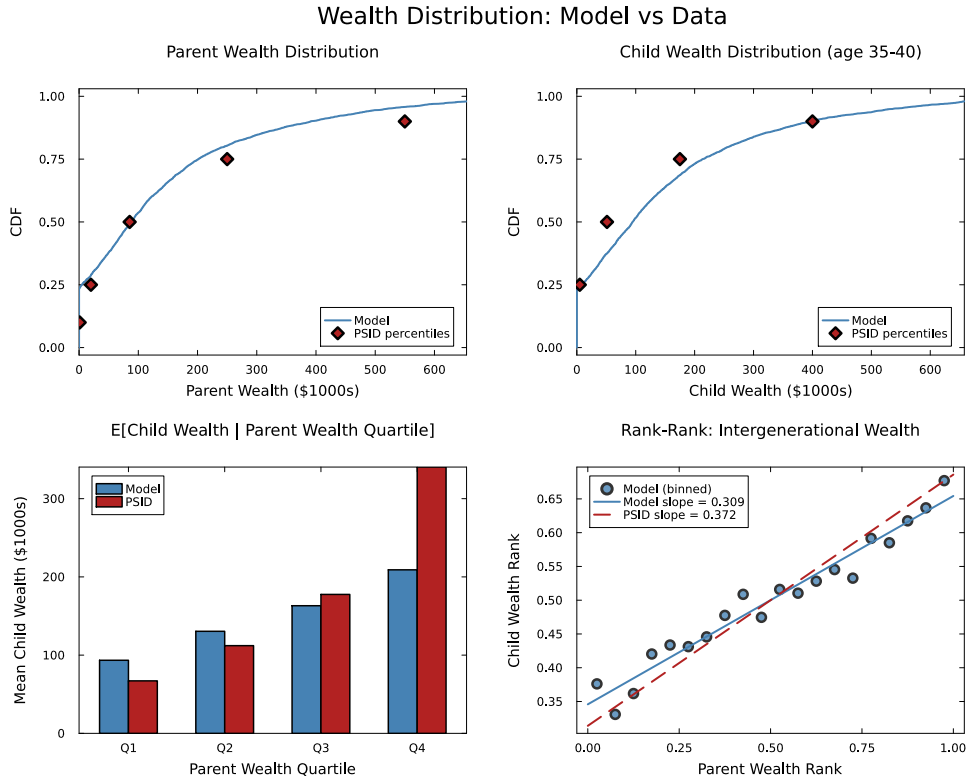


Figure 3. Wealth Distribution: Model vs. Data

Notes: Comparison of the parent wealth distribution in the model (simulated) and the data (PSID). Wealth is measured in thousands of 2016 dollars.

6.2 The Role of Parental Altruism in College Attainment

To quantify the importance of parental altruism, I solve the model with $\eta = 0$: parents make no transfers, so children finance college through their own earnings, loans, and grants alone. Table 11 reports college attendance by child ability quartile in the baseline and in this no-altruism economy. Removing altruism lowers overall attendance from 43% to 35%, and the effect is sharply heterogeneous: attendance falls most at the bottom of the ability distribution (27% to 11% in the lowest quartile, a 61% drop) and by roughly a fifth in the middle quartiles, while at the top it slightly *rises* (53% to 56%). The sign pattern previews the two opposing forces formalized in Section 7.1: for low- and medium-ability children the dominant role of altruism is to finance enrollment—the college gift and anticipated in-college support—so removing it gates them out of college; high-ability children attend largely regardless, so for

them the dominant role of altruism is the cushion it extends to the non-college path, and removing it slightly pushes them in.

Table 11. College Attendance With and Without Parental Altruism

| Ability quartile | Baseline Pr(C) | No altruism ($\eta = 0$) Pr(C) | Change (%) |
|------------------|-------------------|-------------------------------------|---------------|
| Q1 | 0.27 | 0.11 | -61 |
| Q2 | 0.43 | 0.35 | -18 |
| Q3 | 0.47 | 0.38 | -20 |
| Q4 | 0.53 | 0.56 | 6 |
| Overall | 0.43 | 0.35 | -18 |

Notes: College attendance by child ability quartile in the baseline economy and in a counterfactual with parental altruism shut down ($\eta = 0$), so the parent makes no transfers and the child finances college through own earnings, loans, and grants. The final column reports the percentage change relative to the baseline. The model is re-solved at the estimated parameters with the altruism weight set to zero; statistics are computed from $M = 4,000$ simulated dynasties.

This heterogeneity reflects channels that work in opposite directions. On one side, altruistic parents directly subsidize college costs, and they internalize that a child without college will generate larger lifetime transfer costs—for a wealthy parent, the present value of expected transfers to a high-school child can exceed the cost of tuition, making college the more efficient investment. On the other side, the child’s perspective introduces a countervailing force: anticipated parental transfers during and after college subsidize enrollment, but anticipated lifetime transfers also cushion children who forgo college against income risk. The net effect on college attendance is therefore theoretically ambiguous; Section 7 disentangles these forces. I first describe how college reshapes the transfer region in the baseline economy.

6.3 How College Reshapes the Transfer Region

In the baseline economy, parental support to adult children operates through the transfer region \mathcal{T} . In the mechanism, college reduces the child’s exposure to the transfer region by permanently raising expected income through a higher ability gradient ($\lambda_C = 0.797 > \lambda_{HS} = 0.517$) and a steeper life-cycle earnings profile: the higher income level keeps college children’s

resources further from the transfer boundary, so that a given negative shock is less likely to pull them into \mathcal{T} .

Panel (a) of Figure 4 plots the fraction of simulated individuals in the transfer region at each parent age over the post-college overlap window, by education, and Panel C of Table 12 quantifies this exposure, reporting the fraction of time receiving transfers separately for the in-college and post-college periods. The contraction is strong: college children spend 9.0% of the post-college overlap receiving transfers, against 30.1% for high-school children—a reduction of roughly two thirds in transfer-region exposure. During the college years themselves the child is heavily supported (67.6% of the time), the enrollment-subsidy channel of Force 1; after graduation, the permanently higher income level pulls the child away from the free boundary $\partial\mathcal{T}$, despite the student debt—serviced at the estimated borrowing wedge—with which college children exit school. This exposure measure is the model counterpart of the extensive-margin fact of Section 4.3, where college reduces the within-family probability of any transfer by 21% in the HRS; the model’s contraction has the right direction but is considerably stronger than in the data, a comparison taken up in Section 6.4.

Table 12. Transfer Region Exposure and Savings Behavior

| | HS Child | College Child |
|---|----------|---------------|
| <i>Panel A: Child's own savings effort $s_c^{own} = ra_c + y_c - c_c$ (\$k/yr)</i> | | |
| When receiving transfers | -7.67 | -4.41 |
| When not receiving transfers | 7.23 | 18.19 |
| Difference (Samaritan's dilemma) | -14.90 | -22.60 |
| <i>Panel B: Over-consumption vs. autarky $\Delta c_c = c_c^{MPE} - c_c^{aut}$ (\$k/yr)</i> | | |
| In transfer region | 17.53 | 9.53 |
| Outside transfer region | 12.33 | 44.62 |
| <i>Panel C: Fraction of time receiving parental transfers</i> | | |
| During college (ages 18–22) | — | 0.676 |
| After college (ages 22–48) | 0.301 | 0.090 |

Notes: Statistics from simulated dynasties ($M = 4,000$). Panel A reports the child's own savings effort $s_c^{own} = \tilde{r}_c a_c + y_c - c_c$, which excludes the mechanical effect of transfers on wealth accumulation. A negative difference indicates the child saves less of their own resources when receiving transfers (the Samaritan's dilemma). Panel B reports the consumption differential relative to an overlap-period autarky benchmark c_c^{aut} : the child's optimal consumption from a lifecycle problem solved at the same age and wealth level with no parent interaction, using the simulated parent wealth to compute the EFC for financial aid. A negative value indicates the child in the baseline economy consumes less than the autarky benchmark at the same state. Panel C reports the fraction of time the child receives parental transfers, separately for the in-college years (ages 18–22, the enrollment-subsidy channel) and the post-college period (ages 22–48, the insurance channel); the in-college figure applies only to college children.

The same pattern shows up in transfer flows. Panel (b) of Figure 4 shows average post-college transfers by parent wealth quartile and child education, and Table 13 reports the same quantities numerically. College children receive substantially *less* post-college support than their high-school counterparts in every parent wealth quartile: the reduction is about 60% in the bottom two quartiles and 80–85% in the top two, and pooled across quartiles college children receive \$13.7k against \$70.1k per year. The absolute levels rise steeply with parent wealth—in the model, wealthy parents heavily support a high-school child whose income disappoints (\$229k per year on average in the top quartile, against \$45k for a college child)—so the transfer-reducing effect of college is largest in dollar terms precisely where the Samaritan cushion is strongest. The top-quartile level of support to high-school children is implausibly large, the flow counterpart of the over-predicted top-quartile transfer probability

in Panel C of Table 10. The underlying force, however, is exactly the mechanism: a dollar spent on tuition generates a *permanent* increase in the child’s income drift, reducing the probability of entering the transfer region, whereas a dollar transferred directly smooths consumption without changing the drift.

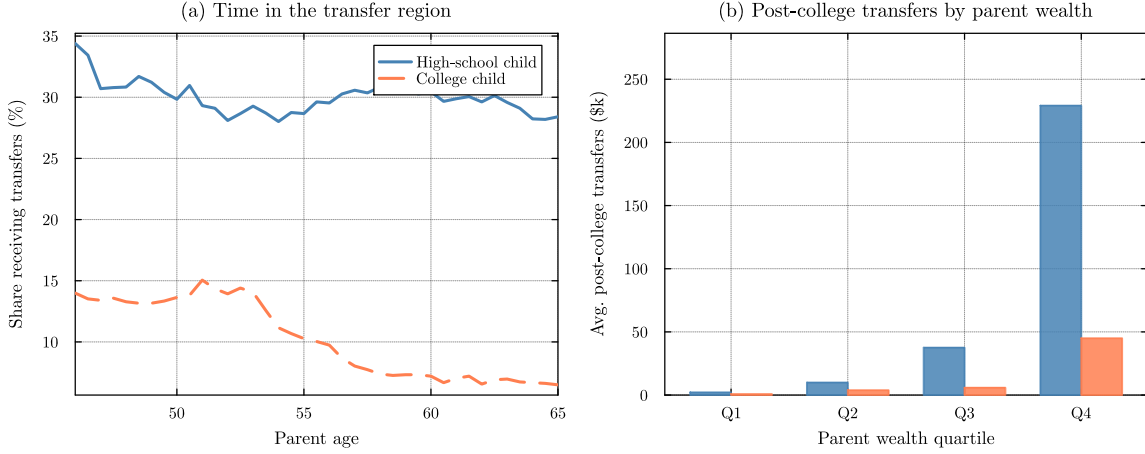


Figure 4. College, Transfer-Region Exposure, and Post-College Transfers

Notes: Simulated dynasties ($M = 4,000$). Panel (a) plots the fraction of children receiving positive transfers ($\tau > 0$) at each parent age over the post-college overlap window—child age 22–48, parent age 46–72—by child education. Panel (b) reports average post-college transfers received by the child (excluding the in-college years), by parent wealth quartile and child education.

