Dynamic treatment effects in Education

Anzheng Li*

October 2026

Abstract

Lifetime welfare calculations often rely on policy effects on educational attainment measured in a study's final observation window—for example, effects on college completion by age X. These estimates leave unresolved whether subsequent attainment widens or narrows the observed gap between the policy and comparison groups, and when that attainment occurs. This paper develops a framework to assess the welfare implications of this uncertainty. I document the timing information in existing studies and model possible continuations of observed education effects. Two deterministic simulations show how different education paths can match the same observed effect yet imply different lifetime earnings values, and how additional observations can narrow this range. I apply the framework to 38 policy calculations to reconstruct ranges for the marginal value of public funds (MVPF) and propose an ordering rule based on these ranges. Under the specified path and valuation assumptions, the same education evidence can support conflicting welfare judgments even when the estimated effects are treated as known. My contribution is a framework for assessing the robustness of welfare comparisons to unobserved education paths and quantifying how estimates from each additional observation window can narrow uncertainty.

1 Introduction

At what age should we measure a policy's effect on educational attainment to assess its lifetime welfare implications? For higher education, projecting lifetime welfare from an effect measured within a finite observation window—for example,

*University of Delaware. I thank the Open Science Pool (OSPool) and the Open Science Grid (OSG) Consortium for providing computing resources and support for this research.

on bachelor's degree (BA) completion by age X—can be misleading if subsequent changes in attainment are overlooked. In the United States, college attainment continues to rise within birth cohorts well into adulthood (Bárány et al. 2023). Moreover, a treatment effect on degree completion measured at a single age may reflect a permanent change in attainment, a temporary difference arising from accelerated or delayed graduation, or both (Scott-Clayton 2011; Angrist et al. 2022). Distinguishing these possibilities matters because welfare depends on both how much additional education a policy generates and when it occurs. Given an earnings mapping and discount rate, the education path shapes the net present value (NPV) of earnings changes and the associated fiscal returns, central components of MVPF calculations (Hendren and Sprung-Keyser 2020).

The empirical challenge is that educational attainment continues to evolve beyond commonly observed ages, while policy studies often report treatment effects for only one age or window. For the 1985–1987 birth cohorts, Figure 1 shows BA attainment, including higher degrees, rising from 29.34 percent at age 25 to 40.80 percent at 35, while approximate Years of Schooling rise from 13.30 to 13.97 years.[1] These cohort profiles document continued accumulation, whereas the implications for a policy effect depend on whether subsequent attainment differs between treated and untreated populations. Evidence on those relative trajectories is limited: in my review of 459 studies, 249 (54.2 percent) report only one estimation time or window, and 132 of the 190 studies reporting BA outcomes (69.5 percent) report only one window for that outcome. This concentration also extends to studies that use education estimates in valuation: of 31 studies projecting lifetime or career-long earnings, fiscal returns, or welfare, 18 report only one window per education series used in the calculation.[2] These estimates inform policy evaluation within their observation windows, while lifetime welfare analysis additionally requires assumptions about the evolution of the treatment–comparison attainment gap at unobserved ages.

This paper asks whether education effects measured within a finite window can rule out conflicting lifetime welfare judgments for the same policy. Using alternative lifetime education paths consistent with reported treatment effects, I

[1]The profiles use full-population American Community Survey (ACS) data for 2001–2022, with equal cohort weights at ages 16–35. Appendix A.2 documents construction and the cohort comparison in Figure A1, which shows larger age-25-to-35 gains in degree attainment and schooling for these cohorts than for the 1975–1977 cohorts.

[2]Appendix A.1 defines the review sample, timing measures, and valuation-sample classification. Figures A2–A4 describe the reported windows and endpoints; Table A1, Panels A and B, reports the full-sample and valuation-sample counts. A window may span several years; these counts describe reported education evidence, not every input used in valuation.

reconstruct MVPF ranges for 38 policy calculations: 37 from the Policy Impacts Library and one based on my earlier paper.[3] Most calculations retain the main valuation assumptions of Hendren and Sprung-Keyser (2020), including returns to education, discounting, and fiscal accounting. Among the 37 cases with an original welfare benchmark, 30 admit paths with positive net public costs and MVPFs on both sides of one while matching the same retained education inputs. The same education evidence can therefore support opposing welfare judgments even when the reported effects are treated as known. Additional comparable observations narrow some of these ranges under the maintained assumptions, showing how more information about education timing can strengthen welfare comparisons. A separate simulation quantifies how much observing education effects at two ages rather than one narrows lifetime earnings NPV ranges and, under a fixed illustrative welfare account, MVPF ranges.[4]

To reconstruct MVPFs from lifetime education paths, I start by defining cumulative educational attainment at each age for a fixed population. For each case whose MVPF I reconstruct, I combine the original study's reported treatment effects with Census–ACS cohort profiles, which provide an external calibration for baseline attainment and empirical restrictions on the paths. I use logistic functions to distinguish permanent changes in attainment from temporary effects due to changes in acquisition speed or timing. Constrained minimum-distance fitting then selects paths that best fit the reported effects subject to the maintained restrictions. Each path is translated into an earnings NPV and an MVPF using the valuation rules specified for that policy calculation. The resulting MVPF ranges show how welfare conclusions depend on the education path. I distinguish this path uncertainty, evaluated while holding reported effects fixed, from uncertainty propagated into welfare values through imprecision in the original estimates.[5]

[3]The Policy Impacts Library is a public database of MVPF estimates for government policies, with links to the underlying empirical studies and details of the welfare calculations. See `https://policyimpacts.org/policy-impacts-library/`. Li (2026) studies how recreational marijuana legalization (RML) affects BA completion among first-time, full-time college entrants, using a staggered difference-in-differences (DID) design. That study does not calculate an MVPF; the present calculation values only the education channel.

[4]Appendix B.4 defines the observation designs and percentage-reduction measure and reports the NPV results and fixed-account MVPF extension. Figure B4 compares all single-age and two-age designs, while Figure 3 and Tables 1–2 illustrate the changes in paths, NPV ranges, and percentage reductions.

[5]Appendix E.1 defines the conditional endpoint summaries, and Appendix E.2 describes the standard errors (SEs) and dependence assumptions used to vary the education inputs and reconstruct their compatible paths. Section 5.3 discusses the resulting welfare uncertainty; Table E1 reports willingness to pay (WTP) and net-cost summaries.

In 24 of the 37 calculations with an original welfare benchmark, the same education inputs admit opposing welfare outcomes. One compatible path implies beneficiary losses and positive net public costs; another implies beneficiary gains and a policy that pays for itself.[6] Among the 12 benchmark calculations with additional comparable education observations, the number with reconstructed MVPFs on both sides of one falls from ten to one under the maintained path restrictions and fitting rule. In the three illustrative scenarios with a fixed welfare account, observing any two distinct ages between 18 and 25 reduces MVPF range widths by 65–99 percent relative to either age alone when the single-age range is finite. Infinite single-age ranges become finite, although all 84 two-age ranges still cross one. Under the proposed ordering rule, the reconstructed ranges also produce a different policy ranking from the original MVPF estimates. In single-parameter comparisons, I find that changes in acquisition speed alone can match the education inputs used in most original welfare calculations. Under the maintained path restrictions, effects interpreted as permanent can therefore also reflect temporary differences in attainment. Allowing inputs to vary under the specified uncertainty exercises yields MVPF summaries spanning one in all 37 benchmark calculations before additional education observations are imposed.

My contribution is a reusable framework for assessing how unobserved education paths affect lifetime welfare conclusions. It connects research on time to degree and dynamic educational choices (Bound et al. 2012; Bárány et al. 2023; Heckman et al. 2016) to the education-based valuations of Hendren and Sprung-Keyser (2020), tracing how permanent and temporary education effects change projected earnings and fiscal returns. Projecting lifetime welfare from finite-window education effects also requires assumptions about unobserved outcomes, linking this paper to research on long-term extrapolation (Athey et al. 2026; Obradović 2026; Fan et al. 2026). My analysis traces the welfare consequences of alternative education paths when reported effects are held fixed and quantifies how much additional comparable observations change the resulting MVPF ranges. Separating this path uncertainty from uncertainty propagated through imprecisely estimated inputs complements research on inference for calibrated parameters and welfare metrics (Cocci and Plagborg-Møller 2025; Vohra 2026). The framework helps researchers assess which policy comparisons depend on unobserved education paths and whether further

[6]Eleven tax calculations require an explicit policy-unit conversion, Early Head Start and Tennessee Promise use disclosed public-data income reconstructions, and algorithmic placement adds a long-run earnings channel; Section 4.1 explains these adjustments. Several calculations share a source study.

progress requires more precise estimates, observations at other ages, or stronger evidence linking education to earnings.

These questions also link the framework to research on lifetime earnings measurement, long-term prediction, and policy choice under incomplete information. Haider and Solon (2006) show that the relationship between current and lifetime earnings varies with age, while Bhuller et al. (2017) document how additional schooling changes the age profile of earnings. These studies examine earnings over the life cycle; the reconstruction here also varies the education path entering the valuation. To forecast beyond experimental follow-up, García et al. (2020) combine experimental and nonexperimental data with economic models to estimate lifetime program benefits and costs. Chen and Ritzwoller (2023) develop efficient estimators combining short-term experimental and longer-term observational outcomes, and Imbens et al. (2025) use the sequential structure of short-term outcomes to address persistent confounding in such data combinations. These approaches clarify how additional data and identifying restrictions support long-term estimates. The present analysis examines which welfare judgments remain unresolved when selected education observations leave room for different attainment continuations. It also relates to sufficient-statistics welfare analysis, which connects reduced-form estimates to welfare through economic restrictions (Chetty 2009; Kleven 2021). Projecting lifetime earnings and fiscal returns from an education effect similarly requires a restriction on its continuation. Finally, Manski (2011) studies policy choice when treatment responses are only partly known and no probability distribution is assigned to that uncertainty. This distinction between what the evidence determines and what a decision rule adds also guides the interpretation of the policy ordering in this paper.

The remaining sections move from the available education evidence to path construction, welfare applications, and their interpretation. Section 2 describes the education calculations in the Policy Impacts Library and the scope of the reconstruction. Section 3 develops the model and simulations. Sections 4 and 5 present the three main applications and the full policy collection. Section 6 discusses assumptions, robustness, and policy implications, and Section 7 concludes.

2 Education Evidence in the Policy Impacts Library

The Policy Impacts Library contains MVPF calculations for more than 260 policy entries, including 85 classified as Education.[7] For 36 of these education calculations, valuation relies primarily on earnings or other non-education outcomes, such as employment or crime. When earnings effects are observed, the standard projection in Hendren and Sprung-Keyser (2020) expresses the estimated earnings difference as a percentage of comparison-group earnings and scales an ACS age–earnings profile to that group's observed income. Beyond the observed period, both the percentage treatment effect and comparison-group earnings relative to the ACS profile are held constant through age 65.[8] The observed and projected earnings effects are discounted and allocated to WTP and net public costs under the applicable tax and transfer rules. This approach requires assumptions about the persistence of earnings effects but does not separately reconstruct the education paths underlying them.

The remaining 49 calculations use education measures, including enrollment, credits, degree attainment, and test scores, as valuation inputs. In the college calculations of Hendren and Sprung-Keyser (2020), source-specific conversions express education effects as increments in Years of Schooling, which are mapped into predicted earnings using an externally estimated return to education. Valuing these effects over a lifetime requires assumptions about the treatment–comparison attainment gap at unobserved ages, even when a complete education path is not explicitly specified. An observed gap may reflect a permanent difference in eventual attainment, a temporary difference in acquisition timing, or both, with different implications for the duration of earnings gains or losses and the associated fiscal returns. Holding fixed the education estimates selected for each reconstruction, I construct alternative attainment paths under stated restrictions and evaluate their welfare implications.

[7]These counts use the Library export of 20 September 2026, which contains 269 entries. They count policy calculations, including separate contrasts within a program; the review in Appendix A.1 counts distinct studies. The two valuation groups below classify the archived entries for this analysis; they are not the Library's official categories. Source: `https://policyimpacts.org/policy-impacts-library/`.

[8]Their standard forecast uses the 2015 ACS and assumes 0.5 percent annual real wage growth; their baseline discount rate is 3 percent. See Section III.A, Appendix Figure I, and Online Appendix I of Hendren and Sprung-Keyser (2020). Projection horizons remain account-specific: their Job Corps baseline values only the 20-year observed period (Section III.E).

I select 37 calculations from the education-input group with U.S. postsecondary outcomes that can be linked to Census–ACS attainment paths. They cover student aid, education tax provisions, tuition, and college funding, as well as Early Head Start, whose retained input is adult college completion. Selection therefore depends on the outcome used in valuation, not the age at which the policy operates. Of the 12 remaining education-input calculations, seven concern non-U.S. populations outside the calibration's scope. The other five use U.S. childhood test scores, whose mapping to later outcomes differs from the attainment model developed here.

Thirty-three of the selected calculations use the schooling-to-earnings mapping of Hendren and Sprung-Keyser (2020). The other four are the Susan Thompson Buffett Foundation (STBF) scholarship, Early Head Start, Tennessee Promise, and algorithmic college placement. For STBF, I retain the authors' nonlinear credential–experience income function. The complete author income inputs are unavailable for Early Head Start and Tennessee Promise, so I reconstruct these inputs from public data and retain the published MVPFs as separate benchmarks. For algorithmic college placement, I retain the source WTP account and add an earnings-related fiscal channel.[9] Within each reconstructed account, I apply the same earnings mapping, discounting, and fiscal rules to every candidate path. The resulting MVPF range measures the sensitivity of welfare valuation to education paths under these assumptions. Where income inputs are reconstructed or a fiscal channel is added, differences from the original MVPF also reflect these adjustments.

3 Education Paths and Lifetime Valuation

A finite observation window leaves the evolution of educational attainment partly unobserved, even when the effects within that window are measured without error. To separate this information problem from sampling uncertainty, I hold a comparison path fixed and construct treated paths consistent with stipulated education effects. In each of the three displayed scenarios, the same age-25 effect admits paths with positive and negative education-channel NPVs under a common valuation rule. Adding the corresponding age-35 effect reduces the width of the numerical candidate range by 92.3–94.2 percent. These comparisons treat the observed effects as exact and hold the baseline and path restrictions fixed. The first experiment varies the education process while observing it at age 25; the

[9]Appendix C gives the common calibration and account-replacement rules; Appendix D documents the source accounts, income mappings, and policy-unit adjustments. Table D1 lists the selected calculations and their calibration populations.

second holds each generating process fixed and varies the observation ages. These experiments distinguish valuation conclusions supported by the observed effects from those that depend on the unobserved schooling path.

3.1 Education Paths, the Baseline, and Path Restrictions

An education effect measured at a particular age describes the attainment difference at that age, while leaving its earlier and later evolution unresolved. For a fixed population in cohort c, let $H_c^0(a)$ and $H_c^1(a)$ denote the mean of a specified educational attainment outcome at age a without and with the policy, respectively. The outcome may record some college, degree attainment, or completed Years of Schooling.[10] The numerical illustration below uses Years of Schooling. For any of these cumulative outcomes, I distinguish the policy effect on the attainment stock from its age-a increment:

$$\Delta H_c(a) \equiv H_c^1(a) - H_c^0(a), \qquad \Delta h_c(a) \equiv \Delta H_c(a) - \Delta H_c(a-1). \tag{1}$$

Cumulative attainment satisfies $H_c^j(a) - H_c^j(a-1) \geq 0$ for $j \in \{0, 1\}$ in a fixed population. This restriction does not imply monotonicity of the treatment effect $\Delta H_c(a)$, because the two potential attainment paths can accumulate at different rates. The increment $\Delta h_c(a)$ may therefore be positive, negative, or zero, according to whether attainment accumulates faster, more slowly, or at the same rate under treatment. In particular, $\Delta h_c(a) < 0$ indicates a decline in the treatment–comparison attainment gap; neither level path need decline. Appendix B.2 documents the observation restrictions and their implementation, and Figure 2 illustrates the two levels, their age-specific gaps, and the change in the gap between adjacent ages.

I represent the treated path with a three-parameter logistic continuation anchored to the comparison level. For a given starting age a_0, write $\tau_c = a_0 - 1$ and $S_c = H_c^0(\tau_c)$. With $\theta_c = (\Delta K_c, \Delta b_c, \Delta d_c)$, define

$$L_c(a; \theta_c) = \Lambda[(b_c + \Delta b_c)(a - d_c - \Delta d_c)],$$

[10] For binary attainment, $H_c^j(a)$ is a population share and $100\Delta H_c(a)$ is a percentage-point effect. Schooling levels and effects are measured in years. Current enrollment is not cumulative and need not rise with age.

where $\Lambda(x) = (1 + e^{-x})^{-1}$. The treated path is

$$H_c^1(a; \theta_c, a_0) = \begin{cases} H_c^0(a), & a \leq \tau_c, \\ S_c + (K_c + \Delta K_c - S_c) \dfrac{L_c(a; \theta_c) - L_c(\tau_c; \theta_c)}{1 - L_c(\tau_c; \theta_c)}, & a > \tau_c \end{cases}. \quad (2)$$

The normalization makes the continuation pass through S_c at τ_c and approach $K_c + \Delta K_c$ at older ages. The first affected observation therefore reflects acquisition since the common starting point; changing parameters does not introduce a level jump at the anchor. The parameter K is measured in attainment units, b in inverse age-years, and d in age-years. They govern the plateau and the speed and timing of accumulation conditional on the anchor. I call an education effect permanent when it changes eventual attainment, so that $\lim_{a\to\infty} \Delta H_c(a) = \Delta K_c \neq 0$. An effect is temporary when the attainment gap eventually vanishes. In the single-parameter exercises, a nonzero K change is permanent, whereas b or d changes with K fixed are temporary.[11] Holding K, b, and the common anchor fixed, a higher d delays subsequent accumulation: attainment is lower at each finite age after the anchor, and any given level between the starting value and plateau is reached later. The delay varies across attainment levels, so an increase in d does not imply an equally sized delay at every point on the path. These are population-path parameters rather than individual graduation dates. Panels A–C of Figure B1 illustrate their roles under the normalization; Appendix B.2 gives the parameter restrictions and timing interpretation.

The choice of a_0 reflects when the policy can first affect the specified attainment outcome for the fixed target population. I assume no policy effect or anticipation through τ_c, which justifies the common starting level in Equation (2). The two paths describe the same population under alternative policy states; the observed treatment and comparison samples need not have identical pre-policy means. In the simulation, the hypothetical policy first affects the population at age 18, making age 17 the anchor. The anchor is therefore determined by exposure timing, not by the age at which the researcher happens to observe an effect.[12]

For the Years of Schooling simulation, $H_c^0(a)$ is the equally weighted average of fitted ACS profiles for the 1985–1987 birth cohorts, using the full-population data

[11] Temporary describes eventual convergence, not disappearance within the valuation horizon. Such an effect may persist beyond age 65 and can have a nonzero lifetime NPV. A permanent effect can also involve changes in acquisition speed or timing.

[12] Observation at 25 does not justify anchoring at 24. The anchor must precede possible effects and anticipation. If exposure ages differ, use a common unaffected age or cohort-specific anchors before aggregating with fixed weights.

at ages 16–35 introduced in Figure 1. Here c indexes the synthetic three-cohort group. Each cohort is fitted separately, and the fixed one-third weights preserve its contribution at every age; values above age 35 are extrapolated.[13] I vary H_c^1 while holding this comparison curve fixed. Appendix A.2 documents the sample and education measures; Appendix B.1, Table B1, and Figure B2 document the construction and fit.

The admissible path class combines the logistic specification with restrictions on attainment levels and annual increments. Both experiments set the first possible effect age to $a_0 = 18$, value earnings through age 65, and extend education paths through age 72. I require completed schooling under treatment to be nondecreasing, to lie between zero and 25 years, and to increase by no more than one schooling year between adjacent ages, including the transition from age 17 to age 18. The total-parameter domain is $K \in [13.34, 15.14]$, $b \in [0.12, 0.65]$, and $d \in [8, 30]$.[14] Let $\mathcal{C}$ collect these maintained conditions. They define which paths enter the comparison; satisfying them does not establish that a particular path is the response to an actual policy.

3.2 What a Single Observed Effect Leaves Unresolved

An education estimate restricts a feature of $\Delta H_c(\cdot)$ but leaves the rest of the path partly unresolved. Let the observation operator $\mathcal{M}$ map a candidate path into the effects represented by a study:

$$\mathbf{m} = \mathcal{M}\{\Delta H_c(\cdot)\}\,. \tag{3}$$

For attainment measured at one age, the relevant entry of $\mathbf{m}$ is $\Delta H_c(a)$. An average effect over an age window $\mathcal{A}$ instead enters as $\sum_{a\in\mathcal{A}} w_a \Delta H_c(a)$, using the source's weights.[15] An effect on attainment acquired between ages a_1 and a_2 restricts $\Delta H_c(a_2) - \Delta H_c(a_1)$. These restrictions preserve the source study's

[13] The reference triple averages parameters, whereas H_c^0 averages fitted curves. Their anchored continuations differ by at most 2.26×10^{-4} Years of Schooling over ages 18–65; zero parameter increments need not imply an exactly zero effect.

[14] The plateau limits are weighted first and 99th percentiles of a historical state–cohort fit bank. Transporting this support to the full-population baseline is a calibration assumption (Appendix B.1). Speed, timing, and annual-increment limits are design restrictions. Binary outcomes require share-consistent bounds instead.

[15] An exact 0.1-year effect at 25 imposes $\Delta H_c(25) = 0.1$. An equally weighted two-percentage-point BA effect at ages 23–24 imposes $[\Delta H_c(23) + \Delta H_c(24)]/2 = 0.02$, not a restriction at the midpoint.

outcome definition and time scale.[16] The source-specific mappings for credits, enrollment, and credentials enter the applications in Section 4. For the fixed comparison path and target population, the compatible path set is

$$\mathcal{P}(\mathbf{m}; a_0, \mathcal{C}) = \left\{ H_c^1(\cdot; \theta_c, a_0) \in \mathcal{C} : \mathcal{M}\left\{ H_c^1(\cdot; \theta_c, a_0) - H_c^0(\cdot) \right\} = \mathbf{m} \right\}. \tag{4}$$

An additional observed effect changes the moment restrictions; changing the parameter domain or the path conditions changes $\mathcal{C}$. With one or two exact age-specific effects, the matching equations generally leave two or one free parameter dimensions. Appendix B.2 gives the algebra and numerical search used to solve these equations directly. I impose no probability distribution over compatible paths or their parameters. The resulting NPV and MVPF candidate ranges describe conditional possibilities, without assigning probabilities to a positive NPV or an MVPF above one.[17]

The first experiment fixes the observation age at 25 and varies the education process generating that observation. Holding two generating parameters at their reference values, I vary the third over its entire conditional feasible interval. Each scenario supplies its own age-25 effect. I then set aside knowledge of the generating parameters and search the common three-parameter class for paths matching that effect. Only one generating parameter varies at a time, but all three candidate parameters may vary: the observer is not assumed to know which parameter the policy changed. Appendix B.4 specifies the finite grids used to span the intervals.

An education effect observed at age 25 leaves open whether the effect is permanent or temporary. Panel A of Figure 3 compares a permanent increase in completed schooling, a temporary gain from faster accumulation, and a temporary loss from delayed accumulation followed by catch-up, from left to right.[18] The orange curves describe these stipulated responses. In the middle column, the early advantage largely disappears as the comparison population catches up; in the right column, the treated population initially falls behind and subsequently catches up. The shaded regions summarize the alternative paths found when only the age-25

[16] Six years after entering college need not mean age 24. Converting follow-up into age requires entry-age information or source-specific averaging. Policy event time also differs from individual follow-up.

[17] Grid spacing does not assign path probabilities. Here the stipulated effects have no sampling error, so candidate ranges are not confidence intervals. Finite search also does not certify global bounds; Appendix B.2 distinguishes these issues.

[18] Only the named parameter moves to its conditional-interval midpoint. Relative to $(13.92, 0.26, 12.56)$, the increments are $(0.32, 0, 0)$, $(0, 3.00 \times 10^{-3}, 0)$, and $(0, 0, 6.44)$. The rule is independent of valuations and does not imply equal-sized policy shocks. Table B2 summarizes the feasible intervals and scenario positions; the replication files retain all scenario parameters.

effect is observed. Within each column, every candidate reproduces that effect and satisfies the same maintained restrictions, yet the subsequent education outcomes can differ. For both positive early effects, the candidates include persistent gains and later reversals of the attainment advantage. For the negative early effect, they include persistent shortfalls, catch-up, and a later attainment advantage. These are possible evolutions of population means; they do not identify particular individuals' graduation or return-to-school decisions. Observing the gap at one age therefore leaves uncertainty about both eventual schooling and how long an advantage or disadvantage lasts. Under the valuation rule introduced next, these differences change the duration of earnings gains or losses and the timing of schooling costs.

3.3 From Education Paths to NPV

Education paths that reproduce the same observed effects can imply different economic values under a common earnings and cost schedule. For college policies without observed earnings effects, Hendren and Sprung-Keyser (2020) project income changes from education outcomes using returns estimated by Zimmerman (2014). The superscripts point and path distinguish the benchmark and its path-based extension. I write both values as functions of the schooling-effect path ΔH_c, holding the observation age a^* fixed. The benchmark uses only the observed gap $\Delta H_c(a^*)$ in completed Years of Schooling per member of the target population.[19] Let $Y_0(a)$ be annual comparison earnings in a common dollar year and r_t^{HSK} the proportional earnings effect per additional schooling year in period t after policy onset. The dollar premium is $\pi^{\text{HSK}}_{\text{Years of Schooling}}(a) = Y_0(a) r^{\text{HSK}}_{a-a_0+1}$. With first possible effect age a_0, terminal valuation age A, and discount rate ρ, write the common discount factor as $D_\rho(a; a_0) = (1+\rho)^{-(a-a_0)}$. The single-effect benchmark is

$$\text{NPV}_c^{\text{point}}(\Delta H_c; a_0, A, \rho) = \sum_{a=a_0}^{A} D_\rho(a; a_0) \pi^{\text{HSK}}_{\text{Years of Schooling}}(a) \Delta H_c(a^*). \tag{5}$$

The observation age a^* stays fixed as the valuation age a runs from a_0 to A. The benchmark therefore assigns the same NPV to all paths with the same observed gap. The retained schedule assigns negative returns in the first seven periods and

[19] Appendix D.2 documents source-specific education conversions; Appendix B.2 distinguishes age-specific and window observations.

positive returns thereafter.[20] Thus, a single observed education difference generates an income profile that varies with age through earnings, returns, and discounting. Its value depends on using that same education input in periods for which the education path is unobserved.

I extend this benchmark by assigning its signed earnings response to the age at which each part of the attainment gap is acquired. Define the signed completion increment

$$q_c(\tau) = \Delta H_c(\tau) - \Delta H_c(\tau - 1). \tag{6}$$

These increments need not be positive and are not probabilities of individual graduation. They record the difference in attainment acquired at each age between the two population paths. The net annual earnings effect and its present value (PV) are

$$\Delta y_c(a) = Y_0(a) \sum_{\tau} R(a, \tau) q_c(\tau), \tag{7}$$

$$\mathrm{NPV}_c^{\mathrm{path}} = \sum_{a=a_0}^{A} D_\rho(a; a_0) \Delta y_c(a). \tag{8}$$

For the common simulation, $R(a, \tau) = r^- \mathbf{1}\{\tau - 7 \leq a < \tau\} + r^+ \mathbf{1}\{a \geq \tau\}$, evaluated only at income ages $a \geq a_0$. Here $r^- = -6.67\%$ and $r^+ = 815/7241$. A single completion increment of size m at 25 reproduces the benchmark's negative response at ages 18–24 and positive response from 25 onward. Two increments sum two shifted responses; a smooth path sums increments across all completion ages. This extension assumes that the earnings-response template applies at different completion ages. It does not assert that Hendren and Sprung-Keyser (2020) observed everyone graduating at 25 or that each additional schooling year entails seven actual years of enrollment.

Early earnings losses and later gains are already in $\Delta y_c(a)$. No separate opportunity-cost charge is subtracted, and annual changes in completed schooling are not set equal to contemporaneous enrollment. Linearity permits signed increments, including comparison-group catch-up. It also assumes that the same

[20] The schedule uses approximately -6.67% in periods 1–7 and 11.26% thereafter, following the conversion of Zimmerman's estimates in the replication files of Hendren and Sprung-Keyser (2020). It is an illustrative benchmark; their policy accounts also use other education inputs and observed earnings.

marginal response can be transported across the completion ages and sizes considered, without additional interactions or general-equilibrium effects. These assumptions concern valuation rather than the identification of the source education estimates.

The first experiment uses $a^* = 25$, income ages 18–65, a 3 percent discount rate to age 18, and 2000 dollars. The common baseline remains the fitted 1985–1987 cohorts. Education paths extend through 72 because income at 65 can include a pre-completion response associated with acquisition up to seven years later. The terminal attainment statistic remains the gap at 65. Path restrictions, earnings profiles, response coefficients, and discounting are held fixed across candidate paths. Appendix B.3 gives the equivalent stock formula and the continuous-age interpretation. Policy accounts then retain their own income, tax, transfer, and education-cost conventions.

Applying the common path valuation to the candidates in the first experiment produces both positive and negative values in all three displayed scenarios. Table 1 separates each scenario's single-effect benchmark, its known generating-path value, and the values of other paths matching the observed effect. The age-25 candidate NPV ranges, in thousands of 2000 dollars, are $[-3.53,\ 44.80]$ for the plateau change, $[-15.43,\ 42.58]$ for the speed change, and $[-26.41,\ 39.18]$ for the timing change. Finding paths with opposite signs establishes that the stipulated education evidence and maintained restrictions do not determine the NPV sign. The generating-path value exceeds the single-effect benchmark in the plateau and timing examples and is below it in the speed example, so the direction of that comparison is not uniform. Table B2 summarizes the full scans; the complete scenario results remain in the replication files.

Figure 4 compares the single-effect benchmark with full-path valuation for a permanent increase in schooling attainment. Only the plateau changes; the speed and timing parameters remain at their reference values.[21] The gap grows from 0.48 years at 25 to 0.58 at 65; Panel D of Figure B1 shows the corresponding levels. Blue bars show discounted education earnings returns, and red bars show the negative pre-completion contributions labeled opportunity cost. Summing both series gives NPVs of \$20,354.29 for the single-effect benchmark and \$25,587.81 for the full path. Both calculations use the same earnings-response template, and the displayed opportunity costs are already included in these NPVs.

[21] The increments are $(\Delta K_c, \Delta b_c, \Delta d_c) = (0.57, 0, 0)$, with ΔK_c chosen to match $\Delta H_c(25) = 0.48$ years. This illustration satisfies the common path restrictions but is outside the two experiment grids.

The distinction between permanent and temporary effects can also be examined with the education inputs used in the policy applications. Figure B3 in the appendix compares each reference earnings NPV with the value obtained by matching its education input through a plateau, speed, or timing change alone. Table B3 reports the corresponding normalized NPVs, including calculations without a comparable pair; Table 5 reports the parameter changes. The other two parameter families remain at their calibrated baseline values. The K specification produces a permanent effect; the b and d specifications produce temporary effects as the paths eventually converge. Both axes use each account's monetary units and the same reference-expenditure normalization. These comparisons retain the applications' own valuation schedules, rather than imposing the simulation's common schedule on every policy. Appendix B.5 explains the construction, unit bridges, and differing matched samples. Differences from the reference projection can reflect both continuation and the source-to-path earnings mapping; they are not estimates of which mechanism actually occurred.

3.4 What Additional Observation Ages Reveal

Additional observations can distinguish paths that fit the same early effect, but the information they provide depends on when schooling is measured. The second experiment holds each generating path fixed and changes the observation ages, retaining the same baseline, path restrictions, and valuation rule. I use five positions spanning each parameter's conditional feasible interval and examine every single age and every pair of distinct ages from 18 through 35. The generating path remains fixed as the observation ages change; its effects at different ages need not be equal. Each added observation restricts the same underlying path set. Appendix B.4 gives the full design; Appendix B.2 describes the numerical procedure for preserving this nesting.

Adding an age-35 effect sharply narrows the NPV ranges for the same three midpoint scenarios. Panel B of Figure 3 retains the generating parameters, baseline, restrictions, valuation rule, and axis scales of Panel A. The candidate NPV ranges in Table 1 become [10.61, 14.33], [−2.92, 0.46], and [−8.74, −3.97] thousand dollars, reductions in width of 92.3, 94.2, and 92.7 percent. The second observation matters because a single schooling gap does not determine how quickly it emerged or what follows it. Timing also matters for value: a year acquired earlier can produce earnings gains over more years but also moves forgone earnings closer to the discount base. In the timing example, the age-35 gap is only −0.04 Years

of Schooling, yet the candidates found have negative NPVs: near-convergence in later schooling need not erase the value lost during earlier delays. The candidates found for the plateau example are positive, whereas the speed example still admits both signs. Its generating speed lies close to the reference value under the one-year acquisition cap, so the stipulated effect is small. These findings do not establish sign identification: the search is finite, and one of the five timing scenarios also admits both signs after adding age 35.

The location of the additional observation matters alongside its availability. Table 2 compares selected designs within each scenario before summarizing across scenarios.[22] Across the three parameter groups, adding age 20 to age 25 reduces the NPV width by a median 33.5–68.1 percent, whereas adding age 24 reduces it by 38.2–74.9 percent. Adding age 30 yields median reductions of 76.0–88.3 percent. The age-24 pair is more informative than the age-20 pair in these summaries despite the shorter distance between observations. It restricts the curve near the first observation, while later ages constrain how much of the schooling gap persists. These comparisons do not imply a general ranking of nearby and distant observations: the ranking depends on the logistic family, anchor, and support. Figure B4 summarizes all observation combinations, showing both interior-scenario medians and maximum widths across all five scenarios; the individual-scenario results remain in the replication files.

The information comparison also extends to MVPF. Appendix B.4 applies a fixed illustrative grant and tax account to every single-age and two-age design. In the three midpoint scenarios, adding the age-24 effect to the age-25 effect reduces the candidate MVPF span by 81.0–96.0 percent, although all three ranges still cross one. Thus, additional information within the same follow-up horizon can narrow lifetime welfare ranges without resolving every welfare judgment. The numerical contraction depends on the fiscal account as well as the paths: ratios are especially sensitive when net public cost approaches zero.

Finite-window evidence can therefore leave economically consequential differences among admissible education paths, while observations at other ages can narrow the differences among the paths found. The exercise conditions on the fitted baseline, logistic family, parameter domain, and valuation calibration. It does not evaluate source-estimate sampling error or sensitivity to alternative schooling functions and earnings-response mappings. The policy applications in Section 4

[22] Reductions are calculated within scenarios relative to age 25 alone, then summarized by the median at the 25, 50, and 75 percent interval positions. Largest widths and sign counts include both boundaries.

add source-specific outcome definitions, timing bridges, observed mechanisms, and full welfare accounting.

4 From Education Paths to Policy Welfare

Calculating the earnings NPV of an education path is only one step toward evaluating a policy: the welfare comparison also depends on transfers, private education payments, and the net cost to the government. Building on the MVPF framework of Hendren and Sprung-Keyser (2020), I incorporate alternative education paths into beneficiaries' WTP and the government's fiscal account. The broader exercise reconstructs education-related welfare calculations for 38 policy cases, whose results appear in Table 3 and Section 5. This section examines the Adams Scholarship, Social Security student benefits, and Tennessee Promise.[23] Under the reference specification, the MVPF is finite and less than one for the Adams Scholarship, finite and greater than one for the Social Security enrollment alternative, and infinite for Tennessee Promise, which pays for itself. For each case, I relate the candidate paths to the study's dynamic evidence and mechanisms, distinguishing the observed education effects from the assumptions needed to project lifetime benefits and fiscal costs.

4.1 Welfare Accounting and the Education-Path Replacement

The MVPF measures beneficiaries' WTP for a policy change per dollar of its net cost to the government. Following Hendren and Sprung-Keyser (2020), write these components as WTP_i and N_i for policy contrast i. Both are PVs for the same population and policy unit. For a positive net cost,

$$\text{MVPF}_i = \frac{\text{WTP}_i}{N_i}, \qquad N_i > 0. \tag{9}$$

WTP is the monetary value of the policy to its beneficiaries under the maintained welfare assumptions. Net fiscal cost includes the initial expenditure and subsequent fiscal effects of behavioral responses; it is not simply the program's budget. For example, higher earnings can increase private disposable income while also reducing

[23] Panels A and B of the collection table scale both monetary components by the same fixed reference program cost; Panel C reports thousands of native account dollars. This section uses each account's stated dollar unit. These display scalings leave MVPF unchanged.

net public cost through higher taxes and lower means-tested transfers. A policy with positive WTP and nonpositive net fiscal cost pays for itself and is assigned an infinite MVPF.[24]

The education-policy accounts can be organized into earnings-related components and the remaining benefits and costs. Use the superscript ref for Type 1 and write

$$\mathrm{WTP}_i^{\mathrm{ref}} = B_i^{\mathrm{ref}} + \mathrm{NPV}_i^{\mathrm{after,ref}} - C_i^{\mathrm{priv,ref}}, \tag{10}$$

where B_i^{ref} contains the value of program transfers and other private benefits included in the source account, $\mathrm{NPV}_i^{\mathrm{after,ref}}$ is the disposable-income component generated by the education-related earnings projection, and $C_i^{\mathrm{priv,ref}}$ is the associated private education payment. The corresponding public account is

$$N_i^{\mathrm{ref}} = G_i^{\mathrm{ref}} + C_i^{\mathrm{pub,ref}} - \mathrm{NPV}_i^{\mathrm{tax,ref}}. \tag{11}$$

Here G_i^{ref} contains direct program expenditure and other net fiscal costs, $C_i^{\mathrm{pub,ref}}$ is additional public education expenditure not already included in G_i^{ref}, and $\mathrm{NPV}_i^{\mathrm{tax,ref}}$ is the earnings-related fiscal return. The superscript tax includes changes in transfers when the source uses a combined tax-and-transfer rate. All components are signed policy-induced changes: a reduction in tuition payments raises WTP, whereas lost tax revenue raises net cost.[25]

I assemble 37 applications with an original welfare calculation, denoted Type 1, from education-related entries in the Policy Impacts Library and their underlying studies. These calculations project lifetime earnings and fiscal returns from education effects measured within finite windows. Most use the reconstruction by Hendren and Sprung-Keyser (2020), as in Adams; STBF and Tennessee Promise use their original authors' welfare calculations (Angrist et al. 2022; Schoonover et al. 2026). The remaining application uses the education effects in Li (2026) to construct an MVPF for the education channel of cannabis legalization; the original paper does not report an MVPF.[26]

[24]Assigning an infinite MVPF when a policy pays for itself is a welfare convention, not the arithmetic ratio at negative cost. If both components are nonpositive, I report them without an assigned MVPF. Hendren and Sprung-Keyser (2020) also apply the marginal framework to finite policy contrasts.

[25]Components vary by policy, and transfers must not be counted twice. Scholarship benefits need not equal expenditure dollar for dollar: Hendren and Sprung-Keyser (2020) distinguish induced students from those whose behavior is unchanged. Appendix C.3 gives the common unit, component-replacement, and accounting checks.

[26]The 20 September 2026 library snapshot and Appendix D supply 23 source-study families. Several tax contrasts can come from one study. Inclusion requires reconstructable education and accounting inputs; the collection is not a census of education policies.

The 38 calculations use a common net-income accounting structure while retaining different source mappings. Thirty-three use programs valued by Hendren and Sprung-Keyser (2020), three use independent authors' forecasts, and two are explicitly limited extensions. I first reconstruct the source monetary components, including earnings, taxes, education payments, WTP, and net cost. I then replace the education inputs, holding the source mapping fixed where it is recoverable. STBF retains its nonlinear credential–experience function. Early Head Start and Tennessee Promise use public-data reconstructions because their full author income inputs are unavailable; their reported Type 1 values remain separate comparison benchmarks. No income multiplier is chosen to force these reconstructions to equal a reported total. Eleven tax cases retain the fixed finite-policy normalization in Appendix D.4.2.

Let $\mathcal{P}_i^j$ contain the relevant schooling, credential, or time-enrolled paths. The common earnings object is

$$\Delta y_i(a) = F_i(a; \mathcal{P}_i^1, z_i) - F_i(a; \mathcal{P}_i^0, z_i), \tag{12}$$

where z_i contains the retained source prices, response coefficients, and population conventions. The path-based extension of the mapping in Hendren and Sprung-Keyser (2020) uses signed attainment increments with source-specific response templates; Social Security indexes its retained enrollment response to entry age. For nonlinear author forecasts, education inputs enter separately, as in the source model. In each case, early earnings losses enter this net-income effect once, without an additional opportunity-cost deduction. With ordinary age discounting, $E_i = \sum_a D_i(a)\Delta y_i(a)$; where the source discounts response phases from different bases, those phase-specific contributions are retained before aggregation (Appendix C.2).

Write T_i for the fiscal earnings component, P_i for private education payments, and C_i for public education costs. For accounts that value disposable earnings, the component replacement is

$$\mathrm{WTP}_i^{\mathrm{path}} = \mathrm{WTP}_i^{\mathrm{ref}} + (E_i^{\mathrm{path}} - E_i^{\mathrm{ref}}) - (T_i^{\mathrm{path}} - T_i^{\mathrm{ref}}) - (P_i^{\mathrm{path}} - P_i^{\mathrm{ref}}), \tag{13}$$

$$N_i^{\mathrm{path}} = N_i^{\mathrm{ref}} + (C_i^{\mathrm{path}} - C_i^{\mathrm{ref}}) - (T_i^{\mathrm{path}} - T_i^{\mathrm{ref}}). \tag{14}$$

The old component is removed before its replacement is added. For these accounts, source unit prices and case-specific timing allocations also determine how private and public education costs vary with the education path. That additional cost-timing bridge is stated explicitly; it is not inferred from the attainment data.

Independent author cases retain their specified non-income components. Early Head Start retains its source convention of gross earnings in WTP, and algorithmic placement retains source WTP while extending only the fiscal earnings channel. These exceptions retain their stated source conventions. Figure D1 displays paired WTP and net cost, so each ratio comes from the same feasible path.

The comparison curves H_i^0 and empirical restrictions come from the Census–ACS calibration library, matched to each application's geographic target and birth cohorts. Adams uses Massachusetts, retaining all races, sexes, and birthplaces in the external population data. Social Security combines a source age-23 control level with a transported national entry shape. Tennessee requires further adjustments because its education effects are conditional on high school (HS) graduation: I match that denominator, the source's prior-year state definition, and its latest-window control means. These adjustments improve correspondence to the reported outcomes without recovering the source design's untreated potential path. Where a comparable control level is unavailable, a statewide population mean remains an external calibration rather than a measured mean for selected scholarship applicants.[27]

Type 1 comprises the original WTP, net fiscal cost, and MVPF calculation. Type 2 constructs paths consistent with the education inputs retained for Type 1; Type 3 preserves the Type 2 restrictions and adds separately reported source windows or dynamics. For accounts based on Hendren and Sprung-Keyser (2020), both reconstructions retain their outcome-conversion rules; author-reported accounts use the case-specific conversions above. Additional information need not be a directly estimated Years of Schooling effect: an earlier enrollment estimate can enter through the retained duration conversion, provided its placement on the schooling path is stated explicitly. The same coefficient may already enter the reference total; Type 3 then retains its separate timing rather than introducing a previously unused estimate. Both reconstructions hold the comparison curve, empirical support, welfare account, and Type 2 constraints fixed.[28]

[27]The library covers 50 states, the District of Columbia (DC), and the nation. Multi-state targets use fixed weights; identification control states need not belong to the target population. Appendix C.1 gives the common calibration and support rules; Table D1 and the records in Appendix D.2 report the application-specific populations and adjustments.

[28]Appendix D.3 records the observation windows, conversions, and any temporal or population mapping. Institutional policy lags and conditional college-year outcomes require such mappings; they are not relabeled as individual follow-up. Splitting a single entry observation into assumed study years alone does not create Type 3, as in Tennessee Pell and Georgia's Helping Outstanding Pupils Educationally (HOPE) scholarship.

The common path-construction rule is constrained minimum distance:

$$\widehat{\Theta}_{i,L} = \underset{\theta_i \in \Theta_{i,L}}{\arg\min} \Big[\widehat{\mathbf{m}}_{i,L} - \mathbf{g}_{i,L}(\theta_i)\Big]' W_{i,L} \Big[\widehat{\mathbf{m}}_{i,L} - \mathbf{g}_{i,L}(\theta_i)\Big], \qquad L = 2, 3 \qquad (15)$$

Here $\widehat{\mathbf{m}}_{i,L}$ contains the fixed source point estimates, and $\mathbf{g}_{i,L}$ applies their age, cohort, and outcome operators to a candidate path. The set $\Theta_{i,L}$ imposes the anchored logistic form, calibrated support, education-level bounds, the one-year annual schooling-acquisition cap, and credential ordering where relevant. Type 3 retains the Type 2 equalities and adds source observations to the fit. $W_{i,L}$ is fixed in advance; the ten minimum-distance Type 3 applications use the identity matrix in their common within-case units, without standard-error bands. If zero residual is attainable, all exact matches minimize the same criterion. Otherwise the retained paths minimize the discrepancies subject to the unchanged Type 2 restrictions, and their residuals are reported. The numerical search retains distinct minimizing solutions found, not every approximately fitting path. Input uncertainty is studied separately. Table D2 identifies the ten best-fit applications, for which the search currently retains one path each; this is selection within the maintained model, not identification of the true path.

For Tennessee, the retained non-income components yield WTP $= w_0 + (1-t)E$ and $N = n_0 - tE$. At positive cost, MVPF exceeds one exactly when $E > n_0 - w_0$. The Adams and Social Security accounts additionally transport education payments, so their welfare cannot be summarized by a fixed threshold in earnings alone. Welfare comparisons across groups also depend on social welfare weights (Hendren and Sprung-Keyser 2020). All ranges refer to verified numerical candidates rather than certified bounds on the continuous feasible set.

4.2 Adams Scholarship

Cohodes and Goodman (2014) use test-score eligibility thresholds to estimate the effects of Massachusetts' Adams Scholarship for the 2005–2006 HS graduating classes. Their Table 4 reports effects on entry into any college of 1.7 percentage points immediately after HS and 0.7 points within two years, and BA completion effects of -1.7 points within four years and -2.5 points within six years; it also reports fourth-year persistence. Hendren and Sprung-Keyser (2020) use the two-year entry and six-year completion estimates, separately for two-year and four-year colleges, to construct -1.80×10^{-2} schooling-equivalent years per eligible student.

This converts effects from different windows into a schooling equivalent rather than using a directly estimated schooling effect. Their Type 1 MVPF is 0.72.

Type 2 places this converted total at age 24, using the Massachusetts ACS comparison profile for the 1987 birth cohort and an unaffected age-17 anchor. Anchored logistic paths satisfy the calibrated support, nondecreasing schooling, and the one-year annual acquisition cap. The source earnings response, tax mapping, and education-cost unit schedules remain fixed; signed path increments change earnings and the timing of private and public education costs. Type 3 preserves the age-24 equality and adds three dated targets: immediate entry at 18, entry within two years at 20, and four-year completion at 22. The source conversion gives 2.60×10^{-2}, 7.00×10^{-3}, and -1.10×10^{-2} Years of Schooling, respectively. These age placements are temporal allocations of converted estimates; the age-20 target adds the timing of coefficients already used in Type 1. Equal-weight minimum distance fits the additional points while keeping the age-24 total exact. Defined Type 2 MVPFs span 0.48 to $+\infty$; the selected Type 3 path yields 0.63, with positive WTP below positive net cost (Figure 5). Its nonzero fitting residuals make the singleton conditional on the model and fitting rule. Appendix D.2 gives the conversion arithmetic, source interpretation, and accounting details.

4.3 Social Security Student Benefits

Dynarski (2003) studies the 1982 withdrawal of college benefits for eligible children of deceased, disabled, or retired workers. The design compares changes in attendance among children with deceased fathers against those with living fathers. I retain the positive benefit-availability contrast. Table 3 reports increases in cumulative full-time college attendance of 21.9 percentage points by age 23 and 22.4 points by age 28, the latter under conservative attrition imputation. These are effects on ever having attended, rather than current enrollment. The enrollment alternative in Hendren and Sprung-Keyser (2020) values the age-28 effect using two Years of Schooling per induced entrant and 3 percent discounting, yielding a Type 1 MVPF of 4.87.

Type 2 matches the age-28 effect exactly with an unaffected age-17 anchor. The comparison curve uses the national ACS 1985–1987 some-college age shape, rescaled to the source's age-23 control mean of 35.2 percent. This transported shape remains an external calibration for full-time entry. Candidates retain donor-based speed and timing restrictions and probability bounds between zero and one. Each signed entry increment receives the source's two-year schooling conversion, entry-indexed

earnings response, and education-cost schedule; fixed program components are unchanged. Type 3 adds the age-23 attendance effect as an exact restriction, holding calibration and valuation fixed. Two observed effects can leave multiple paths within the three-parameter logistic family. The Type 2 MVPF range of 3.79–4.99 narrows to 4.21–4.52, a 74.4 percent contraction in its searched width (Figure 6). All retained candidates have positive WTP and net cost, with MVPF above one. The Type 3 range lies below the reference 4.87: the reference single-entry batch need not belong to the logistic class or match the added age-23 effect. Appendix D.2 documents the source specifications, calibration, and entry-based income and cost mappings.

4.4 Tennessee Promise

Schoonover et al. (2026) study the expansion of tuition-free community college (CC) from a Tennessee county program in 2009 to a statewide program in 2015. Using ACS 2000–2024 data for HS graduates aged 18–24, they report effects for six overlapping age pairs from 18–19 through 23–24. In the latest window, cumulative some college, associate's degree (AA), and BA attainment rise by 3.5, 4.0, and 2.7 percentage points. Their welfare calculation uses Tennessee education returns, a 26 percent tax rate, and 3 percent discounting; under the reported Type 1 account, the policy pays for itself.

Type 2 matches the three latest-window effects as equal-weight age-23–24 averages across the 1992–1998 birth cohorts, with fixed equal cohort weights. The ACS reference matches the HS-graduate denominator, prior-year residence, and source control levels. Both policy states share the age-17 anchor, and credentials remain nested at every age. State-calibrated support and a donor-based limit on subsequent completion restrict continuations, including a BA plateau cap of 45.4 percent. The authors' complete age-by-credential income inputs are unavailable. I therefore reconstruct earnings from public ACS data, assuming two pre-completion study years for some college without a degree and AA-only holders and four for BA holders, zero earnings during study, and earnings differences relative to HS graduates thereafter. Assumptions about participant shares and price years link these paths to the fixed account WTP $= 2{,}124.31 + 0.74E$ and $N = 14{,}792 - 0.26E$, in 2021 dollars. Appendix D.2 documents these assumptions and the discrepancy from the authors' income projection. Comparisons with Type 1 consequently reflect accounting reconstruction as well as education timing; comparisons across candidate paths hold the account fixed.

Type 3 adds the fifteen earlier coefficients to the three latest-window effects. It preserves the latter exactly and fits all eighteen point inputs by equal-weight minimum distance, without SE bands. Defined Type 2 MVPFs span -1.49 to $+\infty$. The retained Type 3 path has earnings NPV of 113.94 thousand dollars and an infinite MVPF, with input root mean squared error (RMSE) of 0.35 percentage points. Figure 7 shows the paths; Table D2 reports all fitted values and residuals. The singleton is conditional on the maintained restrictions and fitting criterion. These comparisons hold reported education effects fixed. Appendix E separately documents the input repetitions, their dependence assumptions, and unsuccessful searches. Section 5 extends the comparison across the policy collection, examining how often alternative paths leave welfare judgments unresolved and how additional observations narrow those differences.

5 Welfare Judgments across Policies and Populations

The welfare implications of an education effect depend on its persistence, the ages at which attainment changes, and uncertainty in the education inputs. I examine whether the observed information establishes an MVPF above or below one, whether the policy pays for itself, and how these conclusions change when additional comparable observations restrict the path. The collection contains 38 calculations based on finite attainment paths, including 37 with Type 1. The tax-input conversions and account adjustments described in Section 4.1 are distinguished throughout the comparison. Single-parameter comparisons distinguish permanent effects on eventual attainment from temporary effects on acquisition speed or timing. Input repetitions examine a separate source of variation under each case's specified uncertainty scheme. Appendix F compares men and women within seven policies, holding terminal education effects and valuation prices fixed.

5.1 Continuation Values and Policy Comparisons

Figure 8 places the Type 2 ranges and available Type 3 exact-match ranges or best-fit values beside the reference MVPFs, ordered by descending Type 1. All 38 calculations use finite attainment paths. Among the 37 with Type 1, 30 admit positive-cost paths on both sides of MVPF one at fixed Type 2 inputs. Of these, 17 belong to the 24-account group displayed separately from the finite tax and special

accounts. Taking the education coefficients as known therefore need not determine whether beneficiary welfare exceeds one dollar per dollar of net public cost.[29]

The uncertainty can concern the signs of the monetary components, rather than only the size of their ratio. In 24 of the 37 calculations with Type 1, including 12 in the 24-account group, one compatible path gives negative WTP and positive public cost, whereas another gives positive WTP and nonpositive public cost. These paired outcomes follow from different projected earnings and tax receipts under the same education restrictions. Fixed transfers remain unchanged. Where recovered, source private and public education-cost schedules are evaluated on each path, so these costs can also change. A gain that persists into high-earning ages can benefit recipients and repay the public expenditure; an initial gap that disappears or reverses can produce a different conclusion. These are possible aggregate continuations, not estimates of individual dropout or return-to-school behavior.

A large MVPF need not represent a large dollar benefit. Numerical checks find positive-cost approaches to zero net cost with a positive limiting WTP in 24 of the 37 calculations, and with a negative limiting WTP in 16. The resulting ratio diverges while both monetary components remain finite. Figure 8 places positive-cost ratio limits and outcomes in which policies pay for themselves on a common axis, with positive infinity at its endpoint; Table 3 retains their accounting distinction. The span does not assert that every intermediate ratio is attained, and a finite searched endpoint does not certify the boundary of the continuous feasible set.[30]

Figure 9 applies a lower-endpoint ordering to all 37 calculations with a reference MVPF. It first compares lower endpoints and then upper endpoints, both descending, using Type 3 where available and Type 2 otherwise. For a best-fit singleton the endpoints coincide; its rank conditions on the fitting rule as well as the path restrictions. Finite endpoints are rounded to two decimals, and equal endpoint pairs share a rank. Some tax calculations retain unbounded positive-cost ratio spans although attainment and both monetary components are bounded. The divergence comes from finite paths approaching zero net cost. No path probabilities

[29]Several tax contrasts come from one study, so the calculation count is not a count of independent studies. The 24-account group spans 20 source-study families, with threshold pairs in 15. MVPF one weights beneficiary dollars and public resources equally; welfare priorities can also depend on distributional weights and financing alternatives (Hendren and Sprung-Keyser 2020).

[30]Appendix D.3 documents complete paired paths, constraint checks, and one-sided ratio sequences. A large finite ratio or a rounded zero cost is insufficient evidence of an infinite limit.

are used to break ties.[31] Both axes increase from small to large ranks, so a point below the diagonal has higher priority under this rule than under Type 1. The ordering expresses a preference over conditional endpoint summaries; it is not a social-welfare or certified worst-case ranking.

Twelve calculations with Type 1 admit Type 3 under the common rule for additional source windows or dynamics. Replacing one calculation's Type 2 by Type 3 improves its displayed priority in ten cases and leaves it unchanged in two, holding the others fixed. This is a consequence of the endpoint-ordering rule, not evidence that information raises true welfare. Additional restrictions or the best-fit criterion can exclude favorable as well as adverse continuations. For the nine best-fit Type 3 applications with Type 1, the displayed contraction includes selection by the maintained loss function. The added restrictions preserve the source windows under explicit outcome and age mappings. For example, Adams combines enrollment and completion coefficients, while Wisconsin retains their joint cohort–follow-up weights. Narrowing therefore depends on these mappings as well as on the additional observations. Appendix D.2 records the inputs and conversions; Figures D2–D9 display H^0 and H^1 for all 38 calculations. No probability distribution is assigned to the paths.

Table 3 reports WTP and net cost alongside MVPF. Its policy-type column distinguishes 14 direct student-aid calculations, 17 education tax-benefit calculations, four tuition or public-college-funding calculations, one college-assessment and course-placement intervention, one early-childhood intervention, and one education channel of another policy. These types describe the intervention, whereas the panels distinguish accounting specifications. Several calculations share a source study, so the category counts are not independent-study counts. Panel A contains the 24 retained comparison cases, including Tennessee's disclosed public-data reconstruction. Panel B reports the eleven finite tax reconstructions, whose source marginal coefficients require the stated finite-policy conversion. Their WTP and net-cost extrema are finite, even when their ratios are not. Panel C distinguishes the public-data Early Head Start account, the algorithmic-placement extension, and the cannabis education channel. Monetary components in Panels A and B use a fixed reference program-cost divisor within each calculation; Panel C uses native thousands. Every MVPF combines the two components of the same path. Very large normalized amounts indicate that the allowed future response can dominate

[31]The tax calculations transport local attendance and accounting responses over a fixed finite schedule change calibrated to one point-estimate credit dollar. Appendix D.4.2 states this additional assumption.

a small reference expenditure; they do not establish that such a response is likely. The main finding is the existence of conflicting welfare judgments within a fixed reconstruction, conditional on both its path restrictions and its input conversion.

5.2 Permanent and Temporary Education Effects

The same finite-window education input can support a permanent or a temporary effect, with different implications for lifetime welfare. I examine this distinction by holding two parameter families at their calibrated comparison values and solving for a change in the third that matches the retained Type 2 education input. A nonzero change in K produces a permanent effect by altering eventual attainment. Changing b or d with K fixed produces a temporary effect: the population acquires education faster or more slowly, but the attainment gap eventually vanishes. The observation operator, empirical support, and welfare account remain unchanged across these three exercises. The comparison makes the continuation assumption explicit without treating the observed input as evidence that one particular parameter changed. Appendix B.5 gives the matching procedure and its relation to the main path construction.[32]

Interpreting the observed education input as a permanent effect tends to preserve the reference welfare comparison. In Panel A of Figure 10, the finite MVPF pairs follow an upward pattern broadly aligned with the 45-degree equality line. Changing K leaves a permanent difference in eventual attainment, so the education gain or loss continues to affect projected earnings beyond the observed window. This resembles a reference calculation that uses the measured education effect to project sustained earnings gains or losses. The alignment is consistent with interpreting the reference valuations in this collection as assigning a permanent effect to the short-window education input. Shared education inputs and related return calibrations contribute to this resemblance, which therefore does not independently establish that the education effect is permanent.[33]

Temporary effects produce a different welfare comparison. In Panels B and C, several finite reconstructed MVPFs cluster near one across a wide range of

[32]The exercise preserves the individual components of each fitted H^0, so a zero change reproduces that baseline exactly. K and d change additively; b changes by a common proportion. STBF permits four outcome-specific changes and Tennessee Promise permits three credential-specific changes in the active family, subject to joint restrictions. This is a diagnostic rather than a decomposition of the main candidate ranges.

[33]Panel A contains 22 finite pairs, seven calculations that move from a finite reference MVPF to an outcome in which the policy pays for itself, and seven in which the policy pays for itself in both accounts. The diagonal denotes equality on the common asinh scale; it is not an estimated regression line.

reference values, rather than following the near-diagonal pattern in Panel A. With eventual attainment fixed, changes in b or d match the education input through faster or slower acquisition, and the attainment gap eventually fades. Figure B3 shows that these temporary effects still generate dispersed earnings NPVs, even though several corresponding MVPFs lie close to one. This pattern reflects how net earnings and education-dependent costs enter the retained welfare accounts. When earnings NPV is small relative to the retained transfers and other non-earnings components, those components largely determine WTP relative to net public cost. Several accounts consequently remain near MVPF one even as their earnings NPVs differ.[34] Temporary effects therefore retain economic value; their influence on MVPF depends on their duration and on the scale and composition of the welfare account under the specified valuation rule. The contrast with Panel A separates two explanations of the same education input: a permanent change in eventual attainment and a temporary change in when education is acquired. Distinguishing these interpretations in a lifetime welfare calculation requires additional evidence or an explicit continuation assumption. Greater precision within the observed window does not by itself resolve that distinction.

Table 5 gives each reference MVPF, its single-parameter reconstruction, and the required parameter change. It also records calculations without a matching solution. The matched samples differ, which limits comparisons of the plotted patterns. Holding two parameters fixed can leave no solution within the empirical support; one, eight, and seventeen calculations have no matched plateau, speed, or timing solution, respectively. Unmatched calculations are reported separately and are not interpreted as zero responses. None of the single-family solutions exactly reproduces all additional point inputs where Type 3 is available. Additional observations can therefore challenge a one-parameter explanation even when the reduced input is matched. This diagnostic is stricter than merely attaining the best-fit criterion used for fitted Type 3 applications. The results do not establish that the original lifetime projections are incorrect or that their authors ignored dynamics. They show how the welfare interpretation of a retained short-window input depends on the continuation assigned to it, while the source studies' observed dynamics and mechanism evidence remain relevant to choosing among those continuations.

[34] Among finite pairs, 7 of 23 temporary speed results and 7 of 15 temporary timing results lie between 0.75 and 1.25. Holding the other components fixed, $W = w_0 + (1-t)E$ and $N = n_0 - tE$ approach the ratio w_0/n_0 as net earnings NPV E approaches zero, provided $n_0 > 0$. Path-dependent education costs can also change this limiting comparison. That ratio need not equal one. This accounting relationship does not imply general insensitivity to temporary effects: the ratio can change sharply near zero net cost.

5.3 From Uncertain Inputs to Paths and Welfare

Uncertainty in an estimated education effect carries through to the education paths that match it and to their welfare values. This section examines that propagation, together with sensitivity to selected accounting inputs, using case-specific input repetitions. Figure 11 displays point-input values in dark colors and input-repetition endpoint summaries in lighter colors. Policy categories and input-method codes appear in Tables 3 and 4. The figure's Type 1 panel reports available source 95 percent MVPF intervals. The Type 2 and Type 3 panels instead summarize the 1,000 prespecified repetitions per calculation, keeping the ACS calibration and specified conversion and valuation rules fixed. The source income-response and education-cost functions stay fixed in form. The 33 calculations based on Hendren and Sprung-Keyser (2020) propagate their paired accounting inputs, including earnings-return draws where available. Their education perturbations can still be conditional or constructed, as identified below. The eleven finite tax displacements remain fixed at their point normalizations. STBF holds its recovered author-model coefficients fixed; Early Head Start and Tennessee hold their public wage schedules and cost or participation conventions fixed. Each summary runs from the 2.5th percentile of candidate minima to the 97.5th percentile of candidate maxima. The reported methods distinguish among propagation of archived joint source/accounting draws, constructed normal or correlated perturbations, and exercises varying economic inputs while holding education fixed.[35]

Among the 37 calculations with Type 1, 37 Type 2 summaries cross MVPF one under their specified repetition schemes. WTP summaries cross zero in 34 calculations and net-cost summaries cross zero in 33; 33 have an infinite upper MVPF endpoint under the convention for policies that pay for themselves or a verified positive-cost limit. Within the 24-account group, the corresponding counts are 24, 22, 20, and 20. These counts describe how the reconstructed endpoints respond to the inputs being varied. They do not measure the chance that a policy harms beneficiaries or pays for itself. The monetary components remain informative when the ratio becomes unstable near zero net cost.[36]

Table 4 identifies each repetition method, endpoint summary, and usable count. Across the collection, all 38,000 scheduled reduced-information repetitions were

[35]Appendix E.1 defines the order-statistic summaries and the inputs varied. These procedures do not share a sampling law or establish confidence coverage for all compatible paths. No probability distribution is assigned to paths within a repetition.

[36]Source intervals and reconstructed summaries can vary different inputs and use different dependence assumptions. Their widths are not a controlled comparison of inferential methods.

evaluated; 38,000 yielded verified paths and 38,000 yielded a defined MVPF. All thirteen Type 3 specifications propagate their complete input vectors through 1,000 paired repetitions, giving 12,980 verified outcomes. Type 2 uses the same retained-input draw as Type 3. The original correlation groups are preserved and extended to added source windows under explicit assumptions, rather than treating an unavailable source covariance as known.[37] Tennessee's constructed normal shocks perturb and refit all eighteen point targets. Two displayed zero SEs receive a 5.00×10^{-4} reporting-precision scale. This is an imposed sensitivity convention, not a recovered SE or a draw from the source permutation distribution.

Path uncertainty remains even if the original education effect is known exactly: different paths can match that effect and imply different welfare values. Estimation uncertainty arises because the original effect is itself estimated imprecisely. Allowing the input to vary changes the compatible paths and, in turn, the welfare range. The two sources can operate together, since each possible input may still admit several continuations. More precise estimates within the same window address estimation uncertainty; comparable observations at additional ages can help resolve path uncertainty. The repetitions reported here trace this propagation under the specified source-draw or input-sensitivity scheme.

6 Discussion

Education effects observed within a finite window can support different lifetime welfare judgments, depending on assumptions about persistence, attainment timing, and valuation. I take the source studies' causal identification as given; the additional assumptions link their estimates to education paths for the target population. Conversions must preserve policy units and observation windows, while assumed schooling durations do not provide new observations. The external comparison profile must also be relevant to the affected population. Matching Census–ACS geography and birth cohorts improves comparability, but population averages may not describe scholarship applicants or marginal students. Historical cohorts' unobserved early attainment is reconstructed using other cohorts' shapes. The

[37] Appendix E.2 records these extensions. Unsuccessful searches are retained without replacement or a claim of infeasibility; usable counts exclude them. Cannabis' six-year-only repetitions are diagnostic, while its primary result uses the three-window Type 3.

common anchor imposes no effect or anticipation before the specified exposure age. Fit to ACS profiles cannot establish these population and exposure assumptions.[38]

Path and valuation assumptions further determine the comparisons. Logistic curves may omit interrupted schooling or reentry, and historical variation need not bound a policy response. The single-parameter exercises distinguish permanent from temporary explanations without identifying which applies. Earnings returns may also depend on completion age (Bárány et al. 2023). The policy reconstructions apply the same income function to treated and comparison education paths and value the difference. For accounts based on Hendren and Sprung-Keyser (2020), signed attainment increments move the source negative-then-positive response template across the corresponding completion ages, or entry ages for the Social Security enrollment alternative. This assumes that the response template applies across ages and is additive in signed increments; it does not identify individuals' additional time enrolled from completed attainment. STBF retains a nonlinear credential–experience model, conditional on recovered sufficient statistics and the subgroup bridge. Early Head Start and Tennessee use disclosed public-data reconstructions where the source income tables are unavailable, so comparisons with their Type 1 benchmarks reflect changes in valuation inputs as well as continuation. Recoverable source education-cost schedules move with the paths; STBF, Early Head Start, and Tennessee retain their stated non-income amounts. Appendix D.6 varies calibration, valuation, and search at fixed source inputs. Using Type 3 where available, positive-cost MVPFs cross one in 21 of 38 calculations, compared with 20 under narrower support. All seven common specifications admit negative-WTP, positive-cost paths alongside positive-WTP, nonpositive-cost paths in 16 calculations. These conditional contrasts neither validate all assumptions nor exclude unsearched paths; wider support preserves feasible paths but can change the minimum-distance solution. The ten fitted Type 3 results report nonzero residuals and one best-fitting path each. This selection is conditional on the logistic family and fixed weights; it is not proof that the underlying educational response or its welfare value is point identified.

Interpreting the welfare ranges also requires distinguishing input uncertainty from the rule used to rank policies. Different repetition schemes and incomplete covariance information prevent a common sampling interpretation, although all repetitions hold calibration and the path family fixed. Endpoint percentiles do not

[38]Appendix C.1 gives the common calibration rules; Table D1 and Appendix D.2 record age support and donor shapes. The selected literature and policy collections also limit the generalizability of the reported frequencies.

establish confidence coverage (Imbens and Manski 2004), and more repetitions can reduce simulation error without resolving these limits or assigning probabilities to paths. The lower-endpoint ordering gives weight to adverse continuations among the paths found. Its ranks depend on search coverage, calibration, and information available for Type 3; they are neither expected-welfare nor certified worst-case rankings. Joint reporting of WTP and net cost distinguishes large benefits from ratios driven by near-zero costs, while outcomes in which policies pay for themselves and paths without an assigned ratio require separate interpretation. Spending recommendations further require distributional weights, a financing comparison, and assumptions about policy expansion. To assess the underlying extrapolations, later education outcomes withheld from path construction could be used to compare logistic predictions with constant effects and flexible monotone alternatives. Observed earnings could test valuation, while comparable age windows and joint coefficient covariances would improve reconstruction and inference. Such evidence can distinguish permanent from temporary effects with different welfare implications.

For future applications, researchers can begin by fitting $H^0(a)$ for the population whose education effects enter the welfare calculation. When the original microdata contain an appropriate control or comparison group, its age-specific attainment can be used to estimate the logistic baseline, with the outcome definition, sample, and weighting aligned with the policy estimand. Comparable external cohort data provide an alternative when suitable control observations are unavailable. Researchers should report treatment effects at as many comparable ages or windows as the data permit, together with control-group attainment and, where feasible, the joint covariance of the reported estimates. Conditional on the fitted $H^0(a)$, constrained minimum-distance fitting combines these effects to construct possible $H^1(a)$ paths under the maintained restrictions. Additional observations help distinguish permanent attainment changes from temporary differences in acquisition and can narrow the resulting welfare range. Researchers can then calculate NPV and MVPF for each pair using an explicit earnings mapping and fiscal account, with input uncertainty examined separately. The resulting lifetime paths remain conditional on the assumptions used to extend them beyond the observed ages.

7 Conclusion

An education effect measured within a finite window leaves two economically distinct questions open: how much of the effect persists and when the remaining attainment

occurs. This paper addresses these questions by constructing education paths that match the education inputs retained for the reference calculation, fitting additional point observations under the stated criterion, and valuing each path under a specified rule. The simulations show how additional observations can distinguish continuations that agree at an earlier age. Single-parameter comparisons separate permanent effects on eventual attainment from temporary effects on acquisition speed or timing. Temporary effects can change lifetime earnings even when the attainment gap eventually vanishes. In the policy applications, verified paths often support different welfare judgments even when the retained education estimates are treated as known. At fixed inputs, 30 of the 37 calculations with Type 1 admit positive-cost MVPFs on both sides of one. More precise estimation within the same window would not remove these particular continuation differences.

The framework provides a reusable tool for researchers to translate education effects measured within finite windows into ranges of lifetime MVPFs under explicit path and valuation assumptions. It combines reported estimates with a baseline education profile from the study's control group or a comparable external population. Researchers can then examine how alternative continuations change welfare conclusions while retaining the valuation rules of the original policy calculation where recoverable. As additional comparable education outcomes become available, the same procedure can assess whether, and by how much, they narrow the range. Researchers can use the framework both to reassess existing welfare calculations and, when designing future empirical studies, to evaluate the value of reporting education effects at additional ages.

The contribution is a method for assessing which welfare judgments remain dependent on unobserved education paths. Compatible paths with different classifications establish that dependence under the stated calibration and valuation rules; they do not establish either path's probability. Input repetitions answer a separate question about imprecision in the retained estimates. The distinction helps identify whether further progress requires more precise inputs, additional observations of attainment, or stronger evidence on the valuation relationship.

The results favor reporting the observed education window and the assumed continuation separately. When additional outcomes are available, their contribution depends on whether they constrain attainment at a new age for a comparable population. Reporting WTP and net fiscal cost alongside MVPF also distinguishes changes in benefits from ratios driven by costs near zero. Longer follow-up, richer education histories, and observed earnings can each improve the assessment, but

they answer different questions. The relevant criterion is whether the available evidence distinguishes paths with different welfare implications.

References

Abraham, Katharine G. and Melissa A. Clark, "Financial Aid and Students' College Decisions: Evidence from the District of Columbia Tuition Assistance Grant Program," *Journal of Human Resources*, 2006, *41* (3), 578–610.

Acton, Riley, "Effects of Reduced Community College Tuition on College Choices and Degree Completion," *Education Finance and Policy*, 2021, *16* (3), 388–417.

Angrist, Joshua D., David H. Autor, and Amanda Pallais, "Marginal Effects of Merit Aid for Low-Income Students," *The Quarterly Journal of Economics*, 2022, *137* (2), 1039–1090.

Athey, Susan, Raj Chetty, Guido W. Imbens, and Hyunseung Kang, "The Surrogate Index: Combining Short-Term Proxies to Estimate Long-Term Treatment Effects More Rapidly and Precisely," *Review of Economic Studies*, 2026, *93* (4), 2284–2312.

Bárány, Zsófia L., Moshe Buchinsky, and Pauline Corblet, "Late Bloomers: The Aggregate Implications of Getting Education Later in Life," NBER Working Paper 31874, National Bureau of Economic Research 2023.

Bartik, Timothy J., Brad Hershbein, and Marta Lachowska, "The Effects of the Kalamazoo Promise Scholarship on College Enrollment and Completion," *Journal of Human Resources*, 2021, *56* (1), 269–310.

Bergman, Peter, Elizabeth Kopko, and Julio Rodríguez, "A Seven-College Experiment Using Algorithms to Track Students: Impacts and Implications for Equity and Fairness," 2026. Author manuscript, June 2026 version.

Bettinger, Eric, "How Financial Aid Affects Persistence," in Caroline M. Hoxby, ed., *College Choices: The Economics of Where to Go, When to Go, and How to Pay for It*, Chicago: University of Chicago Press, 2004, pp. 207–238.

Bhuller, Manudeep, Magne Mogstad, and Kjell G. Salvanes, "Life-Cycle Earnings, Education Premiums, and Internal Rates of Return," *Journal of Labor Economics*, 2017, *35* (4), 993–1030.

Bound, John, Michael F. Lovenheim, and Sarah Turner, "Increasing Time to Baccalaureate Degree in the United States," *Education Finance and Policy*, 2012, *7* (4), 375–424.

Bruce, Donald J. and Celeste K. Carruthers, "Jackpot? The Impact of Lottery Scholarships on Enrollment in Tennessee," *Journal of Urban Economics*, 2014, *81*, 30–44.

Bulman, George B. and Caroline M. Hoxby, "The Returns to the Federal Tax Credits for Higher Education," *Tax Policy and the Economy*, 2015, *29* (1), 13–88.

Carruthers, Celeste K. and Jilleah G. Welch, "Not Whether, but Where? Pell Grants and College Choices," *Journal of Public Economics*, 2019, *172*, 1–19.

Castleman, Benjamin L. and Bridget Terry Long, "Looking beyond Enrollment: The Causal Effect of Need-Based Grants on College Access, Persistence, and Graduation," *Journal of Labor Economics*, 2016, *34* (4), 1023–1073.

Chen, Jiafeng and David M. Ritzwoller, "Semiparametric estimation of long-term treatment effects," *Journal of Econometrics*, 2023, *237* (2), 105545.

Chetty, Raj, "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 2009, *1* (1), 451–488.

Cocci, Matthew D. and Mikkel Plagborg-Møller, "Standard Errors for Calibrated Parameters," *Review of Economic Studies*, 2025, *92* (5), 2952–2978.

Cohodes, Sarah R. and Joshua S. Goodman, "Merit Aid, College Quality, and College Completion: Massachusetts' Adams Scholarship as an In-Kind Subsidy," *American Economic Journal: Applied Economics*, 2014, *6* (4), 251–285.

Cornwell, Christopher, David B. Mustard, and Deepa J. Sridhar, "The Enrollment Effects of Merit-Based Financial Aid: Evidence from Georgia's HOPE Program," *Journal of Labor Economics*, 2006, *24* (4), 761–786.

Deming, David J. and Christopher R. Walters, "The Impact of Price Caps and Spending Cuts on U.S. Postsecondary Attainment," NBER Working Paper 23736, National Bureau of Economic Research 2017.

Denning, Jeffrey T., "College on the Cheap: Consequences of Community College Tuition Reductions," *American Economic Journal: Economic Policy*, 2017, *9* (2), 155–188.

Dynarski, Susan M., "Does Aid Matter? Measuring the Effect of Student Aid on College Attendance and Completion," *American Economic Review*, 2003, *93* (1), 279–288.

Fan, Yanqin, Carlos A. Manzanares, Hyeonseok Park, and Yuan Qi, "A Sensitivity Analysis of the Surrogate Index Approach for Estimating Long-Term Treatment Effects," 2026. Working paper, arXiv:2603.00580v1.

García, Jorge Luis, James J. Heckman, Duncan Ermini Leaf, and María José Prados, "Quantifying the Life-Cycle Benefits of an Influential Early-Childhood Program," *Journal of Political Economy*, 2020, *128* (7), 2502–2541.

Goldrick-Rab, Sara, Robert Kelchen, Douglas N. Harris, and James Benson, "Reducing Income Inequality in Educational Attainment: Experimental Evidence on the Impact of Financial Aid on College Completion," *American Journal of Sociology*, 2016, *121* (6), 1762–1817.

Haider, Steven and Gary Solon, "Life-Cycle Variation in the Association between Current and Lifetime Earnings," *American Economic Review*, 2006, *96* (4), 1308–1320.

Heckman, James J., John Eric Humphries, and Gregory Veramendi, "Dynamic Treatment Effects," *Journal of Econometrics*, 2016, *191* (2), 276–292.

Hendren, Nathaniel and Ben Sprung-Keyser, "A Unified Welfare Analysis of Government Policies," *The Quarterly Journal of Economics*, 2020, *135* (3), 1209–1318.

Hess, Andrew John, "Think of the Children: The Role of Early Childhood Interventions in Long-Term Human Capital Formation." PhD dissertation, University of California, Los Angeles 2025.

Hoxby, Caroline M. and George B. Bulman, "The Effects of the Tax Deduction for Postsecondary Tuition: Implications for Structuring Tax-Based Aid," *Economics of Education Review*, 2016, *51*, 23–60.

Imbens, Guido, Nathan Kallus, Xiaojie Mao, and Yuhao Wang, "Long-term causal inference under persistent confounding via data combination," *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 2025, *87* (2), 362–388.

Imbens, Guido W. and Charles F. Manski, "Confidence Intervals for Partially Identified Parameters," *Econometrica*, 2004, *72* (6), 1845–1857.

Kleven, Henrik J., "Sufficient Statistics Revisited," *Annual Review of Economics*, 2021, *13* (1), 515–538.

Li, Anzheng, "Higher Education or Just Higher? The Unintended Effects of Recreational Marijuana Legalization," 2026. Working paper, September 8. SSRN 7433138.

Long, Bridget Terry, "The Impact of Federal Tax Credits for Higher Education Expenses," in Caroline M. Hoxby, ed., *College Choices: The Economics of Where to Go, When to Go, and How to Pay for It*, Chicago: University of Chicago Press, 2004, pp. 101–168.

Manski, Charles F., "Choosing Treatment Policies Under Ambiguity," *Annual Review of Economics*, 2011, *3* (1), 25–49.

Marx, Benjamin M. and Lesley J. Turner, "Borrowing Trouble? Human Capital Investment with Opt-In Costs and Implications for the Effectiveness of Grant Aid," *American Economic Journal: Applied Economics*, 2018, *10* (2), 163–201.

Obradović, Filip, "Identification of Long-Term Treatment Effects via Temporal Links, Observational, and Experimental Data," 2026. Working paper, revised March 2026, arXiv:2411.04380v2.

Schoonover, Paige, Jonathon Attridge, Celeste K. Carruthers, and Jilleah G. Welch, "A Promise Worth Keeping? Impacts of Free Community College on Degrees and Earnings," NBER Working Paper 35226, National Bureau of Economic Research 2026.

Scott-Clayton, Judith, "On Money and Motivation: A Quasi-Experimental Analysis of Financial Incentives for College Achievement," *Journal of Human Resources*, 2011, *46* (3), 614–646.

Seftor, Neil S. and Sarah E. Turner, "Back to School: Federal Student Aid Policy and Adult College Enrollment," *Journal of Human Resources*, 2002, *37* (2), 336–352.

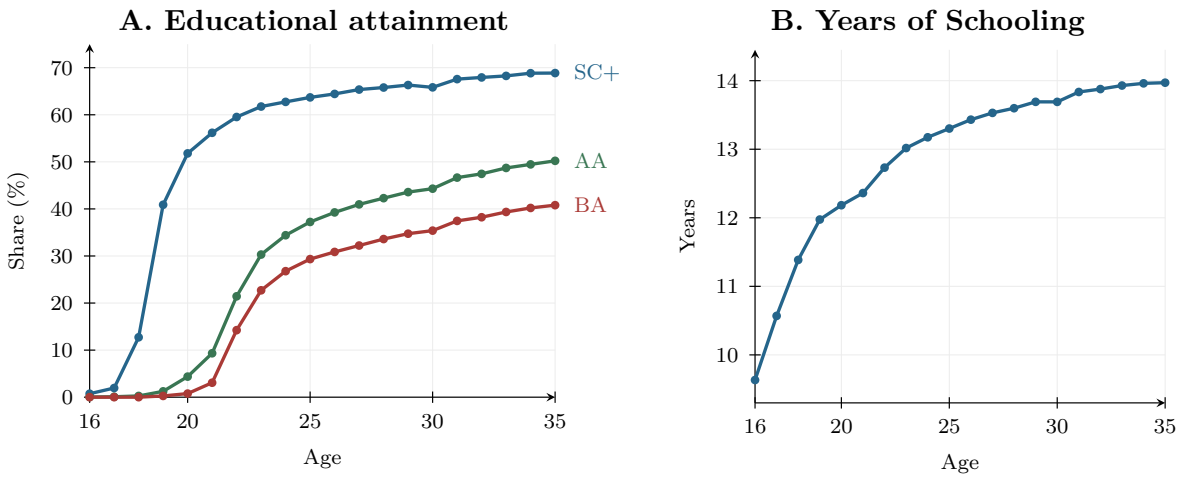
Turner, Nicholas, "The Effect of Tax-Based Federal Student Aid on College Enrollment," *National Tax Journal*, 2011, *64* (3), 839–861.

Vohra, Vedant, "Inference for Welfare Metrics," *Journal of Public Economics*, 2026, *256*, 105603.

Zimmerman, Seth D., "The Returns to College Admission for Academically Marginal Students," *Journal of Labor Economics*, 2014, *32* (4), 711–754.

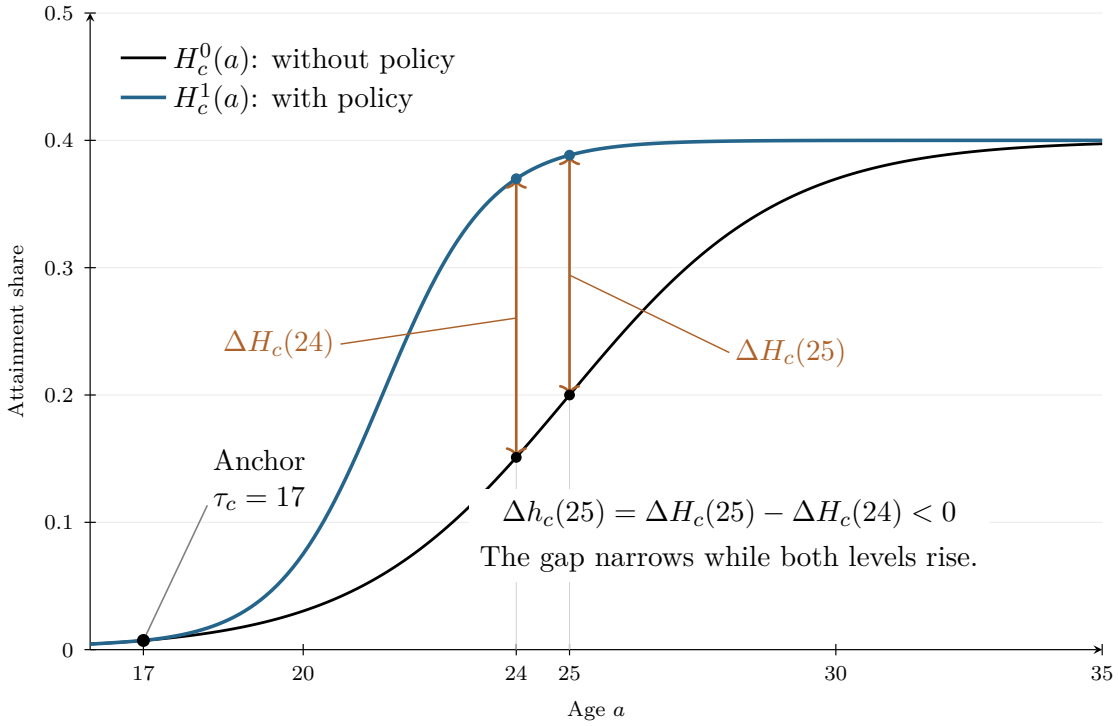
Figures

Figure 1: Education profiles for a fixed set of birth cohorts



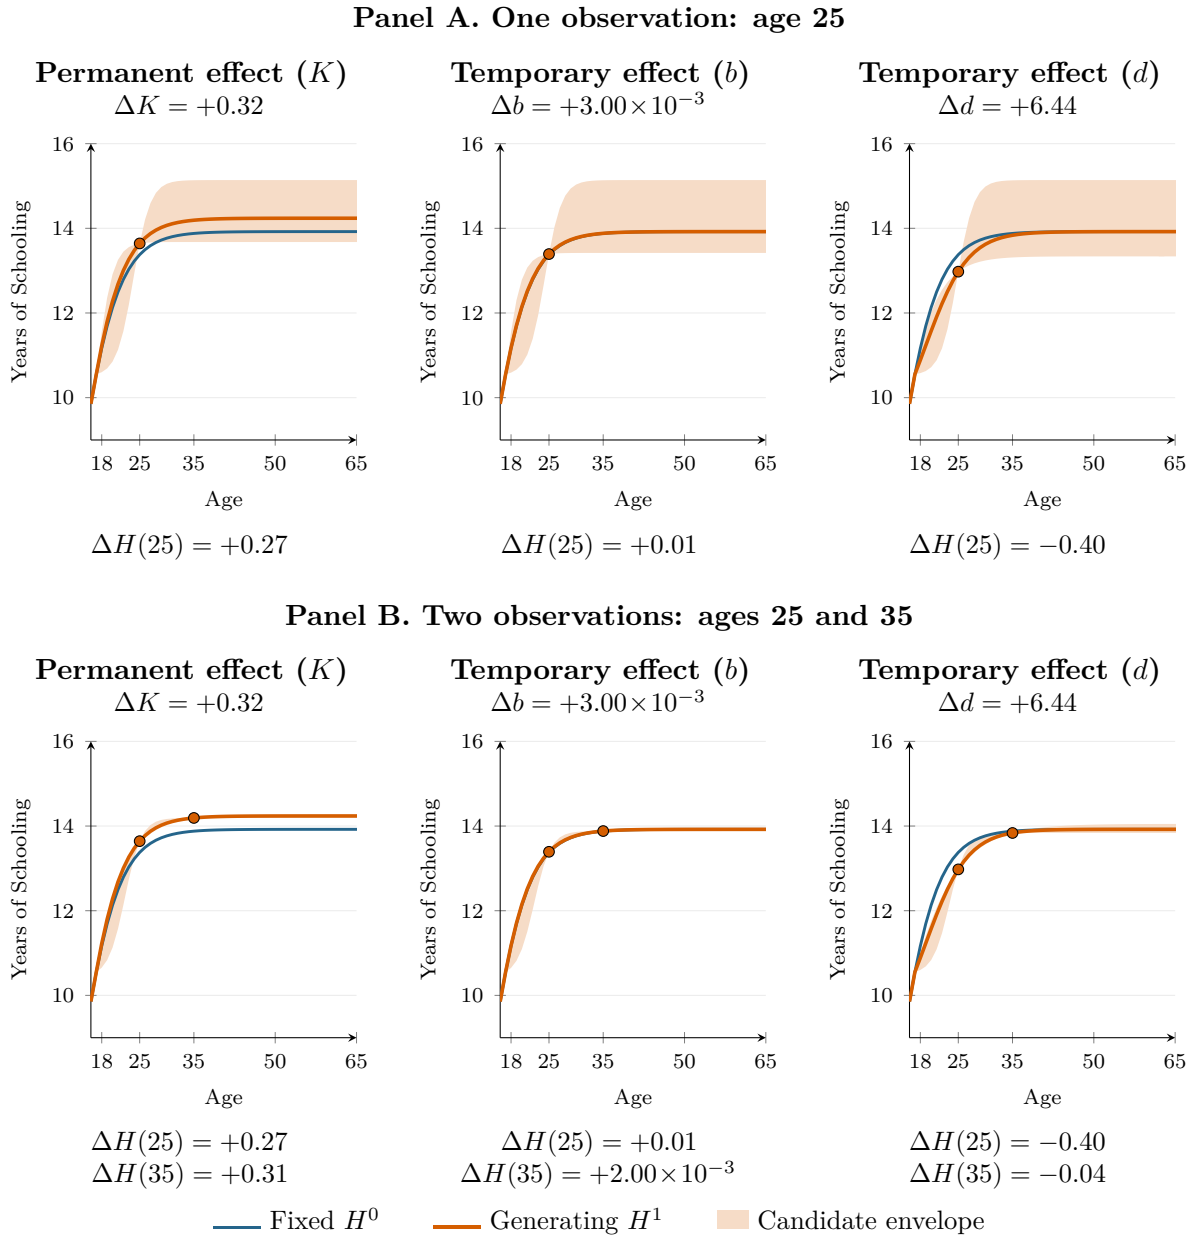
Notes: ACS 2001–2022, birth cohorts 1985–1987, ages 16–35; all sexes, birthplaces, and racial and ethnic groups are included. Person weights apply within cohort–age cells; cohorts receive fixed one-third weights. Repeated cross sections do not fix composition within birth cohorts. SC+ denotes some college, including degree holders, and serves as an attendance proxy; AA and BA denote cumulative degree thresholds. Years of Schooling is measured using the crosswalk for detailed educational attainment (EDUCD). Markers show annual means; lines connect adjacent ages. Appendix A.2 documents these descriptive profiles.
Source: Author's calculations from ACS microdata.