Table 13. Post-College Transfers by Child Education and Parent Wealth

| Parent wealth quartile | High-school child (\$k/yr) | College child (\$k/yr) | Reduction (%) |
|------------------------|----------------------------|------------------------|---------------|
| Q1 | 2.1 | 0.8 | 63 |
| Q2 | 9.9 | 3.8 | 61 |
| Q3 | 37.5 | 5.8 | 85 |
| Q4 | 229.1 | 45.0 | 80 |
| All | 70.1 | 13.7 | 80 |

Notes: Average post-college transfers received by the child (excluding the in-college years), by parent wealth quartile and child education, from $M = 4,000$ simulated dynasties; this is the numerical counterpart of Panel (b) of Figure 4. The reduction is $1 - \bar{\tau}_C / \bar{\tau}_{HS}$, the proportional fall in transfers received by college relative to high-school children; negative values indicate that college children receive more. The “All” row pools across quartiles. Transfers are in thousands of 2016 dollars per year.

6.4 Untargeted Validation

The descriptive patterns of Section 4—the parent consumption release, the lower probability of any transfer, and the weakly larger transfer conditional on one occurring—are not targeted in the SMM objective, which disciplines the college-attendance gradients, the wealth distribution, the pace of decumulation, and the level of parental support (Section 5). They are therefore out-of-sample checks on the mechanism. Table 14 reports each data moment alongside its model counterpart, computed from the simulated panel over the post-college overlap.

The model now matches the direction of the extensive margin—the check most directly tied to the mechanism—but overshoots it, and at the current estimate misses the sign of the other two. The probability of any transfer falls with college in the model as in the data, but by 68% against 21% in the HRS: the transfer-region contraction visible in Panel C of Table 12 is roughly three times too strong. The consumption release currently has the *wrong sign*: parents of college children consume less in the model ($-\$19.8\text{k}$ per year, -14%) where in the data they consume more ($+\$2.5\text{k}$, $+8\%$). The intensive margin likewise: conditional on a transfer occurring, college children receive less in the model, while in the data the conditional amount is weakly larger. The misses are consistent with transfer flows that remain too large at the top of the wealth distribution: the model supports high-school children of wealthy parents at implausible rates (Table 13) and moves substantial wealth to college children around enrollment, so parents of college children—having parted with resources early—consume less thereafter, the reverse of the option-value release in the data, where the consumption gain reflects precautionary resources the parent no longer holds in reserve. Matching the extensive-margin contraction at realistic transfer magnitudes, and with it the sign of the consumption release, is the sharpest diagnostic the estimation must meet.

Table 14. Untargeted Moments: Data and Model

| Untargeted moment (effect of child’s college) | Data | Model |
|--|----------------------|-------|
| <i>Parent consumption release (college – HS, within parent-wealth quartile)</i> | | |
| Level (\$k/yr) | +2.5 | –19.8 |
| Percent | +8% | –14% |
| <i>Transfer extensive margin: Pr(any transfer)</i> | | |
| College effect (%) | –21 | –68 |
| <i>Transfer intensive margin: $\mathbb{E}[\tau \mid \tau > 0]$, college – HS</i> | | |
| Direction | + (weak, $p < .10$) | – |

Notes: Each row reports the effect of the child’s college attainment on a parent-side outcome that is *not* targeted in the SMM objective (Section 5); the data column reproduces the estimates of Section 4 and the model column is computed from $M = 4,000$ simulated dynasties over the post-college overlap, within parent-wealth quartile. The consumption release is the college–HS difference in mean parent consumption (\$k/yr and percent). The extensive margin is the percent change in the probability of any transfer; the model counterpart is the annual transfer probability, the data counterpart the HRS wave probability, so the levels are not directly comparable and only the college effect is. The intensive margin is the college–HS difference in the transfer conditional on a positive transfer; the data moment is measured per event and the model moment as a flow, so only the sign is comparable. Source: simulated model and Section 4.

7 Moral Hazard Decomposition

Parental altruism raises college attendance, but it also generates a fundamental tension: the Samaritan’s dilemma. Because children anticipate transfers when their wealth is low, they under-save relative to autarky, imposing larger and more frequent transfer costs on parents. The effect of this lack of commitment on the college decision is theoretically ambiguous. On one hand, anticipated parental support during the college years effectively subsidizes enrollment, making college more attractive to the child. On the other hand, anticipated lifetime transfers cushion non-college children against income risk, reducing their incentive to invest in human capital. Which force dominates is a quantitative question. This section formalizes both forces, constructs counterfactual economies that isolate them, and reports the decomposition.

7.1 Full Commitment Benchmark and Two Opposing Forces

To isolate the effects of the Samaritan's dilemma, I first characterize the efficient allocation under full commitment. Suppose a benevolent planner commits at $t = 0$ to a state-contingent transfer schedule $\{\tau(a_p, a_c, z, t)\}$ for the entire overlap period. Because both agents pool resources under a single plan, the problem reduces to a single-agent optimization over total household wealth $A \equiv a_p + a_c$:

$$\begin{aligned} \rho V^{\text{FC}}(A, z, t) = \max_{c_{\text{tot}}} \left\{ u_{\text{FC}}(c_{\text{tot}}) + V_A^{\text{FC}} [rA + y_p(t) + y_c(z, t) - c_{\text{tot}}] \right. \\ \left. + \frac{1}{2} \sigma_a^2 A^2 V_{AA}^{\text{FC}} + \mathcal{L}_z V^{\text{FC}} + V_t^{\text{FC}} \right\} \end{aligned} \quad (16)$$

Under CRRA preferences, the efficient allocation equates weighted marginal utilities, $u'(c_p) = \eta u'(c_c)$, so consumption shares are time-invariant: $c_c = \eta^{1/\gamma} c_p$. Total consumption $c_{\text{tot}} = c_p(1 + \eta^{1/\gamma})$ yields an aggregated flow utility $u_{\text{FC}}(c_{\text{tot}}) = \frac{c_{\text{tot}}^{1-\gamma}}{1-\gamma} (1 + \eta^{1/\gamma})^\gamma$, and the planner's HJB reduces to a single-state-variable problem. Full commitment differs from the baseline economy in three ways (Barczyk and Kredler, 2014a,b): it eliminates overconsumption by both agents, it renders the timing of transfers indeterminate, and—most important for the college decision—it provides insurance without moral hazard. The child's college choice under full commitment follows the same probit structure as the baseline (equation 13), with the planner's continuation values replacing the baseline values.

Relative to full commitment, the Samaritan's dilemma creates two opposing forces on the child's college decision. From the parent's perspective, a college-educated child is cheaper to support: college permanently raises the income level, shifting the child away from the transfer boundary. Under full commitment, the parent could internalize this motive directly by conditioning transfers on e_c . In the baseline economy, however, the parent cannot commit: the transfer policy responds only to the current state (a_p, a_c, z) , and the parent transfers whenever the child is sufficiently poor, regardless of the child's education choice.

The child chooses education to maximize own continuation utility, taking the parent's equilibrium transfer policy as given. The continuation values V_C and V_{HS} embed the entire expected path of future transfers. The child does not internalize the cost transfers impose on the parent, generating overconsumption relative to the efficient allocation whenever the

child’s state is near or inside \mathcal{T} . Two forces emerge from this interaction:

Force 1: Anticipated transfers subsidize enrollment. The child starts college with zero assets and reduced income ($\ell_C \cdot y_c$ net of tuition), placing the initial state deep in the transfer region \mathcal{T} . The parent—bound by time consistency—optimally transfers resources that effectively subsidize enrollment costs. This implicit subsidy raises the child’s valuation of the college path relative to full commitment, pushing attendance *above* the efficient level.

Force 2: Anticipated transfers cushion non-college children. Children who forgo college face lower expected income, so they spend substantially more of their adult lives in the transfer region. The parent cannot withhold this support. The anticipation of this lifetime safety net raises the child’s valuation of the high-school path—the child can under-save and over-consume, knowing that parental transfers will materialize when income falls. This cushion reduces the child’s incentive to attend college, pushing attendance *below* the efficient level.

Whether Force 1 or Force 2 dominates depends on ability, parental wealth, and the income process parameters. The decomposition below quantifies these forces; at the current estimate the enrollment subsidy of Force 1 dominates for most families—moral hazard raises attendance throughout the bottom three parent-wealth quartiles—while the non-college cushion of Force 2 dominates in the top wealth quartile, where the anticipated support is largest relative to the return to college.

7.2 Counterfactual Economies and Quantitative Decomposition

To isolate the role of commitment from the role of insurance, I construct two counterfactual economies that eliminate ongoing parent-child interaction entirely, replacing the dynamic transfer game with a one-time lump-sum transfer at $t = 0$. After the initial transfer, the child solves the standard lifecycle problem alone—shutting down both the Samaritan’s dilemma and the parental insurance channel. The two counterfactuals differ in whether the parent gives the same transfer in both education states or an education-contingent menu (τ_C, τ_{HS}) , which separates the pure moral hazard channel from the education-targeting channel.¹³

¹³Appendix H provides a formal derivation of the full commitment benchmark and discusses its relationship to the one-time transfer counterfactuals. The key distinction is that the FC benchmark removes moral hazard

Counterfactual A: One-time college gift. The parent chooses a college-conditional gift τ_C at $t = 0$, paid at enrollment exactly as in the baseline:

$$\begin{aligned} \max_{\tau_C \in [0, a_p]} \Pr(C | \tau_C, \theta) & \left[V_p^{\text{alone}}(a_p - \tau_C) + \eta V^{c, \text{alone}}(\tau_C | C) + \xi \right] \\ & + (1 - \Pr(C | \tau_C, \theta)) \left[V_p^{\text{alone}}(a_p) + \eta V^{c, \text{alone}}(0 | HS) \right] \end{aligned} \quad (17)$$

where $V_p^{\text{alone}}(a)$ is the parent's remaining lifetime utility with no further interaction with the child. The child's education choice follows the same probit rule as the baseline (equation 13), with the child-alone value gain replacing the coupled values:

$$\Pr(C | \tau_C, \theta) = \begin{cases} \Phi\left(\frac{G_0(\tau_C, \theta)/\bar{V} - \kappa_0 - \kappa_\theta \log \theta}{\sigma_\kappa}\right) & \text{if } G_0 > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (18)$$

where $G_0(\tau_C, \theta) \equiv V^{c, \text{alone}}(\tau_C | C) - V^{c, \text{alone}}(0 | HS)$ is the child-alone value gain from college given the gift τ_C . This economy retains the baseline's $t = 0$ decision structure—the college-conditional gift and the paternalistic motive ξ both remain—and removes only the *ongoing* interaction: after the enrollment date, no further transfers flow. The education decision therefore reflects fundamentals—human capital returns net of psychic cost, given the gift—and not the anticipation of state-contingent support. Comparing baseline attendance with this economy isolates the *pure moral hazard effect*: the distortion to college enrollment caused solely by the child's anticipation of future transfers.