Figure 2: Educational attainment, treatment effects, and the pre-policy anchor



Notes: The hypothetical reference is $H_c^0(a) = 0.4\Lambda[0.5(a-25)]$. The treated path uses Equation (2), $(K, b, d) = (0.4, 1, 21.5)$, and the same age-17 level. Arrows mark attainment gaps at 24 and 25. Both levels rise while the gap narrows, giving $\Delta h_c(25) < 0$; both plateaus equal 0.4. These are illustrative attainment shares, not estimated effects.

Figure 3: Compatible schooling paths with one and two observed ages



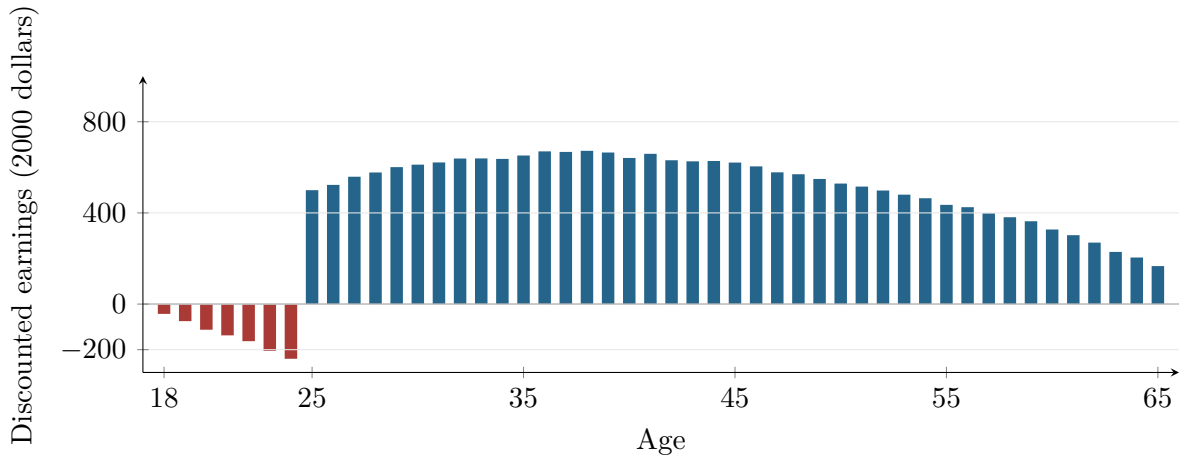
Notes: Panel A matches the age-25 effect; Panel B additionally matches the age-35 effect for the same generating process in each column. Columns vary one generating parameter at a time to the midpoint of its conditional feasible interval; the other two increments are zero. Permanent and temporary refer to the generating path, not all paths in its envelope. Blue lines show the fixed comparison path, and orange-red lines the generating path. Matching shading is the pointwise envelope of verified candidate paths, whose three parameters may all vary. Dots locate the treated levels implied by the observed effects and the fixed baseline. Every treated continuation is anchored to the comparison level at age 17. All six plots use identical age and schooling scales. Valuation results are reported in Table 1. Envelope edges need not form a single feasible path. No probability distribution is imposed on candidate paths, so the shading carries no probability interpretation. The envelopes are not confidence bands or certified global envelopes. Baseline values above age 35 are extrapolated.

Figure 4: Single-effect and full-path valuation with a permanent schooling gain

$\Delta H_c(25) = 0.48$ years — 2000 dollars; 3% discounting to age 18

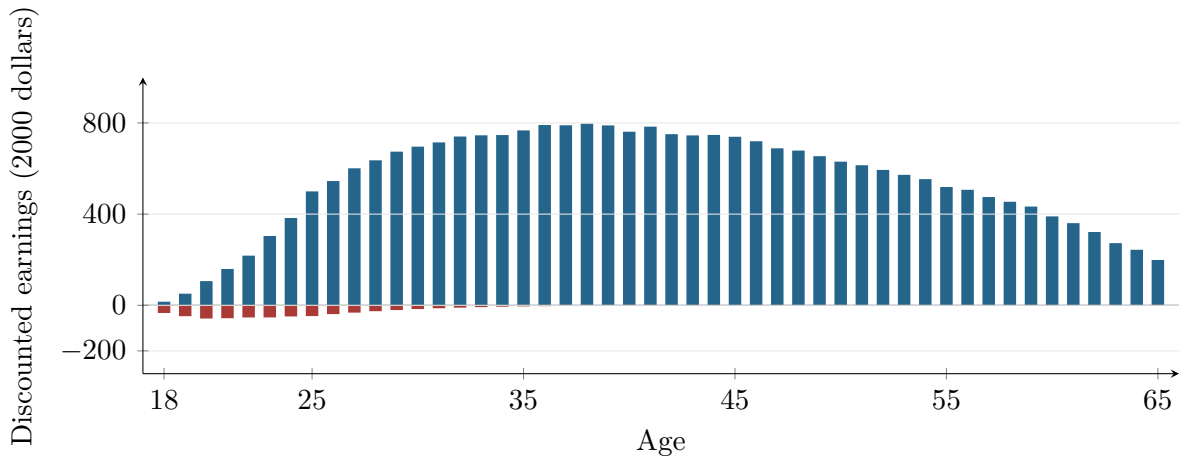
A. Single-effect benchmark — NPV: \$20,354.29

Opportunity cost PV: \$975.19



B. Full-path valuation — NPV: \$25,587.81

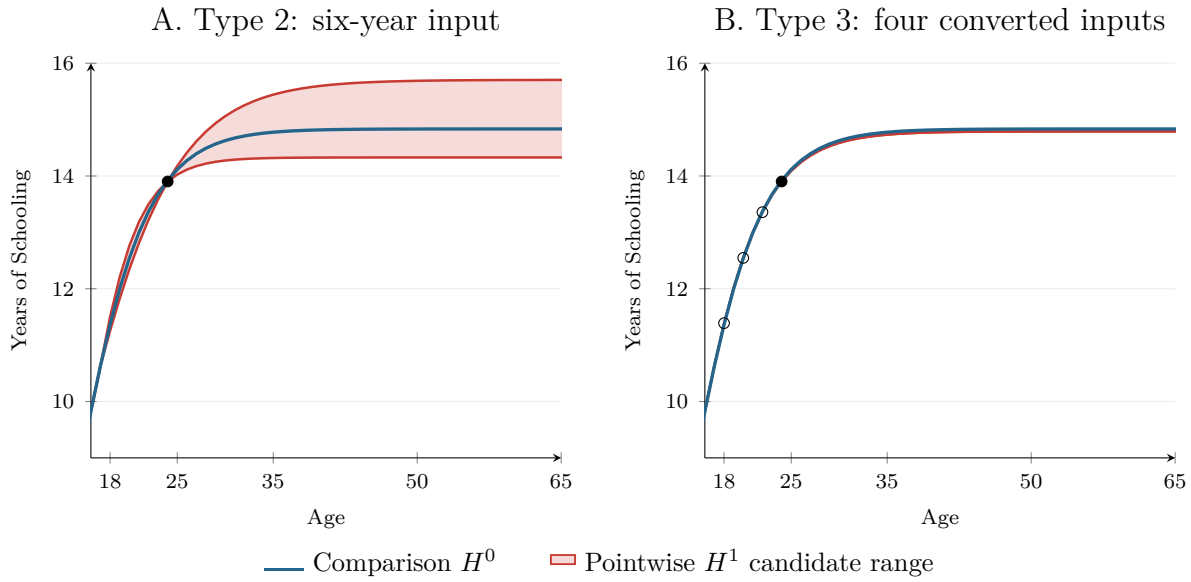
Opportunity cost PV: \$591.49



Education earnings return Opportunity cost

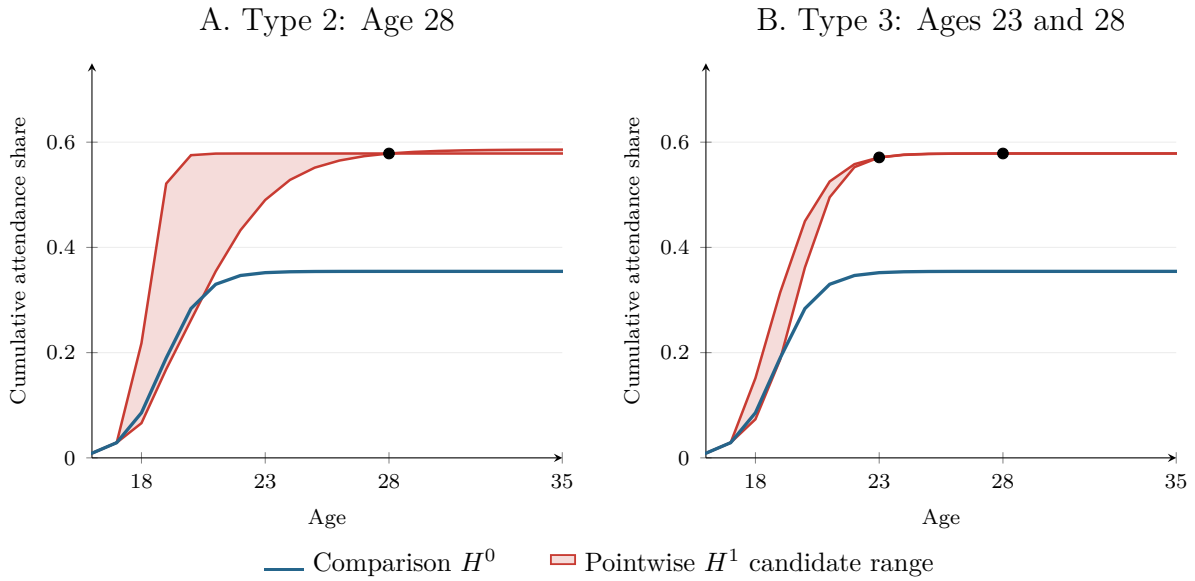
Notes: The full path changes only the plateau: $(\Delta K_c, \Delta b_c, \Delta d_c) = (0.57, 0, 0)$ relative to the registered reference parameters. Its schooling gap is 0.48 years at 25 and 0.58 at 65. Panel D of Figure B1 shows the level paths. Panel A assigns the age-25 effect to one completion batch at 25; Panel B uses the full attainment path. Blue bars show discounted education earnings returns; red bars show negative pre-completion contributions, labeled opportunity cost. The sum of both series is the reported NPV; the opportunity-cost PV is already deducted. Both panels use the same response template, ages 18–65, 2000 dollars, and 3 percent discounting to age 18. The paths share the age-17 anchor. Education paths extend to 72 to preserve the seven-year pre-completion window, with no income response before 18. This illustration is outside the two experiment grids.

Figure 5: Adams: adding three earlier converted education inputs



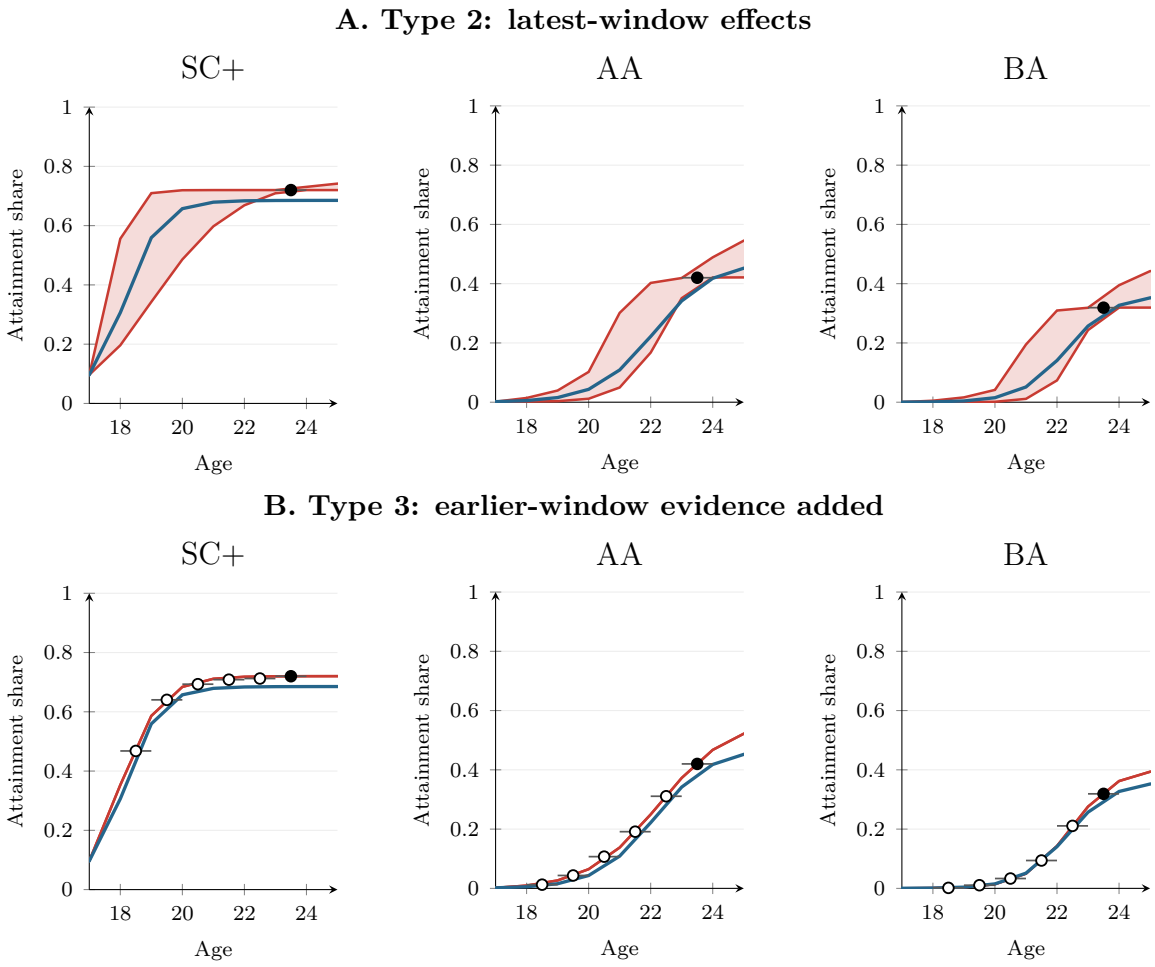
Notes: Type 2 matches -1.80×10^{-2} schooling-equivalent years at HS+6 (model age 24). Panel B fits the additional converted targets 2.60×10^{-2}, 7.00×10^{-3}, and -1.10×10^{-2} at ages 18, 20, and 22 while preserving the age-24 equality. The anchor is 17. Open markers indicate fitted targets; the filled marker indicates the exact restriction. Blue is H^0; red boundaries and shading give the pointwise H^1 candidate range. Its edges may combine different paths and are not confidence limits. Scales are common. Appendix D.2 documents the calibration and allocation.

Figure 6: Social Security: adding the earlier college-attendance observation



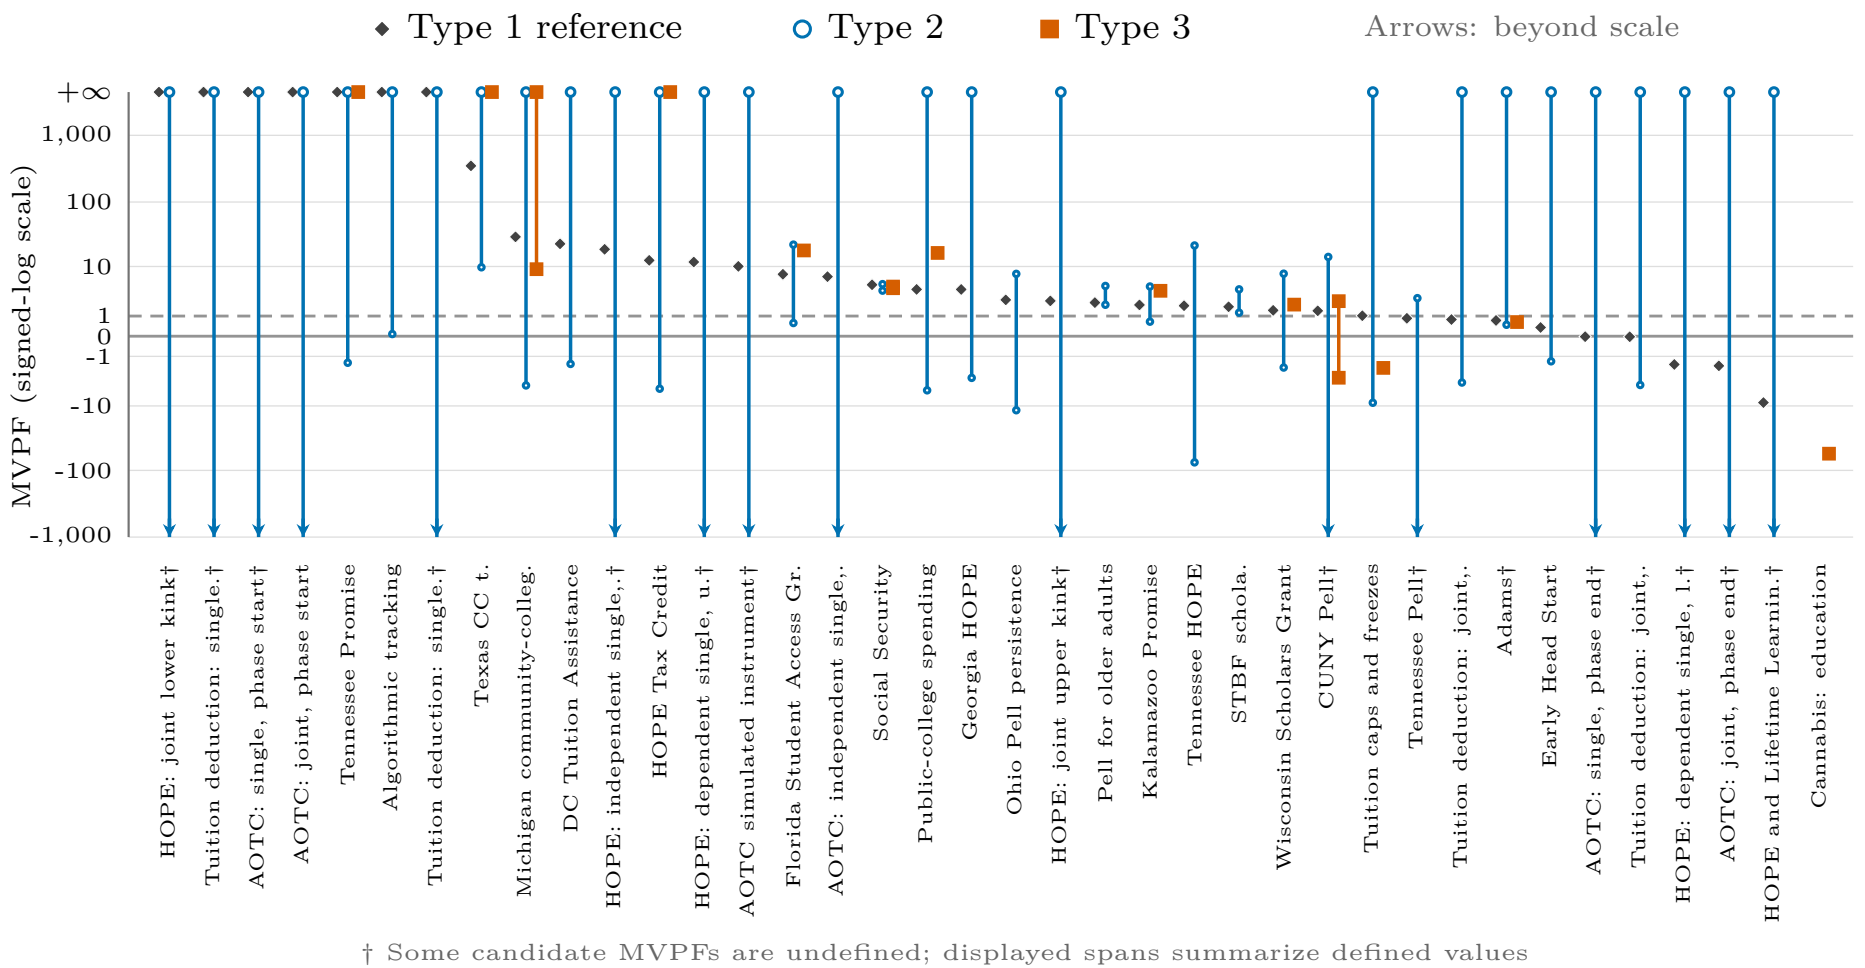
Notes: Type 2 matches the 22.40-percentage-point effect on ever attending college full time by age 28. Type 3 also matches the 21.90-percentage-point effect by age 23, keeping the age-28 equality. Filled markers show H^0 plus the exact source effects. Both panels use the same age-17 anchor, source age-23 control level, transported ACS entry shape, and probability and shape restrictions. The figure shows ages 16–35 for readability; valuation uses ages 18–65. The boundaries may combine different paths and are not confidence limits. Appendix D.2 documents the outcome proxy and calibration.

Figure 7: Tennessee Promise: Type 2 attainment ranges and the Type 3 fitted path



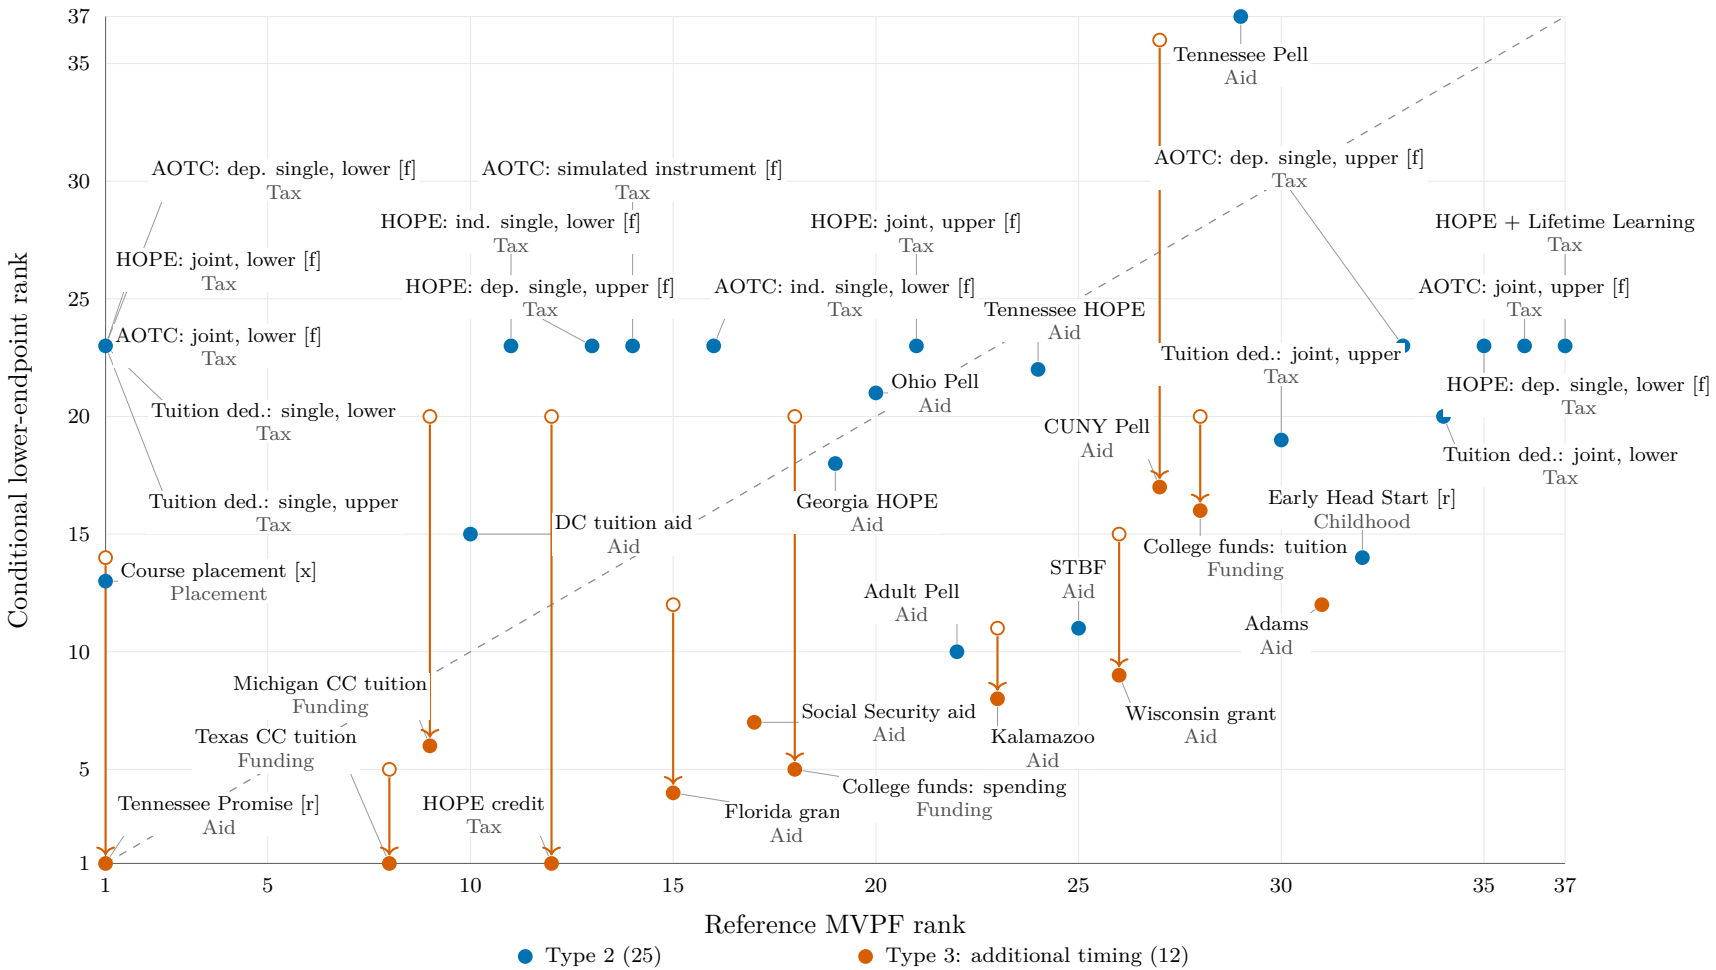
Notes: Blue is H^0; red boundaries and shading give the pointwise H^1 range among joint credential candidates. Filled markers show the exact 23–24 inputs; open markers show the fifteen earlier-window targets fitted in Type 3 from Schoonover et al. (2026). Marker heights equal the window mean of H^0 plus the reported effect. Horizontal bars span the source age windows; their midpoints locate windows, not single-age restrictions. Both panels use the same baseline, age-17 anchor, support, completion caps, and scales. Envelope edges may combine different paths across ages and credentials; they are not confidence bands. Appendix D.2 documents the restrictions; the appendix atlas shows ages through 65.

Figure 8: Education continuation ranges, ordered by the reference MVPF



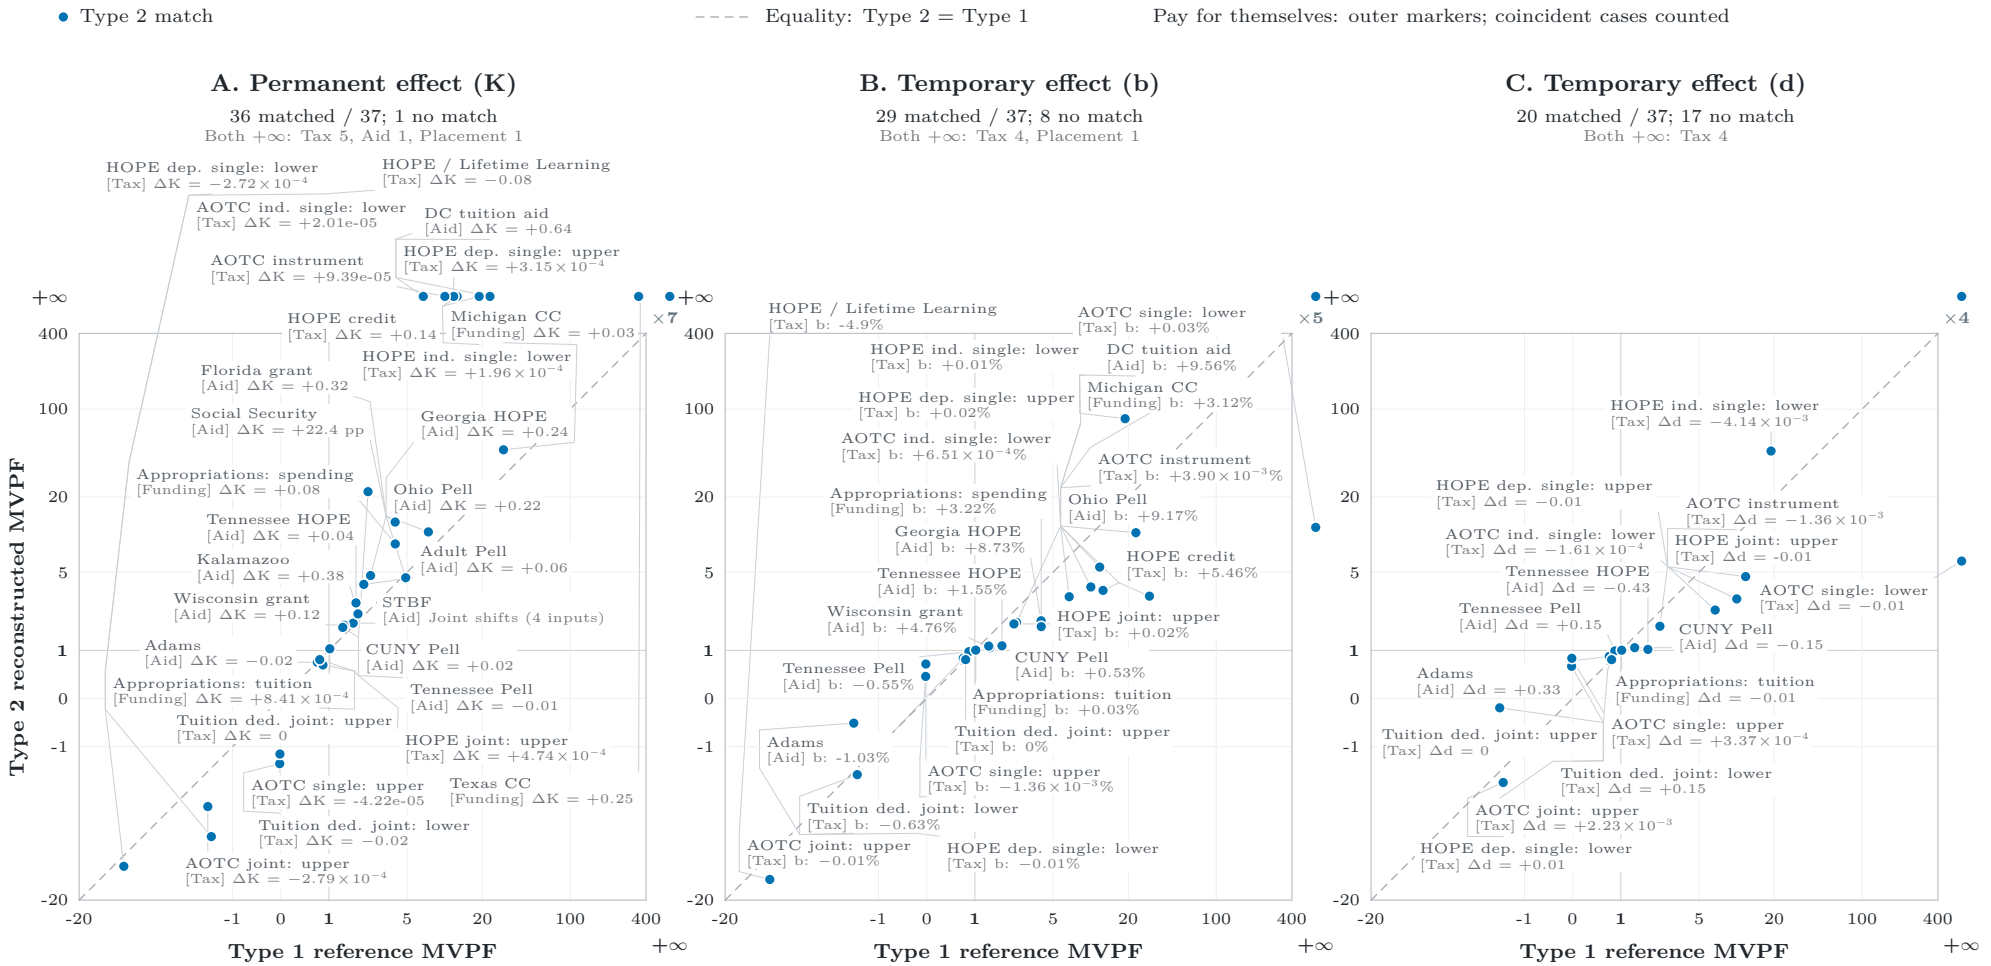
Notes: Policies follow descending Type 1. Gray diamonds denote Type 1; blue circles and orange squares mark Type 2 and Type 3 endpoints. The finite scale is signed-log; zero is solid and one is dashed. Positive-infinity endpoints are connected to their spans, including outcomes in which policies pay for themselves ($W > 0$, $N \leq 0$) and verified positive-cost ratio limits. Arrows denote values beyond the finite plotting scale. A connecting segment does not establish that every intermediate ratio is attained. Daggers mark undefined-ratio paths. Missing Type 3 means no added restriction is admitted under the retained mapping. Cannabis has no whole-policy Type 1 and shows primary Type 3. Ten Type 3 cases have coincident endpoints at one best-fitting path; the other three retain exact-match ranges. These summaries are not confidence intervals or certified bounds. Table 3 gives values, policy categories, and accounting qualifications. Dep./ind. denote dependent/independent filers; lower/upper denote income thresholds.

Figure 9: Type 1 and candidate-based lower-endpoint ranks for 37 policy calculations



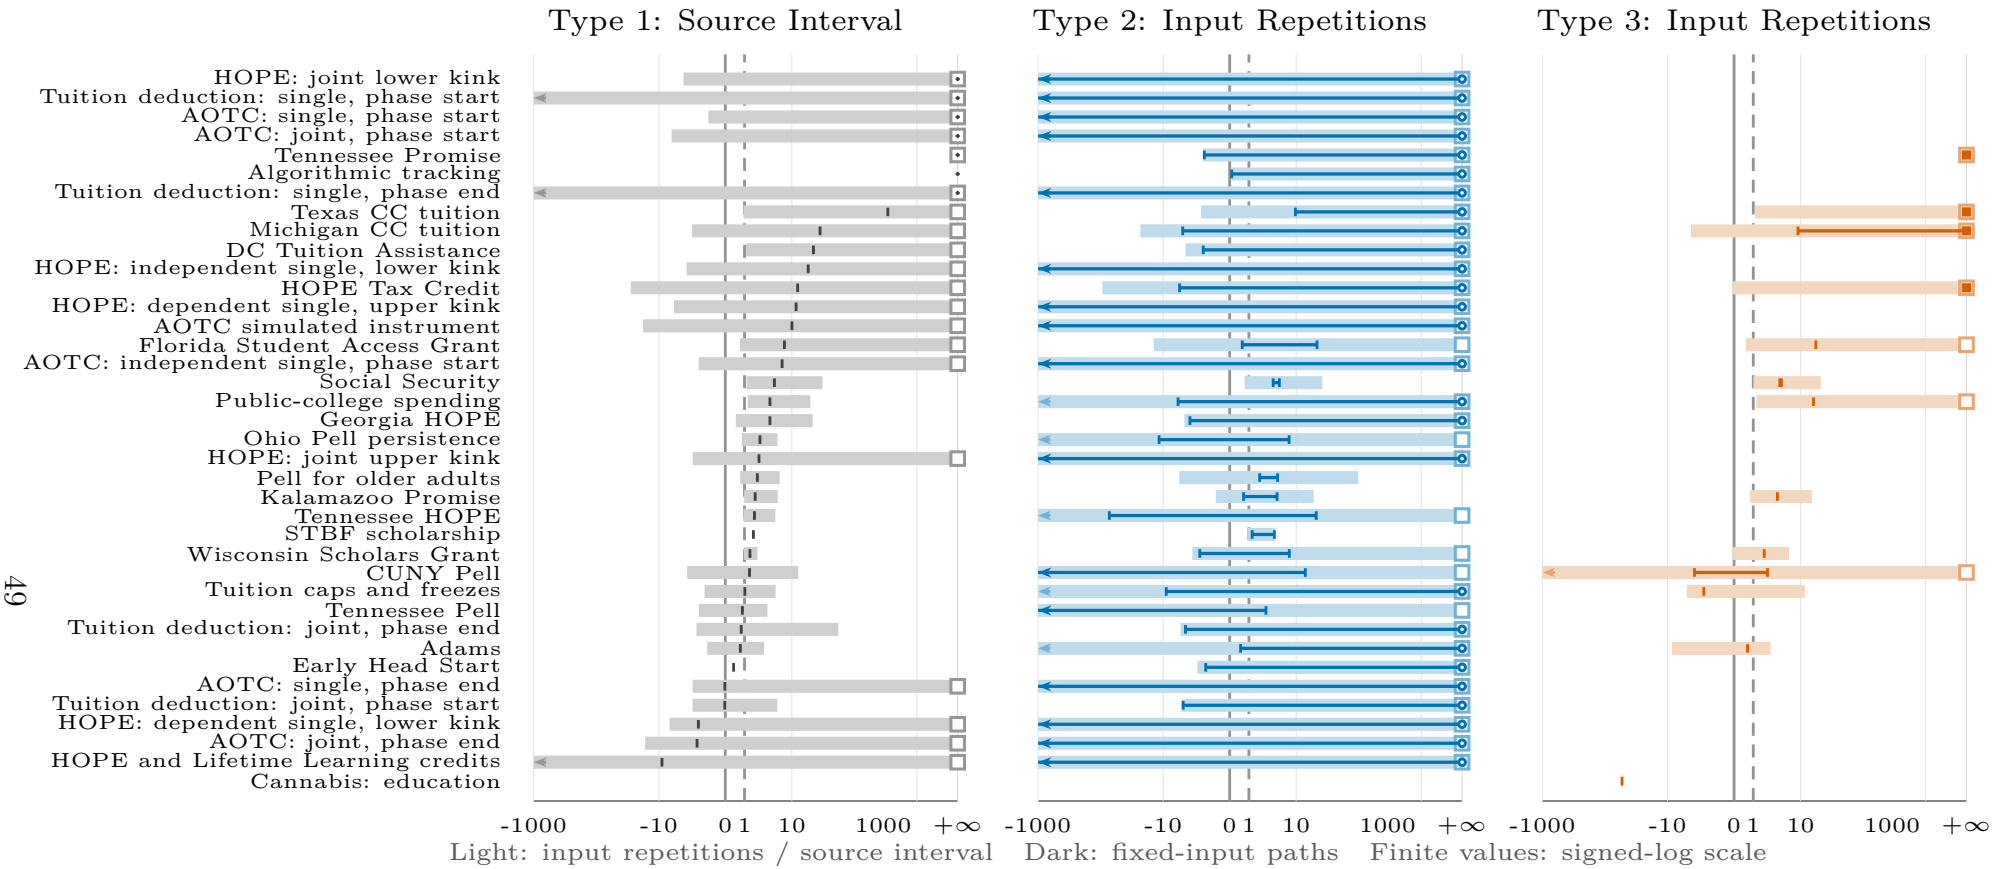
Notes: Higher lower endpoints rank first, followed by higher upper endpoints; admitted Type 3 replaces Type 2. Finite endpoints are rounded to two decimals, and equal pairs share a competition rank. Labels give policy names and Table 3 policy categories; [f], [r], and [x] mark tax conversion, public-data reconstruction, and a fiscal-only earnings extension. Coincident endpoint pairs share a marker. Arrows replace only the indicated policy's Type 2 by Type 3. Thirteen selected specifications contain undefined-ratio paths; those paths, not the calculations, are omitted. Cannabis lacks Type 1. Table D4 gives all ranks. This is a candidate-based display rule, not a social-welfare or certified worst-case ordering.

Figure 10: Reference and reconstructed MVPFs under permanent and temporary effects



Notes: Blue points match Type 2. None of these single-family solutions also matches every additional point input exactly; the main Type 3 fits remain separate. Bracketed labels give the policy types defined in Figure 8. Finite axes have identical asinh scales and square dimensions. Outer markers indicate policies that pay for themselves: $W > 0$, $N \leq 0$. Top markers preserve finite Type 1, right markers preserve finite Type 2, and corner counts identify coincident cases in which the policy pays for itself in both accounts; their policy types appear below the panel counts. Labels give ΔK in Years of Schooling or percentage points for credentials (STBF jointly shifts time and credential inputs), proportional b changes in percent, and Δd in years. The other two parameter families stay fixed. Permanent effects change eventual attainment; temporary effects eventually vanish but can retain lifetime value. Dep./ind. denote dependent/independent filers; lower/upper denote income thresholds. Appendix B.5 gives the method and case-level details.

Figure 11: Point-input ranges, source intervals, and conditional input repetitions



Notes: Dark marks show fixed-input results; pale bands show reported Type 1 intervals or conditional Type 2/Type 3 endpoint summaries. Positive infinity is the right axis endpoint, connected to each span with an infinite upper endpoint. Open squares mark infinite summary endpoints; dark symbols mark infinite fixed-input endpoints. Such endpoints include outcomes in which policies pay for themselves or verified positive-cost limits. An infinity singleton remains a single endpoint symbol. Arrows denote values beyond the finite plotting scale; zero is solid and one is dashed. A connecting segment does not establish that every intermediate ratio is attained. All thirteen Type 3 specifications include their full input vector. Tennessee's normal perturbations are not source permutation draws. Table 4 reports endpoints, usable repetition counts, and input-method codes; Table 3 reports accounting qualifications. Cannabis shows its primary Type 3 only. Reconstructed summaries have no demonstrated coverage for all continuations.

Tables

Table 1: Education valuation under one and two observed ages

Generating process	$\Delta H_c(25)$ (years)	Single-effect NPV	Generating-path NPV	Candidate NPV range	
				Age 25	Ages 25 and 35
Permanent (K)	0.27	11.20	14.08	[−3.53, 44.80]	[10.61, 14.33]
Temporary (b)	0.01	0.60	0.23	[−15.43, 42.58]	[−2.92, 0.46]
Temporary (d)	−0.40	−16.96	−6.07	[−26.41, 39.18]	[−8.74, −3.97]

Notes: NPVs are thousands of 2000 dollars per person. Calculations use unrounded inputs. Each generating scenario changes only the named parameter; the other two increments are zero. The increments are $(0.32, 0, 0)$, $(0, 3.00\times10^{-3}, 0)$, and $(0, 0, 6.44)$, respectively. $\Delta H_c(25)$ is the implied schooling effect. The benchmark uses that effect; the generating-path value uses the full path. Permanent and temporary describe the generating path. All three parameters may vary among candidates matching the stated ages. Ranges summarize verified candidates, not confidence intervals or certified global bounds.