Counterfactual B: Education-contingent transfer menu. The parent commits at $t = 0$ to a transfer schedule that may differ by the child's education choice, choosing a transfer τ_C paid if the child attends college and a transfer τ_{HS} paid if the child chooses high school:

$$\begin{aligned} \max_{(\tau_C, \tau_{HS}) \in [0, a_p]^2} \Pr(C | \tau_C, \tau_{HS}, \theta) & \left[V_p^{\text{alone}}(a_p - \tau_C) + \eta V^{c, \text{alone}}(\tau_C | C) + \xi \right] \\ & + [1 - \Pr(C | \tau_C, \tau_{HS}, \theta)] \left[V_p^{\text{alone}}(a_p - \tau_{HS}) + \eta V^{c, \text{alone}}(\tau_{HS} | HS) \right] \end{aligned} \quad (19)$$

while preserving insurance, whereas the one-time transfer counterfactuals remove both.

$$\Pr(C \mid \tau_C, \tau_{HS}, \theta) = \begin{cases} \Phi\left(\frac{G_m(\tau_C, \tau_{HS}, \theta)/\bar{V} - \kappa_0 - \kappa_\theta \log \theta}{\sigma_\kappa}\right) & \text{if } G_m > 0, \\ 0 & \text{otherwise,} \end{cases} \quad (20)$$

where the value gain is the child's own, evaluated at the education-specific endowments:

$$G_m(\tau_C, \tau_{HS}, \theta) \equiv V^{c,\text{alone}}(\tau_C \mid C) - V^{c,\text{alone}}(\tau_{HS} \mid HS).$$

where $\kappa(\theta, \varepsilon) \geq 0$ is the psychic cost of college (zero for the HS alternative). This menu is the efficient one-time-transfer benchmark: it nests Counterfactual A (the restriction $\tau_{HS} = 0$) and an unconditional one-time transfer (the restriction $\tau_C = \tau_{HS}$), and because it relaxes both restrictions it weakly dominates them in parent welfare. Crucially, the altruistic parent can now support a child in the high-school state directly, so a low-return child is no longer pushed into college merely to capture parental resources. Comparing Counterfactual B with Counterfactual A isolates the *education-targeting effect*: the change in attendance when the parent can direct transfers by education rather than give the same endowment in both states.

Decomposition. Define:

$$\Delta_{\text{MH}}^{\text{pure}}(\theta, a_p) = \Pr(\text{College} \mid \text{baseline}) - \Pr(\text{College} \mid \text{unconditional}) \quad (21)$$

$$\Delta_{\text{targ}}(\theta, a_p) = \Pr(\text{College} \mid \text{unconditional}) - \Pr(\text{College} \mid \text{menu}) \quad (22)$$

$$\Delta_{\text{MH}}^{\text{total}}(\theta, a_p) = \Delta_{\text{MH}}^{\text{pure}} + \Delta_{\text{targ}} \quad (23)$$

A positive $\Delta_{\text{MH}}^{\text{pure}}$ indicates that the baseline economy raises college attendance relative to the no-moral-hazard economy; a negative value indicates that moral hazard depresses college attendance. The education-targeting component Δ_{targ} isolates the change in enrollment when the parent can direct transfers by education: a positive sign indicates that targeting *lowers* attendance, because the parent diverts low-return children out of college and supports them in the high-school state instead.

Results. Table 15 reports the decomposition.

Table 15. Moral Hazard Decomposition of College Attendance

| | Overall | Parent Wealth Quartile | | | |
|--|---------|------------------------|--------|--------|-----------|
| | | Q1 (Low) | Q2 | Q3 | Q4 (High) |
| <i>Panel A: College Attendance Rate (all ability types)</i> | | | | | |
| Baseline | 0.426 | 0.373 | 0.419 | 0.494 | 0.417 |
| No MH (unconditional) | 0.400 | 0.306 | 0.318 | 0.422 | 0.555 |
| No MH (menu) | 0.400 | 0.306 | 0.318 | 0.421 | 0.555 |
| Δ_{MH}^{pure} | +0.026 | +0.067 | +0.101 | +0.072 | -0.138 |
| Δ_{MH}^{total} | +0.026 | +0.067 | +0.101 | +0.073 | -0.138 |
| <i>Panel B1: Pure Moral Hazard Effect (Δ_{MH}^{pure}) by Ability</i> | | | | | |
| Low ability | | -0.056 | +0.019 | +0.031 | -0.067 |
| Medium ability | | +0.170 | +0.210 | +0.117 | -0.128 |
| High ability | | +0.129 | +0.066 | +0.072 | -0.192 |
| <i>Panel B2: Education-Targeting Effect (Δ_{targ}) by Ability</i> | | | | | |
| Low ability | | +0.000 | -0.000 | -0.000 | -0.000 |
| Medium ability | | +0.000 | -0.000 | -0.000 | +0.000 |
| High ability | | +0.000 | +0.000 | +0.000 | +0.000 |
| <i>Panel B3: Total Effect ($\Delta_{MH}^{total} = \Delta_{MH}^{pure} + \Delta_{targ}$) by Ability</i> | | | | | |
| Low ability | | -0.056 | +0.019 | +0.031 | -0.067 |
| Medium ability | | +0.170 | +0.210 | +0.117 | -0.128 |
| High ability | | +0.129 | +0.066 | +0.072 | -0.192 |
| <i>Panel C: Optimal One-Time Transfer τ_0^* (\$000s)</i> | | | | | |
| <i>Unconditional:</i> | | | | | |
| Low ability | | 0.0 | 0.0 | 0.0 | 426.8 |
| Medium ability | | 0.0 | 44.1 | 104.0 | 269.7 |
| High ability | | 0.0 | 66.0 | 145.9 | 242.5 |
| <i>Menu, if college (τ_C):</i> | | | | | |
| Low ability | | 0.0 | 0.0 | 0.0 | 0.0 |
| Medium ability | | 0.0 | 44.1 | 104.0 | 269.7 |
| High ability | | 0.0 | 66.0 | 145.9 | 242.5 |
| <i>Menu, if high school (τ_{HS}):</i> | | | | | |
| Low ability | | 0.0 | 0.0 | 104.0 | 500.9 |
| Medium ability | | 0.0 | 0.0 | 0.0 | 0.0 |
| High ability | | 0.0 | 0.0 | 0.0 | 0.0 |

Notes: Panel A reports average college attendance across ability types under three economies: the baseline, the no-MH economy with an unconditional transfer, and the no-MH economy with the education-contingent menu (τ_C, τ_{HS}). The ‘‘Overall’’ column averages across wealth quartiles. In both no-MH economies, parent and child interact only at $t = 0$; the child solves the lifecycle problem alone thereafter.

$\Delta_{MH}^{pure} = \Pr(\text{College} \mid \text{baseline}) - \Pr(\text{College} \mid \text{unconditional})$ isolates the overconsumption distortion; $\Delta_{MH}^{total} = \Pr(\text{College} \mid \text{baseline}) - \Pr(\text{College} \mid \text{menu})$ adds the education-targeting channel. Panels B1–B3 disaggregate these effects by child ability. Panel C reports the optimal unconditional transfer and the menu transfers τ_C (if college) and τ_{HS} (if high school). Source: simulated model ($M = 4,000$ dynasties).

Aggregate result. Table 15 reports college attendance under the baseline, the unconditional one-time transfer, and the education-contingent menu; Panel A is computed on the simulated population, so its baseline rate coincides with Table 10. The ongoing game *raises* aggregate attendance at the current estimate: $\Delta_{MH}^{\text{pure}} = +2.6$ percentage points (42.6% in the baseline against 40.0% under the one-time transfer), so on average the enrollment subsidy of Force 1 outweighs the high-school cushion of Force 2. The aggregate, however, conceals a sharp reversal across the wealth distribution: moral hazard raises attendance in the bottom three wealth quartiles (+6.7, +10.1, and +7.2 pp) and *lowers* it in the top quartile (−13.8 pp), where the anticipated cushion is largest. The targeting effect is $\Delta_{\text{targ}} = 0$: the optimal menu is essentially indistinguishable from the unconditional transfer (discussed below), leaving a total distortion relative to the efficient menu of $\Delta_{MH}^{\text{total}} = +2.6$ pp. The menu is the appropriate benchmark: a college-or-nothing restriction ($\tau_{HS} = 0$) would mechanically force attendance upward by withholding all support from non-college children, overstating the role of commitment.

Decomposition by ability. Panels B1–B3 of Table 15 disaggregate the effects by child ability and parent wealth. The over-enrollment from the ongoing transfer game is concentrated among medium-ability children of low- and middle-wealth parents: the pure moral-hazard effect reaches +0.17 and +0.21 in the bottom two wealth quartiles, where the anticipated in-college support of Force 1 finances an enrollment that own resources and loans alone would not; high-ability children in those quartiles show the same pattern more weakly (+0.07 to +0.13), and for low-ability children the two forces roughly cancel. In the top wealth quartile every ability group is pushed *out* of college by the cushion, and most strongly the able (−0.19 for high-ability children): the children with the most to gain from enrollment are precisely those for whom a wealthy parent’s anticipated lifetime support is the most attractive outside option. The under-prediction of top-quartile attendance in the model fit (Table 10) is the same phenomenon—removing the ongoing game lifts top-quartile attendance to 55.5%, in line with the 54% in the data. The education-targeting channel (Panel B2) is essentially zero everywhere even though the menu’s policy functions display the disciplining pattern (Panel C): in the upper wealth quartiles the menu reserves the college-state gift for able children ($\tau_C > 0$ for medium and high ability) and grants a high-school-state transfer to

low-ability children (τ_{HS} up to roughly \$500k in the top quartile)—the “I fund you only if you enroll” logic the model is built to express. But a one-time lump raises the child’s value on both branches nearly symmetrically, and in the cells where the menu does redirect support the enrollment probabilities barely move, so the disciplining menu is present in the optimal policy yet quantitatively inert at the current estimate.

This pattern complements the no-altruism comparison in Section 6.2. Shutting down altruism entirely (Table 11) lowers aggregate attendance by eight percentage points, concentrated at the bottom of the ability distribution: the *level* effect of altruism is to finance enrollment. Removing only the ongoing transfer game while preserving a one-time transfer (Table 15) lowers attendance by 2.6 points on average but *raises* it in the top wealth quartile. Read together, the two exercises say that parental support raises college attendance on net, and that even its non-committal form does so for most families, because the in-college subsidy outweighs the cushion; only where parents are wealthy enough that anticipated lifetime support rivals the return to college does the lack of commitment depress enrollment—precisely the families for whom an education-contingent commitment would bind, were it not quantitatively inert at the current estimate.

8 Welfare Analysis

This section quantifies the welfare consequences of the Samaritan’s dilemma—weighing the cost of moral hazard against the insurance value of ongoing transfers—and evaluates a free-college policy with and without commitment.

8.1 Consumption-Equivalent Variation

I measure welfare using the consumption-equivalent variation (CEV): the permanent proportional change in consumption that makes an agent indifferent between the baseline (Samaritan’s dilemma) economy and an alternative. Under CRRA utility $u(c) = c^{1-\gamma}/(1-\gamma)$, CRRA homogeneity implies that scaling all consumption flows by a factor $(1 + \Delta)$ scales lifetime

utility by $(1 + \Delta)^{1-\gamma}$. Inverting gives the CEV

$$\Delta = \left(\frac{V^{\text{alt}}}{V^{\text{base}}} \right)^{\frac{1}{1-\gamma}} - 1, \quad (24)$$

where $\Delta > 0$ means the alternative is welfare-improving.

For the child, $V^c = \int_0^T u(c_c) e^{-\rho t} dt$ and equation (24) applies directly. For the parent, $V^p = \int_0^T [u(c_p) + \eta u(c_c)] e^{-\rho t} dt$ captures both own and altruistic consumption. Because u is homogeneous of degree $1 - \gamma$, scaling both c_p and c_c by $(1 + \Delta)$ scales V^p by $(1 + \Delta)^{1-\gamma}$. The parent CEV therefore measures the permanent household-wide consumption change—own plus the altruistic component—that achieves indifference.

I compute CEV at each grid point (a_p, θ) by evaluating the value functions at the $t = 0$ decision state ($z = 0$), with the college branch at the post-gift state $(a_p - \tau_C^*, a_c = \tau_C^*)$ and the high-school branch at $(a_p, 0)$, integrating over the college-choice psychic-cost shock. I then average over the simulated population: each grid CEV is weighted by the model’s stationary distribution of parent wealth and child ability, using the same quartile sort as the targeted moments, so the welfare averages are comparable to the model-fit and decomposition tables rather than to a uniform grid.

8.2 Welfare Across Commitment Structures

Table 16 reports CEV for three alternative economies relative to the baseline. The first two rows of Panel A correspond to the no-moral-hazard counterfactuals from Section 7: the unconditional one-time transfer eliminates ongoing transfers but gives the child the same endowment in both education states, while the education-contingent menu lets the parent choose separate transfers (τ_C, τ_{HS}) . Both allocations remove the Samaritan’s dilemma by construction—the parent commits to a transfer schedule at $t = 0$ and the child solves an independent problem thereafter. In both, the parent’s own continuation value is the single-agent parent HJB of Appendix J.7, solved with the same consumption floor, bequest motive, and wealth diffusion as the baseline coupled value, so the consumption-equivalent comparison differences like with like rather than against a permanent-income approximation.

Measured by the dynastic (η -weighted) consumption-equivalent variation—the welfare object the parent optimizes—the two one-time-transfer economies are *worse* than the baseline:

the unconditional transfer delivers -7.6% and the education-contingent menu -7.4% , while full commitment delivers $+33.1\%$. The contrast is the central result of this section. Both one-time transfers eliminate the Samaritan’s dilemma, but they do so by severing the ongoing relationship entirely, which also removes the insurance that repeated transfers provide against the child’s income risk; the dynasty is worse off, with the loss carried by the parent’s side of the ledger (-8.6% and -8.3% on the parent’s own value) while the child loses little (-0.6% and -1.3%). The menu adds almost nothing to the unconditional transfer—the welfare counterpart of the inert targeting channel of Section 7—so the disciplining value of conditioning support on schooling does not register in welfare here. Full commitment removes the moral hazard while *preserving* insurance—the planner still moves resources to the child state by state—and is the only economy that raises dynastic welfare. That gain is almost entirely the parent’s—parent-own CEV $+52.9\%$ against the child’s -34.6% —because commitment strips the child of the ability to extract resources by under-saving. With altruism estimated at $\eta = 0.11$, the dynastic measure $V_{\text{own}}^p + \eta V_{\text{own}}^c$ is dominated by the parent’s own utility, so the positive dynastic figure should be read as the value of commitment *to the parent*: it measures the rents the parent recovers by no longer being exploited, not a Pareto improvement.

Table 16. Welfare Analysis: CEV (%) Relative to Baseline

| | Parent (own) CEV (%) | Child CEV (%) | Dynastic (+ η) CEV (%) |
|--|-------------------------|------------------|---------------------------------|
| <i>Panel A: Commitment Structure</i> | | | |
| Baseline (Samaritan's dilemma) | 0.00 | 0.00 | 0.00 |
| Unconditional transfer (no MH) | -8.58 | -0.64 | -7.57 |
| Education-contingent menu (no MH) | -8.30 | -1.30 | -7.40 |
| Full commitment (planner) | 52.89 | -34.64 | 33.08 |
| <i>Panel B: Free College Policy ($\phi = 0$, lump-sum tax)</i> | | | |
| Free college, no commitment | -0.29 | 2.98 | 0.13 |
| Free college, with commitment | -8.95 | 2.48 | -7.53 |
| <i>Panel C: Value of commitment for the free-college policy (CEV points)</i> | | | |
| Value of commitment | -8.66 | -0.50 | -7.66 |

Notes: Consumption-equivalent variation (CEV) measures the permanent proportional change in consumption that achieves indifference between each economy and the baseline. The *Parent (own)* column uses the parent's own-consumption value V_p^{own} ; the *Child* column uses the child's own value; the *Dynastic* column uses the η -weighted dynastic value $V_p^{\text{own}} + \eta V_c^{\text{own}}$ that the parent optimizes and is the welfare-relevant aggregate. All averages are weighted by the model's stationary population distribution of (a_p, θ) . Panel A compares commitment structures holding tuition at its estimated value. Panel B sets $\phi = 0$ and funds the policy with a budget-balancing lump-sum tax T^{LS} (equation 25) solved in general equilibrium for each regime. Panel C reports the value of commitment for the free-college policy—the difference in its CEV gain under commitment versus the time-consistent equilibrium, $\Delta^{\text{cmt}} - \Delta^{\text{no cmt}}$, in CEV points—reported as a signed level rather than a ratio so that it remains well defined when gains are small or negative. Source: model solution at estimated parameters.

The parent would like to implement full commitment—the one allocation that raises dynastic welfare—but cannot, and the reason is precisely the parent's altruism. Committing to hand the child a one-time transfer and then stand back is not time-consistent. Ex ante the parent would like to promise to transfer τ_0 once and never again, inducing the child to self-insure and eliminating the moral hazard. Ex post, however, whenever the child draws a bad shock and approaches the transfer region, the parent—precisely because $\eta > 0$ —strictly prefers to transfer rather than let the child consume little. The child anticipates this and under-saves, so the promise unravels: the only time-consistent (Markov-perfect) policy is the

baseline one, in which the parent transfers whenever the child is sufficiently poor, regardless of education. The commitment allocations are therefore benchmarks, not policies the parent can choose unilaterally. This is why college matters: by permanently raising the child’s income, a tuition payment moves the child away from the transfer region for the remainder of life, achieving through human capital the commitment that cash cannot.

8.3 Free-College Policy

I consider a policy that eliminates tuition ($\phi = 0$) and finances the revenue shortfall through a lump-sum tax T^{LS} on initial parent wealth. No financial aid is needed when tuition is zero, so $\bar{g} = 0$. The tax is set in general equilibrium to balance the government budget: it must cover tuition for the children who enroll, while enrollment itself responds to the tax through the parent’s post-tax wealth. The per-dynasty tax therefore solves the fixed point

$$T^{\text{LS}} = 4\phi \cdot \mathbb{E}[\text{Pr}(C \mid a_p - T^{\text{LS}}, \theta)], \quad (25)$$

where ϕ is the annual tuition defined above, 4ϕ is the four-year cost of tuition the government covers per enrolled child, and the expectation is over the (a_p, θ) distribution. Enrollment, transfers, and value functions are all evaluated at the post-tax wealth $a_p - T^{\text{LS}}$. Because enrollment differs across the two behavioral regimes, the budget-balancing tax is solved separately for each. Parents with $a_p < T^{\text{LS}}$ have post-tax wealth clamped at zero; they still benefit because their child pays no tuition.

I solve the free-college economy under two regimes. Under *no commitment*, the model retains its Markov-perfect equilibrium structure—parents and children interact strategically with $\phi = 0$ —and the Samaritan’s dilemma persists. Under *commitment*, the parent commits to a one-time education-contingent menu (τ_C, τ_{HS}) at $t = 0$ and the child solves independently with $\phi = 0$.

Panel B of Table 16 reports the results. Under no commitment, free college raises enrollment from 42.6% to 52.7% and operates as a transfer from parents to children: the child gains (+3.0%) while the parent bears the lump-sum tax (−0.3%, present value \$13.2k), for a dynastic CEV that is roughly neutral (+0.1%) and modestly positive in the bottom parent-wealth quartile (Table 17). Under commitment, enrollment rises further to 56.0% but the

child gains slightly less (+2.5%) and the parent now bears a heavy burden (−9.0%), so the dynastic change is sharply negative (−7.5%). That burden is not the policy: it is the cost of severing the ongoing relationship—the same loss the one-time-menu commitment carries in Panel A—which here swamps the tuition subsidy. The lesson is the reverse of the usual commitment story: because the parent cannot commit without giving up the insurance value of repeated transfers, free college delivers more to the family precisely *when* the parent cannot commit and the ongoing relationship is preserved.

Panel C reports the *value of commitment* for the free-college policy: the difference between the policy’s dynastic CEV gain under commitment and under the time-consistent (Markov) equilibrium, $\Delta^{\text{cmt}} - \Delta^{\text{no cmt}}$, in consumption-equivalent points. I report it as a signed level rather than a ratio so that it remains interpretable when the underlying gains are small or, as here, negative. The value of commitment is itself *negative* (about −7.7 CEV points on the dynastic measure): committing to the one-time menu is worse than leaving the policy to play out under the ongoing relationship, because the implementable form of commitment in this model forfeits the insurance value of repeated transfers. The parent cannot commit with cash without destroying the insurance the family relies on; the permanent income increase from college, not a transfer schedule, is the commitment technology that matters. The first-best benchmark that would preserve insurance, the full-commitment planner of Panel A, is not implementable by the altruistic parent. The shortfall widens with parental wealth: Table 17 shows the value of commitment falling from about −5.9 CEV points in the bottom parent-wealth quartile to −10.1 in the top, because wealthier families hold more of their consumption insurance in the ongoing relationship and so lose more by replacing it with a fixed schedule.