Table 2: How additional observation ages change the NPV range

Observed ages	Median width ($1,000)	Largest width ($1,000)	Reduction (percent)	Both signs found (of 5 scenarios)
Panel A. Plateau (ΔK)				
25	48.34	65.92	0.0	3/5
20, 25	20.78	23.32	57.0	1/5
24, 25	16.39	18.40	66.1	1/5
25, 30	7.63	8.65	84.2	0/5
25, 35	3.72	4.31	92.3	0/5
Panel B. Speed (Δb)				
25	58.01	72.12	0.0	4/5
20, 25	18.50	22.71	68.1	2/5
24, 25	14.57	21.10	74.9	2/5
25, 30	6.81	12.76	88.3	1/5
25, 35	3.38	7.60	94.2	1/5
Panel C. Timing (Δd)				
25	65.59	91.01	0.0	5/5
20, 25	49.39	49.51	33.5	4/5
24, 25	40.51	47.01	38.2	4/5
25, 30	15.76	38.76	76.0	4/5
25, 35	4.77	13.54	92.7	1/5

Notes: Each panel varies only the named generating parameter; the other two increments are zero. Candidate parameters may all vary. Median widths and percentage reductions use the three interior interval positions (25, 50, and 75 percent). Reductions are calculated within each scenario relative to age 25 alone before taking the median. Largest widths and sign counts use all five positions, including the boundaries. Both signs found requires verified NPV candidates above $0.01 and below −$0.01. A zero count does not establish that such paths are absent from the continuous feasible set. Dollar units are 2000 dollars per person; the ranges are conditional numerical candidates.

Table 3: Reference welfare values, continuation ranges, and Type 3 best fits

Policy	Category	MVPF			WTP			Net fiscal cost		
		Type 1	Type 2	Type 3	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
Panel A. Primary comparisons: WTP and net cost per reference program dollar										
Tuition deduction: single, lower threshold	Tax	$+\infty^{s}$	$[-\infty, -7.40]^{s,\dagger}$	–	5.38	[−345.81, 583.82]	–	−5.10	[−93.18, 46.75]	–
Tennessee Promise^{r}	Aid	$+\infty^{s}$	$[-1.49, +\infty]^{s}$	$+\infty^{s}$	4.92	[−3.29, 10.31]	5.84	−0.68	[−2.57, 2.21]	−1.00
Tuition deduction: single, upper threshold	Tax	$+\infty^{s}$	$[-\infty, -6.33]^{s,\dagger}$	–	1.00	[−236.30, 372.31]	–	−1.13	[−60.20, 37.33]	–
Texas CC tuition	Funding	349.51	$[9.70, +\infty]^{s}$	$+\infty^{s}$	9.31	[5.87, 24.13]	10.87	0.03	[−1.63, 0.61]	−0.08
Michigan CC tuition	Funding	29.46	$[-4.44, +\infty]^{s}$	$[9.03, +\infty]^{s}$	7.80	[−35.00, 96.71]	[5.99, 14.42]	0.26	[−13.56, 7.90]	[−0.59, 0.66]
DC tuition assistance	Aid	22.98	$[-1.60, +\infty]^{s}$	–	7.62	[−2.62, 40.96]	–	0.33	[−2.89, 1.64]	–
HOPE tax credit	Tax	12.58	$[-5.09, +\infty]^{s}$	$+\infty^{s}$	5.23	[−19.03, 59.43]	29.30	0.42	[−6.58, 3.74]	−2.60
Florida access grant	Aid	7.42	[0.57, 22.32]	18.09	7.40	[0.67, 31.22]	23.02	1.00	[1.16, 1.40]	1.27
Social Security student benefits	Aid	4.87	[3.79, 4.99]	[4.21, 4.52]	3.03	[2.54, 3.08]	[2.74, 2.88]	0.62	[0.62, 0.67]	[0.64, 0.65]
College appropriations: spending	Funding	4.00	$[-5.43, +\infty]^{s}$	16.52	3.03	[−18.53, 42.71]	6.53	0.76	[−4.03, 3.41]	0.40
Georgia HOPE	Aid	4.00	$[-3.19, +\infty]^{s}$	–	3.12	[−5.43, 23.20]	–	0.78	[−1.15, 1.70]	–
Ohio Pell grants	Aid	2.49	[−11.75, 7.53]	–	2.88	[−5.34, 27.29]	–	1.16	[0.45, 3.62]	–
Adult Pell introduction	Aid	2.18	[1.95, 4.63]	–	2.71	[1.61, 13.45]	–	1.25	[0.83, 2.91]	–
Kalamazoo Promise	Aid	1.93	[0.65, 4.53]	3.77	1.76	[0.63, 4.12]	3.48	0.91	[0.91, 0.97]	0.92
Tennessee HOPE	Aid	1.86	[−75.64, 21.66]	–	1.89	[−29.05, 44.95]	–	1.02	[0.38, 2.08]	–
STBF scholarship	Aid	1.75	[1.25, 4.01]	–	1.52	[1.19, 2.50]	–	0.87	[0.62, 0.95]	–
Wisconsin Scholars Grant	Aid	1.43	[−1.94, 7.58]	1.96	1.42	[−2.13, 6.60]	1.94	0.99	[0.87, 1.10]	0.99
City University of New York (CUNY) Pell grants	Aid	1.39	$[-\infty, 14.28]^{\dagger}$	[−3.17, 2.34]	1.42	[−27.52, 39.88]	[−2.69, 2.54]	1.02	[−0.14, 2.79]	[0.85, 1.09]
College appropriations: tuition	Funding	1.02	$[-8.86, +\infty]^{s}$	−1.97	1.02	[−22.45, 38.90]	−2.41	1.00	[−1.42, 2.53]	1.22
Tennessee Pell eligibility	Aid	0.85	$[-\infty, 2.70]^{\dagger}$	–	0.79	[−18.13, 27.17]	–	0.93	[−5.34, 10.06]	–
Tuition deduction: joint, upper threshold	Tax	0.77	$[-3.92, +\infty]^{s}$	–	1.00	[−120.89, 191.27]	–	1.29	[−44.89, 30.83]	–
Adams Scholarship	Aid	0.72	$[0.48, +\infty]^{s,\dagger}$	0.63	1.21	[−10.19, 6.55]	1.47	1.68	[−18.52, 13.39]	2.34
Tuition deduction: joint, lower threshold	Tax	−0.02	$[-4.35, +\infty]^{s}$	–	−0.03	[−143.15, 215.89]	–	1.38	[−46.06, 32.89]	–
HOPE and Lifetime Learning credits	Tax	−8.81	$[-\infty, -9.59]^{s,\dagger}$	–	−42.82	[−199.55, 199.21]	–	4.86	[−19.57, 20.78]	–
Panel B. Finite tax reconstructions: per reference program dollar										
HOPE: joint, lower kink^{f}	Tax	$+\infty^{s}$	$[-\infty, +\infty]^{s,\dagger}$	–	18.41	[−21,373.75, 40,380.60]	–	−3.15	[−4,694.42, 2,529.55]	–
American Opportunity Tax Credit (AOTC): dependent single, lower kink^{f}	Tax	$+\infty^{s}$	$[-\infty, +\infty]^{s,\dagger}$	–	23.39	[−40,937.11, 64,893.24]	–	−1.61	[−6,442.87, 3,927.52]	–
AOTC: joint, lower kink^{f}	Tax	$+\infty^{s}$	$[-\infty, +\infty]^{s}$	–	9.96	[−53,630.62, 85,312.41]	–	−1.37	[−10,116.74, 6,349.58]	–

Continued on next page

Table 3 (continued)

Policy	Category	MVPF			WTP			Net fiscal cost		
		Type 1	Type 2	Type 3	Type 1	Type 2	Type 3	Type 1	Type 2	Type 3
HOPE: independent single, lower kink^{f}	Tax	18.86	$[-\infty, +\infty]^{s,\dagger}$	–	5.91	[−53,768.73, 101,434.66]	–	0.31	[−17,374.29, 9,427.43]	–
HOPE: dependent single, upper kink^{f}	Tax	11.83	$[-\infty, +\infty]^{s,\dagger}$	–	4.59	[−17,357.02, 32,615.58]	–	0.39	[−3,366.47, 1,855.80]	–
AOTC: simulated instrument^{f}	Tax	10.05	$[-\infty, +\infty]^{s,\dagger}$	–	5.36	[−82,841.87, 145,584.75]	–	0.53	[−16,160.55, 9,118.53]	–
AOTC: independent single, lower kink^{f}	Tax	6.75	$[-\infty, +\infty]^{s}$	–	2.45	[−179,550.49, 284,893.86]	–	0.36	[−56,780.46, 36,085.36]	–
HOPE: joint, upper kink^{f}	Tax	2.37	$[-\infty, +\infty]^{s,\dagger}$	–	8.21	[−23,368.86, 44,169.47]	–	3.47	[−4,875.76, 2,696.42]	–
AOTC: dependent single, upper kink^{f}	Tax	−0.02	$[-\infty, +\infty]^{s,\dagger}$	–	−0.02	[−42,932.08, 67,915.87]	–	1.12	[−7,231.21, 4,415.06]	–
HOPE: dependent single, lower kink^{f}	Tax	−1.64	$[-\infty, +\infty]^{s,\dagger}$	–	−1.91	[−16,274.53, 30,497.00]	–	1.16	[−2,764.33, 1,548.51]	–
AOTC: joint, upper kink^{f}	Tax	−1.77	$[-\infty, +\infty]^{s,\dagger}$	–	−7.63	[−55,169.57, 87,656.11]	–	4.31	[−10,682.18, 6,711.54]	–
Panel C. Reconstructed or extended accounts and education channel: native thousands										
Algorithmic college placement^{x}	Placement	$+\infty^{s}$	$[0.08, +\infty]^{s}$	–	0.22	0.22	–	−0.09	[−5.55, 2.82]	–
Early Head Start^{r}	Childhood	0.35	$[-1.37, +\infty]^{s}$	–	3.91	[−37.51, 114.60]	–	11.22	[−1.10, 27.35]	–
Cannabis legalization: education channel	Other	–	–	−55.85	–	–	−2.50	–	–	0.04

Policy types: Aid, direct student financial aid; Tax, education tax benefits; Funding, tuition and public-college funding; Placement, college assessment and course placement; Childhood, early childhood services; Other, the education channel of a noneducation policy. Types describe interventions, whereas panels distinguish accounting specifications. *Notes:* Type 1 reproduces source values. A single number denotes coincident endpoints; the ten minimum-distance Type 3 cases each retain one best-fitting path. Brackets denote distinct fixed-input endpoints, including limiting values, not confidence intervals or certified global bounds. WTP and net-cost endpoints need not use the same path; each positive-cost MVPF uses paired components from one complete path. Infinite range endpoints retain verified one-sided limits from finite schooling paths. Superscript s marks positive WTP and nonpositive net fiscal cost (WTP > 0, $N \leq 0$), for which the policy pays for itself and is assigned $+\infty$. On a range, s indicates at least one such path, not all paths. A dagger (†) marks additional verified paths with WTP ≤ 0 and $N \leq 0$. These paths receive no MVPF under the paper's convention and are excluded from the reported MVPF span, but remain included in the WTP and net-cost summaries. Neither symbol denotes statistical significance or a failed search. Superscripts f, r, and x identify the tax conversion, public-data reconstruction, and earnings extension. Panel A divides both monetary components by one fixed reference program-cost component per case. Panel B uses the same bounded logistic family. Source attendance and monetary inputs are transported to a fixed finite schedule change calibrated to one point-estimate credit dollar, retaining the two-year schooling conversion. The conversion and its age placement are assumptions. The first-stage point normalization stays fixed across repetitions. The AOTC simulated-instrument reconstruction also corrects the direct-credit unit while retaining the reported Type 1. Panel C uses thousands of native dollars. Early Head Start uses unscaled public ACS income and a funded-slot cost proxy, in 2000 dollars. Its source price year is unresolved. Tennessee uses public ACS income and an account-implied participant bridge, under a disclosed 2021-dollar convention. Algorithmic placement uses the verified June 2026 source components, but its reconstruction extends fiscal earnings only, keeping source WTP fixed. Their monetary differences cannot be attributed only to continuation. Cannabis has no whole-policy Type 1 and uses its three-horizon Type 3 as the primary result; its last-horizon-only Type 2 remains in the appendix data. A dash denotes an unavailable reference value or an inapplicable computation type; missing Type 3 means no eligible additional source-window restriction. NF denotes an unsuccessful search. Duration allocations alone do not qualify as Type 3. Appendix D.4 documents units, normalization, and account-version qualifications.

Table 4: Reported intervals and conditional input-repetition endpoint summaries

Policy	Type 1: reported 95%	Type 2: endpoints	n_2	Type 3: endpoints	n_3	Scope
HOPE: joint, lower kink	$[-3.51, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
Tuition deduction: single, lower threshold	$[-\infty, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
AOTC: dependent single, lower kink	$[-0.84, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
AOTC: joint, lower kink	$[-5.96, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
Tennessee Promise	$[+\infty, +\infty]$	$[-1.62, +\infty]$	1000	$[+\infty, +\infty]$	1000	P
Algorithmic college placement	–	$[0.07, +\infty]$	1000	–	–	P
Tuition deduction: single, upper threshold	$[-\infty, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
Texas CC tuition	$[0.93, +\infty]$	$[-1.81, +\infty]$	1000	$[1.14, +\infty]$	1000	J
Michigan CC tuition	$[-2.33, +\infty]$	$[-24.02, +\infty]$	1000	$[-3.74, +\infty]$	1000	J
DC tuition assistance	$[1.12, +\infty]$	$[-3.93, +\infty]$	1000	–	–	X
HOPE: independent single, lower kink	$[-3.03, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
HOPE tax credit	$[-29.15, +\infty]$	$[-97.46, +\infty]$	1000	$[-0.04, +\infty]$	1000	J
HOPE: dependent single, upper kink	$[-5.36, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
AOTC: simulated instrument	$[-18.36, +\infty]$	$[-\infty, +\infty]$	1000	–	–	P
Florida access grant	$[0.69, +\infty]$	$[-14.48, +\infty]$	1000	$[0.53, +\infty]$	1000	J
AOTC: independent single, lower kink	$[-1.61, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
Social Security student benefits	[1.16, 32.23]	[0.72, 27.30]	1000	[1.01, 21.68]	980*	J
College appropriations: spending	[1.25, 20.44]	[−1,925.69, $+\infty$]	1000	$[1.23, +\infty]$	1000	J
Georgia HOPE	[0.46, 22.38]	$[-4.11, +\infty]$	1000	–	–	J
Ohio Pell grants	[0.82, 5.54]	$[-\infty, +\infty]$	1000	–	–	J
HOPE: joint, upper kink	$[-2.22, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J
Adult Pell introduction	[0.71, 6.11]	[−5.15, 102.12]	1000	–	–	J
Kalamazoo Promise	[0.97, 5.61]	[−0.65, 19.65]	1000	[0.78, 15.51]	1000	J
Tennessee HOPE	[0.92, 5.08]	$[-\infty, +\infty]$	1000	–	–	J
STBF scholarship	–	[0.94, 4.14]	1000	–	–	P
Wisconsin Scholars Grant	[0.96, 2.18]	$[-2.84, +\infty]$	1000	[−0.04, 6.29]	1000	J
CUNY Pell grants	[−2.95, 12.88]	$[-\infty, +\infty]$	1000	$[-\infty, +\infty]$	922*	J
College appropriations: tuition	[−1.11, 5.13]	$[-\infty, +\infty]$	1000	[−4.48, 11.82]	997*	J
Tennessee Pell eligibility	[−1.59, 3.57]	$[-\infty, +\infty]$	1000	–	–	J
Tuition deduction: joint, upper threshold	[−1.83, 57.84]	$[-4.87, +\infty]$	1000	–	–	J
Adams Scholarship	[−0.92, 3.05]	$[-\infty, +\infty]$	1000	[−8.38, 2.73]	931*	J
Early Head Start	–	$[-2.20, +\infty]$	1000	–	–	P
AOTC: dependent single, upper kink	$[-2.25, +\infty]$	$[-\infty, +\infty]$	1000	–	–	J

Continued on next page

Table 4 (continued)

Policy	Type 1: reported 95%	Type 2: endpoints	n_2	Type 3: endpoints	n_3	Scope
Tuition deduction: joint, lower threshold	$[-2.25,\ 5.52]$	$[-4.67,\ +\infty]$	1000	–	–	J
HOPE: dependent single, lower kink	$[-6.45,\ +\infty]$	$[-\infty,\ +\infty]$	1000	–	–	J
AOTC: joint, upper kink	$[-17.06,\ +\infty]$	$[-\infty,\ +\infty]$	1000	–	–	J
HOPE and Lifetime Learning credits	$[-\infty,\ +\infty]$	$[-\infty,\ +\infty]$	1000	–	–	J
Cannabis legalization: education channel	–	–	–	$[-57.03,\ -54.19]$	1000	R

Notes: Type 1 reports available source 95 percent intervals. Type 2/Type 3 give the 2.5th percentile of candidate minima and the 97.5th percentile of maxima under 1,000 paired input repetitions. Each of the thirteen Type 3 schemes includes its complete input vector. Counts are repetitions with verified paths and a defined ratio; asterisks denote fewer than 1,000, and dashes denote no admitted layer or source interval. These summaries do not establish confidence coverage. Infinite endpoints include outcomes in which policies pay for themselves or verified positive-cost limits. J: archived joint source/accounting draws, extended to added inputs; P: constructed education perturbations; X: economic inputs with education fixed; R: synthetic correlated shocks. All schemes condition on the comparison profile, calibration, and stated dependence. Source intervals can vary different inputs. Appendix E gives details. Cannabis shows only its primary three-window Type 3. Policy types and superscripts follow Table 3.

Table 5: Reference and single-parameter reconstructed MVPFs

	Reference	Permanent (K)		Temporary (b)		Temporary (d)	
Policy	MVPF	ΔK	MVPF	Δb (%)	MVPF	Δd	MVPF
Florida grant	7.42	0.32	10.54	—	—	—	—
DC tuition aid	22.98	0.64	$+\infty$	9.56	10.38	—	—
Adams	0.72	-0.02	0.72	-1.03	0.81	0.33	0.85
STBF	1.75	Joint	1.86	—	—	—	—
Wisconsin grant	1.43	0.12	1.78	4.76	1.07	—	—
Kalamazoo	1.93	0.38	2.25	—	—	—	—
Michigan CC	29.46	0.03	47.45	3.12	3.18	—	—
CUNY Pell	1.39	0.02	1.70	0.53	1.11	-0.15	1.07
Ohio Pell	2.49	0.22	4.70	9.17	1.89	—	—
Tennessee HOPE	1.86	0.04	2.79	1.55	1.12	-0.43	1.03
Tennessee Pell	0.85	-0.01	0.66	-0.55	0.97	0.15	0.99
Georgia HOPE	4.00	0.24	12.63	8.73	1.96	—	—
Social Security	4.87	22.4	4.51	—	—	—	—
Appropriations: spending	4.00	0.08	8.45	3.22	1.73	—	—
Early Head Start	0.35	—	—	—	—	—	—
Texas CC	349.51	0.25	$+\infty$	—	—	—	—
Adult Pell	2.18	0.06	3.98	—	—	—	—
Appropriations: tuition	1.02	8.41×10^{-4}	1.04	0.03	1.01	-0.01	1.00
HOPE credit	12.58	0.14	$+\infty$	5.46	3.55	—	—
HOPE ind. single: lower	18.86	1.96×10^{-4}	$+\infty$	0.01	83.88	-4.14×10^{-3}	46.41
HOPE dep. single: lower	-1.64	-2.72×10^{-4}	-6.21	-0.01	-1.89	0.01	-2.22
HOPE dep. single: upper	11.83	3.15×10^{-4}	$+\infty$	0.02	5.50	-0.01	4.61
HOPE joint: lower	$+\infty$	1.24×10^{-3}	$+\infty$	0.06	$+\infty$	-0.03	$+\infty$
HOPE joint: upper	2.37	4.74×10^{-4}	22.04	0.02	1.83	-0.01	1.74
Tuition ded. single: lower	$+\infty$	0.04	$+\infty$	1.05	$+\infty$	-0.26	$+\infty$
Tuition ded. joint: lower	-0.02	-0.02	-1.49	-0.63	0.42	0.15	0.63
HOPE / Lifetime Learning	-8.81	-0.08	-10.75	-4.90	-13.68	—	—
AOTC single: upper	-0.02	-4.22×10^{-5}	-1.20	-1.36×10^{-3}	0.68	3.37×10^{-4}	0.81
AOTC single: lower	$+\infty$	9.72×10^{-4}	$+\infty$	0.03	11.45	-0.01	6.15
AOTC joint: upper	-1.77	-2.79×10^{-4}	-3.54	-0.01	-0.47	2.23×10^{-3}	-0.17
AOTC joint: lower	$+\infty$	2.98×10^{-4}	$+\infty$	0.01	$+\infty$	-2.38×10^{-3}	$+\infty$
AOTC ind. single: lower	6.75	2.01×10^{-5}	$+\infty$	6.51×10^{-4}	3.15	-1.61×10^{-4}	2.43
AOTC instrument	10.05	9.39×10^{-5}	$+\infty$	3.90×10^{-3}	3.79	-1.36×10^{-3}	3.02
Tennessee Promise	$+\infty$	Joint	$+\infty$	—	—	—	—
College placement	$+\infty$	0.04	$+\infty$	3.41	$+\infty$	—	—
Tuition ded. single: upper	$+\infty$	0.00	$+\infty$	0.00	$+\infty$	0.00	$+\infty$
Tuition ded. joint: upper	0.77	0.00	0.77	0.00	0.77	0.00	0.77

Notes: Each pair varies the named parameter family; the other two remain at H^0 values. K changes eventual attainment; b and d effects eventually vanish. ΔK is in Years of Schooling, except percentage points for Social Security and Early Head Start; Δd is in years and Δb is the proportional speed change in percent. STBF jointly changes time enrolled and three credential plateaus; Tennessee jointly changes three credential plateaus. Their outcome-specific changes are saved in the matched-solution data. No single-family solution matches every additional point input exactly; — denotes no admissible Type 2 match. Here, $+\infty$ denotes cases in which policies pay for themselves ($W > 0$, $N \leq 0$), as in Figure 10. The table includes all 37 reference calculations; cannabis has no whole-policy reference. Appendix B.5 defines the matching procedure. Dep./ind. denote dependent/independent filers; lower/upper denote income thresholds.

Appendix

A Literature Review and Descriptive Education Profiles

A.1 Literature Evidence for Lifetime Valuation

The review contains 459 distinct studies: 447 retained after full-text screening and deduplication of a selected search sample from 12 economics journals over 2000–2025, 11 additional welfare-source studies, and one author manuscript (Li 2026).[39] Inclusion requires an original estimate relating a policy, intervention, or external exposure to a higher-education outcome. Studies reporting only test scores, aspirations, earnings, or HS outcomes are excluded, as are purely descriptive studies and reviews. Versions of the same study count once; there is no maximum follow-up length. The replication register records inclusion decisions, source versions, and evidence locations. The counts are based on the recorded source versions, which can differ from the published articles cited in the main text. The sample is selected, with the additional studies chosen for their relevance to welfare valuation. Within this sample, the outcome groups distinguish some college, AA, and BA, with papers allowed to contribute to several groups. The some college category covers first entry or ever attending college; current enrollment or persistence alone does not qualify. AA retains each source's degree definition, while the overall group also includes credits and other education measures. The timing count distinguishes observation times within a comparable education series: effects reported separately at years 3 and 6 count twice, whereas one estimate pooled over years 3–6 counts once. Alternative specifications, treatment arms, and subgroups at the same horizon do not increase the timing count. Estimates from independent samples or different time origins are not combined into one sequence.

Figure A2 and Table A1, Panel A, report the counts. Of the 459 studies, 249 report one time or window; among 190 BA studies, 132 do so. The BA share is 69.5 percent overall and 69.9 percent in the journal sample. These statistics measure the number of reported windows, each of which may span several years. Figures A3 and A4 relate these counts to observation endpoints, separately for

[39]The journals are the *American Economic Review*, *Quarterly Journal of Economics*, *Journal of Political Economy*, *Review of Economic Studies*, *Review of Economics and Statistics*, *American Economic Journal: Applied Economics*, *American Economic Journal: Economic Policy*, *Journal of Human Resources*, *Journal of Labor Economics*, *Journal of Public Economics*, *Economics of Education Review*, and *Education Finance and Policy*.

attained age, individual follow-up, and policy or cohort event time. Follow-up is measured from a stated origin such as HS exit or first enrollment; six years of follow-up cannot be assigned age 24 without entry-age information. Event time may compare successive cohorts rather than follow the same people. The figures include 366 studies with an exact endpoint on at least one clock; the other 93 remain in the review without imputed coordinates. Bubble areas represent paper counts, and correlations are calculated from individual study records. The overall panels can summarize different education outcomes within a study; the outcome-specific panels restrict the comparison to the recorded category. Both figures describe reporting patterns rather than causal relationships.

Among the reviewed studies, 31 quantitatively project education effects into lifetime or career-long earnings, fiscal returns, benefits relative to costs, or MVPFs. A complete welfare account is not required, but qualitative references to education returns and calculations confined to education costs over a few years do not qualify. Table A1, Panel B, classifies the education information available for these valuations. The count concerns comparable observations of the same education outcome. Different outcomes observed at different deadlines, earnings observations, and successive exposed cohorts do not constitute repeated observations of one attainment path. When several education outcomes enter a valuation, the classification uses the series with the most comparable observation times. Thus a calculation can have several education inputs while each series is observed only once. Eighteen studies report one age, follow-up, or stage-specific window per valued education series, and 12 report multiple times. The remaining study values institutional degree awards indexed by spending-shock time, which does not supply an individual completion deadline. A single pooled window can cover several adult ages, and multiple observations need not extend to final attainment. These counts describe the available education evidence; they do not establish which reported observations enter each original valuation formula. The linked 31-study record preserves outcomes, time origins, reported windows, source versions, and supporting exhibits. This literature count differs from the policy–estimand calculations classified in Section 2, where one study may support several welfare accounts.

A.2 ACS Education Profiles across Ages and Cohorts

The ACS profiles provide descriptive evidence for the introduction and the data underlying the baseline paths in Section 3. Figure 1 uses ACS microdata from 2001–2022 for the 1985–1987 birth cohorts at ages 16–35. Birth cohort is survey

year minus reported integer age. All 60 cohort–age cells are observed, yielding 1,978,860 respondent records from all 50 states and DC, without restrictions on birthplace, race, Hispanic origin, sex, or marijuana policy status. Positive person weights and valid EDUCD are required. Person weights apply within each cell, and each cohort receives one-third weight at every age. Let S_{ca} contain respondents in cohort c at age a, w_i be the Integrated Public Use Microdata Series (IPUMS) person weight, and H_{ik} be respondent i's education measure k. For a three-cohort group C_g, the profile is

$$\overline{H}_{gk}(a) = \frac{1}{3} \sum_{c \in C_g} \frac{\sum_{i \in S_{ca}} w_i H_{ik}}{\sum_{i \in S_{ca}} w_i}. \tag{16}$$

Figure 1 uses $C_g = \{1985, 1986, 1987\}$. Fixed cohort weights hold the mixture of birth cohorts constant; repeated cross sections can still differ in immigration, survival, and other within-cohort composition. The some college proxy includes degree holders and records attainment rather than current enrollment or the date of first entry. AA and BA are cumulative degree thresholds. The EDUCD crosswalk assigns 12 Years of Schooling to HS completion, 12.5 to less than one year of college, 13 to at least one year without a degree, 14 to an AA, 16 to a BA, 18 to a master's degree, 19 to a professional degree, and 20 to a doctorate. Lower categories use the reported grade or documented midpoint, making Years of Schooling an approximation from attainment categories. The crosswalk and cohort–age means are retained in the replication files. Appendix B.1 describes the subsequent logistic fits. Their use as untreated potential-outcome paths requires the calibration assumptions stated in the applications.

Figure A1 compares the 1975–1977, 1980–1982, and 1985–1987 birth cohorts using Census 2000 and ACS 2001–2022, with common sample filters, education measures, and weights.[40] Within this survey window, all three cohorts in each group are observed at ages 25–35, 20–35, and 16–35, respectively; missing earlier ages remain unplotted. The combined sample contains 4,626,027 respondent records in 141 cohort–age cells. All groups are observed at 25 and 35, permitting a common comparison of subsequent accumulation. BA attainment rises by 5.86, 9.73, and 11.46 percentage points across these groups, while approximate Years of Schooling rise by 0.42, 0.63, and 0.67 years. AA gains also increase across groups, whereas

[40] ACS 2001–2005 excludes group quarters, and the 2020 ACS uses experimental pandemic weights; IPUMS retains these in its person-weight variable. These differences limit comparisons across survey years. Common sample filters and fixed cohort weights do not remove them. See the IPUMS sample documentation and the Census guidance on 2020 ACS comparisons.

gains in the college-attendance proxy are not monotone. The comparison therefore concerns later degree attainment and schooling accumulation, rather than a uniform increase in every outcome. These profiles describe changes in cohort means, not individual completion histories or causal effects of birth year. Their implications for a policy effect depend on subsequent attainment differences between treated and comparison populations.

B Education-Path Framework and Simulation

B.1 Comparison Path and Simulation Calibration

The simulations in Section 3 use the 1985–1987 birth cohorts and the sample and education crosswalk documented in Appendix A.2. Let $\overline{\text{Years of Schooling}}_{qa}$ denote the person-weighted mean for cohort q at age a, and let ω_{qa} be the cell population weight normalized over ages 16–35 within that cohort. Each cohort is fitted separately by nonlinear weighted least squares:

$$(\widehat{K}_q, \widehat{b}_q, \widehat{d}_q) = \arg\min_{K,b,d} \sum_{a=16}^{35} \omega_{qa} \left[\overline{\text{Years of Schooling}}_{qa} - K\Lambda\{b(a-d)\}\right]^2. \quad (17)$$

The fitting bounds are $K \in [0.01, 25]$, $b \in [1.00 \times 10^{-3}, 5]$, and $d \in [6, 75]$, distinct from the treated-path search domain. Five starting values yield the same retained objective in each cohort, with at most 3,000 function evaluations per fit and tolerances of 10^{-13}. Table B1 reports the sample sizes, fitted parameters, and weighted RMSEs.

The comparison path averages the fitted curves with fixed one-third weights:

$$H_c^0(a) = \frac{1}{3} \sum_{q=1985}^{1987} \widehat{H}_{0q}(a). \quad (18)$$

Here c denotes the synthetic three-cohort population. Figure B2 compares this curve with the equally weighted observed cohort means. Weighted RMSEs are 0.11–0.12 Years of Schooling, and the largest absolute cell residual is 0.33 years. These are in-sample diagnostics over ages 16–35; values above 35 are extrapolated. The simulation takes the fitted curve as given, without propagating its sampling or extrapolation uncertainty. This descriptive calibration does not identify the untreated path of a particular policy population. The reference triple for $\theta_c = (\Delta K_c, \Delta b_c, \Delta d_c)$ averages the three fitted parameter vectors, giving approximately $(13.92, 0.26, 12.56)$. Its logistic curve differs from the average of fitted curves by at

most 4.12×10^{-4} Years of Schooling over ages 16–65. After normalization to the common age-17 level, the reference-parameter continuation differs from H_c^0 by at most 2.26×10^{-4} Years of Schooling over ages 18–65. Zero parameter increments therefore share the anchor exactly but need not imply an exactly zero effect at later ages.

The simulation's plateau range, $K \in [13.34, 15.14]$, uses the weighted first and 99th percentiles of 716 historical state–birth-cohort logistic fits. The bank covers the 50 states and DC, birth cohorts 1980–1994, and Census 2000 and ACS observations from 2001–2023. Its population comprises U.S.-born non-Hispanic White, non-Hispanic Black, and Hispanic respondents with valid education records, positive person weights, and recorded male or female sex. Profiles are fitted by population-weighted nonlinear least squares; retained fits converge inside the fitting bounds, contain at least 14 age cells spanning at least 13 years, begin by age 20, and reach at least age 29.[41] Transporting this range to the full-population simulation is a calibration assumption because the populations and survey windows differ. It is neither a confidence interval nor a validated bound on policy responses. The speed and timing domains, $b \in [0.12, 0.65]$ and $d \in [8, 30]$, and the annual acquisition cap in equation (20) are design restrictions. The simulation does not impose timing-index bounds or a joint parameter ellipsoid; policy-specific calibration and joint support are described separately in Appendix C.1. Its reported ranges condition on the baseline, support, and valuation rule; alternative support banks or path families have not been evaluated for this simulation.

B.2 Path Restrictions and Matching

The simulations hold the comparison path and observed education effects fixed while varying the unobserved attainment path. Stocks and acquisition flows follow equation (1); the observation operator preserves whether a source reports an age-specific effect, a weighted window average, or acquisition between two ages:

$$\mathcal{M}(\Delta H_c) = \begin{cases} \Delta H_c(a_*), & \text{age-specific stock,} \\ \sum_{a\in\mathcal{A}} w_a \Delta H_c(a), \quad \sum_{a\in\mathcal{A}} w_a = 1, & \text{window average,} \\ \Delta H_c(a_2) - \Delta H_c(a_1), & \text{acquisition over an interval.} \end{cases} \tag{19}$$

[41] Profile weights sum person weights; quantiles interpolate cumulative normalized weights. The limits are 13.34 and 15.14. The bank collectively spans ages 16–43, so plateaus remain fitted asymptotes.

Two separately observed stocks impose two restrictions; their difference imposes one. Follow-up requires a source-specific age mapping. The experiments use exact stocks for the same population at one or two attained ages, holding the comparison path in Appendix B.1 fixed and assuming no policy effect through age 17.

Write total treated parameters as (K, b, d), $S = H_c^0(17)$, $L(a) = \Lambda[b(a-d)]$, and $F(a) = [L(a) - L(17)]/[1 - L(17)]$. The anchored continuation is $H_c^1(a) = S + (K-S)F(a)$ after 17, with $K > S$ and $b > 0$. Candidates satisfy the parameter domain in Section 3.1 and

$$0 \leq H_c^1(a) \leq 25, \qquad 0 \leq H_c^1(a) - H_c^1(a-1) \leq 1, \qquad a = 18, \ldots, 72. \tag{20}$$

The acquisition cap applies to annual schooling increments, including the first post-anchor increment, rather than to the treatment effect or instantaneous derivative. It also holds for the fixed comparison curve. Nondecreasing levels allow either sign of $\Delta h_c(a)$. Holding the anchor fixed gives $\partial_K H_c^1(a) = F(a)$ and $\partial_d H_c^1(a) = -(K-S)b[1-L(a)]F(a) < 0$. The age at which fraction $q \in (0,1)$ of the remaining schooling is attained satisfies

$$a_q = 17 + \frac{1}{b}\log\left(\frac{1 + q\exp[b(d-17)]}{1-q}\right), \qquad 0 < \frac{\partial a_q}{\partial d} < 1. \tag{21}$$

Thus d delays accumulation without producing a uniform horizontal shift or identifying an individual graduation age. The inflection point lies in the continuation only when $d > 17$; the reference value is below 17. Figure B1, Panels A–C, illustrates the parameter roles. The anchor requires an unaffected common age; applications with different exposure ages require an appropriate common anchor or cohort-specific construction.

The constrained minimum-distance criterion in equation (15) has a zero minimum for the stipulated simulation inputs. All fixed positive-definite weight matrices then yield the same exact-match set. With one observation at $a_* > 17$, let $y = H_c^0(a_*) + m_{a_*}$. Conditional on (b, d), the plateau is

$$K = S + \frac{y - S}{F(a_*; b, d)}. \tag{22}$$

The search evaluates a 401×401 grid in (b, d) and refines candidate metric extrema locally. For two observations, define $z_j = H_c^0(a_j) + m_{a_j} - S$, $t_j = \exp[-b(a_j - 17)]$,

and $r = (z_2/z_1)(1-t_1)/(1-t_2)$. Exact matching gives

$$q = \frac{r-1}{t_1 - rt_2}, \qquad d = 17 + \frac{\log q}{b}, \qquad K = S + z_1 \frac{1+qt_1}{1-t_1}. \tag{23}$$

The resulting one-dimensional search uses 1,601 values of b, boundary root finding, and local refinement. It requires $q > 0$, finite parameters, and all path restrictions; the generating scenarios satisfy $0 < z_1 < z_2$. Including these generating scenarios verifies that the candidate sets are nonempty. Exact simulation matching is distinct from the policy Type 3 specifications that require constrained best fits with reported nonzero residuals.

Acceptance requires moment residuals no larger than 10^{-8} Years of Schooling and constraint violations no larger than 10^{-9}. Near-boundary proposals are moved toward a verified seed, rematched, and checked without relaxing these tolerances. Every two-age candidate is inserted into both component single-age libraries before recomputing extrema. All scenario–design calculations passed the reported feasibility and nesting checks; the largest retained moment residual was 3.55×10^{-15} Years of Schooling, and independent endpoint recalculation differed by at most 2.36×10^{-10} dollars. These checks do not include a grid-refinement stability exercise. Complete search settings and diagnostics remain in the replication files. Verified feasible outputs are retained as candidates regardless of optimizer stopping flags; they do not certify global extrema or search completeness. Pointwise envelopes can switch candidate identity across ages and need not themselves be feasible paths. No path probabilities or confidence levels are assigned. Supplementary Figures S1–S3 show the individual-scenario NPV heatmaps.

B.3 Earnings Valuation

The common valuation uses comparison earnings $Y_0(a)$ in 2000 dollars, values income at ages 18–65, and applies the discount factor $D(a) = 1.03^{-(a-18)}$. The signed returns are $r^- = -6.67\%$ and $r^+ = 815/7241$, retaining the conversion of Zimmerman (2014) used by Hendren and Sprung-Keyser (2020). For

$q(\tau) = \Delta H(\tau) - \Delta H(\tau - 1)$ and $\Delta H(17) = 0$,

$$R(a, \tau) = r^{-}\mathbf{1}\{\tau - 7 \leq a < \tau\} + r^{+}\mathbf{1}\{a \geq \tau\}, \tag{24}$$

$$\Delta y(a) = Y_0(a) \sum_{\tau=18}^{72} R(a, \tau) q(\tau) = Y_0(a) \left[r^{+} \Delta H(a) + r^{-} \{\Delta H(a + 7) - \Delta H(a)\} \right], \tag{25}$$

$$E^{\text{path}} = \sum_{a=18}^{65} D(a) \Delta y(a). \tag{26}$$

The policy-onset restriction excludes earnings responses before 18. Education extends through 72 because acquisition up to seven years after 65 can affect earnings before 65; the terminal schooling statistic remains $\Delta H(65)$. A step $\Delta H(a) = m\mathbf{1}\{a \geq 25\}$ reproduces the single-effect benchmark in equation (5); multiple steps sum shifted responses. This equivalence is an accounting construction, not evidence that everyone graduates at 25.[42]

The mapping assumes that the earnings response can be transported linearly across acquisition ages. Signed increments represent differences between population acquisition flows, not graduation probabilities. The seven-year negative-response window is a valuation template rather than seven additional years of enrollment. Early losses and later gains are already included, so no further opportunity-cost charge is deducted. Tuition, education resources, and transfers enter the separate welfare account. Figure 4 uses a plateau-only change $(\Delta K, \Delta b, \Delta d) = (0.57, 0, 0)$ matching $\Delta H(25) = 0.48$ years; Figure B1, Panel D, shows its levels. Its annual increments satisfy the cap, and both colored series enter the displayed NPV. The applications retain their source-specific earnings functions and monetary conventions, documented in Appendix C.

B.4 Simulation Design and Information Gains

The first experiment varies the generating process while observing it at age 25. For each parameter, the other two are fixed at the mean cohort parameters in Appendix B.1, and the feasible interval is the intersection of the parameter domain and all path restrictions. Table B2 reports these conditional intervals. The generating value is $p_j = L_j + u(U_j - L_j)$, with $u = 0, 0.01, \ldots, 1$, plus the reference value: 102 scenarios per parameter and 306 rows. The three reference rows repeat

[42] The continuous-age counterpart is $\Delta y(a) = Y_0(a) \int R(a, \tau)\, d\Delta H(\tau)$, including both smooth changes and discrete batches. Computation uses annual sums; the education observations do not identify the response kernel.

the same generating path. Only the generating path is restricted in this way; all three parameters may vary when matching the resulting effect. The second experiment uses five positions, $u \in \{0, 0.25, 0.5, 0.75, 1\}$, for each parameter and varies the observed ages. Its 15 scenarios each have 18 single-age designs and 153 distinct age pairs over 18–35, giving 2,565 comparisons. The age-25 designs overlap the first experiment, leaving 2,856 distinct scenario–design identifiers. Midpoint increments are 0.32, 3.00×10^{-3}, and 6.44 in their respective units. The acquisition cap binds at the speed interval's upper end, placing its midpoint near the reference speed. Scenario positions are chosen independently of welfare values.

For the verified candidate library $\widehat{\mathcal{P}}_s(D)$ of scenario s and observation design D, define

$$W_s(D) = \max_{h\in\widehat{\mathcal{P}}_s(D)} E^{\text{path}}(h - H_c^0) - \min_{h\in\widehat{\mathcal{P}}_s(D)} E^{\text{path}}(h - H_c^0). \tag{27}$$

Reductions relative to age 25 alone are $1 - W_s(D)/W_s(\{25\})$, computed within scenarios before aggregation; a zero baseline width makes this ratio undefined. Medians use the three interior positions, while largest widths and sign counts include both boundaries. Both signs require a candidate below $-\$0.01$ and another above \$0.01. The libraries also retain terminal schooling and the sum of annual stock gaps; the latter has schooling-year–age-year units. Endpoints of different metrics need not correspond to the same path. These summaries weight deterministic scenarios equally and do not estimate the prevalence of policy mechanisms.

Table B2 reports the complete age-25 scans. The midpoint comparisons in Figure 3 use the same generating paths before and after adding age 35, reducing NPV widths by 92.3–94.2 percent. Candidate values then lie above zero for the plateau example and below zero for the timing example, while the speed example retains both signs. Across all five scenarios, one speed and one timing scenario still admit both signs. Table 2 compares selected designs; Figure B4 reports every age combination, including maximum widths not fully tabulated. Adding age 24 is more informative than adding age 20 in the reported interior-scenario summaries despite the shorter separation. This pattern depends on the maintained curve family and support and does not establish an optimal observation design with noisy inputs.

To separate education information from differences in fiscal accounts, the MVPF extension applies the same hypothetical grant to every second-experiment design. The grant's PV is $B = \$10{,}000$ per person in 2000 dollars at age 18, recipients value it dollar-for-dollar, and the earnings tax rate is $t = 0.20$. All other non-earnings

components are zero. For earnings NPV E, WTP is $B + (1-t)E$ and net public cost is $B - tE$, giving

$$V(E) = \begin{cases} \dfrac{10{,}000 + 0.80E}{10{,}000 - 0.20E}, & E < 50{,}000, \\ +\infty, & E \geq 50{,}000. \end{cases} \tag{28}$$

The infinite value denotes beneficiary gains with nonpositive public cost. The positive-cost branch is strictly increasing, so earnings endpoints determine MVPF endpoints. A set with finite and infinite values has infinite span; a set containing only infinite values has zero dispersion under this convention. An infinite span does not establish that every intervening finite ratio is attainable. Proportional reductions require a positive finite initial span. Table B4 reports selected midpoint comparisons, and Figure B5 reports the median and maximum spans for all age combinations. Near zero net public cost, the nonlinear ratio can make percentage contractions differ substantially from those of NPV.

Restricting observations to ages 18–25 gives 28 pairs in each midpoint scenario, or 84 designs. Comparing each pair with either constituent age yields 168 comparisons. In the 49 with a positive finite single-age span, the MVPF span falls by 65.1–98.7 percent; the other 119 replace an infinite single-age span with a finite two-age span, for which no conventional percentage reduction is defined. All 84 two-age ranges still cross one. Adding age 35 to age 25 instead puts the plateau midpoint range above one and the timing range below one, while the speed range still crosses one. Across the full grid, 2,324 designs have only positive-cost candidates, 203 combine these with candidates under which the policy pays for itself, and 38 contain only the latter. These results condition on the grant, tax rate, baseline, path family, support, and earnings mapping. They measure the effect of additional exact observations, not sampling-error reductions or confidence coverage. Complete scenario results and individual heatmaps remain in the replication materials.