Table 17. Free-College Welfare and the Value of Commitment by Parent Wealth Quartile

| Wealth Q | No Commitment | | With Commitment | | Value of comm. (pp) |
|----------|---------------|-------|-----------------|-------|---------------------|
| | Dynastic | Child | Dynastic | Child | |
| Q1 | 0.96 | 3.73 | -4.98 | 3.15 | -5.94 |
| Q2 | -0.16 | 3.73 | -6.99 | 3.01 | -6.83 |
| Q3 | -0.30 | 3.00 | -8.09 | 2.14 | -7.79 |
| Q4 | 0.01 | 1.44 | -10.07 | 1.62 | -10.07 |

Notes: Consumption-equivalent variation (CEV, %) of the free-college policy relative to the baseline, by parent wealth quartile, under the no-commitment (Markov) equilibrium and under the education-contingent menu commitment. The *Dynastic* columns use the η -weighted value the parent optimizes; the *Child* columns the child's own value. The final column is the value of commitment for the policy—the dynastic CEV gap, $\Delta^{\text{cmt}} - \Delta^{\text{no cmt}}$, in CEV points. All averages are weighted by the model's stationary population distribution of (a_p, θ) within each parent-wealth quartile. Source: model solution at estimated parameters.

9 Conclusion

This paper argues that college serves as a commitment device against the Samaritan's dilemma. In matched parent-child data, parent consumption is higher when children attend college and transfers are less frequent though conditionally larger, with the consumption gain far exceeding the fall in realized transfers—the signature of the mechanism: college shrinks the transfer region, releasing the precautionary resources parents hold to stand ready to support their children.

The net effect of the Samaritan's dilemma on college attendance is theoretically ambiguous—the lack of commitment creates an enrollment subsidy (Force 1) that raises attendance and a transfer cushion (Force 2) that depresses it—and which force dominates, and for whom, is a quantitative question. The structural decomposition makes that question answerable: it separates the pure moral-hazard effect from the education-targeting effect and traces both across the joint distribution of ability and parental wealth, where the balance of the two forces differs systematically across families. Whether eliminating the friction improves welfare depends on what replaces it: the ongoing transfer relationship that distorts the college margin also insures children against income risk, so severing it with a one-time transfer can

leave the dynasty worse off, whereas a commitment that preserves insurance raises welfare. The welfare effects of a free-college policy likewise hinge on whether parents can commit. The same altruism parameter η that generates the Samaritan's dilemma also drives parents to invest in college as a way to mitigate it.

The framework carries implications for the design of college subsidies and financial aid. Because tuition policy operates on a margin embedded in an ongoing family relationship, its incidence is not confined to students: lowering the price of college reallocates resources within the dynasty, and how the gains divide between parent and child depends on the transfer relationship the policy leaves in place. The welfare effects of financial aid therefore depend on how policies interact with private transfers: grants that crowd out parental contributions may be less effective at raising enrollment than those that complement family support, while policies that expand parents' capacity to transfer without encouraging education may deepen the Samaritan's dilemma.

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.
- Altonji, Joseph G, Fumio Hayashi, and Laurence J Kotlikoff**, “Is the extended family altruistically linked? direct tests using micro data,” *The American Economic Review*, 1992, *82*, 1177–1198.
- Ameriks, John, Joseph Briggs, Andrew Caplin, and Mi Luo**, “Inter-generational Transfers and Precautionary Saving,” 2016 Meeting Paper 1616, Society for Economic Dynamics 2016.
- Attanasio, Orazio, Costas Meghir, and Corina Mommaerts**, “Insurance in Extended Family Networks,” Working Paper 21059, National Bureau of Economic Research 2018.
- 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 —, “Evaluating Long-Term-Care Policy Options, Taking the Family Seriously*,” *The Review of Economic Studies*, April 2018, *85* (2), 766–809.
- 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.
- Becker, Gary S and Nigel Tomes**, “An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility,” *Journal of Political Economy*, 1979, *87* (6), 1153–1189.
- and —, “Human Capital and the Rise and Fall of Families,” *Journal of Labor Economics*, 1986, *4* (3, Part 2), S1–S39.
- Belley, Philippe and Lance Lochner**, “The changing role of family income and ability in determining educational achievement,” *Journal of Human capital*, 2007, *1* (1), 37–89.
- Bernheim, B. Douglas, Andrei Shleifer, and Lawrence H. Summers**, “The Strategic Bequest Motive,” *Journal of Political Economy*, 1985, *93* (6), 1045–1076.
- Boar, Corina**, “Dynastic Precautionary Savings,” *The Review of Economic Studies*, 2021, *88* (6), 2735–2765.

- Brandsaas, Eirik E**, “Illiquid Homeownership and the Bank of Mom and Dad,” Finance and Economics Discussion Series, Board of Governors of the Federal Reserve System 2025.
- Brown, M., J. Karl Scholz, and A. Seshadri**, “A New Test of Borrowing Constraints for Education,” *The Review of Economic Studies*, April 2012, 79 (2), 511–538.
- Bruce, Neil and Michael Waldman**, “The Rotten-Kid Theorem Meets the Samaritan’s Dilemma,” *The Quarterly Journal of Economics*, 1990, 105 (1), 155–165.
- Cameron, Stephen V and James J Heckman**, “Life cycle schooling and dynamic selection bias: Models and evidence for five cohorts of American males,” *Journal of Political Economy*, 1998, 106 (2), 262–333.
- Caucutt, Elizabeth M and Lance Lochner**, “Early and late human capital investments, borrowing constraints, and the family,” *Journal of Political Economy*, 2020, 128 (3), 1065–1147.
- Charles, Kerwin Kofi and Erik Hurst**, “The Correlation of Wealth Across Generations,” *Journal of Political Economy*, 2003, 111 (6), 1155–1182.
- Cunha, Flavio, James Heckman, and Salvador Navarro**, “Separating uncertainty from heterogeneity in life cycle earnings,” *oxford Economic papers*, 2005, 57 (2), 191–261.
- Daruich, Diego and Julian Kozlowski**, “Explaining Intergenerational Mobility: The Role of Fertility and Family Transfers,” *Review of Economic Dynamics*, 2020, 36, 220–245.
- Duffie, Darrell and Kenneth J Singleton**, “Simulated Moments Estimation of Markov Models of Asset Prices,” *Econometrica*, 1993, 61 (4), 929–952.
- Fu, Chao, Ami Ko, and Chamna Yoon**, “Intergenerational Time Transfers: Altruism and Strategic Motives within the Family,” Working Paper, University of Wisconsin–Madison 2026.
- Hayashi, Fumio, Joseph Altonji, and Laurence Kotlikoff**, “Risk-Sharing between and within Families,” *Econometrica: Journal of the Econometric Society*, 1996, pp. 261–294.
- Heckman, James J and Stefano Mosso**, “The economics of human development and social mobility,” *Annu. Rev. Econ.*, 2014, 6 (1), 689–733.
- , **Lance J Lochner, and Petra E Todd**, “Earnings functions, rates of return and treatment effects: The Mincer equation and beyond,” *Handbook of the Economics of Education*, 2006, 1, 307–458.
- Hotz, V Joseph, Emily E Wiemers, Joshua Rasmussen, and Kate Maxwell Koegel**, “The role of parental wealth and income in financing children’s college attendance and its consequences,” Technical Report, National Bureau of Economic Research 2018.
- Hurd, Michael D.**, “Research on the Elderly: Economic Status, Retirement, and Consumption and Saving,” *Journal of Economic Literature*, 1990, 28 (2), 565–637.

- Kaplan, Greg**, “Moving back home: Insurance against labor market risk,” *Journal of Political Economy*, 2012, 120 (3), 446–512.
- Keane, Michael P and Kenneth I Wolpin**, “The effect of parental transfers and borrowing constraints on educational attainment,” *International Economic Review*, 2001, 42 (4), 1051–1103.
- Laitner, John**, “Bequests, gifts, and social security,” *The Review of Economic Studies*, 1988, 55 (2), 275–299.
- Lee, Bong-Soo and Beth Fisher Ingram**, “Simulation Estimation of Time-Series Models,” *Journal of Econometrics*, 1991, 47 (2–3), 197–205.
- Lee, Sang Yoon (Tim) and Ananth Seshadri**, “On the Intergenerational Transmission of Economic Status,” *Journal of Political Economy*, April 2019, 127 (2), 855–921.
- Lindbeck, Assar and Jörgen W Weibull**, “Altruism and time consistency: The economics of fait accompli,” *Journal of Political Economy*, 1988, 96 (6), 1165–1182.
- Lochner, Lance and Alexander Monge-Naranjo**, “The Nature of Credit Constraints and Human Capital,” *American Economic Review*, 2011, 101 (6), 2487–2529.
- Lockwood, Lee M.**, “Incidental Bequests and the Choice to Self-Insure Late-Life Risks,” *American Economic Review*, September 2018, 108 (9), 2513–2550.
- Love, David A., Michael G. Palumbo, and Paul A. Smith**, “The Trajectory of Wealth in Retirement,” *Journal of Public Economics*, 2009, 93 (1-2), 191–208.
- McFadden, Daniel**, “A Method of Simulated Moments for Estimation of Discrete Response Models Without Numerical Integration,” *Econometrica*, 1989, 57 (5), 995–1026.
- McGarry, Kathleen**, “Inter vivos transfers and intended bequests,” *Journal of Public Economics*, 1999, 73 (3), 321–351.
- , “Dynamic aspects of family transfers,” *Journal of Public Economics*, 2016, 137, 1–13.
- Nardi, Mariacristina De**, “Wealth Inequality and Intergenerational Links,” *The Review of Economic Studies*, 2004, 71 (3), 743–768.
- **and Fang Yang**, “Bequests and heterogeneity in retirement wealth,” *European Economic Review*, November 2014, 72, 182–196.
- , **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.
- Nishiyama, S**, “Bequests, Inter Vivos Transfers, and Wealth Distribution,” *Review of Economic Dynamics*, October 2002, 5 (4), 892–931.
- Pakes, Ariel and David Pollard**, “Simulation and the Asymptotics of Optimization Estimators,” *Econometrica*, 1989, 57 (5), 1027–1057.

Pfeffer, Fabian T. and Alexandra Killewald, “Intergenerational Wealth Mobility and Racial Inequality,” *Socius*, 2018, *4*, 1–2.

Restuccia, Diego and Carlos Urrutia, “Intergenerational Persistence of Earnings: The Role of Early and College Education,” *American Economic Review*, November 2004, *94* (5), 1354–1378.

A PSID Sample Construction

This appendix describes the construction of the analysis sample from the Panel Study of Income Dynamics (PSID). The sample combines four PSID data products: the Family Interview File, the Individual File, the Transition into Adulthood Supplement, and the Family Identification Mapping System. I use the thirteen biennial waves covering 1999 through 2023. The exact survey variable codes underlying each measure are documented in the replication package.

A.1 Data Sources

Family Interview File. The core source is the PSID family-level interview file, which provides household-level information on demographics, income, wealth, employment, and housing tenure for each survey wave from 1999 to 2023. A small number of supplementary measures that are not part of the main interview module—principally liquid savings and detailed industry classifications—are drawn from a companion family-level release and merged by household and wave.

Individual File. The PSID individual file provides person-level information for every sample member, including age, relationship to the household head, position within the family unit, and—beginning in 2013—years of completed education.

Transition into Adulthood Supplement. This supplement covers young adults aged roughly 18 to 28 in the 2015 and 2017 waves. I use it to confirm college enrollment directly, complementing the education-based measures available in the main files.

Family Identification Mapping System. This product provides the intergenerational links between children and their parents. Each child record carries identifiers for up to four parent types—biological father, biological mother, adoptive father, and adoptive mother—together with a generation marker that distinguishes original 1968 sample members from their children, grandchildren, and later descendants. The target sample consists of third-generation and later children (born roughly 1975–2005, reaching college age between 1993 and 2023) matched to their parents.

A.2 Variables Used in the Analysis

Table A.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 A.1. Measures Drawn from the PSID

| Measure | Waves | Source |
|--|-----------|---------------|
| <i>Panel A: Demographics</i> | | |
| 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 |
| <i>Panel B: Income</i> | | |
| 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 |
| <i>Panel C: Wealth</i> | | |
| 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 |
| <i>Panel D: Employment</i> | | |
| 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 |
| <i>Panel E: Consumption</i> | | |
| 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 |

continued on next page

Table A.1 (continued)

| Measure | Waves | Source |
|----------------------------------|-----------|-----------------|
| <i>Panel F: Individual-Level</i> | | |
| 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).

A.3 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.

A.4 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.

A.5 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.

A.6 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.

A.7 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.

A.8 Merge Attrition

Table [A.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 A.2. PSID Sample Construction and Merge Attrition

| Stage | Observations | Unique IDs | % Retained |
|--|--------------|------------|------------|
| <i>Panel A: Building the Analysis Sample</i> | | | |
| Individual file (all persons \times 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 \times year | 71 896 | 10 492 | — |
| After sample restrictions | 62 320 | 10 145 | 86.7 |
| <i>Panel B: Parent Identification Source</i> | | | |
| Linked via biological father | 297 284 | | |
| Linked via adoptive father | 4 114 | | |
| Linked via biological mother | 56 856 | | |
| No identifiable parent (dropped) | 0 | | |
| <i>Panel C: Merge Losses</i> | | | |
| 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 | 6 566 | | |

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.

B 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.

B.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.

B.2 Variables Used in the Analysis

Table [B.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 B.1. Measures Drawn from the NLSY97

| Measure | Years | Source |
|--|------------|------------------|
| <i>Panel A: Demographics and Ability</i> | | |
| 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 |
| <i>Panel B: Family Background</i> | | |
| 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 |
| <i>Panel C: College Enrollment and Costs</i> | | |
| 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 |
| <i>Panel D: Education Attainment</i> | | |
| Highest grade attended | 1998–2003 | Annual |
| Highest degree ever received | Cumulative | Cumulative |
| Withdrawal from college | Annual | Schooling |
| <i>Panel E: Income Process</i> | | |
| Annual labor income | 1997–2011 | Income |
| Annual hours worked (all jobs) | 1997–2007 | Employment |
| <i>Panel F: Parental Transfers</i> | | |
| 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.

B.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.

B.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 Section 3.7.

C 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.

C.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.

C.2 Variables Used in the Analysis

Table C.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 C.1. Measures Drawn from the HRS

| Measure | Coverage | Source |
|---|----------------|----------|
| <i>Panel A: Respondent Demographics</i> | | |
| 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 |
| <i>Panel B: Household Wealth</i> | | |
| 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 C.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 |
| <i>Panel C: Income</i> | | |
| 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 |
| <i>Panel D: Bequest Expectations</i> | | |
| 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 |
| <i>Panel E: Child Information</i> | | |
| 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 |
| <i>Panel F: Inter-Vivos Transfers</i> | | |
| 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 |
| <i>Panel G: College Support</i> | | |
| 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 |
| <i>Panel H: Consumption</i> | | |
| Total household consumption | 2001–2017 | CAMS |
| Non-durable spending | 2001–2017 | CAMS |
| Housing expenditure | 2001–2017 | CAMS |
| <i>Panel I: Bequests</i> | | |
| 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.

C.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.

C.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.

D Income Process: Discrete-to-Continuous-Time Mapping and Life-Cycle Profiles

This appendix derives the mapping from the discrete-time AR(1) income process estimated by [Abbott et al. \(2019\)](#) to the continuous-time Ornstein-Uhlenbeck (OU) specification used in the model, and presents the implied life-cycle income profiles.

D.1 Mapping Derivation

The discrete-time income shock process is an AR(1) at annual frequency:

$$z_{t+1}^e = \rho_e^{\text{ann}} z_t^e + \eta_{t+1}^e, \quad \eta_{t+1}^e \sim N(0, (\sigma_\eta^e)^2). \quad (26)$$

The continuous-time counterpart is the OU process:

$$dz^e = -\kappa_e z^e dt + \sigma_e dW, \quad (27)$$

where $\kappa_e > 0$ is the mean-reversion rate and σ_e is the diffusion coefficient.

Mean-reversion rate. The exact discretization of (27) at unit time intervals ($\Delta = 1$ year) gives the conditional mean $\mathbb{E}[z_{t+1}^e | z_t^e] = e^{-\kappa_e} z_t^e$. Matching this with the AR(1) coefficient ρ_e^{ann} yields:

$$e^{-\kappa_e} = \rho_e^{\text{ann}} \implies \kappa_e = -\ln(\rho_e^{\text{ann}}). \quad (28)$$

Diffusion coefficient. The conditional variance of the discretized OU process is:

$$\text{Var}(z_{t+1}^e | z_t^e) = \frac{\sigma_e^2}{2\kappa_e} (1 - e^{-2\kappa_e}). \quad (29)$$

The conditional variance of the AR(1) is simply $(\sigma_\eta^e)^2$. Equating the two and using $e^{-\kappa_e} = \rho_e^{\text{ann}}$:

$$\begin{aligned} (\sigma_\eta^e)^2 &= \frac{\sigma_e^2}{2\kappa_e} (1 - (\rho_e^{\text{ann}})^2) \\ \implies \sigma_e &= \sigma_\eta^e \cdot \sqrt{\frac{2\kappa_e}{1 - (\rho_e^{\text{ann}})^2}}. \end{aligned} \quad (30)$$

This mapping preserves both the autocorrelation and the stationary variance $\text{Var}(z^e) = \sigma_e^2/(2\kappa_e)$ of the process. Note that the stationary variance of the OU process equals the unconditional variance of the AR(1), $(\sigma_\eta^e)^2/(1 - (\rho_e^{\text{ann}})^2)$, confirming internal consistency.

Numerical values. Applying (28)–(30) to the male estimates from [Abbott et al. \(2019\)](#) (Table 3.2): for high-school graduates the annual persistence is $\rho_{HS}^{\text{ann}} = 0.952$ with innovation variance $(\sigma_\eta^{HS})^2 = 0.017$, giving $\kappa_{HS} = -\ln(0.952) = 0.049$ and $\sigma_{HS} = \sqrt{2\kappa_{HS}(\sigma_\eta^{HS})^2/(1 - (\rho_{HS}^{\text{ann}})^2)} = 0.134$; for college graduates $\rho_C^{\text{ann}} = 0.966$ with $(\sigma_\eta^C)^2 = 0.017$, giving $\kappa_C = 0.035$ and $\sigma_C = 0.133$. Both the mean-reversion rates and the diffusions are set externally to these AGMV-implied values rather than estimated. Because college earnings are more persistent ($\rho_C > \rho_{HS}$) while the innovation variance is common, the implied stationary variance of the persistent shock, $\sigma_e^2/(2\kappa_e) = (\sigma_\eta^e)^2/(1 - (\rho_e^{\text{ann}})^2)$, is *higher* for college graduates: 0.254 (standard deviation 0.51) versus 0.182 for high school (standard deviation 0.43). The model therefore does not generate the transfer-region contraction through a reduction in income risk; the operative channel is the income level (Section 3).

D.2 Life-Cycle Profiles

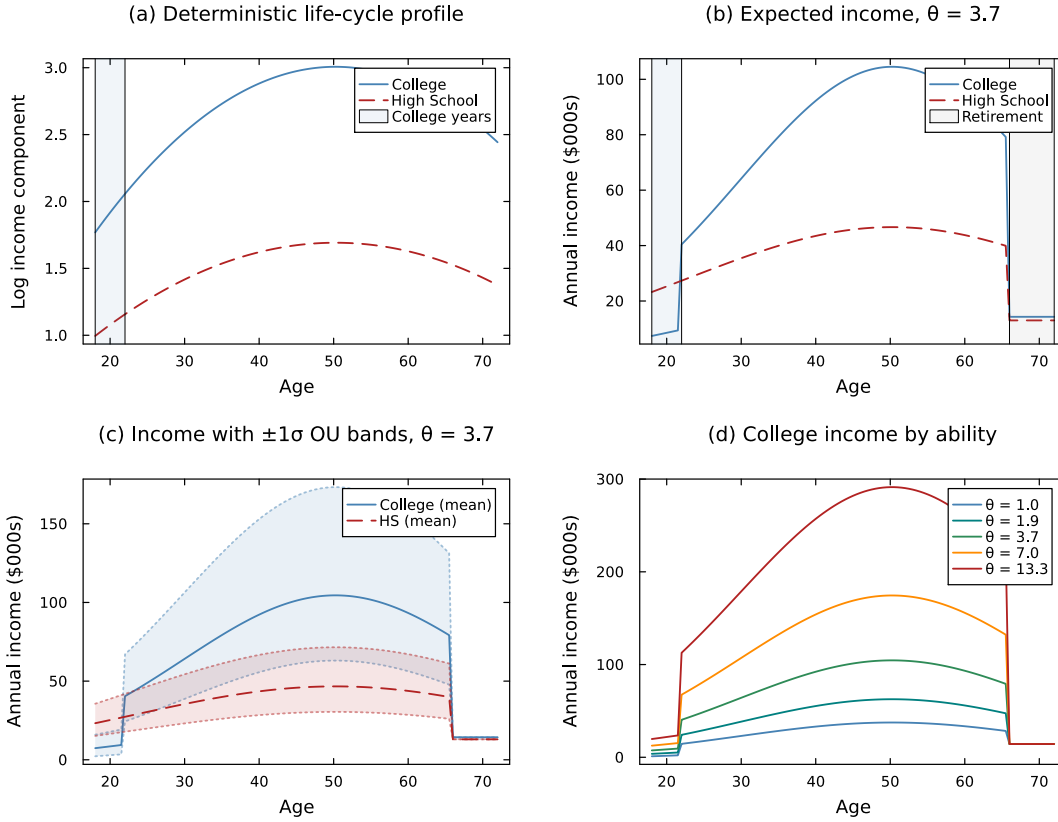


Figure D.1. Income Process: Life-Cycle Profiles by Education

Notes: Panel (a): deterministic age-income component $\gamma(t, e_c) = \gamma_{1,e} t - \gamma_{2,e} t^2$ from [Abbott et al. \(2019\)](#). Panel (b): expected annual income $y_c(t) = w \cdot \theta^{\lambda_e} \cdot \exp(\gamma(t, e_c))$ for median ability ($\theta = 3.9$); college workers earn $\ell_C = 56\%$ of full-time income during ages 18–22 and pay net tuition. Panel (c): ± 1 stationary standard deviation of the OU income shock ($\sigma_e / \sqrt{2\kappa_e}$) around expected income. Panel (d): college income profiles for low ($\theta = 1.5$), medium ($\theta = 3.9$), and high ($\theta = 10.2$) ability workers, illustrating the complementarity between ability and education ($\lambda_C > \lambda_{HS}$).