B.5 Single-Parameter Policy Comparisons

The policy diagnostic holds the reported Type 2 education inputs fixed and varies one parameter family while holding the other two fixed. Unlike the simulations, it preserves each application's observation operator, geography, cohort calibration, policy dose, and valuation account. Let $z_{ij}^0 = (K_{ij}^0, \log b_{ij}^0, d_{ij}^0)$ denote baseline component j of policy i, with fixed weight ω_{ij}. For family $q \in \{K, b, d\}$, the

construction and matching set are

$$z_{ij}^{1,q}(t) = z_{ij}^{0} + te_q, \qquad \mathcal{T}_{iq} = \left\{ t : H_i^{1,q}(\cdot; t) \in \mathcal{C}_i,\ \mathcal{M}_i\{\Delta H_i^q(\cdot; t)\} = \mathbf{m}_i^{\text{Type 2}} \right\}. \quad (29)$$

The anchored component paths are aggregated with their original weights, reproducing H_i^0 at $t = 0$. Changes in K and d are additive; speed changes multiply b by $\exp(t)$. Plateau changes are permanent, whereas speed and timing changes preserve eventual attainment. The conditional domain retains level and acquisition bounds, timing support, and the calibrated parameter ellipsoid; Social Security uses its shape-only ellipse. Scanning, root bracketing, and tangency checks retain matches with residuals below 10^{-8}, without SE bands. These component-preserving shifts need not be coordinate slices of the main common-parameter search.

STBF permits four outcome-specific shifts within the active family, jointly matching time enrolled, any college, AA, and BA. Tennessee Promise permits three credential-specific shifts for some college, AA, and BA. Credential ordering, subgroup feasibility, and applicable late-completion restrictions remain imposed. Only the plateau family yields jointly admissible matches for these two cases. None of the single-family solutions exactly matches every additional Type 3 point input; this diagnostic does not replace the main constrained best fits. An absent solution means that the specified one-family explanation fails within its conditional support, not that the effect is zero or that all timing responses are excluded.

Figure B3 and Table B3 compare matched earnings NPVs with their source-account benchmarks, dividing both by the same positive reference expenditure after reconciling monetary and discount-base units. The paired samples contain 34, 28, and 20 calculations for K, b, and d. Algorithmic placement has no source lifetime-earnings projection; Early Head Start and Tennessee lack verified common monetary years for paired dollar comparisons. Their dimensionless MVPF comparisons remain available. Cannabis has no whole-policy benchmark. The table retains all 37 reference calculations, including unpaired cases; parameter changes appear in Table 5. Each earnings comparison retains its application's signed response or author income function, so deviations from the benchmark can reflect both continuation and valuation mapping. They do not identify the realized mechanism. The replication files retain all parameter sweeps, matching checks, and unit reconciliations.

C Common Calibration and Welfare Accounting

The implementation combines externally calibrated education paths with the valuation rules of each source account. This appendix documents the common calibration, income mapping, and component-replacement procedures underlying Section 4.1. Appendix D.2 records the inputs and implementation of all 38 applications, including the three discussed in Section 4; Appendix E treats uncertainty in the reported inputs.

C.1 Common Calibration and Path Restrictions

The comparison profiles use the unrestricted IPUMS Census–ACS extract for 2000–2022, covering the 50 states, DC, and the national population without race, sex, or birthplace restrictions. The 2000 component is Census sample 200001, followed by annual ACS samples. Person weights form outcome means within geographic–birth-cohort–age cells, with birth cohort defined as survey year minus integer age. Geographic and cohort weights are fixed across ages within each application. The resulting profiles describe repeated cross sections; migration and other composition changes can remain within cells. The library contains 10,920 fits across 52 geographic units, 35 birth cohorts, and six outcome definitions. Table D1 records the population, observed age support, and calibration mode for each application.

Logistic profiles are fitted to person-weighted age means, using ages 16–35 where observed support permits. State shapes are regularized toward their national counterparts. When young-adult observations are insufficient, the profile borrows a national shape and uses the available observations to calibrate its level; the registry identifies the donor and calibration ages. A published control mean replaces an external level only when its population, outcome, and clock can be reconciled with the retained effect. Even then, matching a sample mean does not recover the untreated prediction from the source regression. The external profile and any population adjustment remain calibration assumptions, held fixed across Type 2 and Type 3.

Empirical parameter support is constructed before the policy-path search. Write $x = (K, \log b, d)^\top$. For each outcome, the donor set contains the 1983–1987 cohorts in the 51 state-level regions. The covariance for a state target averages its within-donor covariance and the pooled-state covariance with equal weights. National targets use the pooled covariance; multi-state targets average state covariances using their fixed population weights. Centering this dispersion at the target

reference's weighted parameter location gives

$$(x - \bar{x}_i)^\top \Sigma_i^{-1} (x - \bar{x}_i) \leq 11.34. \tag{30}$$

The cutoff uses a 99-percent chi-square scale as an empirical calibration convention. It is not a confidence region for a treatment response. A raw logistic profile and its anchored normalization have the same parameter triple, so no separate post-anchor refit is needed; such refits are retained as diagnostics. Timing limits use the anchored completion ages defined in equation (21). For completion fractions 0.1 and 0.9, the 0.5th and 99.5th percentiles of pooled donor deviations are transported around the corresponding target reference. The wider sensitivity doubles Σ_i, multiplies the timing deviations by $\sqrt{2}$, and retains the verified main candidates. Outcome-specific exceptions are recorded in Appendix D.2.

Candidates share an unaffected anchor with the comparison path. Schooling levels are nondecreasing and gain at most one year per annual step; probability levels remain in $[0, 1]$, with credential ordering where required. Equation (15) fits the fixed point inputs subject to these conditions. Type 2 uses Sobol and 97×97 regular shape grids, with the retained moment determining the plateau conditional on shape. Exact additional matches use bracketing and root refinement. The best-fit specifications use up to 64 diverse feasible starts and constrained optimization, followed by least-squares polishing only when every original restriction still holds. Distinct minimizing solutions are retained rather than repeated outputs from the same basin. Annual paths, input equalities, and paired welfare components are independently reconstructed, and verified Type 3 solutions enter the paired Type 2 bank. These checks establish feasibility of the saved solutions, not global optimality or completeness. Appendix D.3 and the replication records document fit diagnostics and search settings.

C.2 From Education Paths to Earnings

The observation operator preserves whether the source reports a stock, a window average, or acquisition over an interval. Converting enrollment, credits, or credentials into Years of Schooling is a separate operation: it retains the source's duration and incremental-credential conventions without counting the same education twice. Assigned ages and population weights are stated with each input in Appendix D.2. Neither the simulation's age-17 anchor nor its 1985–1987 cohorts are imposed on all applications. The common earnings object is the difference between the source income function evaluated at the policy and comparison paths, as in equation (12).

The retained prices, returns, and population conventions are fixed within each path comparison.

For accounts using a linear schooling-to-earnings response, signed increments of the education-effect path enter through the source response kernel:

$$e_i^{\text{path}}(a) = Y_{0i}(a) \sum_{\tau} R_i(a, \tau) \Big[\Delta H_i(\tau) - \Delta H_i(\tau - 1) \Big]. \tag{31}$$

The event indexing τ, phase lengths, age-income bases, and conversion factors follow the source account. A signed batch is a difference in population acquisition flows; it need not represent a nonnegative group of individual graduates. The illustrative seven-year negative-response template in Appendix B.3 does not override a source program's own income clock or overlapping response phases. An author forecast that is nonlinear in credentials and experience is instead evaluated using its original vector of education inputs. Where complete author inputs are unavailable, the case record distinguishes the public-data reconstruction from the original benchmark. Early earnings losses enter the net-income response once, without a separate schooling-flow opportunity-cost deduction.

Let $s \in \{\text{ref}, \text{path}\}$ index the source and path accounts, and let $\mathcal{A}_i$ contain the valued income ages. With source discount rate ρ_i and base age $a_{d,i}$, write $D_i(a) = (1 + \rho_i)^{-(a - a_{d,i})}$. Annual earnings effects must use the same price year. Under ordinary age discounting,

$$\text{NPV}_i^{\text{earn},s} = \sum_{a \in \mathcal{A}_i} D_i(a) e_i^s(a). \tag{32}$$

When source response phases b use different discount bases, their contributions remain separate: $E_i = \sum_a \sum_b D_{ib}(a) e_{ib}(a)$. The nominal income curve is $\sum_b e_{ib}(a)$, whereas its discounted counterpart is $\sum_b D_{ib}(a) e_{ib}(a)$, whose age sum equals the implemented NPV. A common discount factor is used only when it reproduces those phase conventions. The source-block records retain these distinctions without aligning totals through a final scaling factor.

For a linear fiscal schedule, let $\vartheta_i(a)$ be the share of an earnings change accruing to government through taxes and reduced transfers. The private and fiscal income components are

$$\text{NPV}_i^{\text{after},s} = \sum_{a \in \mathcal{A}_i} D_i(a)[1 - \vartheta_i(a)] e_i^s(a), \tag{33}$$

$$\text{NPV}_i^{\text{tax},s} = \sum_{a \in \mathcal{A}_i} D_i(a) \vartheta_i(a) e_i^s(a). \tag{34}$$

Source phase-specific tax rates remain attached to the corresponding phase contributions. A single average rate generally cannot preserve both components across alternative paths. Nonlinear tax or transfer rules require the source mapping and, where relevant, distributional inputs. Likewise, a source account that values gross earnings in WTP or changes only fiscal earnings retains that convention; the disposable-income formula is not imposed on every application.

C.3 Welfare Accounts and Component Replacement

All components within a welfare account refer to the same population, policy contrast, monetary year, and discount convention. Inputs expressed per eligible person, per recipient, or per marginal unit of aid require explicit conversions before they can be combined. Multiplying every component by the same positive constant preserves the MVPF; combining different treatment margins does not. Cross-policy monetary levels require common dollar units, although MVPF ratios are dimensionless. The fixed private and public components in equations (10) and (11) can include transfers and other measured benefits or costs. A scholarship already counted as a private benefit cannot also reduce the tuition charge to generate a second benefit, and public education spending must exclude expenditures already in the program component.

The reconstruction first recovers source WTP and net cost from their recorded components. Equations (13) and (14) then remove each replaced component at its reference value before adding its path value. Reinserting the original components must recover the original totals. This identity checks the accounting without requiring the reference continuation to belong to the logistic family. Missing components are recovered from documented source information or explicit calibrations, rather than set to zero. Published benchmarks remain separately identified when reconstructed income inputs, policy units, or price conventions differ. In those applications, a comparison with Type 1 reflects these differences as well as the education continuation; comparisons among candidate paths hold the reconstructed account fixed.

Recoverable private-payment and public-education schedules are transported with the source schooling equivalents and event timing. Applying their per-unit schedules to signed batches reproduces source costs at the reference batch and permits costs to vary with timing and quantity. This transport holds the institution mix and marginal prices fixed; attainment alone does not determine actual enrollment, sector choice, or aid eligibility. Other accounts retain their

stated non-income components and replace only the specified earnings channel. Such comparisons hold grant duration, tuition, and education resources fixed where the source implementation requires it. Appendix D records the replacement scope and the source-derived or reconstructed inputs for every application.

Interpreting income as WTP retains the source welfare assumptions, including whether the earnings response is a human-capital return rather than compensation for additional effort. Unmeasured consumption benefits, externalities, or behavioral costs require separate valuation. Each reported MVPF pairs WTP and net cost from the same complete candidate path; marginal component extrema cannot be combined into a new ratio. Positive WTP with nonpositive net cost receives the convention defined in Section 4.1. Large finite ratios and rounded zero denominators alone do not establish an infinite bound. Appendix D.3 documents the candidate comparisons, while input uncertainty remains the separate exercise in Appendix E.

D Full Policy Collection and Case Registry

This appendix documents the 38 policy calculations in Section 5: their source inputs, calibration, welfare comparisons, and ordering rule. The case records cover all 38 applications, with technical details for the three main examples supplementing Section 4. Appendix C gives the common calibration and accounting procedures. Appendix E treats input uncertainty; the replication materials supply full-precision observation operators and complete paths.

D.1 Cases, Source Families, and Calibration

Table D1 lists all 38 calculations. The collection retains 34 earlier calculations and adds algorithmic college placement, two further tuition-deduction contrasts, and the cannabis education account. Several calculations can come from one paper. The source-account cases represent 22 research families, with one further family for the author cannabis study. The grouping combines related Bulman–Hoxby versions and tax-schedule calculations, the spending and tuition calculations associated with Deming–Walters, and the Kalamazoo Promise research versions. Source author names are recorded separately. A family-level witness count records whether at least one member supplies the specified paths; it is not an estimate of a policy-population frequency.

The geographic registry selects the source's treated population rather than combining it mechanically with control states. Birth cohorts and mixture weights

remain fixed across candidates. Geographic matching improves comparability, but an ACS mean is not an identified untreated trajectory for a source study's selected population. The source-control status records that distinction, including Tennessee's HS-graduate denominator and control-level calibration. Education conversion, age mapping, the anchored logistic continuation, empirical support, and transport of welfare prices remain maintained assumptions. Reproduction of a source coefficient or account is distinct from validation of those assumptions. The accompanying replication materials record each moment's outcome, reconstruction, target, available SE, and nonzero age–cohort weights, together with the comparison profiles, anchors, valuation coefficients, and sources. A converted coefficient is not relabeled as observed completed schooling. In particular, current enrollment and accumulated education are different outcomes; the source-specific bridge remains part of each reconstruction. Full-precision inputs permit the matching equation to be reproduced without inferring an age from a figure or replacing a window by its midpoint.

D.2 Education Inputs and Observation Windows

The following records document all 38 calculations: source estimates and windows, path-input conversions, welfare accounts, and results. The three main examples supplement Section 4. Unless stated otherwise, Type 1 and its education conversion follow Hendren and Sprung-Keyser (2020) and the associated Policy Impacts calculation. A probability change of 0.01 is one percentage point; conversions to schooling equivalents and assigned observation ages remain separate assumptions. Calculations retain full precision. Small defining inputs use percentages or scientific notation so that two-decimal displays do not erase differences across windows. Numerical equalities with rounded displays refer to the archived unrounded inputs. Repetition summaries use the 2.5th percentile of minima and 97.5th percentile of maxima from 1,000 paired inputs, conditional on verified paths and defined ratios. Appendix E describes their dependence assumptions and limitations; these summaries are not validated confidence regions. Source Type 1 intervals retain their original inferential meaning. Appendix C gives the common procedures.

Florida Student Access Grant

Castleman and Long (2016) use a regression discontinuity (RD) design to estimate how Florida Student Access Grant eligibility affects enrollment, credits, and completion. Outcomes extend from the first semester through seven years after HS.

Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 7.42. The education input is the 4.37-credit effect on four-year cumulative credits in the original study's Table 4, column (5). Dividing by 24 gives 1.82×10^{-1} schooling-equivalent years. With first exposure assigned to age 18 and the unaffected anchor at 17, Type 2 retains $\Delta H(22) = 1.82\times10^{-1}$. Type 2 yields a searched positive-cost MVPF span of $[0.57, 22.32]$, crossing one; these endpoints bracket Type 1. Type 3 adds Table 4, columns (1)–(4): 0.36, 1.10, 2.67, and 3.85 credits. The same conversion gives fixed targets of 1.48×10^{-2}, 4.58×10^{-2}, 1.11×10^{-1}, and 1.60×10^{-1} at assigned ages 18.5, 19, 20, and 21. These are modeled age placements of the source windows; the half-year target is evaluated on the continuous curve. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 yields a positive-cost MVPF of 18.09, above one. This fitted value lies above Type 1; the narrower span relative to Type 2 also reflects the fitting criterion. Figure D2, Florida access grant panel, shows age-65 gaps of $[-0.10, 1.07]$ under Type 2 and 0.72 years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-14.48, +\infty]$; Type 3 $[0.53, +\infty]$. The reported Type 1 95 percent interval is $[0.69, +\infty]$.

DC Tuition Assistance Grant

Abraham and Clark (2006) examine whether the DC Tuition Assistance Grant increases college enrollment, comparing enrollment changes for DC graduates with those in comparison jurisdictions. The retained outcome is first-time freshman enrollment among recent HS graduates, rather than subsequent completion. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 22.98. The original study's published Table 7a, Total row for all schools, final difference column, supplies the 8.90×10^{-2} enrollment effect used in the reference calculation. Its two-year duration assumption gives $2(8.90\times10^{-2}) = 1.78\times10^{-1}$ schooling-equivalent years. Type 2 places this converted total at age 20, imposing $\Delta H(20) = 1.78\times10^{-1}$ with an age-17 anchor. Age 20 is an assumed endpoint of the duration conversion, not a directly observed schooling age for every student. Type 2 yields a searched positive-cost MVPF span of $[-1.60, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: assigning 8.90×10^{-2} years to an earlier age would split the same assumed duration rather than add a distinct follow-up estimate. The source's enrollment comparisons therefore establish the input, but do not determine its subsequent persistence. Figure D2, DC tuition assistance panel, shows the Type 2

paths; the age-65 gap is $[-0.36, 1.27]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-3.93, +\infty]$. The reported Type 1 95 percent interval is $[1.12, +\infty]$.

Adams Scholarship

The policy contrast is scholarship eligibility, not the separate effect among students induced to switch colleges. Eligibility shifts students toward participating public colleges. Table 4 of Cohodes and Goodman (2014) reports cumulative entry and completion by each window, alongside fourth-year persistence. The authors interpret similar declines in persistence and graduation as evidence favoring dropout over pure delay within the observed period. A modeled continuation that later catches up extends beyond this evidence; it does not contradict the observed mechanism. The implementation of Section 4.2 retains Table 4, columns (3) and (5), of Cohodes and Goodman (2014) and the education conversion in Hendren and Sprung-Keyser (2020). Two-year entry, four-year entry, two-year completion, and BA completion receive incremental schooling multipliers $(1, 2, 0.5, 1)$. The two-year entry and six-year completion windows give $7.00{\times}10^{-3}+2(0)+0.5(0)-2.50{\times}10^{-2} = -1.80\times10^{-2}$ Years of Schooling per eligible student. This combines outcomes observed at different horizons. The external comparison is the full-population Massachusetts profile for cohort 1987, with an unaffected age-17 anchor; it is not the untreated mean of the local RD population. Type 2 places the converted total at age 24:

$$\Delta H(24) = -1.80\times10^{-2}. \tag{35}$$

Type 3 retains this equality and fits the additional targets

$$\Delta H(18) = 2(9.00\times10^{-3}) + 8.00\times10^{-3} = 2.60\times10^{-2}, \tag{36}$$

$$\Delta H(20) = 2(0) + 7.00\times10^{-3} = 7.00\times10^{-3}, \tag{37}$$

$$\Delta H(22) = 7.00\times10^{-3} - 1.70\times10^{-2} + 0.5(-2.00\times10^{-3}) = -1.10\times10^{-2}. \tag{38}$$

These are converted point targets, not exact restrictions on the selected Type 3 fit. The age-20 input preserves the timing of entry coefficients already in the reference total; ages 18 and 22 add immediate-entry and four-year-completion evidence. Carrying the enrollment contribution forward follows the source's additive valuation, with completion adding only incremental schooling. The assigned ages are reconstruction assumptions, and fourth-year persistence is not imposed as a

schooling-stock target. Table D3 records the source coefficients and SEs; Table D2 reports the fitted discrepancies.

The source reference batch reproduces earnings NPV, taxes, private payments, and public education costs separately in 2004 dollars (Table D5): earnings NPV is $-\$1,204.04$, WTP is \$1,364.23, and net cost is \$1,896.35. The retained Type 2 earnings range is -27.52–42.21 thousand dollars, compared with -2.74 thousand for the selected Type 3 path. With E, T, P, C defined in Section 4.1, retain

$$w_{\text{fixed}} = W^{\text{ref}} - E^{\text{ref}} + T^{\text{ref}} + P^{\text{ref}}, \tag{39}$$

$$n_{\text{fixed}} = N^{\text{ref}} + T^{\text{ref}} - C^{\text{ref}}. \tag{40}$$

Candidates use $W = w_{\text{fixed}} + E - T - P$ and $N = n_{\text{fixed}} + C - T$. Both education-payment schedules move with the signed schooling-equivalent batches, retaining the source institution mix and unit costs even when the converted effect is negative. All kernels already use 2004 dollars; no second 188.9/172.2 conversion is applied. Input repetitions preserve and extend the source's two-year and four-year error groups as specified in Appendix E.2.

STBF Scholarship

The deposited workbooks and earnings programs of Angrist et al. (2022) supply the 2012–2014 four-year strata's pooled six-year effects: time enrolled $m_T = 0.36$, cumulative credential shares $m_C = 7.16 \times 10^{-3}$, $m_A = 0.05$, and $m_B = 0.08$, with SEs $(0.04, 2.96 \times 10^{-3}, 0.02, 0.02)$. Each has a separate anchored path satisfying

$$\mathcal{M}_o\{\Delta H_o\} = m_o, \qquad o \in \{T, C, A, B\}. \tag{41}$$

The graduating cohorts receive equal weights; graduation ages 17, 18, and 19 receive external weights $(6.52\%, 57.59\%, 35.89\%)$, placing the six-year inputs at ages 23–25 and anchors at 16–18. These are calibration weights rather than verified trial entry-age frequencies. Credential levels satisfy $H^s_C \geq H^s_A \geq H^s_B$. Time enrolled uses an external schooling-shaped timing index, not an observed level of completed schooling. Four outcomes at one horizon do not supply a second observation date, so no Type 3 is admitted.

For race–sex group g, the retained author Poisson mean-input function is

$$F_g(a; \Delta p_g, \Delta T_g) = \exp\Big[\alpha_g + (B_g + \beta_g' \Delta p_g(a)) \mathbf{1}\{x_g(a) > 0\} + f_g(x_g(a))\Big], \tag{42}$$

$$x_g(a) = \max\{a - 18 - \bar{T}_g - \Delta T_g(a), 0\}. \tag{43}$$

Here f_g is the source quartic with its original scaling, $\Delta p_g = (\Delta C_g - \Delta A_g, \Delta A_g - \Delta B_g, \Delta B_g)$ contains exclusive credential changes, and $B_g = \beta_g' \bar{p}_g$. The forecast evaluates mean inputs rather than integrating over an unobserved individual distribution. Source coefficients, income ages 18–65, 3 percent discounting to 18, and 2019 dollars are retained. The source's 3 and 5 percent PVs are used to recover sufficient statistics $(\bar{T}_g, B_g)$ conditional on its coefficients and effects; baseline credential means remain underidentified.[43] Every candidate must admit one fixed baseline credential-share vector per group that remains valid at all ages; a feasibility witness is not treated as an author mean.

Because pooled and subgroup controlled effects are not exactly aggregation-consistent, each pooled timing index is transported using

$$\Delta H_{go}(a) = m_{go} \Delta H_o(a)/m_o.$$

The group effect ratios remain fixed, without claiming that the pooled path equals weighted group means at every age. Annual income is $\sum_g w_g[F_g(a; \Delta p_g, \Delta T_g) - F_g(a; 0, 0)]$. Evaluating the function at the source effects recovers the workbook aggregate of \$21,129.23; the printed \$21,150 is retained separately, with no scaling correction. The account is $W = 32{,}250 + 0.8E$ and $N = 32{,}250 - 0.2E$. Printed Type 1 is 1.75; Type 2 earnings NPVs are 7.60–60.60 thousand dollars and MVPFs are 1.25–4.01. Figure D9 displays the four joint paths, distinguishing the external time index from the author's mean enrollment time. The 1,000 repetitions apply one common normal shock to all four marginal SEs and repeat the eight-dimensional joint search, holding income coefficients, recovered statistics, effect ratios, and external calibration fixed (Appendix E.2).

Wisconsin Scholars Grant

Goldrick-Rab et al. (2016) randomize Wisconsin Scholars Grants across three entering cohorts and report semester credits, cumulative credits, and completion. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 1.43. The reference input combines Table 5's 0.9-credit third-year effect for cohort 1 with its pooled 2.1-credit effect for cohort 2 after three years and cohort 3 after two years. Weights of 31.46% and 68.54% and division by 24 yield $7.18{\times}10^{-2}$ Years of Schooling. With age-17 anchors, Type 2 imposes $31.46\%\Delta H_{1990}(21) + 34.27\%\Delta H_{1991}(21) +$

[43] In source group order, the recovered pairs are $(1.91, 0.28)$, $(2.02, 0.32)$, $(1.57, 0.24)$, and $(1.62, 0.22)$. Conditional on the control total of 1,924 and exported aggregate PVs, recovered integer weights are proportional to $(522, 798, 242, 362)$. These recoveries are distinguished from directly supplied coefficients.

$34.27\%\Delta H_{1992}(20) = 7.18\times10^{-2}$. Equal splitting of the pooled cohorts and entry at 18 are calibration assumptions; cohort 3's unobserved third year receives no weight. Type 2 yields a searched positive-cost MVPF span of $[-1.94, 7.58]$, crossing zero and one; these endpoints bracket Type 1. Type 3 separately retains both cumulative estimates and converts Table 5's annual credit contributions: $(0.5, 0.2, 0.1)/24$ for cohort 1, $(0.4, 1.0)/24$ for pooled cohorts 2–3, and $0.5/24$ for cohort 2's third year. The first-year contribution restricts $\Delta H(19) - \Delta H(17)$; later contributions restrict the next annual differences. Reported semester and cumulative coefficients do not satisfy exact adding-up identities, so the fit also reflects this source aggregation. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 yields a positive-cost MVPF of 1.96, above one. This fitted value lies above Type 1; the narrower span relative to Type 2 also reflects the fitting criterion. Figure D2, Wisconsin Scholars Grant panel, shows age-65 gaps of $[-0.54, 1.01]$ under Type 2 and 0.15 years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-2.84, +\infty]$; Type 3 $[-0.04, 6.29]$. The reported Type 1 95 percent interval is $[0.96, 2.18]$.

Kalamazoo Promise

Bartik et al. (2021) estimate how Kalamazoo Promise eligibility affects enrollment, attempted credits, and credentials. Their credit results cover two, three, and four years after HS; later degree outcomes use different follow-up samples. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 1.93. The original study's Appendix Table A.7, Panel C, column (1), supplies a fourth-year credit effect of 6.56, retained as 6.56 in the reference calculation. Dividing by 24 gives 0.27 schooling-equivalent years. Assuming entry at 18, Type 2 imposes $\Delta H(22) = 0.27$ and anchors the paths at 17. Type 2 yields a searched positive-cost MVPF span of $[0.65, 4.53]$, crossing one; these endpoints bracket Type 1. Type 3 adds the same table's two- and three-year credit effects, 3.24 and 4.31. Division by 24 gives targets $\Delta H(20) = 0.13$ and $\Delta H(21) = 0.18$, while the fourth-year Type 2 input remains exact. Treating these follow-ups as a common response requires comparability across the available cohorts. Enrollment and credentials are not additionally counted as attempted credits. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 yields a positive-cost MVPF of 3.77, above one. This fitted value lies above Type 1; the narrower span relative to Type 2 also reflects the fitting criterion. Figure D2, Kalamazoo Promise panel, shows age-65 gaps of $[-0.12, 1.06]$ under Type 2 and 0.83 years under Type 3; Type 2 permits

eventual reversal. Paired input repetitions give Type 2 $[-0.65, 19.65]$; Type 3 $[0.78, 15.51]$. The reported Type 1 95 percent interval is $[0.97, 5.61]$. The released income program overlaps a positive response beginning at 18 with a negative response at 18–24 and discounts the positive block from age 11. The reconstruction transports those phase conventions rather than substituting the illustrative post-25 template.

Michigan CC tuition

Acton (2021) estimates responses to a $1,000 reduction in Michigan CC tuition. Enrollment within one year is observed for 2009–2011 graduates; credentials are measured in 2017, six to eight years after HS. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 29.46. The reference uses the original study's Table 4, Panel B, columns (1)–(4), enrollment effects $(3.60\times10^{-2}, -2.10\times10^{-2}, -4.00\times10^{-3}, -3.00\times10^{-3})$, and Table 5, columns (4)–(6), credential effects $(-2.00\times10^{-3}, 3.00\times10^{-3}, 1.10\times10^{-2})$. Schooling weights $(1, 1, 1, 2, 0.5, 1, 2)$ give 2.90×10^{-2} years. Type 2 places this total at representative age 25, $\Delta H(25) = 2.90\times10^{-2}$, with anchor 17. Type 2 yields a searched positive-cost MVPF span of $[-4.44, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Type 3 separately uses the early enrollment contribution, $3.60\times10^{-2} - 2.10\times10^{-2} - 4.00\times10^{-3} + 2(-3.00\times10^{-3}) = 5.00\times10^{-3}$. It imposes $[\Delta H(18) + \Delta H(19)]/2 = 5.00\times10^{-3}$, preserving Type 2. The equal-weight age window is an additional temporal allocation of the converted enrollment input, not a source estimate of completed schooling at either age. It uses an earlier source window rather than inventing a second observation from the same enrollment result. Type 3 yields a searched positive-cost MVPF span of $[9.03, +\infty]$ and separate paths under which the policy pays for itself, entirely above one among the paths found. Figure D2, Michigan CC tuition panel, shows age-65 gaps of $[-0.22, 0.57]$ under Type 2 and $[0.01, 0.06]$ years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-24.02, +\infty]$; Type 3 $[-3.74, +\infty]$. The reported Type 1 95 percent interval is $[-2.33, +\infty]$.

CUNY Pell grants

Marx and Turner (2018) estimate how an additional $1,000 of Pell aid affects borrowing and academic attainment among CUNY entrants in 2007–2011. Their first-year-student results include current-year and three-year credit outcomes. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 1.39. The

original study's Table 8, Panel C, column (2), gives the 2.23×10^{-1}-credit three-year effect retained for lifetime valuation. Dividing by 24 gives 9.29×10^{-3} schooling-equivalent years. With entry assigned to 18 and anchor 17, Type 2 imposes $\Delta H(21) = 9.29\times10^{-3}$. Type 2 yields a searched positive-cost MVPF span of $[-\infty, 14.28]$, crossing zero and one; these endpoints bracket Type 1. Type 3 adds the first-year entrants' current-year effect, 2.12×10^{-1} credits, as $\Delta H(19) = 2.12\times10^{-1}/24 = 8.83\times10^{-3}$. Both converted points are matched exactly. The first- and third-year regressions have 38,100 and 32,271 observations, so this construction assumes comparability of their available samples. The separate 0.49-credit estimate for returning students is excluded because it concerns a different risk set. Type 3 yields a searched positive-cost MVPF span of $[-3.17, 2.34]$, crossing zero and one. Figure D3, CUNY Pell grants panel, shows age-65 gaps of $[-0.70, 1.05]$ under Type 2 and $[-0.10, 0.04]$ years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-\infty, +\infty]$; Type 3 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-2.95, 12.88]$.

Ohio Pell grants

Bettinger (2004) examines whether Pell aid prevents withdrawal between the 1999–2000 and 2000–2001 academic years in Ohio. The retained estimate concerns students already enrolled in four-year colleges. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 2.49. The original study's Table 5.3, column (4), reports that \$1,000 in imputed Pell aid reduces withdrawal by 6.40 percentage points. The reference reverses the sign and assigns one additional schooling year to persistence, obtaining 0.06 expected Years of Schooling. Type 2 imposes $\Delta H(20) = 0.06$. The modeled first affected year begins at 19, so age 18 is the unaffected anchor. This timing reflects a persistence intervention for existing college students, rather than an initial-entry intervention. Type 2 yields a searched positive-cost MVPF span of $[-11.75, 7.53]$, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted. Other columns change the sample or specification for the same transition and do not provide a later attainment observation for these students. Consequently, the positive short-window persistence effect leaves both its duration and later catch-up to the maintained path restrictions. Figure D3, Ohio Pell grants panel, shows the Type 2 paths; the age-65 gap is $[-0.28, 1.07]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[0.82, 5.54]$. The source

earnings program has six early-response years and distinct early and later discount bases; both are retained in the path mapping.

Tennessee HOPE scholarship

Bruce and Carruthers (2014) study Tennessee HOPE eligibility and post-HS college choice among 2006–2009 graduates. Their discontinuity comparison observes enrollment choices rather than a common later degree-completion endpoint. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 1.86. The original study's Table 3 reports a -0.02 effect on two-year enrollment in the retained eligibility comparison for students with low grade point averages (GPAs). The reference uses a one-year schooling contrast for the sector response and reverses its sign, giving 0.02 schooling-equivalent years. Type 2 imposes $\Delta H(20) = 0.02$, with first exposure assigned to 18 and anchor 17. The assigned endpoint locates the reference duration conversion; it is not an observed age common to all graduates. Type 2 yields a searched positive-cost MVPF span of $[-75.64, 21.66]$, crossing zero and one; these endpoints bracket Type 1. No additional comparable attainment window is admitted to Type 3. Sector-specific choices and robustness specifications are not later observations, and the requirement to enroll within 16 months is an eligibility rule, not a common follow-up period. Type 2 therefore examines continuations of the retained sector-switching input without rescaling it into a recipient effect. Figure D3, Tennessee HOPE panel, shows the Type 2 paths; the age-65 gap is $[-0.68, 1.09]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[0.92, 5.08]$.

Tennessee Pell eligibility

Carruthers and Welch (2019) estimate effects on some college in the initial fall at the Pell eligibility cutoff among 2006–2009 Tennessee graduates. The original study's Table 4, column (1), distinguishes enrollment by sector and location. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 0.85. The retained four-year probability effects sum to $P_4 = -6.24 \times 10^{-3}$. Two-year and grouped for-profit effects sum to $P_2 = 5.19 \times 10^{-3}$. The reference assigns two Years of Schooling to four-year entry and one to two-year entry, giving $2P_4 + P_2 = -7.28 \times 10^{-3}$. Type 2 matches this total in the second post-HS year, at ages 19–21 for graduation ages 17–19. Their weights are 6.52%, 57.59%, and 35.89%; graduation cohorts receive equal weights. Each component is anchored one year before graduation, at 16, 17,

or 18. Type 2 yields a searched positive-cost MVPF span of $[-\infty, 2.70]$, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted. The sum $P_4 + P_2 = -1.05 \times 10^{-3}$ uses the same initial-entry coefficients; imposing it at an earlier age would allocate the assumed duration rather than add an observed window. The age and cohort weights are calibration assumptions, not additional causal estimates. Figure D3, Tennessee Pell eligibility panel, shows the Type 2 paths; the age-65 gap is $[-0.68, 1.03]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-1.59, 3.57]$. The released source program gives an MVPF of 0.84, distinct from the retained library Type 1 of 0.85; their vintages remain separately recorded.

Georgia HOPE scholarship

Cornwell et al. (2006) estimate Georgia HOPE effects on freshman enrollment in a 1988–1997 state-year panel. The outcome compares successive entering classes rather than following one class through completion. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 4.00. The published study's Table 3, Southern Regional Education Board (SREB) control columns (4), (6), and (10), supplies log effects 0.09, 0.13, and −0.05 for public four-year, private four-year, and public two-year enrollment. Table 4, column (5), contributes 560 fewer out-of-state entrants. The reference applies $\exp(\beta) - 1$, sector counts, and recipient normalization, obtaining four- and two-year probability responses of 5.60×10^{-2} and -1.47×10^{-2}. The duration conversion gives $2(5.60 \times 10^{-2}) - 1.47 \times 10^{-2} = 9.73 \times 10^{-2}$. Type 2 imposes this schooling gap at age 20, equally weighting 1975–1979 cohorts with anchor 17. Type 2 yields a searched positive-cost MVPF span of $[-3.19, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: a first-year value of 0.04 would recombine the same sector estimates, not introduce another attainment window. The continuation comparison therefore remains conditional on both the reference enrollment conversion and its representative age placement. Figure D3, Georgia HOPE panel, shows the Type 2 paths; the age-65 gap is $[-0.36, 1.16]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-4.11, +\infty]$. The reported Type 1 95 percent interval is $[0.46, 22.38]$. The income mapping retains the source sector mixture, including the zero early-response component for CC enrollment.

Social Security Student Benefits

Section 4.3 uses the full-time-attendance column of Table 3 in Dynarski (2003): 21.90 percentage points (SE 10.20 points) by age 23 and 22.40 points (SE 10.60 points) by age 28. The latter is the attrition-imputation lower-bound specification; the label refers to missing-outcome treatment, not a confidence limit. The source compares the two ages to assess acceleration of attendance. Their similar effects suggest a lasting gain over the observed period, with the reported SEs limiting stronger conclusions. Alternative age-28 estimates of 24.80 and 25.60 percentage points change missing-outcome treatment, not observation age. The enrollment alternative in Hendren and Sprung-Keyser (2020) converts each induced entrant into two Years of Schooling. Type 2 matches $\Delta H(28) = 22.40\%$ and Type 3 also matches $\Delta H(23) = 21.90\%$ exactly; the separate completed-schooling specification, with an input of 0.68 years and reported MVPF of 4.86, is excluded. The comparison level at 23 is 35.20%, the post-termination deceased-father mean in the source's Table 1. It is a sample mean rather than the controlled DID untreated prediction. The age shape uses equally weighted national ACS 1985–1987 some-college profiles, rescaled to this mean; cohort 1963 remains the application index. Transport from some-college attainment to full-time attendance is an additional assumption. The historical 1963 library profile observes ages 37–59 and borrows a 1978 shape observed at 22–35; its backward-extrapolated entry support permits an age-28 effect of only about 0.08. The application instead uses the observed-youth shape, the logical plateau bound $0 \leq K \leq 1$, and the marginal $(\log b, d)$ covariance from the 1983–1987 state donors with cutoff 9.21. The cutoff is a calibration convention. Anchored timing restrictions follow Appendix C.1; the baseline and restrictions are fixed before matching and shared across Type 2 and Type 3.

Writing $q(\tau) = \Delta H(\tau) - \Delta H(\tau - 1)$, the entry-indexed response is

$$\Delta y(a) = 2 \sum_{\tau} q(\tau) \left[r_{-} Y_{-}(a) \mathbf{1}\{\tau \leq a < \tau + 7\} + r_{+} Y_{+}(a) \mathbf{1}\{a \geq \tau + 7\} \right], \qquad (44)$$

with source returns $r_{-} = -6.67\%$ and $r_{+} = 11.26\%$, source income bases Y_{-}, Y_{+}, income ages 18–65, and 3 percent discounting to 18. Early and later tax rates are 3.00×10^{-3} and 0.20. Negative batches reverse the source response; seven years describes its earnings template rather than time actually enrolled.

Let $d_2 = 1 + 1/1.03$, $B = 6700$, $s_0 = 35.20\%$, $m = 22.40\%$, $p = 2968.01$, and $c = 6466.03$, retaining archived precision in calculation. Fixed private and program components are $w_0 = B s_0 d_2 \simeq 4648.11$ and $n_0 = B(s_0 + m) d_2 \simeq 7606.00$. The

source's distinct private and public duration conventions give

$$P = \sum_{\tau} q(\tau) 1.03^{-(\tau-18)}[-2(B-p)], \qquad C = \sum_{\tau} q(\tau) 1.03^{-(\tau-18)}[(c-p)d_2]. \quad (45)$$

With $W = w_0 + E - T - P$ and $N = n_0 + C - T$, a reference batch $q(18) = m$ gives $(E, T, P, C) = (21163.75, 4411.05, -1671.93, 1544.29)$ and reproduces WTP of \$23,072.74 and net cost of \$4,739.23 in 2000 dollars per eligible person (Table D5), without scaling. Type 2 earnings NPVs range from 17.09 to 21.59 thousand dollars; Type 3 narrows them to 18.71–19.87 thousand dollars. The fixed program amount does not respond to alternative grant eligibility or duration. Input repetitions preserve source attendance and return draws; the age-23 coefficient receives the same standardized shock as age 28, imposing perfect cross-window dependence as documented in Appendix E.2.

College appropriations: spending

Deming and Walters (2017) estimate public-college enrollment responses to institutional spending using 1990–2013 institution-year data. Their current-year and next-three-year coefficients describe policy time, not individual follow-up. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 4.00. The retained elasticities are 3.04×10^{-1}, 7.96×10^{-1}, 8.45×10^{-1}, and 8.30×10^{-1}, corresponding to the original study's spending panel across four lag columns (Table 3, Panel A, in the available draft; Table A6 in the reference version). The reference divides their sum by 2.02 and multiplies by 1,000/29,856.23, giving 4.60×10^{-2} Years of Schooling. Assigning the four contributions to ages 18–21, Type 2 imposes $\Delta H(21) = 4.60 \times 10^{-2}$ with anchor 17. Type 2 yields a searched positive-cost MVPF span of $[-5.43, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Type 3 fits the separate annual targets $(5.04 \times 10^{-3}, 1.32 \times 10^{-2}, 1.40 \times 10^{-2}, 1.38 \times 10^{-2})$ for $\Delta H(a) - \Delta H(a-1)$ at $a = 18, \ldots, 21$, preserving their exact total. This requires transporting institutional policy-time responses to a representative age path. All coefficients already enter Type 1's sum; Type 3 adds their sequence, rather than previously unused coefficient values. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 yields a positive-cost MVPF of 16.52, above one. This fitted value lies above Type 1; the narrower span relative to Type 2 also reflects the fitting criterion. Figure D3, College appropriations: spending panel, shows age-65 gaps of $[-0.50, 1.16]$ under Type 2 and 0.12 years under Type 3; Type 2

permits eventual reversal. Paired input repetitions give Type 2 $[-1{,}925.69, +\infty]$; Type 3 $[1.23, +\infty]$. The reported Type 1 95 percent interval is $[1.25, 20.44]$.

Early Head Start

Hess (2025) estimates how access to Early Head Start affects later educational attainment. The source pools adult outcomes in linked Census–ACS data and distinguishes population intent-to-treat effects from participant effects. Hess calculates an MVPF (Type 1) of 0.35. The original study's Table 1.2, full-exposure All/F row and College Grad ATT column, supplies an average treatment effect on the treated (ATT) of 0.17 on BA attainment; the corresponding population effect is 1.20%. The reconstruction uses 0.17 directly as a probability gap, without converting it into Years of Schooling. Type 2 imposes a weighted $\Delta H_c(22) = 0.17$ across 1989–1999 cohorts. Age 22 is a representative endpoint, not a common source follow-up. Because exposure occurs in childhood, no college-age equality anchor is imposed. Both BA profiles remain zero through 19 under the credential model's structural restriction. Type 2 yields a searched positive-cost MVPF span of $[-1.37, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: partial-exposure, sex-specific, and some-college estimates change the target or outcome. The reconstructed public ACS earnings and funded-slot cost account is an extension of the source valuation, with an unresolved original price-year comparison. Figure D8, Early Head Start panel, shows the Type 2 paths; the age-65 gap is $[-0.06, 0.27]$ probability units, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-2.20, +\infty]$.

Texas CC tuition

Denning (2017) estimates how Texas CC district access and lower tuition affect 1994–2005 HS graduates. Enrollment is observed in the graduation year and each of the following six years; BA completion is observed after six and eight years. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 349.51. The reference sums the original study's Table 6 graduation-year and Table 7 subsequent enrollment effects across CCs and universities, giving 0.21 Years of Schooling. Online Appendix Table A4 adds $2(7.80\times10^{-3} - 3.30\times10^{-3}) = 9.00\times10^{-3}$ for later BA completion. Type 2 retains $\Delta H(26) = 0.22$, with graduation assigned to 18 and anchor 17. Type 2 yields a searched positive-cost MVPF span of $[9.70, +\infty]$ and separate paths under which the policy pays for itself, entirely above one among

the paths found; these endpoints bracket Type 1. Type 3 separately fits cumulative targets $(0.03, 0.08, 0.11, 0.15, 0.17, 0.19, 0.21)$ at ages 18–24, keeping the age-26 total exact. These sums preserve the seven reported annual enrollment responses under the reference conversion. The later completion difference supplies the final increment; no year-seven enrollment response is imputed. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 selects a path under which the policy pays for itself. Figure D4, Texas CC tuition panel, shows age-65 gaps of $[-4.00 \times 10^{-3}, 0.80]$ under Type 2 and 0.22 years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-1.81, +\infty]$; Type 3 $[1.14, +\infty]$. The reported Type 1 95 percent interval is $[0.93, +\infty]$.