E Parent Wealth Portfolio

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

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

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

F Parent Consumption Detail

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

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

G Additional Robustness: Parent Fixed Effects

As additional robustness for the consumption stylized fact (Section 4.2), 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 G.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 Section 4.2 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 G.1. Parent Consumption and Child College Status: Parent Fixed Effects (PSID)

| | (1) | (2) |
|----------------------|---------------------|--------------------|
| | OLS + Income Q FE | Parent FE |
| College (BA+) | 0.068*** (0.018) | 0.027** (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. * $p < .10$, ** $p < .05$, *** $p < .01$.

H Full Commitment Benchmark

This appendix provides additional details on the full commitment benchmark introduced in Section 7.1 of the main text.

Under full commitment, a benevolent planner commits at $t = 0$ to a state-contingent transfer schedule $\{\tau(a_p, a_c, z, t)\}$ for the entire overlap (parent ages 42–72), solving a single-agent problem over total household wealth $A \equiv a_p + a_c$, whose law of motion carries the same multiplicative wealth diffusion $\sigma_a A dB$ as the decentralized economy (equation 16 in the main text). With CRRA preferences, the efficient allocation equates weighted marginal utilities:

$$u'(c_p) = \eta u'(c_c), \quad \text{so that} \quad c_c = \eta^{1/\gamma} c_p \quad (31)$$

Consumption shares are time-invariant. Substituting yields:

$$u_{\text{FC}}(c_{\text{tot}}) = \frac{c_{\text{tot}}^{1-\gamma}}{1-\gamma} \left(1 + \eta^{1/\gamma}\right)^\gamma \quad (32)$$

Relationship to the one-time transfer counterfactuals. The full commitment benchmark and the one-time transfer counterfactuals (Section 7.2) remove the Samaritan’s dilemma through different mechanisms. Full commitment eliminates moral hazard while preserving state-contingent insurance: the planner adjusts transfers in response to income shocks throughout the overlap period. The one-time transfer counterfactuals eliminate both moral hazard and ongoing insurance: the child receives a lump sum at $t = 0$ and is subsequently on their own. The difference between FC and the one-time transfer economies therefore isolates the *insurance-removal effect*—the change in college attendance and welfare attributable to shutting down ongoing parental support, holding the absence of moral hazard fixed.

Comparing FC with the baseline economy yields the pure moral hazard effect holding insurance constant. As discussed in Section 7.1, this comparison reveals two opposing forces: anticipated transfers subsidize enrollment (raising college attendance above the efficient level) but also cushion non-college children (lowering attendance below the efficient level). The net sign depends on the child’s position in the ability–wealth distribution.

I Equilibrium Properties

This appendix derives the key properties of the Markov-Perfect Equilibrium in the continuous-time coupled problem.

I.1 Optimality Conditions

Following [Barczyk and Kredler \(2014a\)](#), the equilibrium partitions the state space into three regions: a *self-sufficient* region, where the possibility of transfers is remote; an *overconsumption* region, where no transfer currently flows ($\tau^* = 0$) but the state is drifting toward the transfer region; and the *transfer* region, where $\tau^* > 0$. The child's first-order condition $u'(c_c^*) = V_{a_c}^c$ holds throughout. In the self-sufficient region, the child's continuous-time Euler equation takes the standard form $\dot{c}_c/c_c = (r - \rho)/\gamma$. In the overconsumption region, the child anticipates that accumulating wealth will reduce future parental support, so the effective return on saving is $r + \partial\tau^*/\partial a_c$ and the Euler equation becomes:

$$\frac{\dot{c}_c}{c_c} = \frac{1}{\gamma} \left(r - \rho + \frac{\partial\tau^*}{\partial a_c} \right), \quad \frac{\partial\tau^*}{\partial a_c} < 0. \quad (33)$$

This anticipation wedge is the Samaritan's dilemma in continuous time: it lowers the child's return on saving below r and induces over-consumption *before* any transfer flows. Inside the transfer region, by contrast, the parent equalizes the marginal values of own and child wealth ($V_{a_p}^p = V_{a_c}^p$) and acts as a family dictator, setting the child's consumption directly; the child's own Euler equation is therefore not the operative condition there ([Barczyk and Kredler, 2014a](#)).

The parent's first-order condition gives $u'(c_p^*) = V_{a_p}^p$. In the no-transfer region:

$$\frac{\dot{c}_p}{c_p} = \frac{1}{\gamma} (r - \rho) \quad (34)$$

In the transfer region, $V_{a_p}^p = V_{a_c}^p$: the parent equalizes the marginal values of own and child wealth and makes a positive gift $\tau > 0$. The parent funds the gift out of its own resources but continues to optimize its wealth, so its drift $\dot{a}_p = r a_p + y_p - c_p - \tau$ need not be zero—the parent makes a *bounded* transfer (it lifts the child only toward its own optimum) and keeps saving.

I.2 Transfer Region Characterization

The complementary-slackness condition (3) partitions the state space into:

1. **No-transfer region** $\mathcal{N} = \{(a_p, a_c, z, t) : V_{a_p}^p > V_{a_c}^p\}$: Parent and child accumulate wealth independently.
2. **Transfer region** $\mathcal{T} = \{(a_p, a_c, z, t) : V_{a_p}^p = V_{a_c}^p\}$: The parent makes a positive gift that brings the child toward the altruistic target $c_c = \eta^{1/\gamma} c_p$, bounded by the child's own optimum, $\tau^* = \max\{0, \min(c_c^{\text{opt}} - R_c, \eta^{1/\gamma} c_p - R_c)\}$ with $R_c = \tilde{r}_c a_c + y_c$. The parent funds the gift from its resources but continues to save ($\dot{a}_p = r a_p + y_p - c_p - \tau^*$).

The no-transfer region \mathcal{N} subdivides into a *self-sufficient* subregion, where the wealth distribution is balanced and future transfers are remote, and an *overconsumption* subregion bordering \mathcal{T} , where no transfer currently flows but the state is drifting toward \mathcal{T} and the anticipation wedge of (33) is active. The boundary of \mathcal{T} is a free boundary determined endogenously. Parents transfer when the child is relatively poor (low a_c), faces negative income shocks (low z), or when the parent is wealthy (high a_p). The region shrinks as the parent approaches death ($t \rightarrow T$).

Lump-sum exercise rule. The same partition governs the lump-sum channel of Section 3.3.4. At a Poisson arrival, the parent solves (5); the first-order condition for an interior lump is

$$V_{a_c}^p(a_p - \tau_L^*, a_c + \tau_L^*, z; t) = V_{a_p}^p(a_p - \tau_L^*, a_c + \tau_L^*, z; t),$$

so $\tau_L^* > 0$ if and only if $V_{a_c}^p > V_{a_p}^p$ at the pre-transfer state—the state lies strictly inside \mathcal{T} —and the optimal lump moves the dynasty onto the free boundary $\partial\mathcal{T}$ along the constant-total-wealth line $a_p + a_c$. In \mathcal{N} and on $\partial\mathcal{T}$ the opportunity expires unused ($\tau_L^* = 0$). With bounded flow transfers, income and wealth shocks can carry the state strictly into \mathcal{T} between arrivals, which is what gives the lump policy positive mass: the flow props up the child's consumption while the dynasty waits for an opportunity to rebalance the stocks. Anticipated lumps also feed the child's saving incentive: the jump term in (2) adds an expected-rebalancing component to the child's continuation value that is decreasing in own wealth, reinforcing the anticipation wedge of (33)—the Samaritan's dilemma extends to the lump margin.

I.3 The Samaritan’s Dilemma

Both agents over-consume relative to the commitment solution. The child’s distortion arises in the overconsumption region: anticipating that accumulating wealth would move the state out of \mathcal{T} and eliminate support, the child faces a reduced effective return on saving. This marginal disincentive to save is captured by $\partial\tau^*/\partial a_c$ in (33), and it operates *before* transfers flow, not inside \mathcal{T} itself, where the parent dictates the allocation.

The parent’s behavior is also shaped by the lack of commitment: unable to commit to withholding support, the parent transfers whenever the child is sufficiently constrained, and the child anticipates this. Unlike a reflecting-barrier formulation, the parent does not exhaust its wealth in \mathcal{T} ; it makes a bounded gift—lifting the child only toward its own optimum—and continues to save. The distortion therefore operates through the size and likelihood of transfers, and through the child’s reduced self-insurance, rather than through a degenerate parent saving rule.

The continuous-time formulation offers three advantages over discrete-time alternatives. First, no transfer carries within-period commitment: the flow is revisable at every instant, and the lump-sum channel operates only at exogenous arrivals, so a discrete transfer is an action taken when an opportunity strikes rather than a promise held for the length of a period. Second, the free boundary provides a sharp characterization of where moral hazard operates—for both the flow and the lump margins. Third, the multiplicative wealth diffusion $\sigma_a a dB$ (Section 3) renders the value functions C^2 almost everywhere and the free boundary smooth, so the Markov-perfect equilibrium exists in pure strategies and is computed without the non-smoothness—the “blast from the past”—that afflicts deterministic and discrete-time altruism models (Barczyk and Kredler, 2014a, 2021); the stochastic timing of lump opportunities preserves this smoothness, where deterministic transfer dates would reintroduce anticipation kinks.