Adult Pell Grant introduction

Seftor and Turner (2002) estimate how the introduction of federal student aid changes current enrollment among adults ages 22–35. Their repeated-cross-section comparison does not follow individual entrants through later attainment. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 2.18. The original study's Table 2, final panel, columns (1) and (3), supplies enrollment effects of 1.50 percentage points for men and 1.40 points for women. Weights proportional to 47,407 and 62,405 observations yield approximately 1.44×10^{-2}. The reference's two-year duration gives 2.89×10^{-2} Years of Schooling. Type 2 imposes $\Delta H(30) = 2.89 \times 10^{-2}$ with anchor 27, representing exposure beginning at 28. The source's broad adult age range does not identify these specific model ages; they locate the assumed acquisition period in the representative adult path. Type 2 yields a searched positive-cost MVPF span of $[1.95, 4.63]$, entirely above one among the paths found; these endpoints bracket Type 1. No Type 3 is admitted. Sex and age contrasts concern different populations, and splitting the two-year duration would add a timing assumption without another follow-up. The reconstruction consequently holds the converted adult-enrollment input fixed while allowing its later education gap to vary. Figure D4, Adult Pell introduction panel, shows the Type 2 paths; the age-65 gap is $[-0.01, 0.21]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-5.15, 102.12]$. The reported Type 1 95 percent interval is $[0.71, 6.11]$. The source adult-entry income clock and 1973-dollar account are retained. The external 1947 cohort profile is level-calibrated at ages 53–65 using the 1978 national shape observed at ages 22–35; its early ages are backward reconstructions.

College appropriations: tuition

Deming and Walters (2017) also estimate enrollment responses to tuition changes in institution-year data from 1990–2013. Four coefficients cover the current policy year and three subsequent years. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 1.02. The original tuition panel's four lag coefficients are -1.60×10^{-2}, 6.60×10^{-2}, 3.10×10^{-2}, and -7.30×10^{-2} (the available draft's Table 3; the reference version's Table A6). The reference divides their sum by 2.02 and multiplies by 1,000/8,599.86, yielding 4.61×10^{-4} years under its released tuition-cut sign convention. Type 2 imposes $\Delta H(21) = 4.61\times10^{-4}$, with anchor 17 and modeled acquisition ages 18–21. Type 2 yields a searched positive-cost MVPF span of $[-8.86, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Type 3 fits annual changes at those four ages to $(-9.21\times10^{-4}, 3.80\times10^{-3}, 1.78\times10^{-3}, -4.20\times10^{-3})$ while retaining the exact total. The source coefficients already contribute to Type 1; their separate sequence supplies Type 3's extra restriction. Mapping institutional policy time into individual ages is additional to the inherited schooling conversion, and does not establish a longitudinal treatment effect for a fixed entering cohort. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 yields a positive-cost MVPF of -1.97, implying beneficiary losses. This fitted value lies below Type 1; the narrower span relative to Type 2 also reflects the fitting criterion. Figure D4, College appropriations: tuition panel, shows age-65 gaps of $[-0.57, 1.05]$ under Type 2 and -0.09 years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-\infty, +\infty]$; Type 3 $[-4.48, 11.82]$. The reported Type 1 95 percent interval is $[-1.11, 5.13]$.

HOPE tax credit

Turner (2011) estimates tax-aid effects on first-year full-time enrollment and second-year continuation among college-age youth. The second response concerns students who entered their first year in the preceding academic year. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 12.58. The original study's Table 4, columns (1)–(2), Subsidy row, provides probability responses of 2.35×10^{-3} and 2.10×10^{-3} per \$100. The reference combines a \$1,104 subsidy with one year for entry and 1.5 years for continuation: $(1{,}104/100)[2.35\times10^{-3}+1.5(2.10\times10^{-3})] = 6.07\times10^{-2}$. With anchor 17, Type 2 imposes $\Delta H(20.5) = 6.07\times10^{-2}$. The fractional endpoint represents the assumed duration and is evaluated continuously. Type 2 yields a searched positive-cost MVPF span of $[-5.09, +\infty]$ and separate paths under which

the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Type 3 adds $\Delta H(19) = 2.59 \times 10^{-2}$ and $\Delta H(20) = 4.91 \times 10^{-2}$. These respectively allocate the first-year contribution and one additional year of the continuation contribution; the remaining half-year completes Type 2's total. This construction preserves the reference conversion while adding an explicit age allocation and a bridge between the two source risk sets. It does not claim that the source directly estimates three schooling stocks. The constrained fit uses the additional targets while preserving the Type 2 equality. Type 3 selects a path under which the policy pays for itself. Figure D4, HOPE tax credit panel, shows age-65 gaps of $[-0.35, 1.02]$ under Type 2 and 0.44 years under Type 3; Type 2 permits eventual reversal. Paired input repetitions give Type 2 $[-97.46, +\infty]$; Type 3 $[-0.04, +\infty]$. The reported Type 1 95 percent interval is $[-29.15, +\infty]$.

HOPE credit: independent single filers, phase-out start

Bulman and Hoxby (2015) estimate the HOPE attendance response at the phase-out start for independent single filers. The observation is current attendance among ages 19–20 in 2007; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 18.86. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient 3×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 1.16×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 5.19 \times 10^{-5}$ schooling-equivalent years. With anchor 17, Type 2 imposes

$$\{[\Delta H_{1987}(21) - \Delta H_{1987}(19)] + [\Delta H_{1988}(20) - \Delta H_{1988}(18)]\}/2 = 5.19 \times 10^{-5}.$$

These assigned acquisition windows implement the duration conversion; they are distinct from the survey-age window. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D4, HOPE: independent single, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.48, 1.00]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-3.03, +\infty]$.

HOPE credit: dependent single filers, phase-out start

Bulman and Hoxby (2015) estimate the HOPE attendance response at the phase-out start for dependent single filers. The observation is current attendance among ages 19–20 in 2007; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of −1.64. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient -5×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 1.39×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = -7.21 \times 10^{-5}$ schooling-equivalent years. With anchor 17, Type 2 imposes $\{[\Delta H_{1987}(21) - \Delta H_{1987}(19)] + [\Delta H_{1988}(20) - \Delta H_{1988}(18)]\}/2 = -7.21 \times 10^{-5}$. These assigned acquisition windows implement the duration conversion; they are distinct from the survey-age window. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D4, HOPE: dependent single, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.48, 1.00]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-6.45, +\infty]$.

HOPE credit: dependent single filers, phase-out end

Bulman and Hoxby (2015) estimate the HOPE attendance response at the phase-out end for dependent single filers. The observation is current attendance among ages 19–20 in 2007; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 11.83. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient 6×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 1.44×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 8.33 \times 10^{-5}$ schooling-equivalent years. With anchor 17, Type 2 imposes $\{[\Delta H_{1987}(21) - \Delta H_{1987}(19)] + [\Delta H_{1988}(20) - \Delta H_{1988}(18)]\}/2 = 8.33 \times 10^{-5}$. These assigned acquisition windows implement the duration conversion; they are distinct from the survey-age window. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing

zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D5, HOPE: dependent single, upper kink panel, shows the Type 2 paths; the age-65 gap is $[-0.47, 1.00]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-5.36, +\infty]$.

HOPE credit: joint filers, phase-out start

Bulman and Hoxby (2015) estimate the HOPE attendance response at the phase-out start for joint filers. The observation is current attendance among ages 19–20 in 2007; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of $+\infty$. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient 12×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 7.31×10^{-2}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 3.28 \times 10^{-4}$ schooling-equivalent years. With anchor 17, Type 2 imposes $\{[\Delta H_{1987}(21) - \Delta H_{1987}(19)] + [\Delta H_{1988}(20) - \Delta H_{1988}(18)]\}/2 = 3.28 \times 10^{-4}$. These assigned acquisition windows implement the duration conversion; they are distinct from the survey-age window. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; the policy also pays for itself under Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D5, HOPE: joint, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.47, 1.00]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-3.51, +\infty]$.

HOPE credit: joint filers, phase-out end

Bulman and Hoxby (2015) estimate the HOPE attendance response at the phase-out end for joint filers. The observation is current attendance among ages 19–20 in 2007; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 2.37. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope

coefficient 4×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 6.38×10^{-2}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 1.25 \times 10^{-4}$ schooling-equivalent years. With anchor 17, Type 2 imposes $\{[\Delta H_{1987}(21) - \Delta H_{1987}(19)] + [\Delta H_{1988}(20) - \Delta H_{1988}(18)]\}/2 = 1.25 \times 10^{-4}$. These assigned acquisition windows implement the duration conversion; they are distinct from the survey-age window. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D5, HOPE: joint, upper kink panel, shows the Type 2 paths; the age-65 gap is $[-0.47, 1.00]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-2.22, +\infty]$.

Tuition deduction: single filers, phase-start

Hoxby and Bulman (2016) estimate attendance responses around federal tuition-deduction thresholds. This contrast concerns single filers at the deduction's start. The retained outcome is attendance in the first full tax year after expected on-time HS graduation. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of $+\infty$. The original study's Table 5, column (5), baseline specification, supplies the probability effect 0.50%. The reference's two-year enrollment conversion gives 0.01 schooling-equivalent years. Type 2 imposes $[\Delta H(18) + \Delta H(19)]/2 = 0.01$ with anchor 17. Equal weights at ages 18 and 19 are an additional age mapping of the tax-year window; they are neither recovered source weights nor two independently estimated schooling effects. Type 2 yields a searched positive-cost MVPF span of $[-\infty, -7.40]$ and separate paths under which the policy pays for itself, entirely negative among the positive-cost paths; the policy also pays for itself under Type 1. Table 7 contains additional attendance windows around expected graduation. They are not admitted as nested Type 3 because their specification and population have not been reconciled with the retained Table 5 input under a common schooling conversion. The absence of Type 3 therefore does not mean the paper lacks dynamic results. Figure D5, Tuition deduction: single, lower threshold panel, shows the Type 2 paths; the age-65 gap is $[-0.83, 1.55]$ years, allowing either sign of the

late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-\infty, +\infty]$.

Tuition deduction: joint filers, phase-start

Hoxby and Bulman (2016) estimate attendance responses around federal tuition-deduction thresholds. This contrast concerns joint filers at the deduction's start. The retained outcome is attendance in the first full tax year after expected on-time HS graduation. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of -0.02. The original study's Table 5, column (2), baseline specification, supplies the probability effect -3.00×10^{-3}. The reference's two-year enrollment conversion gives -6.00×10^{-3} schooling-equivalent years. Type 2 imposes $[\Delta H(18) + \Delta H(19)]/2 = -6.00\times10^{-3}$ with anchor 17. Equal weights at ages 18 and 19 are an additional age mapping of the tax-year window; they are neither recovered source weights nor two independently estimated schooling effects. Type 2 yields a searched positive-cost MVPF span of $[-4.35, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Table 7 contains additional attendance windows around expected graduation. They are not admitted as nested Type 3 because their specification and population have not been reconciled with the retained Table 5 input under a common schooling conversion. The absence of Type 3 therefore does not mean the paper lacks dynamic results. Figure D5, Tuition deduction: joint, lower threshold panel, shows the Type 2 paths; the age-65 gap is $[-0.88, 1.47]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-4.67, +\infty]$. The reported Type 1 95 percent interval is $[-2.25, 5.52]$.

HOPE and Lifetime Learning tax credits

Long (2004) studies the introduction of HOPE and Lifetime Learning credits using repeated cross sections. The retained result concerns current enrollment among respondents ages 18–24, rather than cumulative attainment at a single age. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of -8.81. The original study's Table 3.11, Panel A, column (2), reports an odds ratio of 0.93; Table 3.10 supplies the baseline enrollment rate of 46.30%. The reference approximation $(0.93 - 1)0.46$ and a two-year duration give -0.06 Years of Schooling. This is the inherited approximation, not the exact odds-to-probability transformation. Type 2 imposes $\Delta H(25) = -0.06$ with an age-17 anchor. Age 25 aligns the converted schooling input with the start of the reference account's long-run earnings returns.

This is an assumed age mapping, not a directly observed age-25 schooling effect; subsequent attainment remains free to change. Type 2 yields a searched positive-cost MVPF span of $[-\infty, -9.59]$ and separate paths under which the policy pays for itself, entirely negative among the positive-cost paths; its positive-cost values fall below Type 1. No Type 3 is admitted. The age-18–19 subsample has a separate odds ratio, but its corresponding baseline enrollment probability has not been established. Reusing the broader sample's baseline would introduce an unverified conversion, so that subgroup is not turned into an additional schooling target. Figure D5, HOPE and Lifetime Learning credits panel, shows the Type 2 paths; the age-65 gap is $[-0.36, 0.50]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-\infty, +\infty]$.

AOTC: dependent single filers, phase-out end

Bulman and Hoxby (2015) estimate the AOTC attendance response at the phase-out end for dependent single filers. The observation is current attendance among ages 19–20 in 2011; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of -0.02. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient -2×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 2.16×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = -1.85 \times 10^{-5}$ schooling-equivalent years. With anchor 17, Type 2 imposes $[\Delta H_{1991}(20) + \Delta H_{1992}(19)]/2 = -1.85 \times 10^{-5}$. Equal age–cohort weights and placement of the duration-equivalent input at the observed ages are explicit reconstruction assumptions. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D6, AOTC: dependent single, upper kink panel, shows the Type 2 paths; the age-65 gap is $[-0.81, 1.43]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-2.25, +\infty]$.

AOTC: dependent single filers, phase-out start

Bulman and Hoxby (2015) estimate the AOTC attendance response at the phase-out start for dependent single filers. The observation is current attendance among ages 19–20 in 2011; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of $+\infty$. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient 43×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 2.02×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 4.26 \times 10^{-4}$ schooling-equivalent years. With anchor 17, Type 2 imposes $[\Delta H_{1991}(20) + \Delta H_{1992}(19)]/2 = 4.26 \times 10^{-4}$. Equal age–cohort weights and placement of the duration-equivalent input at the observed ages are explicit reconstruction assumptions. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; the policy also pays for itself under Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D6, AOTC: dependent single, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.80, 1.43]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-0.84, +\infty]$.

AOTC: joint filers, phase-out end

Bulman and Hoxby (2015) estimate the AOTC attendance response at the phase-out end for joint filers. The observation is current attendance among ages 19–20 in 2011; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of -1.77. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient -7×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 1.15×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = -1.22 \times 10^{-4}$ schooling-equivalent years. With anchor 17, Type 2 imposes $[\Delta H_{1991}(20) + \Delta H_{1992}(19)]/2 = -1.22 \times 10^{-4}$. Equal age–cohort weights and placement of the duration-equivalent input at the observed ages are explicit reconstruction assumptions. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing

zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D6, AOTC: joint, upper kink panel, shows the Type 2 paths; the age-65 gap is $[-0.81, 1.43]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-17.06, +\infty]$.

AOTC: joint filers, phase-out start

Bulman and Hoxby (2015) estimate the AOTC attendance response at the phase-out start for joint filers. The observation is current attendance among ages 19–20 in 2011; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of $+\infty$. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope coefficient 7×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 1.07×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 1.30 \times 10^{-4}$ schooling-equivalent years. With anchor 17, Type 2 imposes $[\Delta H_{1991}(20) + \Delta H_{1992}(19)]/2 = 1.30 \times 10^{-4}$. Equal age–cohort weights and placement of the duration-equivalent input at the observed ages are explicit reconstruction assumptions. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; the policy also pays for itself under Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D6, AOTC: joint, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.81, 1.43]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-5.96, +\infty]$.

AOTC: independent single filers, phase-out start

Bulman and Hoxby (2015) estimate the AOTC attendance response at the phase-out start for independent single filers. The observation is current attendance among ages 19–20 in 2011; it is not a later follow-up of one entering class. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 6.75. The original study's Table 14, relevant filer row and first Attendance column, supplies the signed slope

coefficient 1×10^{-6} per income-margin dollar. Its matching Table 13 credit first stage is 2.27×10^{-1}. A fixed schedule displacement, normalized to a one-dollar credit increase at the point estimate, and the reference two-year duration give $2\beta/\varphi = 8.82 \times 10^{-6}$ schooling-equivalent years. With anchor 17, Type 2 imposes $[\Delta H_{1991}(20) + \Delta H_{1992}(19)]/2 = 8.82 \times 10^{-6}$. Equal age–cohort weights and placement of the duration-equivalent input at the observed ages are explicit reconstruction assumptions. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. No Type 3 is admitted: the source reports one attendance window for this contrast, and splitting its assumed duration would not add an observation. Appendix D.4.2 separates the coefficient, fixed policy unit, and dollar account. Figure D6, AOTC: independent single, lower kink panel, shows the Type 2 paths; the age-65 gap is $[-0.81, 1.43]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-1.61, +\infty]$.

AOTC: simulated instrument

Bulman and Hoxby (2015) use a simulated instrument to estimate how education tax credits affect attendance. The original study's Table 16 reports separate age-19 and age-22 responses over 2006–2011. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 10.05. The reference retains the age-19 coefficient, 1.4×10^{-5} per actual credit dollar, from the first Attendance column. Multiplying by the assumed two-year duration gives 2.80×10^{-5} schooling-equivalent years for a one-dollar policy change. Type 2 imposes the equally weighted $\Delta H_c(20) - \Delta H_c(18) = 2.80 \times 10^{-5}$ restriction across 1987–1992 cohorts, with anchor 17. The positive age-18 gap is not mechanically set to zero. The accounting transfer is one dollar because this is already a second-stage coefficient; Appendix D.4.2 documents that unit correction. Type 2 yields a searched positive-cost MVPF span of $[-\infty, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. The age-22 coefficient concerns different cohorts, so no longitudinal Type 3 is admitted for the retained age-19 population. Type 1 remains the source MVPF; its comparison with Type 2 therefore includes the unit correction as well as the continuation model. Figure D6, AOTC: simulated instrument panel, shows the Type 2 paths; the age-65 gap is $[-0.59, 1.16]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-18.36, +\infty]$.

Tennessee Promise

The attainment regressions in Schoonover et al. (2026) define exposure by eligibility at age 17 and control for state and survey year. The authors allow for completion after age 24 and qualify the projected education returns. Their private-benefit account includes tuition relief for students who would have enrolled without the policy and after-tax earnings for induced students; projected taxes offset scholarship and public education expenditure. The reconstruction retains the reported Type 1 benchmark while documenting the income and participant-unit assumptions below. The reconstruction in Section 4.4 uses fixed equal weights over cohorts 1992–1998, which reach 17 during the county-to-statewide expansion in 2009–2015. These model cohorts do not delimit the source's full regression sample. Tennessee ACS profiles match the HS-graduate denominator and prior-year residence definition in Schoonover et al. (2026). Each profile is rescaled to its source age-23–24 control mean, and treated and comparison paths share the age-17 level. For cumulative some college, AA, and BA, Type 2 imposes

$$\frac{\Delta H_j(23) + \Delta H_j(24)}{2} = m_j, \qquad (m_C, m_A, m_B) = (3.50\%, 4.00\%, 2.70\%). \tag{46}$$

Equal age weights approximate the unavailable regression operator. At every age, both policy states satisfy $H_B^s \leq H_A^s \leq H_C^s$. Each credential has its own anchored parameter triple and state-centered support. The joint Sobol search is supplemented by combinations near earnings extremes and earlier-window fits; separate credential extrema are never combined into a single candidate.

Late completion is bounded relative to those who have not attained the outcome. For donor control level $\bar{H}_{jc}$ at ages 23–24,

$$R_j = \max_{c \in \{1983,\ldots,1987\}} \frac{K_{jc} - \bar{H}_{jc}}{1 - \bar{H}_{jc}}, \tag{47}$$

including the target reference's fraction if larger. The treated restriction is

$$K_j^1 \leq \bar{H}_j^1 + (1 - \bar{H}_j^1) R_j. \tag{48}$$

Donors use the same HS-graduate denominator, and their plateaus are extrapolated. Tables D3 and D2 report levels, completion caps, and the eighteen source point targets. Type 3 preserves the three latest-window equalities and fits the fifteen earlier coefficients by equal-weight least squares in probability units, without SE or permutation bands. Its 42 joint starts are supplemented by separate credential fits.

The table reports every prediction and residual; the maximum absolute residual is 0.67 percentage points. The selected joint path also enters Type 2. Appendix D.6 varies the completion allowance separately from parameter support.

Because the author's age-by-credential income table could not be recovered, the reconstruction uses public ACS 2011–2021 mean wages for employed, out-of-school Tennessee residents under the prior-residence sample, in 2021 dollars. For exclusive category $j \in \{SC, AA, BA\}$, the kernel $K_j(a, \tau)$ equals $-Y_{HS}(a)$ during d_j study years before completion, $Y_j(a) - Y_{HS}(a)$ from completion onward, and zero otherwise or before policy onset. Durations are $(2, 2, 4)$, approximating zero earnings while studying; they do not reproduce the source's 2.3-year participation and five-year BA assumptions as annual paths. Signed increments q_j of $(\Delta C - \Delta A, \Delta A - \Delta B, \Delta B)$ give

$$E = \frac{1}{p} \sum_{a=18}^{64} 1.03^{-(a-18)} \sum_{j} \sum_{\tau} K_j(a, \tau) q_j(\tau). \tag{49}$$

Education extends through 68 for the four-year pre-completion response; graduate degrees receive the BA earnings schedule. Negative exclusive increments represent differences between credential distributions.

The participant-unit bridge uses grant $g = 2711$, private remainder $w_0 = 2124.31$, and entry effect 5.40%: $u = w_0/g$ and $p = 5.40\%/(1 - u) \simeq 0.25$. The interpretation of u as the uninduced participant share is maintained; p is account-implied rather than observed take-up. With earnings scale one, source-rounded exclusive endpoints yield \$59,881.13 and the model's nested endpoint vector yields \$69,261.65, compared with the source-implied \$95,557.69. No factor aligns these totals. The fixed non-income amounts adopt the 2021-dollar convention because their original price vintage is unavailable. Source totals imply

$$w_0 = 72{,}837 - 0.74(24{,}845/0.26) \simeq 2{,}124.31, \tag{50}$$

$$n_0 = -10{,}053 + 24{,}845 = 14{,}792. \tag{51}$$

The one-dollar difference from the displayed source program total of \$14,793 reflects published rounding. Candidates use $W = w_0 + 0.74E$ and $N = n_0 - 0.26E$, retaining unrounded remainders and fixed grant, tuition, and public-resource components. The Type 2 earnings range is -68.71–203.21 thousand dollars. At positive net cost, MVPF exceeds one when $E > \$12,667.69$; positive WTP and nonpositive cost arise when $E \geq \$56,892.31$. These restrictions and the reconstructed wage, unit, and price conventions distinguish the comparison with Type 1 from a change in continuation alone. Source clustered SEs enter only the separate normal-input

repetitions in Appendix E.2; those repetitions do not reproduce the author's permutation distribution.

Algorithmic college placement

Bergman et al. (2026) randomize algorithmic rather than test-only course placement at seven New York CCs. Entering cohorts have unequal follow-up through 2018–2019. Bergman and coauthors calculate an MVPF (Type 1) of $+\infty$. The June 2026 study's Table 3, column (3), reports 5.62×10^{-1} additional college credits for the pooled sample. The reconstruction uses $5.62\times10^{-1}/24 = 2.34\times10^{-2}$ schooling-equivalent years. The authors' own MVPF uses placement benefits and costs while omitting uncertain lifetime earnings; the reconstruction adds the fiscal earnings channel but preserves their private valuation. Type 2 imposes $\Delta H(23) = 2.34\times10^{-2}$ with anchor 20 and first possible exposure at 21. These are representative ages; the pooled coefficient is not a common two-year follow-up. Type 2 yields a searched positive-cost MVPF span of $[0.08, +\infty]$ and separate paths under which the policy pays for itself, crossing one; the policy also pays for itself under Type 1. No Type 3 is admitted. The 1.72-credit result for the more mature fall-2016 cohort changes the population relative to pooled Type 2, and first-term course-placement outcomes are not total earned-credit stocks. Thus, this remains an extension of the source welfare account, not a comparison that changes only the education continuation. Figure D7, Algorithmic college placement panel, shows the Type 2 paths; the age-65 gap is $[-0.31, 0.61]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[0.07, +\infty]$.

Tuition deduction: single filers, phase-end

Hoxby and Bulman (2016) estimate attendance responses around federal tuition-deduction thresholds. This contrast concerns single filers at the deduction's end. The retained outcome is attendance in the first full tax year after expected on-time HS graduation. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of $+\infty$. The original study's Table 5, column (6), baseline specification, supplies the probability effect 0. The reference's two-year enrollment conversion gives 0 schooling-equivalent years. Type 2 imposes $[\Delta H(18) + \Delta H(19)]/2 = 0$ with anchor 17. Equal weights at ages 18 and 19 are an additional age mapping of the tax-year window; they are neither recovered source weights nor two independently estimated schooling effects. Type 2 yields a searched positive-cost MVPF span of $[-\infty, -6.33]$ and separate paths under which the policy pays for itself, entirely negative among

the positive-cost paths; the policy also pays for itself under Type 1. Table 7 contains additional attendance windows around expected graduation. They are not admitted as nested Type 3 because their specification and population have not been reconciled with the retained Table 5 input under a common schooling conversion. The absence of Type 3 therefore does not mean the paper lacks dynamic results. Figure D7, Tuition deduction: single, upper threshold panel, shows the Type 2 paths; the age-65 gap is $[-0.86, 1.50]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-\infty, +\infty]$. The reported Type 1 95 percent interval is $[-\infty, +\infty]$.

Tuition deduction: joint filers, phase-end

Hoxby and Bulman (2016) estimate attendance responses around federal tuition-deduction thresholds. This contrast concerns joint filers at the deduction's end. The retained outcome is attendance in the first full tax year after expected on-time HS graduation. Hendren and Sprung-Keyser (2020) calculate an MVPF (Type 1) of 0.77. The original study's Table 5, column (3), baseline specification, supplies the probability effect 0. The reference's two-year enrollment conversion gives 0 schooling-equivalent years. Type 2 imposes $[\Delta H(18) + \Delta H(19)]/2 = 0$ with anchor 17. Equal weights at ages 18 and 19 are an additional age mapping of the tax-year window; they are neither recovered source weights nor two independently estimated schooling effects. Type 2 yields a searched positive-cost MVPF span of $[-3.92, +\infty]$ and separate paths under which the policy pays for itself, crossing zero and one; these endpoints bracket Type 1. Table 7 contains additional attendance windows around expected graduation. They are not admitted as nested Type 3 because their specification and population have not been reconciled with the retained Table 5 input under a common schooling conversion. The absence of Type 3 therefore does not mean the paper lacks dynamic results. Figure D7, Tuition deduction: joint, upper threshold panel, shows the Type 2 paths; the age-65 gap is $[-0.86, 1.50]$ years, allowing either sign of the late-life effect. Paired input repetitions give Type 2 $[-4.87, +\infty]$. The reported Type 1 95 percent interval is $[-1.83, 57.84]$.

Cannabis legalization: education channel

Li (2026) estimates recreational-cannabis legalization effects on BA completion among first-time, full-time entrants, using 2006–2017 entering cohorts and post-entry policy exposure. The original manuscript's Table 3 reports effects of -1.60, -1.10, and -1.80 percentage points on completion within four, five, and six years.

There is no original MVPF (Type 1). The education-only account assigns two additional Years of Schooling per BA completion, yielding targets -3.20×10^{-2}, -2.20×10^{-2}, and -3.60×10^{-2}. The published reconstruction is Type 3 only. It anchors at 17, assigns entry to 18, and places the targets at ages 22, 23, and 24, equally weighting 1988–1999 calibration cohorts. The age-24 target is exact; the earlier two enter constrained minimum distance. This common onset is a modeling assumption despite heterogeneous actual exposure after entry. All three comparable BA windows enter the primary reconstruction. The two-year-institution outcome concerns another entrant population and is excluded. A six-year-only Type 2 is retained solely as a computational diagnostic and is not treated as a second substantive result. Type 3 yields a positive-cost MVPF of -55.85, implying beneficiary losses. Figure D7, Cannabis legalization: education channel panel, shows the Type 3 paths; the age-65 gap is -0.10 years, preserving a negative late-life gap. Paired input repetitions give Type 3 $[-57.03, -54.19]$.

D.3 Additional Observations and Numerical Verification

Additional observations constrain a common path only when their population, policy contrast, education outcome, and clock can be reconciled with Type 2. Table D1 distinguishes reported timing from its use in reconstruction. Among the 37 calculations with Type 1, 17 have one verified window, five have two, and eleven have at least three; four use institutional or college-stage clocks. The cannabis application reports three completion horizons. Each reported age interval counts once, including overlapping two-year windows, without implying independent information; additional credentials or specifications at the same window do not increase the count. Type 3 retains source windows or dynamic responses under the stated conversion and population mappings. Splitting one enrollment coefficient across assumed study years supplies no additional observation. The case records explain why some studies with several reported windows remain Type 2 only. Thirteen calculations admit Type 3. Michigan, CUNY, and Social Security student benefits have verified exact-match banks; the other ten retain one best-fitting path each, with nonzero additional-input residuals. All fits preserve the Type 2 equalities and use fixed source point estimates with equal weights. Table D2 reports the targets, fitted discrepancies, and welfare values. A fitted singleton is a conditional selection, not evidence of point identification or absence of sampling uncertainty. Every retained Type 3 path also enters Type 2, so the recorded component ranges

are nested. Appendix E separately treats 1,000 joint input repetitions and their dependence assumptions; source SEs neither relax nor weight the point-input fits.

Welfare comparisons pair components from the same complete path. A positive-cost crossing of MVPF one requires paths with $N > 0$ and, respectively, $W < N$ and $W > N$. The stronger comparison between beneficiary losses and a policy that pays for itself pairs $W < 0, N > 0$ with $W > 0, N \leq 0$. Table D5 gives selected paths and their paired components. The separate extrema in Table 3 cannot be combined into a ratio. Numerical checks evaluate each saved path through age 72, or 68 for Tennessee, while valuing income over its registered horizon. They verify anchors, monotonicity, credential nesting, observation operators, empirical support, and welfare accounts. Shape grids are supplemented by multiple-start optimization and searches for zero WTP, WTP equal to cost, zero net cost, and selected ratios. Decreasing positive-cost targets check one-sided ratio limits at zero-cost roots with nonzero WTP. These procedures also apply to the registered input repetitions. They verify the reported candidates without certifying global endpoints, connectedness, or infeasibility when a search fails. Figures D2–D9 display H^0 and H^1 for all 38 calculations. Black denotes the comparison path; blue and red denote pointwise Type 2 and Type 3 treated ranges. Coincident fitted boundaries represent a retained singleton. Envelope edges can switch candidate identity across ages and need not constitute feasible paths. Neither these envelopes nor spans of welfare values imply that every intermediate path or value has been verified. No probability distribution is assigned to the continuation class; complete paths and observation operators remain in the replication materials.

D.4 Monetary Units and Source-Specific Accounts

For each calculation, Table 3 divides WTP and net cost by the same positive reference program-cost component wherever its provenance is established. This changes neither MVPF nor comparisons across paths within that account. The divisor follows the source policy unit, which may be a marginal tax-credit change rather than a program participant; a large normalized component therefore need not imply a large dollar benefit. Appendix C gives the common unit and replacement rules. Table D5 reports the main examples' reference accounts and selected candidate components; Supplementary Figure S4 retains the native-unit spans. Reference ratios and their source components remain separately recorded when rounding or reconstruction prevents exact agreement.

Early Head Start requires an explicit distinction between the library account and the public-data reconstruction. The library reports earnings NPV of 21,944, WTP of 3,906, and net cost of 11,216, with fiscal components 24,257+1,540−10,478−4,103 = 11,216. The dissertation's net cost of approximately 12,755 appears to count the 1,540 CC term twice; its price year remains unresolved. The reconstruction uses person-weighted wages, including zeros, for BA holders and for HS graduates or General Educational Development (GED) credential holders in the 2017–2019 national ACS, converted to 2000 dollars. It values the retained ATT on BA attainment of 0.17 through its path, with age 22 as an assumed placement of the pooled adult outcome. No college-age negative income coefficient is imposed on this childhood intervention. Program cost is the discounted three-year value of fiscal-year (FY) 2001–2008 statutory allocations per funded slot, 28,158.29, an average-cost proxy. Fixed private education payments of 11,463.05 and a public college-cost change of −5,680.05 imply $W = E - 11{,}463.05$ and $N = 22{,}478.24 - \vartheta E$, where $\vartheta = 18.70\%$. The source convention uses gross earnings in WTP while also valuing fiscal revenue. College prices require a stated 2021 publication-price convention before conversion to 2000 dollars. A constant degree gap and an age-22 step give earnings NPVs of 78,683.28 and 77,910.64, respectively, compared with the source 21,944; no multiplier aligns these totals. Input repetitions condition on the reconstructed schedules. Algorithmic placement instead retains source WTP of 220 and net cost of −90 as its reference, adding only the fiscal earnings channel. These comparisons with Type 1 include the disclosed account differences as well as continuation.

D.4.1 Cannabis Education Account

The education inputs and age mapping for cannabis are documented in Appendix D.2. Its calibration uses 14 adopting states and the 1988–1999 birth cohorts, with equal cohort weights and fixed state weights from the age-18 population; these are external calibration weights rather than recovered weights from the causal estimator. Let $E_i = \sum_a D_i(a)\Delta y_i(a)$ be net education earnings, including the signed pre- and post-completion response, and let F_i be the discounted schooling-equivalent acquisition flow priced by the education-resource schedule. With tax revenue $T_i = t_i E_i$, the account is

$$W_i^{\text{edu}} = E_i - T_i - p_i F_i, \tag{52}$$

$$N_i^{\text{edu}} = g_i F_i - T_i, \tag{53}$$

where p_i is private payment and g_i is public education-resource cost net of that payment. The schedule uses $t_i = 18.90\%$, $p_i = 5{,}997.02$, and $g_i = 8{,}557.30$, in 2000 dollars, without a separate forgone-work charge. Prices come from the 2015 public four-year-college inputs in the replication files of Hendren and Sprung-Keyser (2020), aggregated by full-time-equivalent enrollment within states, deflated, and combined using the fixed state weights. Pricing one schooling-equivalent increment as one full-time-equivalent education year and treating average expenditure as marginal cost are calibration assumptions; attainment does not identify actual enrollment or time out of work. The retained Type 3 path gives earnings NPV $-\$3,635.92$, WTP $-\$2,498.50$, and positive net education cost \$44.73. Tax transfers cancel in the resource comparison, $W_i^{\text{edu}} - N_i^{\text{edu}} = E_i - (p_i + g_i)F_i$. Appendix D.6 refits alternative sector prices, state weights, support, and income responses; Appendix E.2 describes the input perturbations. Cannabis excise revenue, health, crime, enforcement, and consumption value remain outside this education-channel account.

D.4.2 Finite Tax-Policy Experiments

The eleven HOPE/AOTC reconstructions convert local education responses into finite policy changes. Table 14 of Bulman and Hoxby (2015) reports attendance–income slope changes at credit kinks for ages 19–20, in 2007 for HOPE and 2011 for AOTC; Table 13 supplies the credit first stages. For signed attendance coefficient β_i and credit first-stage point estimate $\varphi_i > 0$, fix the schedule displacement $q_i = 1/\varphi_i$, giving a one-dollar credit increase at the point estimate. Transporting the local education and accounting responses linearly over this displacement and retaining the source two-year duration gives

$$m_i^{\text{Years of Schooling}} = 2\beta_i q_i = \frac{2\beta_i}{\varphi_i}. \tag{54}$$

This is a schooling-equivalent calibration, not an observed completed-schooling effect. The Table 16 simulated-instrument coefficient is already per actual credit dollar and uses $q_i = 1$.

Table D1 gives the source-age cohorts. Equal cohort-cell weights are maintained because the source regression weights are unavailable. Under the HOPE acquisition convention, enrollment at age a contributes in years a and $a+1$, so its converted total restricts $\Delta H_c(a+1) - \Delta H_c(a-1)$. The joint input is

$$\tfrac{1}{2}[\Delta H_{i,1987}(21) - \Delta H_{i,1987}(19)] + \tfrac{1}{2}[\Delta H_{i,1988}(20) - \Delta H_{i,1988}(18)] = m_i^{\text{Years of Schooling}}. \tag{55}$$

The AOTC kink convention instead places the duration-equivalent total in the observed age cells, $[\Delta H_{i,1991}(20) + \Delta H_{i,1992}(19)]/2 = m_i^{\text{Years of Schooling}}$. The simulated-instrument operator averages $\Delta H_c(20) - \Delta H_c(18)$ across its six cohorts observed at age 19 during 2006–2011. These temporal allocations are maintained assumptions; a second allocated acquisition year does not create Type 3. The simulated instrument's age-22 estimate concerns another cohort composition. All paths retain the age-17 anchor and common level and support restrictions, with a common treated parameter triple across the cohort-specific baselines. The independent-filer income clock beginning at 25 remains a valuation assumption.

For the ten kink contrasts, fixed non-earnings components receive the same displacement, $w_{0i}^F = q_i w_{0i}^{\text{source}}$ and $n_{0i}^F = q_i n_{0i}^{\text{source}}$, while per-schooling-year valuation schedules retain their source units. The cohort-average gap enters

$$W_i^F(h) = w_{0i}^F + \sum_a w_i(a)\overline{\Delta H}_i(a), \tag{56}$$

with $N_i^F(h) = n_{0i}^F + \sum_a n_i(a)\overline{\Delta H}_i(a)$. Reference WTP, net cost, and the display divisor receive the same normalization, leaving Type 1 ratios and intervals unchanged. The simulated-instrument account separately corrects the direct credit payment from approximately 0.38 to one dollar: at each repetition it adds $1 - \text{tax_value}$ to cost and $(1-\text{tax_value})\min(1-\beta_i, 1)$ to non-earnings WTP. Its Type 1 comparison therefore includes this correction.

The point normalization q_i stays fixed in all 1,000 joint input repetitions, including five negative first-stage draws. Those draws are retained without dividing education by the new first stage; their credit effects need not remain one dollar. Simulated-instrument repetitions are a constructed reference-input exercise, not draws from the source sampling distribution. Currency conversions preserve the physical policy experiment, whereas changing its dose changes the attainment restrictions. Section D.6 checks credit changes of 0.50, 1, and 2 with education inputs and non-earnings components scaled together. Finite WTP and cost can still yield infinite ratio limits as positive cost approaches zero; these differ from candidates in which the policy pays for itself. Source coefficients, operator weights, and full-precision paths remain in the replication records.

D.5 Candidate-Based Lower-Endpoint Ordering

Figure 9 and Table D4 rank the 37 calculations with Type 1, selecting Type 3 for twelve and Type 2 for the other twenty-five. Cannabis has no whole-policy

reference. For the defined MVPF endpoints $(\underline{M}_i, \overline{M}_i)$, including audited one-sided limits and the convention for policies that pay for themselves, let r_2 round finite values to two decimals and leave infinities unchanged. The ordering is ascending in

$$\Big(-r_2(\underline{M}_i), -r_2(\overline{M}_i)\Big). \tag{57}$$

Equal pairs share a competition rank, one plus the number of strictly preferred calculations. Reference ranks preserve Type 1 ties. Three- and four-decimal rounding leave the current ranks unchanged. An open marker substitutes that calculation's Type 2 set while holding the others fixed. Added information improves ten displayed ranks and leaves two unchanged under this rule; it does not imply higher true welfare. Thirteen calculations contain nonpositive-WTP, nonpositive-cost paths. Only these unassigned-ratio paths are excluded, while their components remain reported. Best-fit ranks condition on the fitting rule, and all endpoints condition on the finite search. The ordering assigns no probabilities to paths and is neither a social-welfare ranking nor a certified maximin rule. Account extensions and reconstructed valuations remain identified in the table.

D.6 Point-Input Sensitivity across All Applications

Table D6 reports the common and case-specific checks for all 38 calculations, selecting Type 3 for thirteen, including cannabis, and Type 2 for twenty-five. Narrower and wider support multiply the empirical covariance by 0.5 and 2 and timing deviations by their square roots. The geographic alternative uses national profiles for the same cohorts; cases already using national profiles retain their baseline calibration. Tennessee retains its HS-graduate denominator and source-control matching in the common checks. Income alternatives scale net earnings and associated tax effects by 0.75 and 1.25, including negative and positive contributions, while holding transfers and education-resource schedules fixed. STBF's multiplier is applied after evaluating its nonlinear author model; Early Head Start and Tennessee use their disclosed public reconstructions. Each specification preserves the fixed education inputs and common logical restrictions, reoptimizing Type 2 and exact-match Type 3 ranges and refitting the ten minimum-distance Type 3 selections. The denser scalar search uses 2^{16} shape pairs, supplemented joint pools, and additional starts. Wider support need not expand a minimizer set, even though it expands feasible paths.

Positive-cost MVPFs cross one in 21 main selected sets, compared with 20 under narrower support, 22 under wider support, 20 with national calibration,

and 21 under either income multiplier or the denser search. Main and denser searches find paths under which the policy pays for itself in 24 sets. Nineteen main sets contain both a negative-WTP, positive-cost path and a positive-WTP, nonpositive-cost path; sixteen retain both types in all seven common specifications. These are classifications of verified candidates, with several calculations sharing a source study. Fitted singletons reflect the loss criterion as well as additional observations, and the searches do not certify global bounds or uniqueness.

The table's case-specific panel retains STBF's printed 0.36-year input and removal of its author-model experience delay; Adams's zero negative-response template; Tennessee's alternative control matching and late-completion allowance; credit doses of 0.50 and 2; Early Head Start's FY2001 funded-slot cost; and cannabis education-cost, state-weight, and income-response alternatives. Tax-dose checks scale education and fixed monetary components jointly, without asserting invariant endpoints. Across the stated checks, the positive-cost MVPFs found remain above one for STBF and Social Security, below one for Adams, and Tennessee continues to pay for itself in the reconstructed account. STBF's denser-search lower endpoint changes from 1.25 to 1.24, illustrating numerical sensitivity. All comparisons preserve paired monetary components and the registered account differences. They vary assumptions one at a time and remain separate from the input-uncertainty exercise in Appendix E.

E Input Uncertainty and Conditional Endpoint Summaries

E.1 Conditional Endpoint Summaries and Usable Repetitions

The input-uncertainty exercise uses 1,000 prescribed repetitions for each of the 38 calculations and 1,000 paired repetitions for each of the thirteen Type 3 specifications. The six-year-only cannabis calculation is a diagnostic; its substantive result uses Type 3. Comparison profiles, empirical support, education conversions, age mappings, and the functional form of valuation remain fixed. For repetition r, let $\underline{M}_{i\ell}^{(r)}$ and $\overline{M}_{i\ell}^{(r)}$ be the lowest and highest defined MVPFs among verified paths for calculation i and reconstruction ℓ. A fitted singleton has equal endpoints. The

reported summary is

$$\left[q_{2.5\%}\left(\{\underline{M}_{i\ell}^{(r)}\}_{r\in\mathcal{R}_{i\ell}}\right), q_{97.5\%}\left(\{\overline{M}_{i\ell}^{(r)}\}_{r\in\mathcal{R}_{i\ell}}\right)\right], \tag{58}$$

where $\mathcal{R}_{i\ell}$ contains repetitions with verified paths and a defined ratio. Table 4 reports these summaries and usable counts; Table E1 separately summarizes WTP and net cost using all repetitions with verified monetary components. Every ratio pairs components from the same path and repetition. Positive-cost limits at zero net cost remain distinct from paths under which a policy pays for itself, following Section 4.1.[44]

All 38,000 reduced-information repetitions produced verified paths and defined MVPFs. Of 13,000 Type 3 repetitions, 12,980 produced verified paths; the 20 unsuccessful searches concern Social Security, whose endpoint percentiles therefore use 980 repetitions. Unsuccessful searches are not replaced or interpreted as proof of infeasibility. The point-input fitting rules and incorporation of verified Type 3 paths into the paired Type 2 set follow Appendix D.3. These percentiles are conditional summaries: finite searches and imposed input-dependence assumptions do not establish confidence coverage for the continuous path class, and external ACS calibration is not resampled.

Reported Type 1 intervals are available for 34 calculations, using the Policy Impacts Library snapshot of 20 September 2026. Social Security uses the enrollment-alternative interval $[1.16, 32.23]$; Tennessee's $[+\infty, +\infty]$ comes from Schoonover et al. (2026), Section 5.2.2. Comparable reported intervals were not recovered for STBF, Early Head Start, or algorithmic placement; cannabis has no original whole-policy MVPF. Differences from reconstructed summaries can reflect the inputs varied, accounting conventions, dependence assumptions, and coverage targets.

E.2 Input Perturbations and Dependence

The 33 accounts based on Hendren and Sprung-Keyser (2020) retain their paired source and accounting draws, including earnings-return inputs where available. Constructed education perturbations are used for STBF, Early Head Start, algorithmic placement, the simulated-instrument input, and Tennessee. DC tuition assistance varies accounting inputs while holding its education effect fixed because

[44] Quantiles use the inverse empirical distribution, without interpolation between finite and infinite values. With 1,000 usable repetitions, the endpoints are the 25th minimum and 975th maximum order statistics.

a corresponding source SE is unavailable. The finite-policy tax normalization remains fixed across repetitions. Original correlation groups, signs, and marginal SEs are preserved; additional Type 3 coefficients extend these groups according to their outcomes, samples, and windows. These extensions impose dependence where source covariances are unavailable. Type 2 uses the same retained-input draw as Type 3.

Florida's five cumulative-credit inputs take the form $\widehat{m}_j + SE_j Z^{(r)}$, sharing the fourth-year shock across horizons. Earlier Kalamazoo credits, CUNY first-year credits, and Social Security age-23 attendance likewise share their retained endpoint shocks. Wisconsin has independent groups for cohort 1 and pooled cohorts 2–3, with added annual coefficients sharing the corresponding group shock; scales constructed from marginal semester SEs remain imposed upper bounds where covariance is unavailable. Adams retains the source four-year and two-year outcome groups and extends them to immediate entry and four-year completion. Other source-based Type 3 inputs preserve their registered primitive draws through the conversion and age operators. STBF uses one common normal shock across its four marginal education SEs, holding recovered author income coefficients and its stated aggregation bridge fixed (Appendix D.2).

Tennessee uses one normal shock across the six windows within each cumulative credential, with three independent credential groups. All eighteen targets are perturbed and refitted while the three perturbed latest-window inputs remain exact. Marginal scales use displayed clustered SEs; two BA coefficients whose SEs are printed as zero receive a 5.00×10^{-4} reporting-precision floor. This imposed scale and the normal draws do not recover the source permutation distribution. Cannabis extends the registered 15-cluster synthetic Rademacher shock for the six-year effect to all three BA horizons, scaled by their marginal SEs. These are constructed input perturbations, not recovered cluster influence functions or a new microdata bootstrap of Li (2026); they condition on the education account in Appendix D.4.1.

Perfect positive dependence within a group need not produce conservative welfare ranges. With equal marginal scales it removes uncertainty from coefficient differences, which can determine education timing. The summaries therefore remain conditional on the specified grouping; recovering joint source estimates would allow that grouping to change without altering the point-input framework.

F Policy-Specific Male and Female Education Paths

F.1 Subgroup Inputs and Calibration

Published male and female degree effects support within-policy comparisons for seven interventions. Table F1 reports the inputs, available SEs, and standardized education NPVs when terminal effects may vary. The subgroup outcomes and populations can differ from the pooled calculations in Appendix D.2. These exercises retain their own estimands and do not replace the pooled welfare accounts. Adams uses four-year BA completion from Cohodes and Goodman (2014), Table 7, Panel D, column 5, for the 2005–2008 HS classes, rather than the pooled account's 2005–2006 six-year sample. STBF uses six-year BA effects for four-year applicants from Angrist et al. (2022), Online Appendix Figure A4. Earlier outcomes concern changing samples and are not combined into subgroup Type 3. Wisconsin uses the first cohort's administrative-sample four-year BA outcome and treatment-by-male interaction in Goldrick-Rab et al. (2016), Table 7, rather than the smaller financial-aid sample or the pooled credit outcome. Its covariance information is insufficient to recover the male subgroup SE.[45] Michigan uses the BA effect of in-district residence in Acton (2021), Table 7, Panel B, column 7, rather than a per-dollar tuition slope. The sample covers the 2009–2011 HS classes, excludes degrees earned in HS, and observes completion through 2017; enrollment results use different cohorts. Early Head Start uses the full-exposure college-graduation ATT in Hess (2025), Table 1.2, rather than the population ITT. Its pooled adult outcome is assigned age 22. Tennessee uses the textual BA effects at ages 23–24 and earlier windows in Figure 4F of Schoonover et al. (2026). Kalamazoo uses the within-district IV estimates of six-year BA completion in Bartik et al. (2021), Table 8, rather than the between-district estimand or the pooled attempted-credit input.[46]

The Census–ACS calibration follows Appendix C.1, with means fitted separately by sex and cohort and fixed cohort weights shared across sexes within a policy. College-aid references use HS graduates; Early Head Start uses the full

[45] STBF's reported gender contrast is 1.00×10^{-3} with an SE of 5.00×10^{-3}; Wisconsin's treatment-by-male interaction rounds to zero. Identical displayed subgroup inputs therefore need not be identical underlying estimates.

[46] The source tests for Adams, Michigan, and Kalamazoo do not reject equality of male and female coefficients. Tennessee reports significance at 10 percent for women but not men; this does not test equality between the effects.

population and national 1989–1999 cohorts. Geography is current residence except for Tennessee's prior-year residence, whose 1992–1998 mixture and HS-graduate denominator follow the main case. Adams maps birth cohorts 1987–1990 to completion at 22; Wisconsin uses birth cohort 1990 and age 22. STBF retains the pooled case's 2012–2014 entry cohorts and entry-age mixture, placing six-year effects at ages 23, 24, and 25; subgroup entry-age distributions were not recovered. Michigan uses birth cohorts 1991–1993 with equal weights at ages 26, 25, and 24. Kalamazoo uses cohorts 1988–1990 and age 24, rescaling state sex-specific shapes to pre-policy eligible-group means of 29.50% and 30.40%. These geographic and age mappings remain population-transport assumptions.[47]

Anchored continuations obey the probability-level restrictions in Appendix C.1. College-aid anchors follow calibrated entry ages; Early Head Start assumes zero BA attainment through age 19, without ruling out earlier childhood-treatment effects. For subgroup $s \in \{m, f\}$,

$$\mathcal{M}_i\{H_s^1 - H_s^0\} = \widehat{\delta}_{is}, \tag{59}$$

where the operator preserves the component ages and weights above. The geographic support covariance is shared across sexes and centered at each reference. A common cap limits the fraction without a BA at the observed window who can complete one later. It takes the largest residual completion share across both subgroup references and their 1983–1987 sex-specific donor curves, evaluated with the source window weights, and is fixed before search and identical across sexes.

Tennessee additionally screens its five earlier subgroup moments against intervals recovered from Figure 4F, retaining the textual latest-window estimates exactly.[48] Both endpoint-only and screened results are saved. This auxiliary graphical-interval sensitivity is separate from the pooled Type 3 fit to fixed Table 2 points; the intervals are neither normal SEs nor a joint 95 percent confidence region. Other policies receive no extra subgroup moment from allocating one coefficient across assumed years or combining different outcomes or samples.

[47]The Census 2000 and ACS 2001–2022 extract retains valid person weights and education codes without race or birthplace restrictions. Available means through age 35 enter the national-shape shrinkage rule; recent cohorts borrow a same-sex shape and fit their level. Adams, STBF, Wisconsin, Michigan, and Early Head Start use 96 successful reference and donor fits, with maximum weighted RMSE of about 7.5 percentage points. Supplementary Figures S5–S7 retain the seven policies' paths.

[48]The extraction retains image coordinates and scale, rounds interval limits outward to 1.00×10^{-3}, adds a 2.00×10^{-3} graphical-resolution allowance, and includes marker pixels extending beyond short intervals.

F.2 Common-Terminal Comparisons and Results

To separate acquisition timing from eventual attainment, I value BA effects using the common net-income mapping in equations (25)–(26), with two schooling-equivalent years per BA increment relative to some college:

$$V_{is}^{\mathrm{BA}} = 2E^{\mathrm{path}}\left(\Delta H_{is}^{\mathrm{BA}}\right). \tag{60}$$

Valuation prices are common across sexes and policies, in 2000 dollars discounted at 3 percent to age 18. Education extends through 72 to include the seven-year pre-completion response in earnings through 65; early losses enter once. This standardization excludes sex-specific returns, other credentials, transfers, and fiscal costs, so it is an education-channel NPV rather than a full subgroup MVPF. Table F2 further imposes

$$\Delta H_{im}^{\mathrm{BA}}(65) = \Delta H_{if}^{\mathrm{BA}}(65) = \eta, \tag{61}$$

using $\eta \in \{-0.15, -0.10, -0.05, 0\}$ for Adams and $\eta \in \{0.02, 0.05, 0.10, 0.15\}$ elsewhere. Each female-minus-male NPV comparison pairs complete paths at the same η, prices, and exposure-age weights. The table reports every grid value and verified path count. Its common-reference sensitivity replaces both reference curves by their equally weighted mean within each component, retaining the inputs, age weights, support, residual cap, and terminal grid.

Six policies yield verified pairs under separate reference curves at one or more grid values. Adams, STBF, Wisconsin, and Michigan permit both signs of the NPV difference at an individual fixed terminal effect; Early Head Start and Tennessee have positive differences among the saved pairs. At a five-percentage-point terminal effect, Tennessee's female NPV exceeds its male NPV by approximately \$46–\$277; the common-reference comparison has the same direction. Kalamazoo has no verified pair on the separate-reference grid, but its common-reference comparison supports a ten-point terminal effect with higher female NPV. STBF and Wisconsin permit both NPV orderings under the common reference despite identical rounded inputs. Table F2 also records unmatched grid values, including Wisconsin's absence of a verified female path at five points under separate references. The comparisons condition on published subgroup point estimates, their rounding, external calibration, and common valuation prices. Pooled-policy input repetitions do not supply subgroup confidence intervals. Finite searches verify possible path pairs without establishing global bounds, probabilities, or incompatibility when

no pair is found. The resulting differences do not identify sex-specific education or welfare effects beyond the source estimands, nor do they establish marriage, childbearing, or other mechanisms behind attainment timing.

Table A1: Reported education windows in the literature and in lifetime valuation

A. Reported effect-estimation times or windows

	Full sample				Journal sample	
Outcome	Eligible	Coded	One	One (%)	One/coded	One (%)
All education outcomes	459	459	249	54.2	244/447	54.6
Some college	237	236	193	81.8	189/231	81.8
AA	40	40	32	80.0	32/39	82.1
BA	190	190	132	69.5	130/186	69.9

B. Timing of the valued education outcome in 31 studies

Time measure	One window	Multiple windows	Studies
Attained age	5	1	6
Individual follow-up	9	11	20
Education stage only	4	0	4
Institutional event time	—	—	1
Total	18	12	31

Notes: Panel A counts studies in overlapping outcome groups; one some college study lacks a usable timing count. Journal results exclude the 11 additional welfare-source studies and one author manuscript. One estimate may pool several years. Panel B counts comparable observations of the education outcome used in valuation; when several outcomes enter, it uses the series with the most comparable times. The institutional-flow study is separate from individual attainment windows. Counts describe reported evidence, not every input used in valuation. Appendix A.1 gives the definitions.

Table B1: ACS cohort support and comparison-path fits

Birth cohort	Records	Ages	K	b	d	RMSE
1985	630,409	16–35	13.89	0.25	12.51	0.11
1986	659,595	16–35	13.90	0.26	12.62	0.11
1987	688,856	16–35	13.97	0.26	12.57	0.12

Notes: The selected ACS cohort–age cells occur in 2001–2022. Each cohort contributes all 20 ages. Fits minimize weighted squared differences from person-weighted cell means, using normalized cell population weights within each cohort. The comparison curve averages the three fitted curves with fixed one-third weights. RMSE is in schooling years. Parameter estimates and fit diagnostics are descriptive inputs, not causal policy effects.

Table B2: Scenario design and complete age-25 scans

Parameter	Feasible interval	Age-25 effect interval (Years of Schooling)	Candidate NPV width			Both signs found
			Minimum	Median	Maximum	
K	$[13.34, 15.14]$	$[-0.49, 1.02]$	21.92	48.62	65.92	56/102
b	$[0.12, 0.40]$	$[-1.13, 0.39]$	43.83	58.27	72.12	91/102
d	$[8.00, 30.00]$	$[-2.19, 0.08]$	55.49	65.54	91.01	102/102

Notes: The reference parameters are $(K, b, d) = (13.92, 0.26, 12.56)$. For interval $[L, U]$, the generating value is $L + u(U - L)$ and the other two parameters remain fixed. The age-25 scans use $u = 0, 0.01, \ldots, 1$ plus the reference value (102 scenarios per row); the observation-design experiment uses $u = 0, 0.25, 0.5, 0.75, 1$ (15 scenarios in total). All three candidate parameters may vary during matching. Widths are in thousands of 2000 dollars per person. Both signs requires candidates below $-\$0.01$ and above \$0.01. The full scenario parameters and results remain in the replication files; counts are not path probabilities.

Table B3: Reference and single-parameter reconstructed earnings NPVs

Policy	Reference	Permanent (K)	Temporary (b)	Temporary (d)
Florida grant	8.63	15.52	—	—
DC tuition aid	8.95	30.68	7.27	—
Adams	−1.07	−1.43	−0.37	−0.33
STBF	0.66	0.74	—	—
Wisconsin grant	0.59	1.05	0.14	—
Kalamazoo	1.10	1.55	—	—
Michigan CC	9.45	11.49	3.41	—
CUNY Pell	0.56	1.04	0.17	0.12
Ohio Pell	2.53	8.88	1.46	—
Tennessee HOPE	1.31	2.77	0.35	0.23
Tennessee Pell	−0.30	−0.62	−0.08	−0.05
Georgia HOPE	2.81	7.03	1.20	—
Social Security	2.78	2.61	—	—
Appropriations: spending	2.70	5.16	0.89	—
Early Head Start	n.a.	—	—	—
Texas CC	11.16	13.99	—	—
Adult Pell	2.47	6.47	—	—
Appropriations: tuition	0.03	0.05	9.55×10^{-3}	6.55×10^{-3}
HOPE credit	5.95	14.27	2.91	—
HOPE ind. single: lower	6.44	36.06	8.04	8.21
HOPE dep. single: lower	−4.14	−16.49	−3.44	−3.48
HOPE dep. single: upper	5.06	20.18	4.20	4.27
HOPE joint: lower	23.75	94.72	19.73	20.04
HOPE joint: upper	9.72	38.76	8.08	8.20
Tuition ded. single: lower	6.88	27.53	3.22	1.80
Tuition ded. joint: lower	−1.66	−6.62	−0.79	−0.44
HOPE / Lifetime Learning	−57.97	−77.41	−25.65	—
AOTC single: upper	−1.40	−3.42	−0.41	−0.25
AOTC single: lower	31.04	75.53	9.02	5.54
AOTC joint: upper	−11.54	−28.09	−3.36	−2.06
AOTC joint: lower	12.01	29.24	3.49	2.14
AOTC ind. single: lower	1.86	6.37	0.82	0.50
AOTC instrument	5.97	21.26	3.52	2.94
Tennessee Promise	6.46	n.c.	—	—
College placement	n.a.	n.c.	n.c.	—
Tuition ded. single: upper	0.00	0.00	0.00	0.00
Tuition ded. joint: upper	0.00	0.00	0.00	0.00

Notes: NPVs are divided by the same positive reference expenditure within each policy, after reconciling account units. Each column varies one parameter family, holding the other two fixed. — denotes no admissible Type 2 match; n.a., no reference projection or normalizer; n.c., no comparable monetary pair. Early Head Start and Tennessee Promise lack verified common monetary years. Figure B3 plots the comparable pairs; Table 5 reports parameter changes. Appendix B.5 describes joint inputs and exclusions. Dep./ind. denote dependent/independent filers; lower/upper denote income thresholds.

Table B4: Additional observation ages and MVPF ranges under a fixed grant account

Observed ages	MVPF range	Width reduction (%)	
		NPV	MVPF
A. Permanent effect (K)			
25	$[0.67, 44.09]$	—	—
24 and 25	$[0.96, 3.35]$	66.1	94.5
25 and 35	$[2.35, 3.01]$	92.3	98.5
B. Temporary effect (b)			
25	$[-0.18, 29.70]$	—	—
24 and 25	$[-0.01, 1.20]$	74.9	96.0
25 and 35	$[0.72, 1.05]$	94.2	98.9
C. Temporary effect (d)			
25	$[-0.73, 19.11]$	—	—
24 and 25	$[-0.70, 3.07]$	38.2	81.0
25 and 35	$[0.26, 0.63]$	92.7	98.1

Notes: The three midpoint scenarios use a \$10,000 grant and a 20 percent earnings tax rate, with no other non-earnings components. Dollars are valued at age 18 in 2000 prices. Reductions use unrounded endpoints relative to age 25 alone; all displayed candidates have positive net public cost. Across all 84 two-age designs within 18–25, comparison with each constituent age gives 49 finite-baseline reductions of 65.1–98.7 percent and 119 transitions from infinite to finite spans. All 84 two-age ranges still cross one. Figure B5 includes the full age grid and maximum spans. Ranges describe numerical candidates, not confidence intervals or certified bounds.

Table D1: Policy populations, education inputs, and reported timing

Policy	Geography / cohorts	Calibration support / mode	Path input	Reported windows	Type 3
Florida access grant	Florida 1982	18–35 state	Years of Schooling	3+: First-semester and years 1–4 credits	Fit (1)
DC tuition assistance	DC 1981	19–35 state	Years of Schooling	1: First college enrollment within 12 months of HS	–
Adams Scholarship	Massachusetts 1987	16–35 state	Years of Schooling	3+: Entry, persistence, and BA through year 6	Fit (1)
STBF scholarship	Nebraska 1993–1997	16–29 borrowed/state donors 1993–1994	T, C, A, B	3+: Semester enrollment; BA in years 4–6	–
Wisconsin Scholars Grant	Wisconsin 1990–1992	16–32 state	Years of Schooling	3+: Semester credits; cumulative years 2/3	Fit (1)
Kalamazoo Promise	Michigan 1988	16–34 state	Years of Schooling	3+: Credits in years 2–4; degrees in years 4/6	Fit (1)
Michigan CC tuition	Michigan 1992	16–30 state	Years of Schooling	2: Entry within a year; later credentials	Exact (49)
CUNY Pell grants	New York 1991	16–31 state	Years of Schooling	2: First-year and three-year credits	Exact (50)
Ohio Pell grants	Ohio 1981	19–35 state	Years of Schooling	1: Entry to next-year persistence	–
Tennessee HOPE	Tennessee 1989	16–33 state	Years of Schooling	1: First post-HS college enrollment	–
Tennessee Pell eligibility	Tennessee 1987–1992	16–35 state	Years of Schooling	1: Initial-fall college-sector choice	–
Georgia HOPE	Georgia 1975–1979	21–35 borrowed/state donors 1978–1979	Years of Schooling	Other: State-year freshman flows	–
Social Security student benefits	National 1963; shape 1985–1987	16–35 (donors); level at 23 source level donors 1985–1987	Entry	2: Attainment at ages 23 and 28	Exact (56)
College appropriations: spending	National 1983	17–35 national	Years of Schooling	Other: Institution enrollment by policy lag	Fit (1)
Early Head Start	National 1989–1999	16–33 borrowed/national donors 1989–1994	BA	1: Pooled adult attainment	–
Texas CC tuition	Texas 1987	16–35 state	Years of Schooling	3+: Annual enrollment at HS+0 through HS+6	Fit (1)
Adult Pell introduction	National 1947	53–65 borrowed donors 1978	Years of Schooling	1: Pooled current adult enrollment	–
College appropriations: tuition	National 1983	17–35 national	Years of Schooling	Other: Institution enrollment by policy lag	Fit (1)
HOPE tax credit	National 1981	19–35 national	Years of Schooling	Other: College-grade and pooled enrollment	Fit (1)
HOPE: independent single, lower kink	National 1987–1988	16–35 national	Years of Schooling	1: Contemporaneous credit-kink response	–
HOPE: dependent single, lower kink	National 1987–1988	16–35 national	Years of Schooling	1: Contemporaneous credit-kink response	–
HOPE: dependent single, upper kink	National 1987–1988	16–35 national	Years of Schooling	1: Contemporaneous credit-kink response	–
HOPE: joint, lower kink	National 1987–1988	16–35 national	Years of Schooling	1: Contemporaneous credit-kink response	–
HOPE: joint, upper kink	National 1987–1988	16–35 national	Years of Schooling	1: Contemporaneous credit-kink response	–
Tuition deduction: single, lower threshold	National 1988	16–34 national	Years of Schooling	3+: Years 0–4 from expected HS graduation	–

Table D1 (continued)

Policy	Geography / cohorts	Calibration support / mode	Path input	Reported windows	Type 3
Tuition deduction: joint, lower threshold	National 1988	16–34 national	Years of Schooling	3+: Years 0–4 from expected HS graduation	–
HOPE and Lifetime Learning credits	National 1980	20–35 national	Years of Schooling	1: Pooled current college enrollment	–
AOTC: dependent single, upper kink	National 1991–1992	16–31 national	Years of Schooling	1: Contemporaneous credit-kink response	–
AOTC: dependent single, lower kink	National 1991–1992	16–31 national	Years of Schooling	1: Contemporaneous credit-kink response	–
AOTC: joint, upper kink	National 1991–1992	16–31 national	Years of Schooling	1: Contemporaneous credit-kink response	–
AOTC: joint, lower kink	National 1991–1992	16–31 national	Years of Schooling	1: Contemporaneous credit-kink response	–
AOTC: independent single, lower kink	National 1991–1992	16–31 national	Years of Schooling	1: Contemporaneous credit-kink response	–
AOTC: simulated instrument	National 1987–1992	16–35 national	Years of Schooling	2: Attendance at ages 19 and 22	–
Tennessee Promise	Tennessee 1992–1998	18–30; source-level adjustment borrowed/state donors 1992–1994	Some college, AA, BA	3+: Six age windows, 18–19 to 23–24	Fit (1)
Algorithmic college placement	New York 1995	16–27 borrowed donors 1994	Years of Schooling	2: Initial term; pooled and oldest-cohort follow-up	–
Tuition deduction: single, upper threshold	National 1988	16–34 national	Years of Schooling	3+: Years 0–4 from expected HS graduation	–
Tuition deduction: joint, upper threshold	National 1988	16–34 national	Years of Schooling	3+: Years 0–4 from expected HS graduation	–
Cannabis legalization: education channel	RML 14 states 1988–1999	16–34 borrowed/state donors 1988–1994	Years of Schooling	3+: BA completion within 4, 5, 6 years	Fit (1)

Notes: Windows count reported ages or follow-up periods, not coefficients or independent observations; Other denotes institutional or college-stage clocks. Geography and cohorts identify the external calibration. Support gives the observed calibration ages across component cohorts, not a complete age range for every cohort. National denotes a direct national fit; state denotes shrinkage toward its national shape; borrowed denotes a donor shape with a level calibrated for the target cohort; source level denotes a published control mean combined with a donor shape. Mixed modes apply to different component cohorts. Shape donors match the listed cohorts unless separately stated. Tennessee uses HS graduates and source control levels. These are calibration rules, not additional causal observations. Type 3 preserves Type 2 and adds comparable source timing under the mappings in the case records. Fit denotes a minimum-distance selection; Exact denotes an exact-match bank. Parentheses count saved paths, not all possible solutions. A dash means no additional source restriction is admitted; splitting an assumed duration does not create Type 3. Cannabis values education alone and has no Type 1. STBF jointly uses a time-enrolled index and three credentials; its T, C, A, and B inputs are defined in its case record. Full operator weights and source-control qualifications remain in the replication registry.

Table D2: Type 3 minimum-distance fits and point-input diagnostics

Panel A: Fit summaries

Policy	Targets	RMSE	Max. error	NPV	MVPF
Florida grant	5	0.02	0.02	26.46	18.09
Adams Scholarship	4	0.01	0.01	-2.74	0.63
Wisconsin grant	9	0.02	0.03	8.09	1.96
Kalamazoo Promise	3	0.01	0.02	57.29	3.77
College spending	5	9.72×10^{-4}	1.75×10^{-3}	7.57	16.52
Texas tuition	8	3.88×10^{-3}	0.01	13.12	$+\infty$
Tuition caps	5	1.90×10^{-3}	2.47×10^{-3}	-4.60	-1.97
HOPE tax credit	3	5.35×10^{-4}	8.36×10^{-4}	16.19	$+\infty$
Tennessee Promise	18	0.35	0.67	113.94	$+\infty$
Cannabis: education	3	0.01	0.01	-3.64	-55.85

Panel B: Schooling-equivalent inputs (Years of Schooling)

Policy	Source input	Target	Source SE	Fitted	Residual	Exact
Florida grant	Year 1 cumulative	4.58×10^{-2}	–	5.53×10^{-2}	$+9.56\times10^{-3}$	–
	Year 2 cumulative	1.11×10^{-1}	–	9.38×10^{-2}	-1.73×10^{-2}	–
	Year 3 cumulative	1.60×10^{-1}	–	1.37×10^{-1}	-2.37×10^{-2}	–
	Year 4 cumulative	1.82×10^{-1}	–	1.82×10^{-1}	0.00	Yes
	First semester	1.48×10^{-2}	–	3.84×10^{-2}	$+2.36\times10^{-2}$	–
Adams Scholarship	Immediate entry (age 18)	2.60×10^{-2}	–	1.18×10^{-2}	-1.42×10^{-2}	–
	Entry within two years (age 20)	7.00×10^{-3}	–	1.06×10^{-2}	$+3.58\times10^{-3}$	–
	Four-year completion (age 22)	-1.10×10^{-2}	–	-3.65×10^{-3}	$+7.35\times10^{-3}$	–
	Six-year completion (age 24)	-1.80×10^{-2}	–	-1.80×10^{-2}	0.00	Yes
Wisconsin grant	Cohort 1 cumulative	3.75×10^{-2}	–	6.41×10^{-2}	$+2.66\times10^{-2}$	–
	Cohort 1: year 1 flow	2.08×10^{-2}	–	1.80×10^{-2}	-2.83×10^{-3}	–
	Cohort 1: year 2 flow	8.33×10^{-3}	–	2.23×10^{-2}	$+1.39\times10^{-2}$	–
	Cohort 1: year 3 flow	4.17×10^{-3}	–	2.39×10^{-2}	$+1.97\times10^{-2}$	–
	Cohorts 2–3: year 1 flow	1.67×10^{-2}	–	4.19×10^{-2}	$+2.53\times10^{-2}$	–
	Cohorts 2–3: year 2 flow	4.17×10^{-2}	–	2.37×10^{-2}	-1.80×10^{-2}	–
	Cohort 2: year 3 flow	2.08×10^{-2}	–	1.93×10^{-2}	-1.53×10^{-3}	–
	Cohorts 2–3 cumulative	8.75×10^{-2}	–	7.53×10^{-2}	-1.22×10^{-2}	–
	Retained weighted aggregate	7.18×10^{-2}	–	7.18×10^{-2}	0.00	Yes
Kalamazoo Promise	Year 4 cumulative	2.73×10^{-1}	–	2.73×10^{-1}	0.00	Yes
	Year 2 cumulative	1.35×10^{-1}	–	1.24×10^{-1}	-1.03×10^{-2}	–
	Year 3 cumulative	1.80×10^{-1}	–	1.97×10^{-1}	$+1.73\times10^{-2}$	–
College spending	Retained four-year sum	4.60×10^{-2}	–	4.60×10^{-2}	0.00	Yes
	Lag 0 contribution	5.04×10^{-3}	–	6.12×10^{-3}	$+1.07\times10^{-3}$	–
	Lag 1 contribution	1.32×10^{-2}	–	1.14×10^{-2}	-1.75×10^{-3}	–
	Lag 2 contribution	1.40×10^{-2}	–	1.40×10^{-2}	-3.87×10^{-5}	–
	Lag 3 contribution	1.38×10^{-2}	–	1.45×10^{-2}	$+7.14\times10^{-4}$	–
Texas tuition	Model age 18 cumulative	3.16×10^{-2}	–	3.88×10^{-2}	$+7.29\times10^{-3}$	–
	Model age 19 cumulative	7.76×10^{-2}	–	7.83×10^{-2}	$+7.14\times10^{-4}$	–
	Model age 20 cumulative	1.13×10^{-1}	–	1.14×10^{-1}	$+1.54\times10^{-3}$	–
	Model age 21 cumulative	1.46×10^{-1}	–	1.45×10^{-1}	-7.79×10^{-4}	–
	Model age 22 cumulative	1.72×10^{-1}	–	1.69×10^{-1}	-2.75×10^{-3}	–
	Model age 23 cumulative	1.87×10^{-1}	–	1.88×10^{-1}	$+9.38\times10^{-4}$	–
	Model age 24 cumulative	2.09×10^{-1}	–	2.02×10^{-1}	-7.44×10^{-3}	–
	Model age 26 cumulative	2.18×10^{-1}	–	2.18×10^{-1}	0.00	Yes
Tuition caps	Retained four-year sum	4.61×10^{-4}	–	4.61×10^{-4}	0.00	Yes
	Lag 0 contribution	-9.21×10^{-4}	–	1.00×10^{-3}	$+1.92\times10^{-3}$	–
	Lag 1 contribution	3.80×10^{-3}	–	1.33×10^{-3}	-2.47×10^{-3}	–
	Lag 2 contribution	1.78×10^{-3}	–	4.44×10^{-5}	-1.74×10^{-3}	–
	Lag 3 contribution	-4.20×10^{-3}	–	-1.92×10^{-3}	$+2.28\times10^{-3}$	–
HOPE tax credit	Retained duration total (age 20.5)	6.07×10^{-2}	–	6.07×10^{-2}	0.00	Yes

Table D2 (continued): point-input diagnostics

Policy / outcome	Input / age window	Target	Source SE	Fitted	Residual	Exact
	First-year entry (age 19)	2.59×10^{-2}	–	2.63×10^{-2}	$+3.98\times10^{-4}$	–
	Continuation (age 20)	4.91×10^{-2}	–	4.83×10^{-2}	-8.36×10^{-4}	–
Cannabis: education	Four-year completion bridge	-3.20×10^{-2}	–	-2.72×10^{-2}	$+4.77\times10^{-3}$	–
	Five-year completion bridge	-2.20×10^{-2}	–	-3.11×10^{-2}	-9.14×10^{-3}	–
	Six-year completion bridge	-3.60×10^{-2}	–	-3.60×10^{-2}	0.00	Yes

Panel C: Tennessee Promise (percentage points)

Outcome	Age window	Target	Source SE	Fitted	Residual	Exact
SC+	18–19	3.50	0.40	3.68	+0.18	–
SC+	19–20	3.20	0.40	2.68	−0.52	–
SC+	20–21	2.50	0.40	2.95	+0.45	–
SC+	21–22	2.70	0.30	3.34	+0.64	–
SC+	22–23	2.80	0.40	3.47	+0.67	–
SC+	23–24	3.50	0.40	3.50	0.00	Yes
AA	18–19	0.20	0.10	0.76	+0.56	–
AA	19–20	1.40	0.20	1.63	+0.23	–
AA	20–21	3.10	0.30	2.52	−0.58	–
AA	21–22	2.60	0.30	2.84	+0.24	–
AA	22–23	2.90	0.40	2.93	+0.03	–
AA	23–24	4.00	0.40	4.00	0.00	Yes
BA	18–19	−0.10	0.05	−0.03	+0.07	–
BA	19–20	0.10	0.05	−0.08	−0.18	–
BA	20–21	1.00×10^{-3}	0.10	−0.14	−0.14	–
BA	21–22	−0.20	0.20	0.03	+0.23	–
BA	22–23	1.20	0.30	1.06	−0.14	–
BA	23–24	2.70	0.40	2.70	0.00	Yes

Notes: Each of the ten searches retains one best-fitting path with nonzero added-input discrepancies and exact Type 2 constraints (Exact = Yes). Residual equals fitted minus target. RMSE includes all encoded targets, including exact inputs. Errors are Years of Schooling in Panel B and percentage points for Tennessee; NPV is thousands of case-specific dollars. The fixed identity weight matrix uses no SE bands or precision weights. Dashes in Panel B do not assert that source SEs are unavailable; they are not repeated here. Wisconsin retains the source semester/cumulative estimates and cohort windows without imposing an unsupported adding-up identity. Tennessee preserves the three latest-window effects and fits fifteen earlier targets from Schoonover et al. (2026), Table 2; its source SEs enter only the separate input repetitions. Two SEs printed as zero in the source receive a 5.00×10^{-4} probability-scale floor in that exercise. The fitted values are conditional selections, not certified unique or global solutions. Source-age and conversion assumptions remain in the case records; calculations use unrounded inputs.

Table D3: Main-example source inputs and population calibration

Panel A: Adams Scholarship

Education input	Effect (pp)	SE (pp)	Source clock
Four-year enrollment, immediate	0.90	0.80	HS+0
Two-year enrollment, immediate	0.80	0.50	HS+0
Four-year enrollment, within two years	0.00	0.70	HS+2
Two-year enrollment, within two years	0.70	0.50	HS+2
BA completion, within four years	−1.70	1.00	HS+4
Two-year completion, within four years	−0.20	0.30	HS+4
BA completion, within six years	−2.50	0.90	HS+6
Two-year completion, within six years	0.00	0.40	HS+6

Panel B: Social Security student benefits

Education input	Effect (pp)	SE (pp)	Reconstruction role
Ever attended by age 28	22.40	10.60	Exact Type 2; attrition lower bound
Ever attended by age 23	21.90	10.20	Exact additional Type 3

Panel C: Tennessee Promise calibration and continuation limits

Quantity	Some college	AA	BA
Source control mean, ages 23–24	68.50	38.00	29.20
Matched treated mean, ages 23–24	72.00	42.00	31.90
Remaining-completion fraction cap	14.02	26.35	19.78
Implied treated plateau ceiling	75.93	57.28	45.37
Largest saved treated plateau	75.27	57.28	45.37

Notes: Effects and SEs in Panels A and B are percentage points (pp). Probability levels and remaining-completion fractions in Panel C are percentages. Calculations use unrounded inputs. Adams uses Cohodes and Goodman (2014), Table 4, columns (3) and (5); HS+t denotes years after high-school graduation. Its converted schooling targets and fitted residuals appear in Table D2. Assigned model ages are separate assumptions, and converted-input uncertainty follows the grouped source draws. Social Security uses Dynarski (2003), Table 3, full-time attendance; its age-28 lower bound concerns attrition imputation. Tennessee conditions on HS graduation and prior-year state residence. Source-matched levels and donor-based continuation limits are maintained calibration restrictions; plateau ceilings are not estimated causal bounds. Its eighteen education targets, source SEs, and residuals appear in Table D2. STBF retains the joint inputs in Appendix D.2.

Table D4: Conditional welfare ranks

Policy	Type 1	Own Type 2	Selected	Type
HOPE: joint, lower kink [f]†	1	23	23	Type 2
Tuition deduction: single, lower threshold†	1	23	23	Type 2
AOTC: dependent single, lower kink [f]†	1	23	23	Type 2
AOTC: joint, lower kink [f]	1	23	23	Type 2
Tennessee Promise [r]	1	14	1	Type 3
Algorithmic college placement [x]	1	13	13	Type 2
Tuition deduction: single, upper threshold†	1	23	23	Type 2
Texas CC tuition	8	5	1	Type 3
Michigan CC tuition	9	20	6	Type 3
DC tuition assistance	10	15	15	Type 2
HOPE: independent single, lower kink [f]†	11	23	23	Type 2
HOPE tax credit	12	20	1	Type 3
HOPE: dependent single, upper kink [f]†	13	23	23	Type 2
AOTC: simulated instrument [f]†	14	23	23	Type 2
Florida access grant	15	12	4	Type 3
AOTC: independent single, lower kink [f]	16	23	23	Type 2
Social Security student benefits	17	7	7	Type 3
College appropriations: spending	18	20	5	Type 3
Georgia HOPE	19	18	18	Type 2
Ohio Pell grants	20	21	21	Type 2
HOPE: joint, upper kink [f]†	21	23	23	Type 2
Adult Pell introduction	22	10	10	Type 2
Kalamazoo Promise	23	11	8	Type 3
Tennessee HOPE	24	22	22	Type 2
STBF scholarship	25	11	11	Type 2
Wisconsin Scholars Grant	26	15	9	Type 3
CUNY Pell grants	27	36	17	Type 3
College appropriations: tuition	28	20	16	Type 3
Tennessee Pell eligibility†	29	37	37	Type 2
Tuition deduction: joint, upper threshold	30	19	19	Type 2
Adams Scholarship	31	12	12	Type 3
Early Head Start [r]	32	14	14	Type 2
AOTC: dependent single, upper kink [f]†	33	23	23	Type 2
Tuition deduction: joint, lower threshold	34	20	20	Type 2
HOPE: dependent single, lower kink [f]†	35	23	23	Type 2
AOTC: joint, upper kink [f]†	36	23	23	Type 2
HOPE and Lifetime Learning credits†	37	23	23	Type 2

Notes: All 37 calculations with a reference MVPF (Type 1) are retained. Own Type 2 replaces only that calculation's selected set by Type 2. Selected uses Type 3 where admitted and Type 2 otherwise. [f] marks a finite tax experiment with a local-linear input conversion, [r] a public-data reconstructed account, and [x] a fiscal-only earnings extension. † marks at least one unassigned-ratio path; those paths, not the calculations, are omitted from ranking. Finite endpoints are rounded to two decimals before lower-endpoint-first ordering, and tied endpoint pairs share a competition rank. Table 3 reports the endpoints; Table D1 and the case profiles document source observation windows. Cannabis has no whole-policy Type 1. The restricted 24-case ordering and rounding sensitivities remain in the replication materials. These are conditional display ranks, not social-welfare priorities or certified global bounds.

Table D5: Reference welfare accounts and selected complete paths

Panel A: Reference accounts (case-specific dollars)

Component	Adams	Social Security	Tennessee Promise
Net earnings NPV	-1,204.04	21,163.75	95,557.69
Reference private nonincome component	2,329.97	6,320.04	2,124.31
Reference pretax public component	1,658.06	9,150.29	14,792.00
Type 1 WTP	1,364.23	23,072.74	72,837.00
Type 1 net cost	1,896.35	4,739.23	-10,053.00
Type 1 MVPF	0.72	4.87	$+\infty$
Dollar year	2004	2000	2021 convention

Panel B: Selected paths (WTP and net cost in thousands of case-specific dollars)

Path	Outcome	K	b	d	$\Delta H(65)$	WTP	Net cost	MVPF
A1	Years of Schooling	14.33	0.35	14.11	-0.51	7.27	15.10	0.48
A2	Years of Schooling	15.70	0.18	12.78	0.87	-11.50	-20.90	–
A3	Years of Schooling	14.79	0.25	12.70	-0.04	1.66	2.64	0.63
S1	Entry	0.58	0.58	19.95	0.23	19.34	5.10	3.79
S2	Entry	0.58	2.74	18.20	0.22	23.46	4.70	4.99
S3	Entry	0.58	1.28	19.62	0.22	20.81	4.95	4.21
S4	Entry	0.58	1.02	18.61	0.22	21.89	4.85	4.52
T1	SC+	0.73	1.06	19.08	0.04	-48.72	32.66	-1.49
	AA	0.42	2.26	21.22	-0.05			
	BA	0.32	3.00	21.62	-0.04			
T2	SC+	0.72	3.00	17.59	0.03	152.50	-38.04	$+\infty$
	AA	0.57	0.77	22.11	0.10			
	BA	0.45	0.90	22.48	0.09			
T3	SC+	0.72	1.44	17.89	0.04	86.44	-14.83	$+\infty$
	AA	0.57	0.88	22.26	0.10			
	BA	0.41	1.36	22.45	0.04			

Notes: Panel A preserves the source accounts. Non-income components are reference-batch values; Adams and Social Security transport education payments with the paths. Social Security uses the enrollment alternative of Hendren and Sprung-Keyser (2020), with two schooling years per entrant and phase-specific tax rates of 3.00×10^{-3} and 0.20. Tennessee retains its source-implied earnings NPV separately from the public ACS reconstruction and participant-unit bridge. Panel B pairs components from each complete candidate; rows sharing a label form one joint path. A1–A2, S1–S2, and T1–T2 minimize/maximize Type 2 earnings NPV; S3–S4 do so for Social Security Type 3. A3 and T3 are the respective Type 3 minimum-distance fits, each shown once. All candidates pass the registered restrictions. Age 65 is extrapolated; Tennessee values earnings only through 64. Parameters are rounded for display, income includes early losses once, and monetary vintages differ. STBF remains documented in its case record.

Table D6: Point-input sensitivity: numerical comparisons

Policy	Set	Main	Narrow	Wide	National	Earnings 75%	Earnings 125%	Denser
Direct student financial aid								
Florida access grant	Type 3	18.09	18.25	17.88	17.64	5.99	$_^{S}$	18.09
DC tuition assistance	Type 2	$[-1.60, > 10^6]^{S}$	$[6.67, > 10^6]^{S}$	$[-5.01, > 10^6]^{S}$	$[1.04, > 10^6]^{S}$	$[-1.30, > 10^6]^{S}$	$[-1.77, > 10^6]^{S}$	$[-1.60, > 10^6]^{S}$
Adams Scholarship	Type 3	0.63	0.62	0.63	0.63	0.88	0.40	0.63
STBF scholarship	Type 2	[1.25, 4.01]	[1.27, 2.78]	[1.24, 6.55]	[1.28, 5.69]	[1.18, 2.98]	[1.31, 5.43]	[1.24, 4.01]
Wisconsin Scholars Grant	Type 3	1.96	2.09	1.77	1.95	1.60	2.38	1.96
Kalamazoo Promise	Type 3	3.77	3.62	3.84	4.29	2.60	5.43	3.77
CUNY Pell grants	Type 3	[−3.17, 2.34]	[−2.01, 1.79]	[−5.19, 3.02]	[−2.96, 2.60]	[−3.01, 1.78]	[−3.26, 3.03]	[−3.17, 2.34]
Ohio Pell grants	Type 2	[−11.75, 7.53]	[−1.73, 7.08]	$[< -10^6, 7.83]$	[−8.76, 7.49]	$[< -10^6, 3.73]$	[−7.52, 19.58]	[−11.75, 7.53]
Tennessee HOPE	Type 2	[−75.64, 21.66]	[−34.06, 18.34]	[−371.25, 24.53]	[−77.66, 23.24]	$[< -10^6, 6.48]$	$[-15.01, > 10^6]^{S}$	[−75.64, 21.66]
Tennessee Pell eligibility	Type 2	$[< -10^6, 2.70]$	$[< -10^6, 2.66]$	$[< -10^6, 2.72]$	$[< -10^6, 2.70]$	$[< -10^6, 1.66]$	$[< -10^6, 4.25]$	$[< -10^6, 2.70]$
Georgia HOPE	Type 2	$[-3.19, > 10^6]^{S}$	$[-1.36, > 10^6]^{S}$	$[-4.78, > 10^6]^{S}$	$[-3.50, > 10^6]^{S}$	[−3.05, 56.42]	$[-3.28, > 10^6]^{S}$	$[-3.19, > 10^6]^{S}$
Social Security benefits	Type 3	[4.21, 4.52]	[4.27, 4.48]	[4.11, 4.56]	[4.21, 4.52]	[2.88, 3.04]	[6.23, 6.81]	[4.21, 4.52]
Adult Pell introduction	Type 2	[1.95, 4.63]	[2.87, 4.59]	[−0.50, 4.64]	[1.95, 4.63]	[1.64, 2.73]	[2.28, 8.20]	[1.95, 4.63]
Tennessee Promise	Type 3	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$
Education tax benefits								
HOPE tax credit	Type 3	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$
HOPE: indep. single, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-1000.00, > 10^6]^{S}$	$[-52.08, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
HOPE: dep. single, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-2{,}697, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
HOPE: dep. single, end	Type 2	$[< -10^6, > 10^6]^{S}$	$[-443.66, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-220.34, > 10^6]^{S}$	$[-161.93, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
HOPE: joint, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-15.87, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
HOPE: joint, end	Type 2	$[< -10^6, > 10^6]^{S}$	$[-21.92, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, 1{,}000]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
Tuition ded.: single, start	Type 2	$[< -10^6, -7.40]^{S}$	$[< -10^6, -7.65]^{S}$	$[< -10^6, -7.26]^{S}$	$[< -10^6, -7.40]^{S}$	$[< -10^6, -13.36]^{S}$	$[< -10^6, -5.99]^{S}$	$[< -10^6, -7.40]^{S}$
Tuition ded.: joint, start	Type 2	$[-4.35, > 10^6]^{S}$	$[-4.28, > 10^6]^{S}$	$[-4.41, > 10^6]^{S}$	$[-4.35, > 10^6]^{S}$	$[-4.38, > 10^6]^{S}$	$[-4.34, > 10^6]^{S}$	$[-4.35, > 10^6]^{S}$
HOPE and Lifetime Learning	Type 2	$[< -10^6, -9.59]^{S}$	$[< -10^6, -9.79]^{S}$	$[< -10^6, -9.33]^{S}$	$[< -10^6, -9.59]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, -7.02]^{S}$	$[< -10^6, -9.59]^{S}$
AOTC: single, end	Type 2	$[< -10^6, > 10^6]^{S}$	$[-116.46, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-188.71, > 10^6]^{S}$	$[-95.11, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
AOTC: single, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-1{,}240, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
AOTC: joint, end	Type 2	$[< -10^6, > 10^6]^{S}$	$[-8.79, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-12.62, > 10^6]^{S}$	$[-300.32, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
AOTC: joint, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, -13.39]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
AOTC: indep. single, start	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$
AOTC: simulated instrument	Type 2	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[< -10^6, > 10^6]^{S}$	$[-372.79, > 10^6]^{S}$	$[< -10^6, 36.36]^{S}$	$[< -10^6, > 10^6]^{S}$
Tuition ded.: single, end	Type 2	$[< -10^6, -6.33]^{S}$	$[< -10^6, -6.41]^{S}$	$[< -10^6, -6.27]^{S}$	$[< -10^6, -6.33]^{S}$	$[< -10^6, -8.17]^{S}$	$[< -10^6, -5.68]^{S}$	$[< -10^6, -6.33]^{S}$
Tuition ded.: joint, end	Type 2	$[-3.92, > 10^6]^{S}$	$[-3.84, > 10^6]^{S}$	$[-3.98, > 10^6]^{S}$	$[-3.92, > 10^6]^{S}$	$[-3.72, > 10^6]^{S}$	$[-4.03, > 10^6]^{S}$	$[-3.92, > 10^6]^{S}$
Tuition and public-college funding								
Michigan CC tuition	Type 3	$[9.03, > 10^6]^{S}$	$[16.60, > 10^6]^{S}$	$[3.81, > 10^6]^{S}$	$[3.89, > 10^6]^{S}$	[4.70, 26.12]	$[21.35, > 10^6]^{S}$	$[9.03, > 10^6]^{S}$
College spending	Type 3	16.52	25.71	9.27	16.52	6.78	189.51	16.52
Texas CC tuition	Type 3	$_^{S}$	$_^{S}$	$_^{S}$	$_^{S}$	14.33	$_^{S}$	$_^{S}$
College tuition caps	Type 3	−1.97	−1.36	−2.82	−1.97	−1.53	−2.27	−1.97
College assessment and course placement								
Algorithmic placement	Type 2	$[0.08, > 10^6]^{S}$	$[0.12, > 10^6]^{S}$	$[0.06, > 10^6]^{S}$	$[0.08, > 10^6]^{S}$	$[0.10, > 10^6]^{S}$	$[0.06, > 10^6]^{S}$	$[0.08, > 10^6]^{S}$
Early childhood services								
Early Head Start	Type 2	$[-1.37, > 10^6]^{S}$	[−0.48, 13.31]	$[-1.60, > 10^6]^{S}$	$[-1.37, > 10^6]^{S}$	[−1.19, 17.32]	$[-1.54, > 10^6]^{S}$	$[-1.37, > 10^6]^{S}$
Education channel of other policies								
Cannabis: education channel	Type 3	−55.85	−53.65	−58.94	−47.77	–	−14.94	−55.85

Notes: Brackets report positive-cost MVPF spans among verified candidates; equal endpoints appear once. Superscript S marks a separate path under which the policy pays for itself. $> 10^6$ and $< -10^6$ are attained finite values, not certified infinite bounds. Narrow/Wide multiply support covariance by 0.5/2 and timing deviations by the square roots; National changes the external population. Earnings columns scale net income and associated taxes, holding transfers and education resources fixed. Each column retains the same point inputs and refits its information rule.