I.4 College and Parental Influence

The college decision at $t = 0$ is forward-looking: the child compares lifetime values under each education path, incorporating the entire future trajectory of transfers. Altruistic parents affect the decision through two channels. First, they provide generous transfers during the

college years (when child income is low), effectively subsidizing attendance—the transfer region \mathcal{T} is larger when $e_c = C$ during the early years. Second, the continuation value under $e_c = C$ incorporates higher future income, which feeds back into the transfer policy: wealthier children require fewer transfers, freeing parent resources. The net effect generates a positive gradient of college enrollment in parent wealth.

J Solution Algorithm

The model is solved numerically 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.

J.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 $e \in \{\text{HS}, \text{C}\}$ 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 $g > 0$:

$$a_i = a_{\min} + a_{\max} \cdot \frac{(1+g)^{i-1} - 1}{(1+g)^{N-1} - 1}, \quad i = 1, \dots, N. \quad (35)$$

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 $\underline{a}_c = -25$, slightly below the cumulative student borrowing limit $\bar{b}_0 = 23$ (Section 3.7) 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:

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

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).

J.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:

$$\begin{aligned} \rho V^{c,\text{alone}}(a, z, t) = \max_{c \geq 0} \left\{ u(c) + V_a^{c,\text{alone}} [r a + y(t, z) - c] + V_z^{c,\text{alone}}(-\kappa z) \right. \\ \left. + \frac{1}{2} \sigma^2 V_{zz}^{c,\text{alone}} + \frac{1}{2} \sigma_a^2 a^2 V_{aa}^{c,\text{alone}} + V_t^{c,\text{alone}} \right\}. \end{aligned} \quad (36)$$

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 $s^{\text{fwd}} = r a + y - c^{\text{fwd}}$ and $s^{\text{bwd}} = r a + y - c^{\text{bwd}}$. The upwind scheme selects:

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

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:

$$\left[\left(\frac{1}{\Delta t} + \rho \right) \mathbf{I} - \mathbb{A}^n \right] \mathbf{V}^n = \mathbf{u}^n + \frac{1}{\Delta t} \mathbf{V}^{n+1}, \quad (37)$$

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:

$$\begin{aligned}
(\mathbb{A}_z)_{k,k+1} &= \frac{\sigma^2/2}{\Delta z^2} + \frac{\max(-\kappa z_k, 0)}{\Delta z}, \\
(\mathbb{A}_z)_{k,k-1} &= \frac{\sigma^2/2}{\Delta z^2} + \frac{\max(\kappa z_k, 0)}{\Delta z}, \\
(\mathbb{A}_z)_{k,k} &= -(\mathbb{A}_z)_{k,k+1} - (\mathbb{A}_z)_{k,k-1},
\end{aligned} \tag{38}$$

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

$$\begin{aligned}
(\mathbb{A}_a)_{i,i+1} &= D_i \frac{2}{\Delta_i(\Delta_i + \Delta_{i-1})}, \\
(\mathbb{A}_a)_{i,i-1} &= D_i \frac{2}{\Delta_{i-1}(\Delta_i + \Delta_{i-1})}, \\
(\mathbb{A}_a)_{i,i} &= -(\mathbb{A}_a)_{i,i+1} - (\mathbb{A}_a)_{i,i-1}.
\end{aligned} \tag{39}$$

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.

J.3 Stage 2: Coupled Parent-Child HJB

During the overlap period, the parent and child interact strategically. The parent's HJB equation is given by (1), with the transfer decision characterized by the complementarity condition (3). 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 (equation 11). 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:

$$c_p = (V_{a_p}^p)^{-1/\gamma}, \quad c_c^{\text{opt}} = (V_{a_c}^c)^{-1/\gamma},$$

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

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:

$$\tau = \max\left\{0, \min\left(c_c^{\text{opt}} - R_c, \eta^{1/\gamma} c_p - R_c\right)\right\}. \quad (40)$$

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:

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

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 Section 3.9.

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:

$$\mathbf{B} \mathbf{V}^{p,n} = \mathbf{u}^p + \frac{1}{\Delta t} \mathbf{V}^{p,n+1} + \mathbf{J}^{p,n}, \quad (41)$$

$$\mathbf{B} \mathbf{V}^{c,n} = \mathbf{u}^c + \frac{1}{\Delta t} \mathbf{V}^{c,n+1} + \mathbf{J}^{c,n}, \quad (42)$$

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 (Section 3.3.4) 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 *understates* off-node values while $\tau_L = 0$ is evaluated exactly, so discretization error biases against spurious lumps. The jump contributions to (41)–(42) are then

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

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. \(2022\)](#).

This procedure is repeated for each combination of child ability θ , child education $e_c \in \{\text{HS}, \text{C}\}$, parent education $e_p \in \{\text{HS}, \text{C}\}$, 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 $a_p \geq 0$.

J.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, their midpoints, and the cap itself, all bounded by $\min(a_p, \bar{\tau}_C)$ where $\bar{\tau}_C = 4 \phi_{\text{net}}(y_p^0, \theta)$ is the four-year net tuition bill (Section 3.6); 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 probit (13) and the parent’s objective (14)—including the paternalistic motive ξ —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 $\Phi((G/\bar{V} - \kappa_0 - \kappa_\theta \log \theta)/\sigma_\kappa)$ is available in closed form, so no numerical integration is required. In the simulation, each dynasty resolves 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.

J.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 $z_0 \sim \mathcal{N}(\zeta_{pe} \mathbf{1}\{e_p = C\}, \sigma_e^2/2\kappa_e + \sigma_{z_0, e_c}^2)$: 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} (Section 3).

College decision. For each dynasty m in cohort g , we interpolate the college probability $\Pr(C \mid a_p^{m,g}, \theta^{m,g})$ 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 $n = 1, \dots, N_{\text{steps}}$:

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, $a_p \text{ -- } \tau_L^*$, $a_c \text{ + } \tau_L^*$, *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; the model carries no transitory shock, $\sigma_{\varepsilon, e} = 0$, as in [Abbott et al., 2019](#)).

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

$$\begin{aligned} a_p^{n+1} &= a_p^n + (r a_p^n + y_p - c_p - \tau) \Delta t + \sigma_a a_p^n \sqrt{\Delta t} \xi_p^n, \\ a_c^{n+1} &= a_c^n + (\tilde{r}_c a_c^n + y_c - c_c + \tau) \Delta t + \sigma_a a_c^n \sqrt{\Delta t} \xi_c^n, \end{aligned} \quad (43)$$

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.

J.6 Stage 5: Moment Computation and SMM Estimation

The model is estimated by Simulated Method of Moments. It computes 42 moments—38 of which are targeted—organized into seven blocks. All moments are computed from a settled cohort (the second-to-last generation) to avoid transient dynamics from initial conditions.

Moment blocks.

1. **College attendance** (16 moments, weight 3): enrollment rates by parent-wealth quartile \times ability quartile.
2. **Income** (1 targeted, 2 untargeted): HS/college income ratio (targeted); cross-sectional income standard deviations for HS and college workers (computed but untargeted, since the income process is fixed externally).

3. **Transfers** (5 moments): the mean enrollment gift *conditional on attendance*, matching the earmarked NLSY97 aid question asked only of enrollees (in-college consumption flows are excluded to mirror the data concept); and the post-college transfer probability by parent-wealth quartile, computed over two-year windows from child age 26 in which cumulated flow transfers plus any lump-sum transfer reach the \$500 HRS reporting threshold.
4. **Wealth** (11 targeted, 1 untargeted): wealth-to-income ratio, intergenerational rank-rank slope, conditional child wealth $\mathbb{E}[a_c \mid \text{parent wealth } Q]$ (4 moments), median wealth for parents and children, variance of $\text{IHS}(a_c)$ for children, and two quantile-based dispersion measures of parent wealth (the p90/p50 ratio and the fraction with net worth below \$25,000); the variance of $\text{IHS}(a_p)$ for parents is computed but untargeted (unmatchable under the parent constraint $a_p \geq 0$).
5. **Decumulation** (1 moment): parent wealth decumulation rate from age 60 to 70.
6. **College by parent income** (4 moments): college graduation rate by parent-income quartile.
7. **Negative net worth** (1 moment): fraction of children aged 35–40 with negative net worth.

Loss function. The objective uses block-diagonal weighting so that each group contributes proportionally:

$$\mathcal{L}(\boldsymbol{\vartheta}) = \sum_g \frac{w_g}{n_g} \sum_{m \in g} \left(\frac{m_m(\boldsymbol{\vartheta}) - \hat{m}_m}{\hat{m}_m} \right)^2, \quad (44)$$

where n_g is the number of moments in block g and w_g is the group weight (college moments receive $w_g = 3$; all other groups receive $w_g = 1$).

Optimization. We minimize \mathcal{L} over the 14 free parameters using a parallel $(\mu/\mu_w, \lambda)$ -CMA-ES optimizer restricted to the free (non-fixed) subspace. 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.

J.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: $\Delta_{\text{MH}}^{\text{FC}} = \Pr(\text{C} \mid \text{baseline}) - \Pr(\text{C} \mid \text{FC})$. This full-commitment comparison is distinct from the main-text decomposition, which uses the one-time-transfer economies ($\Delta_{\text{MH}}^{\text{pure}}$ and Δ_{targ} in equations 21–23).

No-moral-hazard (unconditional). The parent makes a one-time lump-sum transfer τ_0 at $t = 0$, after which the child solves the child-alone problem for their entire remaining life. The child receives τ_0 regardless of education choice. The parent’s problem is:

$$\begin{aligned} \max_{\tau_0 \in [0, a_p]} V_p^{\text{alone}}(a_p - \tau_0) \\ + \eta \cdot \mathbb{E}_{\varepsilon_\kappa} \left[\max \left\{ V^{c, \text{alone}}(\tau_0 \mid \text{C}) - \kappa(\theta, \varepsilon_\kappa) + \varepsilon_C, V^{c, \text{alone}}(\tau_0 \mid \text{HS}) + \varepsilon_{\text{HS}} \right\} \right]. \end{aligned}$$

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 Section 8 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 ($\tau_C = \tau_{\text{HS}}$) 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:

$$\underbrace{\Pr(C \mid \text{BL}) - \Pr(C \mid \text{menu})}_{\Delta_{\text{MH}}^{\text{total}}} = \underbrace{\Pr(C \mid \text{BL}) - \Pr(C \mid \text{uncond.})}_{\Delta_{\text{MH}}^{\text{pure}}} + \underbrace{\Pr(C \mid \text{uncond.}) - \Pr(C \mid \text{menu})}_{\Delta_{\text{targ}}}. \quad (45)$$

The signs are consistent with the main text (Section 7): a positive $\Delta_{\text{MH}}^{\text{pure}}$ indicates that baseline moral hazard raises college attendance relative to the no-moral-hazard economy.