Table D6: Point-input sensitivity: numerical comparisons (continued)

Policy	Specific change	Set	WTP (1,000)	Net cost (1,000)	Positive-cost MVPF
Adams Scholarship	Zero negative income-response template	Type 3	1.48	2.64	0.56
STBF scholarship	No author-model experience delay	Type 2	$[43.52, 85.86]$	$[18.85, 29.43]$	$[1.48, 4.56]$
STBF scholarship	Printed time-enrolled input (0.36)	Type 2	$[38.39, 81.03]$	$[20.05, 30.71]$	$[1.25, 4.04]$
Tennessee Promise	Pooled late-completion allowance	Type 3	86.44	-14.83	$_^{S}$
Tennessee Promise	No source-control level scaling	Type 3	86.49	-14.85	$_^{S}$
HOPE: indep. single, start	Credit dose: 0.50 dollars	Type 2	$[-53.78, 101.42]$	$[-17.37, 9.43]$	$[< -10^6, > 10^6]^{S}$
HOPE: indep. single, start	Credit dose: 2 dollars	Type 2	$[-53.74, 101.47]$	$[-17.38, 9.42]$	$[< -10^6, > 10^6]^{S}$
HOPE: dep. single, start	Credit dose: 0.50 dollars	Type 2	$[-16.27, 30.50]$	$[-2.77, 1.55]$	$[< -10^6, > 10^6]^{S}$
HOPE: dep. single, start	Credit dose: 2 dollars	Type 2	$[-16.28, 30.48]$	$[-2.76, 1.55]$	$[< -10^6, > 10^6]^{S}$
HOPE: dep. single, end	Credit dose: 0.50 dollars	Type 2	$[-17.36, 32.61]$	$[-3.37, 1.86]$	$[-83.48, > 10^6]^{S}$
HOPE: dep. single, end	Credit dose: 2 dollars	Type 2	$[-17.34, 32.63]$	$[-3.37, 1.86]$	$[-124.32, > 10^6]^{S}$
HOPE: joint, start	Credit dose: 0.50 dollars	Type 2	$[-21.40, 40.34]$	$[-4.69, 2.53]$	$[-38.45, > 10^6]^{S}$
HOPE: joint, start	Credit dose: 2 dollars	Type 2	$[-21.31, 40.46]$	$[-4.70, 2.52]$	$[-22.42, > 10^6]^{S}$
HOPE: joint, end	Credit dose: 0.50 dollars	Type 2	$[-23.38, 44.15]$	$[-4.88, 2.70]$	$[< -10^6, > 10^6]^{S}$
HOPE: joint, end	Credit dose: 2 dollars	Type 2	$[-23.34, 44.20]$	$[-4.87, 2.70]$	$[-42.16, > 10^6]^{S}$
AOTC: single, end	Credit dose: 0.50 dollars	Type 2	$[-42.93, 67.92]$	$[-7.23, 4.41]$	$[-387.50, > 10^6]^{S}$
AOTC: single, end	Credit dose: 2 dollars	Type 2	$[-42.93, 67.91]$	$[-7.23, 4.42]$	$[-112.99, > 10^6]^{S}$
AOTC: single, start	Credit dose: 0.50 dollars	Type 2	$[-40.96, 64.86]$	$[-6.44, 3.93]$	$[< -10^6, 64.57]^{S}$
AOTC: single, start	Credit dose: 2 dollars	Type 2	$[-40.89, 64.97]$	$[-6.45, 3.92]$	$[< -10^6, > 10^6]^{S}$
AOTC: joint, end	Credit dose: 0.50 dollars	Type 2	$[-55.16, 87.67]$	$[-10.69, 6.71]$	$[-20.42, > 10^6]^{S}$
AOTC: joint, end	Credit dose: 2 dollars	Type 2	$[-55.18, 87.63]$	$[-10.68, 6.72]$	$[-8.22, > 10^6]^{S}$
AOTC: joint, start	Credit dose: 0.50 dollars	Type 2	$[-53.64, 85.30]$	$[-10.11, 6.35]$	$[< -10^6, > 10^6]^{S}$
AOTC: joint, start	Credit dose: 2 dollars	Type 2	$[-53.61, 85.34]$	$[-10.12, 6.35]$	$[< -10^6, > 10^6]^{S}$
AOTC: indep. single, start	Credit dose: 0.50 dollars	Type 2	$[-179.55, 284.89]$	$[-56.78, 36.09]$	$[< -10^6, > 10^6]^{S}$
AOTC: indep. single, start	Credit dose: 2 dollars	Type 2	$[-179.55, 284.90]$	$[-56.78, 36.09]$	$[< -10^6, > 10^6]^{S}$
AOTC: simulated instrument	Credit dose: 0.50 dollars	Type 2	$[-31.48, 55.32]$	$[-6.14, 3.47]$	$[< -10^6, > 10^6]^{S}$
AOTC: simulated instrument	Credit dose: 2 dollars	Type 2	$[-31.47, 55.33]$	$[-6.14, 3.46]$	$[-19.32, > 10^6]^{S}$
Early Head Start	FY2001 funded-slot cost proxy	Type 2	$[-37.51, 114.60]$	$[-0.52, 27.93]$	$[-1.34, > 10^6]^{S}$
Cannabis: education channel	All-public-college costs	Type 3	-2.68	0.09	-30.29
Cannabis: education channel	CC costs	Type 3	-2.91	0.13	-22.12
Cannabis: education channel	Equal state weights	Type 3	-2.94	0.13	-21.99
Cannabis: education channel	Zero negative income-response template	Type 3	-2.71	0.09	-28.68

Notes: Monetary components use registered account units; Early Head Start and Tennessee are public-data reconstructions. Component extrema can come from different paths and must not be combined into a ratio. Tax-dose checks change education inputs and fixed account components together. Full-precision NPV, components, and Type 2/Type 3 records remain in the data.

Table E1: WTP and net fiscal cost: point paths and input repetitions

Policy	Type	WTP: point	WTP: repetitions	Net cost: point	Net cost: repetitions
Florida grant	Type 2	[0.60, 27.81]	[−9.49, 51.26]	[1.04, 1.25]	[−4.05, 7.08]
Florida grant	Type 3	[20.50, 20.50]	[2.55, 38.86]	[1.13, 1.13]	[−2.83, 5.43]
DC tuition aid	Type 2	[−3.08, 48.18]	[−10.07, 86.68]	[−3.40, 1.93]	[−12.14, 6.54]
Adams	Type 2	[−11.50, 7.40]	[−787.98, 730.47]	[−20.90, 15.11]	[−228.12, 219.90]
Adams	Type 3	[1.66, 1.66]	[−16.18, 21.72]	[2.64, 2.64]	[−2.92, 7.92]
STBF	Type 2	[38.33, 80.73]	[30.65, 82.05]	[20.13, 30.73]	[19.80, 32.61]
Wisconsin grant	Type 2	[−13.53, 42.02]	[−30.51, 70.02]	[5.54, 6.98]	[−1.28, 13.18]
Wisconsin grant	Type 3	[12.37, 12.37]	[−0.27, 31.07]	[6.29, 6.29]	[4.37, 8.35]
Kalamazoo	Type 2	[11.13, 72.60]	[−13.20, 118.03]	[16.04, 17.15]	[5.87, 26.88]
Kalamazoo	Type 3	[61.25, 61.25]	[17.13, 112.08]	[16.26, 16.26]	[6.92, 25.29]
Michigan CC tuition	Type 2	[−7.88, 21.76]	[−18.38, 42.64]	[−3.05, 1.78]	[−8.15, 4.05]
Michigan CC tuition	Type 3	[1.35, 3.24]	[−1.48, 7.72]	[−0.13, 0.15]	[−0.99, 0.60]
CUNY Pell	Type 2	[−27.52, 39.88]	[−52.36, 70.73]	[−0.14, 2.79]	[−7.07, 11.30]
CUNY Pell	Type 3	[−2.69, 2.54]	[−24.21, 25.42]	[0.85, 1.09]	[−1.79, 3.38]
Ohio Pell	Type 2	[−5.34, 27.29]	[−10.63, 48.18]	[0.45, 3.62]	[−1.50, 9.35]
Tennessee HOPE	Type 2	[−28.18, 43.61]	[−49.85, 76.69]	[0.37, 2.01]	[−6.75, 11.07]
Tennessee Pell	Type 2	[−25.65, 38.45]	[−51.60, 70.85]	[−7.56, 14.24]	[−56.66, 50.68]
Georgia HOPE	Type 2	[−10.03, 42.88]	[−22.65, 84.37]	[−2.13, 3.14]	[−11.96, 8.22]
Social Security aid	Type 2	[19.34, 23.46]	[4.66, 45.25]	[4.70, 5.10]	[1.48, 8.64]
Social Security aid	Type 3	[20.81, 21.89]	[6.77, 42.48]*	[4.85, 4.95]	[1.82, 8.63]*
Appropriations: spending	Type 2	[−19.38, 44.67]	[−36.11, 79.55]	[−4.21, 3.57]	[−12.35, 7.37]
Appropriations: spending	Type 3	[6.83, 6.83]	[1.68, 13.68]	[0.41, 0.41]	[−0.84, 1.49]
Early Head Start	Type 2	[−37.51, 114.60]	[−76.09, 114.87]	[−1.10, 27.35]	[−1.15, 34.50]
Texas CC tuition	Type 2	[5.88, 24.15]	[−2.81, 46.37]	[−1.63, 0.61]	[−6.55, 3.89]
Texas CC tuition	Type 3	[10.88, 10.88]	[1.99, 22.77]	[−0.08, −0.08]	[−2.37, 2.04]
Adult Pell	Type 2	[0.14, 1.18]	[−0.23, 2.31]	[0.07, 0.26]	[0.01, 0.51]
Appropriations: tuition	Type 2	[−22.46, 38.91]	[−42.92, 72.50]	[−1.42, 2.54]	[−11.24, 8.42]
Appropriations: tuition	Type 3	[−2.41, −2.41]	[−12.43, 5.32]	[1.22, 1.22]	[0.38, 2.99]
HOPE credit	Type 2	[−8.27, 25.82]	[−17.03, 48.24]	[−2.86, 1.62]	[−8.35, 3.73]
HOPE credit	Type 3	[12.73, 12.73]	[−0.08, 32.72]	[−1.13, −1.13]	[−5.17, 1.55]
HOPE: ind. single, lower	Type 2	[−53.77, 101.43]	[−97.27, 184.42]	[−17.37, 9.43]	[−38.00, 19.98]
HOPE: dep. single, lower	Type 2	[−16.27, 30.50]	[−30.46, 57.47]	[−2.76, 1.55]	[−9.38, 5.44]
HOPE: dep. single, upper	Type 2	[−17.36, 32.62]	[−32.35, 61.24]	[−3.37, 1.86]	[−10.39, 5.82]
HOPE: joint, lower	Type 2	[−21.37, 40.38]	[−39.52, 74.85]	[−4.69, 2.53]	[−12.63, 6.67]
HOPE: joint, upper	Type 2	[−23.37, 44.17]	[−42.83, 81.41]	[−4.88, 2.70]	[−13.73, 7.18]
Tuition deduction: single, lower	Type 2	[−31.21, 52.69]	[−60.65, 100.97]	[−8.41, 4.22]	[−20.55, 11.27]
Tuition deduction: joint, lower	Type 2	[−41.01, 61.84]	[−80.26, 122.56]	[−13.19, 9.42]	[−27.35, 18.53]

Continued on next page

Table E1 (continued)

Policy	Type	WTP: point	WTP: repetitions	Net cost: point	Net cost: repetitions
HOPE + Lifetime Learning	Type 2	[−9.75, 9.74]	[−18.77, 21.36]	[−0.96, 1.02]	[−3.34, 3.05]
AOTC: dep. single, upper	Type 2	[−42.93, 67.92]	[−78.48, 124.90]	[−7.23, 4.42]	[−20.14, 12.27]
AOTC: dep. single, lower	Type 2	[−40.94, 64.89]	[−74.94, 119.28]	[−6.44, 3.93]	[−18.85, 11.59]
AOTC: joint, upper	Type 2	[−55.17, 87.66]	[−99.34, 158.36]	[−10.68, 6.71]	[−26.88, 16.38]
AOTC: joint, lower	Type 2	[−53.63, 85.31]	[−96.67, 154.28]	[−10.12, 6.35]	[−25.91, 15.77]
AOTC: ind. single, lower	Type 2	[−179.55, 284.89]	[−317.82, 506.61]	[−56.78, 36.09]	[−111.37, 68.42]
AOTC: simulated instrument	Type 2	[−31.48, 55.32]	[−57.80, 102.28]	[−6.14, 3.47]	[−16.84, 9.45]
Tennessee Promise	Type 2	[−48.72, 152.50]	[−58.19, 160.81]	[−38.04, 32.66]	[−40.97, 35.97]
Tennessee Promise	Type 3	[86.44, 86.44]	[70.60, 100.77]	[−14.83, −14.83]	[−19.92, −9.30]
Course placement	Type 2	[0.22, 0.22]	[0.22, 0.22]	[−5.55, 2.82]	[−6.11, 3.20]
Tuition deduction: single, upper	Type 2	[−34.67, 54.63]	[−67.26, 106.70]	[−8.83, 5.48]	[−20.99, 13.09]
Tuition deduction: joint, upper	Type 2	[−40.59, 64.21]	[−80.83, 128.11]	[−15.07, 10.35]	[−29.69, 19.53]
Cannabis education channel	Type 3	[−2.50, −2.50]	[−2.70, −2.28]	[0.04, 0.04]	[0.04, 0.05]

Notes: Components use thousands of each case-specific account unit; monetary years and policy doses differ. Point entries give candidate extrema. Repetition entries give the 2.5th percentile of minima and 97.5th percentile of maxima under 1,000 prescribed paired inputs. Stars mark fewer than 1,000 usable repetitions; Table 4 reports ratio counts. WTP and cost extrema may come from different paths and must not be divided to form an MVPF. All thirteen admitted Type 3 specifications have complete prescribed input schemes, with inherited and extended correlation groups documented in Appendix E.2. These conditional summaries do not establish confidence coverage or path probabilities. Cannabis reports its education-only Type 3.

Table F1: Published male and female degree effects and standardized education NPVs

Policy	Sex	Effect (pp)	SE (pp)	Paths	NPV span ($1,000)
Adams Scholarship	Men	-0.8	0.9	2,325	[−15.36, 1.21]
	Women	-1.4	1.0	1,913	[−16.26, −0.61]
STBF scholarship	Men	8.0	–	5,173	[4.22, 12.63]
	Women	8.0	–	5,232	[4.42, 11.95]
Wisconsin Scholars Grant	Men	4.7	–	2,145	[−14.34, 8.28]
	Women	4.7	–	2,108	[−12.21, 3.46]
Kalamazoo Promise	Men	4.1	7.0	5,642	[−3.64, 6.69]
	Women	13.2	7.7	5,436	[8.46, 16.53]
Michigan CC tuition	Men	1.7	1.0	7,737	[−1.42, 6.35]
	Women	1.9	1.1	7,040	[−2.22, 4.80]
Early Head Start	Men	13.0	4.0	1,063	[−5.47, 12.34]
	Women	21.2	5.0	946	[−1.17, 13.96]
Tennessee Promise	Men	1.2	–	2,888	[−2.13, 10.52]
	Women	4.1	–	902	[−1.71, 7.09]

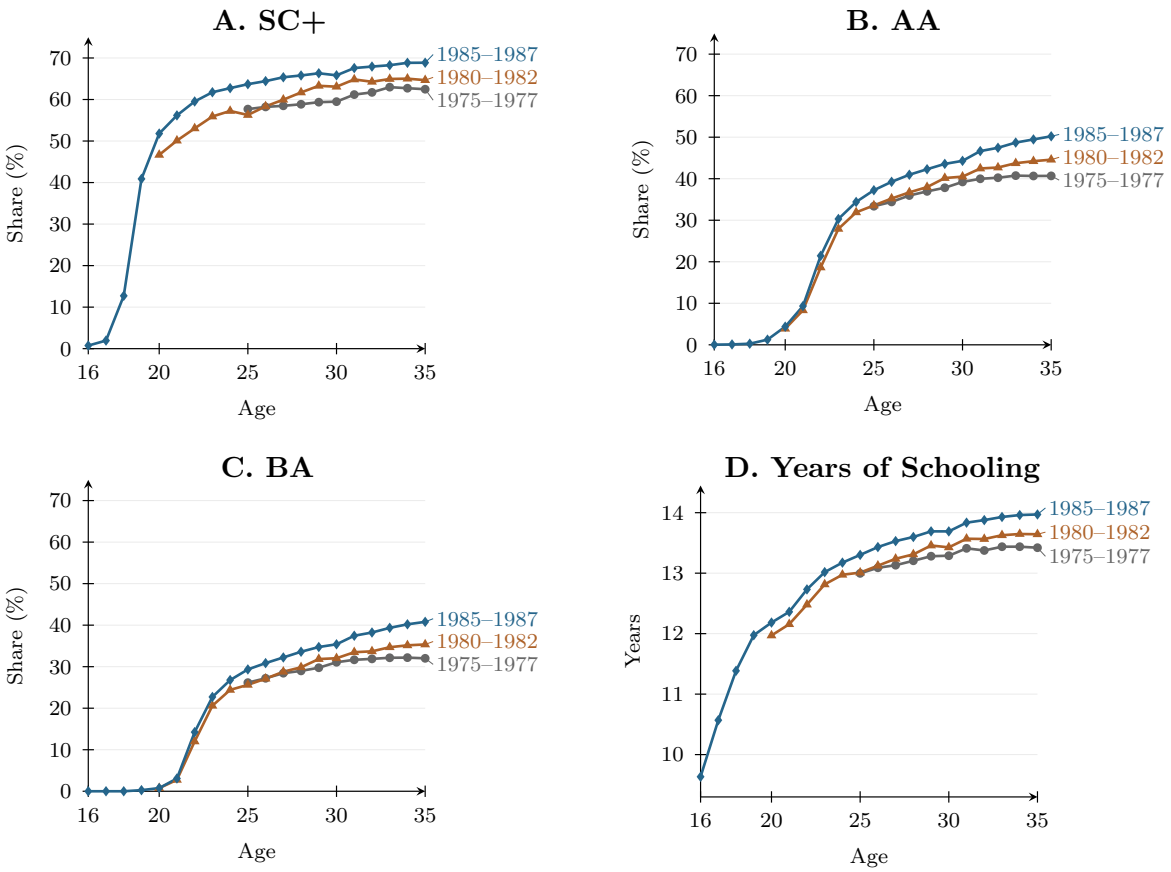
Notes: Effects are the source studies' subgroup point estimates, in percentage points. Source SEs are reported where available; a dash means that a subgroup SE was not recovered. Tennessee retains a separate auxiliary subgroup interval sensitivity. STBF and Wisconsin use rounded published coefficients; their identical displayed effects do not establish exact equality. The NPV spans allow the terminal effects to vary and therefore combine timing with eventual attainment. They summarize verified finite candidate banks, with no probabilities assigned to paths. All accounts use a common two-schooling-year equivalent per BA increment and the Section 3 net-earnings response, in 2000 dollars discounted to age 18 at 3 percent. No other education changes, transfers or fiscal costs are included. These are education-channel values, not total-policy MVPFs. Source tables, populations, age mappings and estimands are described in Appendix F.1.

Table F2: Male and female education NPVs at the same terminal effect

		Sex-specific reference curves		Common reference curve	
Policy	η (pp)	Paths (M/F)	$V_f - V_m$ ($)	Paths (M/F)	$V_f - V_m$ ($)
Adams Scholarship	-15	32/64	[−449, −249]	43/42	[−101, 35]
	-10	62/64	[−539, −31]	44/46	[−252, 179]
	-5	93/54	[−675, 332]	69/75	[−468, 384]
	0	122/0	–	95/0	–
STBF scholarship	2	0/0	–	0/0	–
	5	71/72	[−585, 733]	72/72	[−664, 664]
	10	66/70	[−626, 559]	67/67	[−587, 587]
	15	0/0	–	0/0	–
Wisconsin Scholars Grant	2	74/83	[−95, 863]	84/84	[−425, 425]
	5	90/0	–	98/98	[−576, 576]
	10	0/0	–	0/0	–
	15	0/0	–	0/0	–
Kalamazoo Promise	2	25/0	–	24/0	–
	5	26/0	–	26/0	–
	10	0/19	–	25/20	[475, 1,279]
	15	0/18	–	0/19	–
Michigan CC tuition	2	63/65	[−602, 404]	63/64	[−474, 527]
	5	58/60	[−513, 317]	59/60	[−391, 442]
	10	0/0	–	0/0	–
	15	0/0	–	0/0	–
Early Head Start	2	45/43	[37, 208]	49/33	[210, 356]
	5	45/51	[35, 194]	50/44	[210, 347]
	10	41/59	[29, 231]	48/52	[217, 383]
	15	0/0	–	0/53	–
Tennessee Promise	2	21/11	[39, 320]	13/14	[58, 316]
	5	22/11	[46, 277]	13/16	[85, 312]
	10	21/0	–	0/17	–
	15	0/0	–	0/0	–

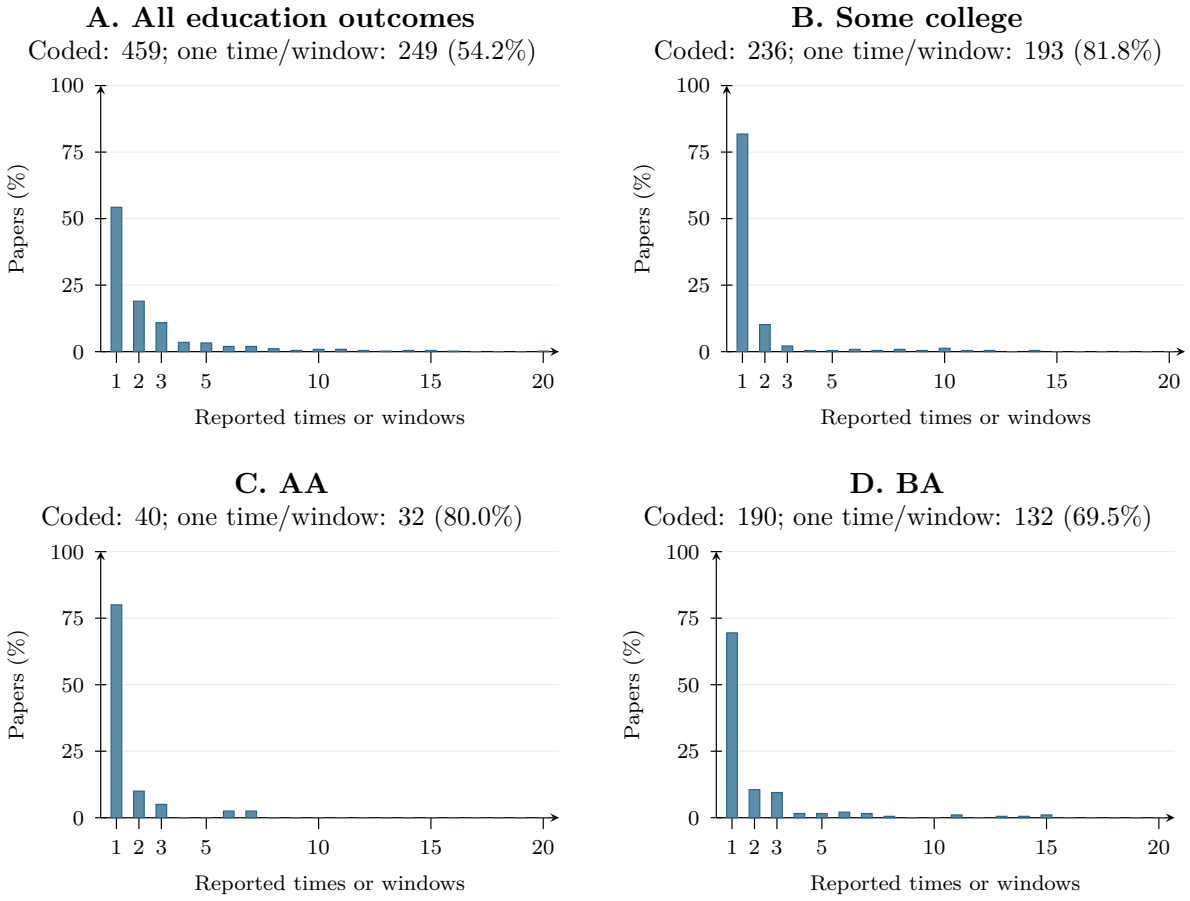
Notes: Each pair matches its own male and female source effects and the same age-65 BA effect η. Prices and exposure-age weights are identical across sexes within a policy. Negative η represents a common terminal reduction, as allowed for Adams. The common-reference sensitivity replaces both reference curves by their equally weighted mean, retaining the same inputs and grid. Entries count verified paths, not observations or repetitions. The displayed dollar span runs from the smallest to the largest female-minus-male NPV difference among the retained pairs. A dash denotes no verified pair at that grid value; finite search failure does not prove incompatibility. Spans are numerical summaries, not confidence intervals or certified global bounds. Tennessee also retains a separate auxiliary subgroup interval sensitivity. Policy populations and estimands differ, so these are within-policy comparisons rather than a cross-policy welfare ranking.

Figure A1: Education profiles across birth cohorts, ages 16–35



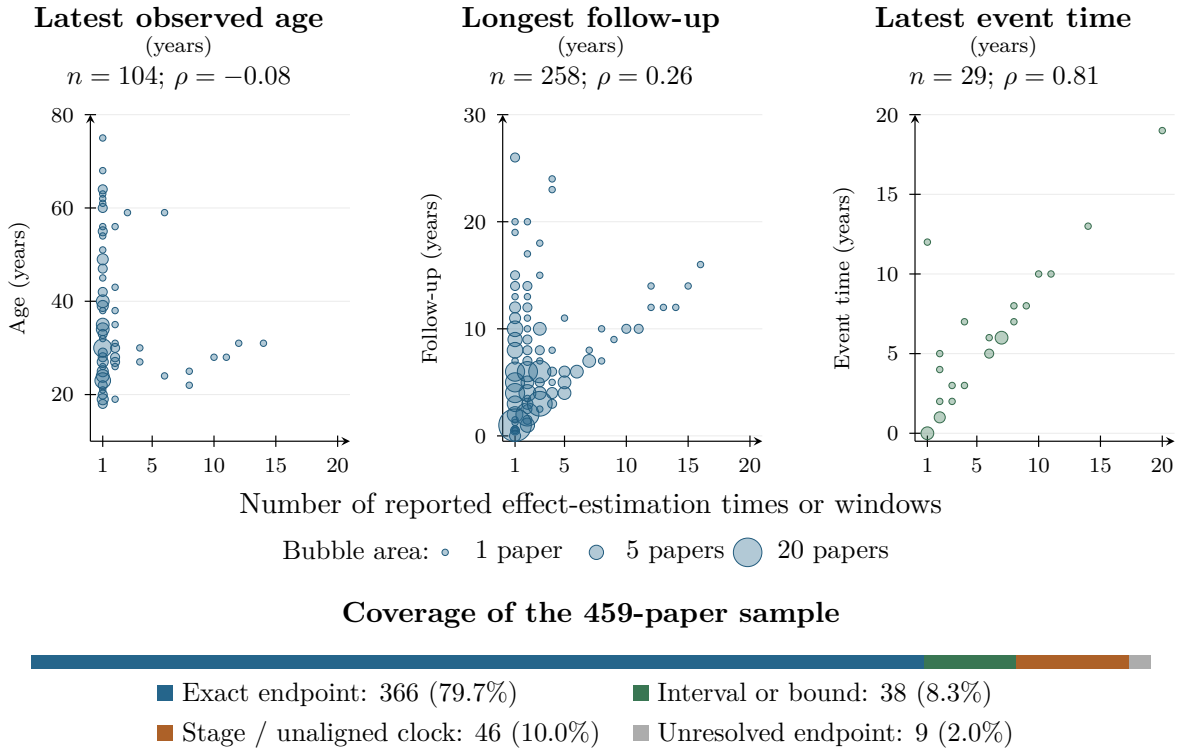
Notes: Labels at the right identify birth cohorts. Markers show observed annual education levels; lines connect adjacent observed ages. With fixed one-third cohort weights, Census 2000 and ACS 2001–2022 support ages 25–35 for 1975–1977, ages 20–35 for 1980–1982, and ages 16–35 for 1985–1987. Missing earlier ages remain unplotted. Person weights apply within cohort–age cells; all sexes, birthplaces, and racial and ethnic groups are included. Attainment panels share percentage scales; schooling uses years. SC+, AA, and BA are cumulative attainment measures. Appendix A.2 gives construction details.

Figure A2: Number of reported effect-estimation times or windows



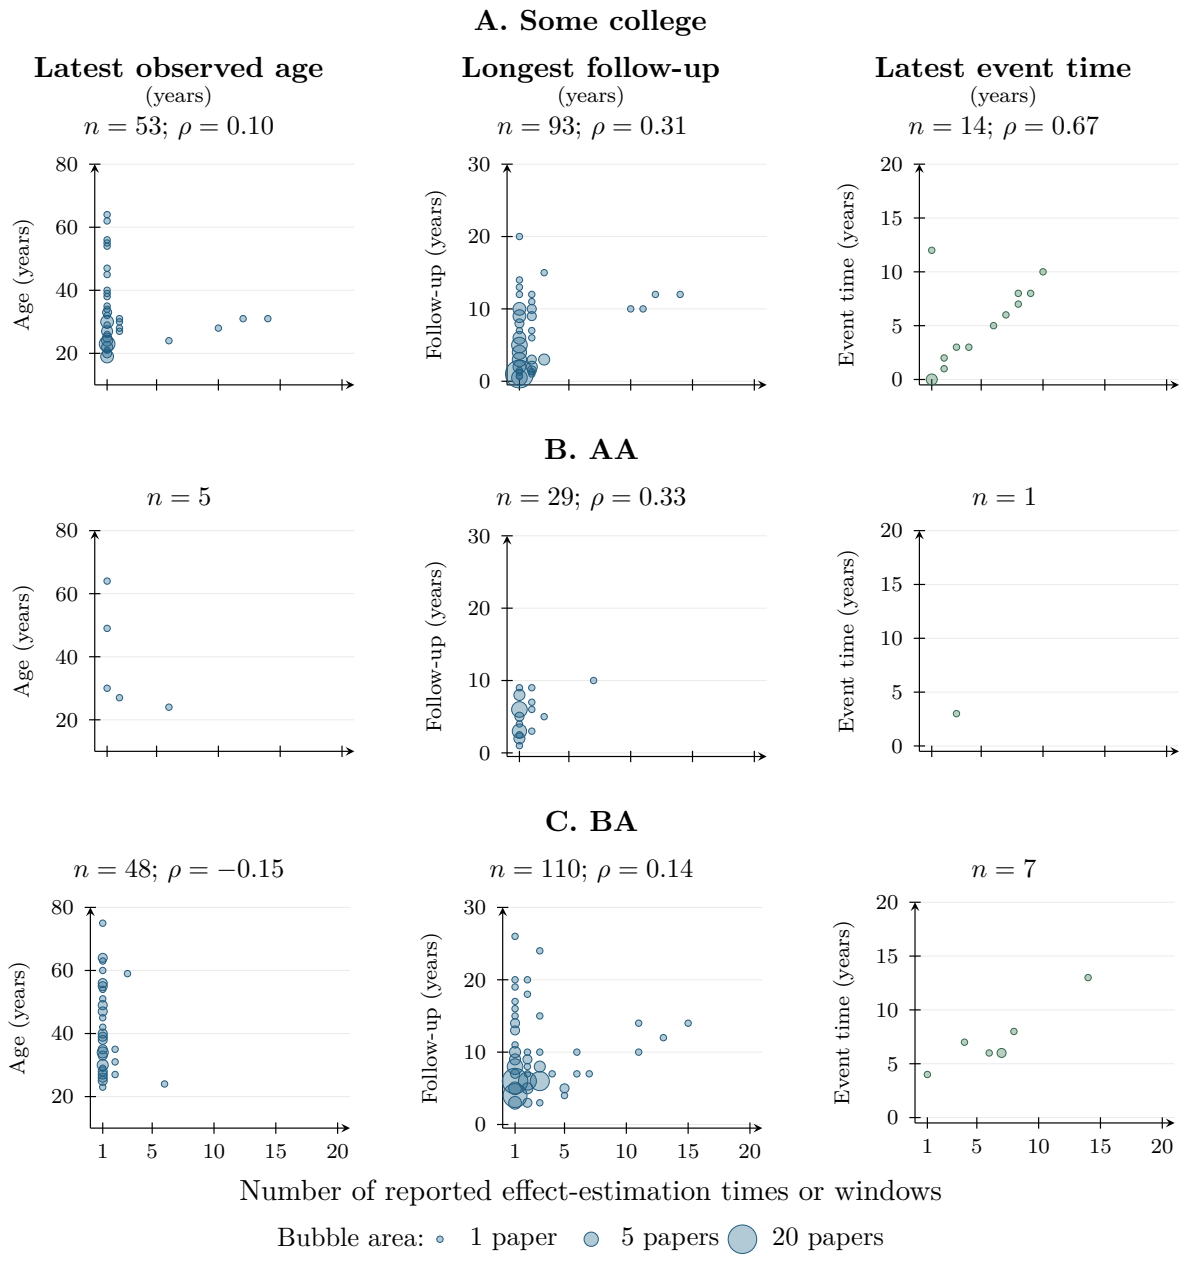
Notes: Percentages use papers with a usable count in each overlapping outcome group; one paper reporting some college is excluded. A point or pooled window counts once; separately reported horizons count separately. Counts do not measure follow-up length or effect persistence. Axes are common. Table A1 gives counts; Appendix A.1 gives coding rules.

Figure A3: Estimation times and observation endpoints: all education outcomes



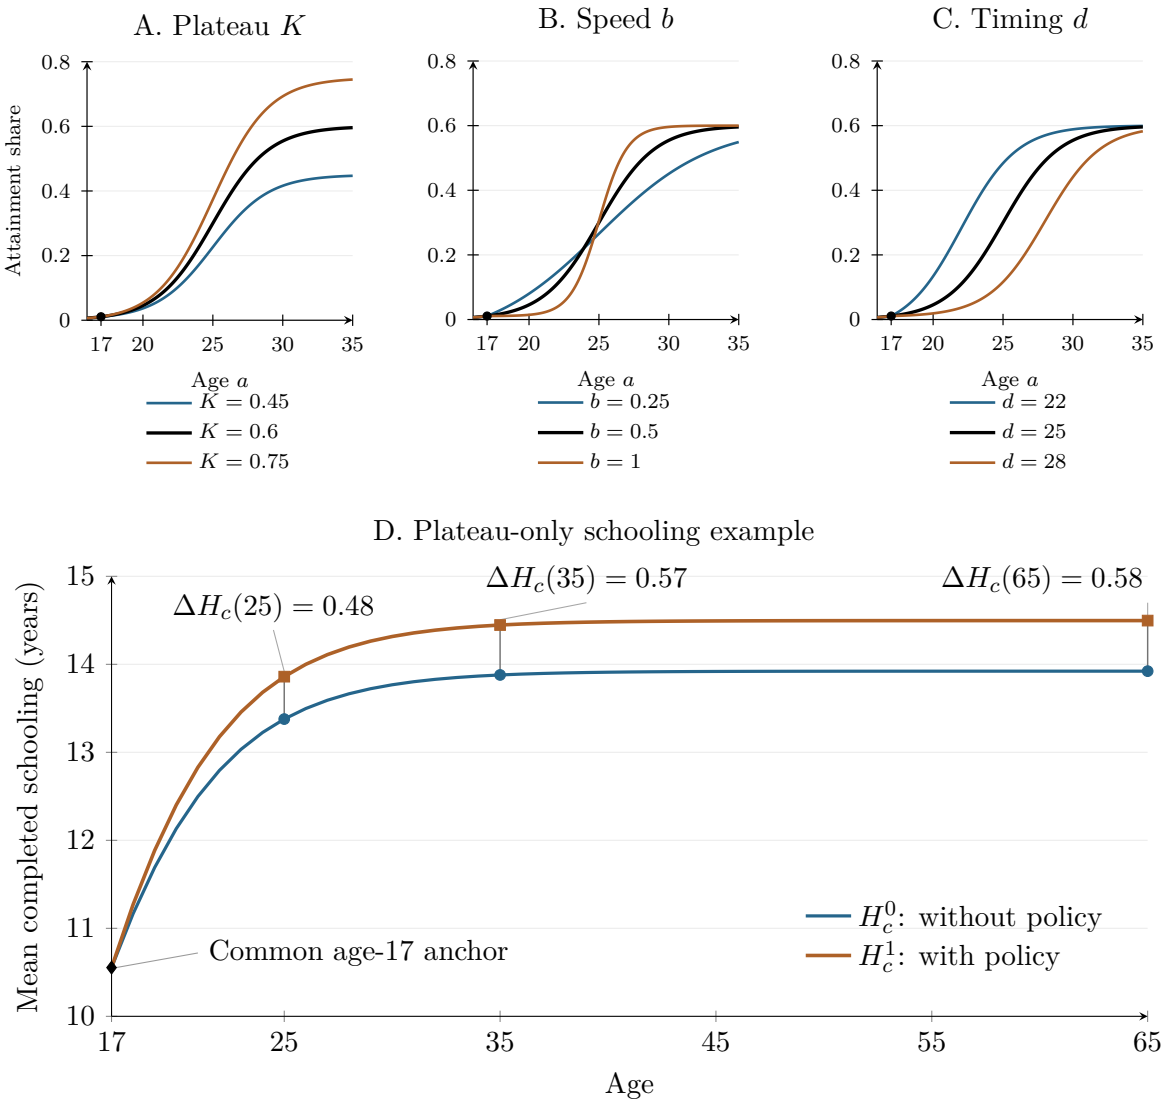
Notes: Bubble area counts papers at each coordinate. Panels separate attained age, individual follow-up, and policy event time, covering 366 distinct papers with overlap across clocks. The coverage bar includes all 459 papers; 93 lack an exact coordinate and receive no imputed endpoint. Correlations are paper-level Spearman correlations. Overall counts and endpoints may summarize different outcomes. Appendix A.1 gives coding rules.

Figure A4: Estimation times and observation endpoints by education outcome



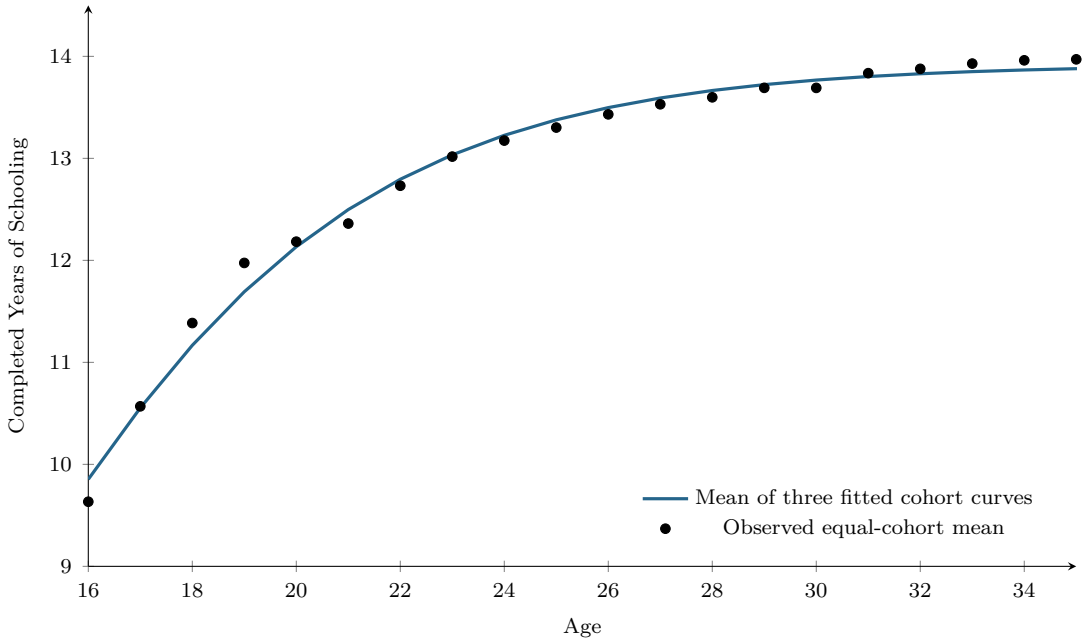
Notes: Rows show some college, AA, and BA outcomes; columns retain separate clocks. Scales are common within clocks; bubble areas share one count scale. Coverage is 146 papers on some college, 34 on AA, and 158 on BA, with overlap. Spearman correlations require at least ten papers and variation in both coordinates. AA preserves each source's exact- or at-least-degree definition.

Figure B1: Parameter roles and the schooling paths underlying Figure 4



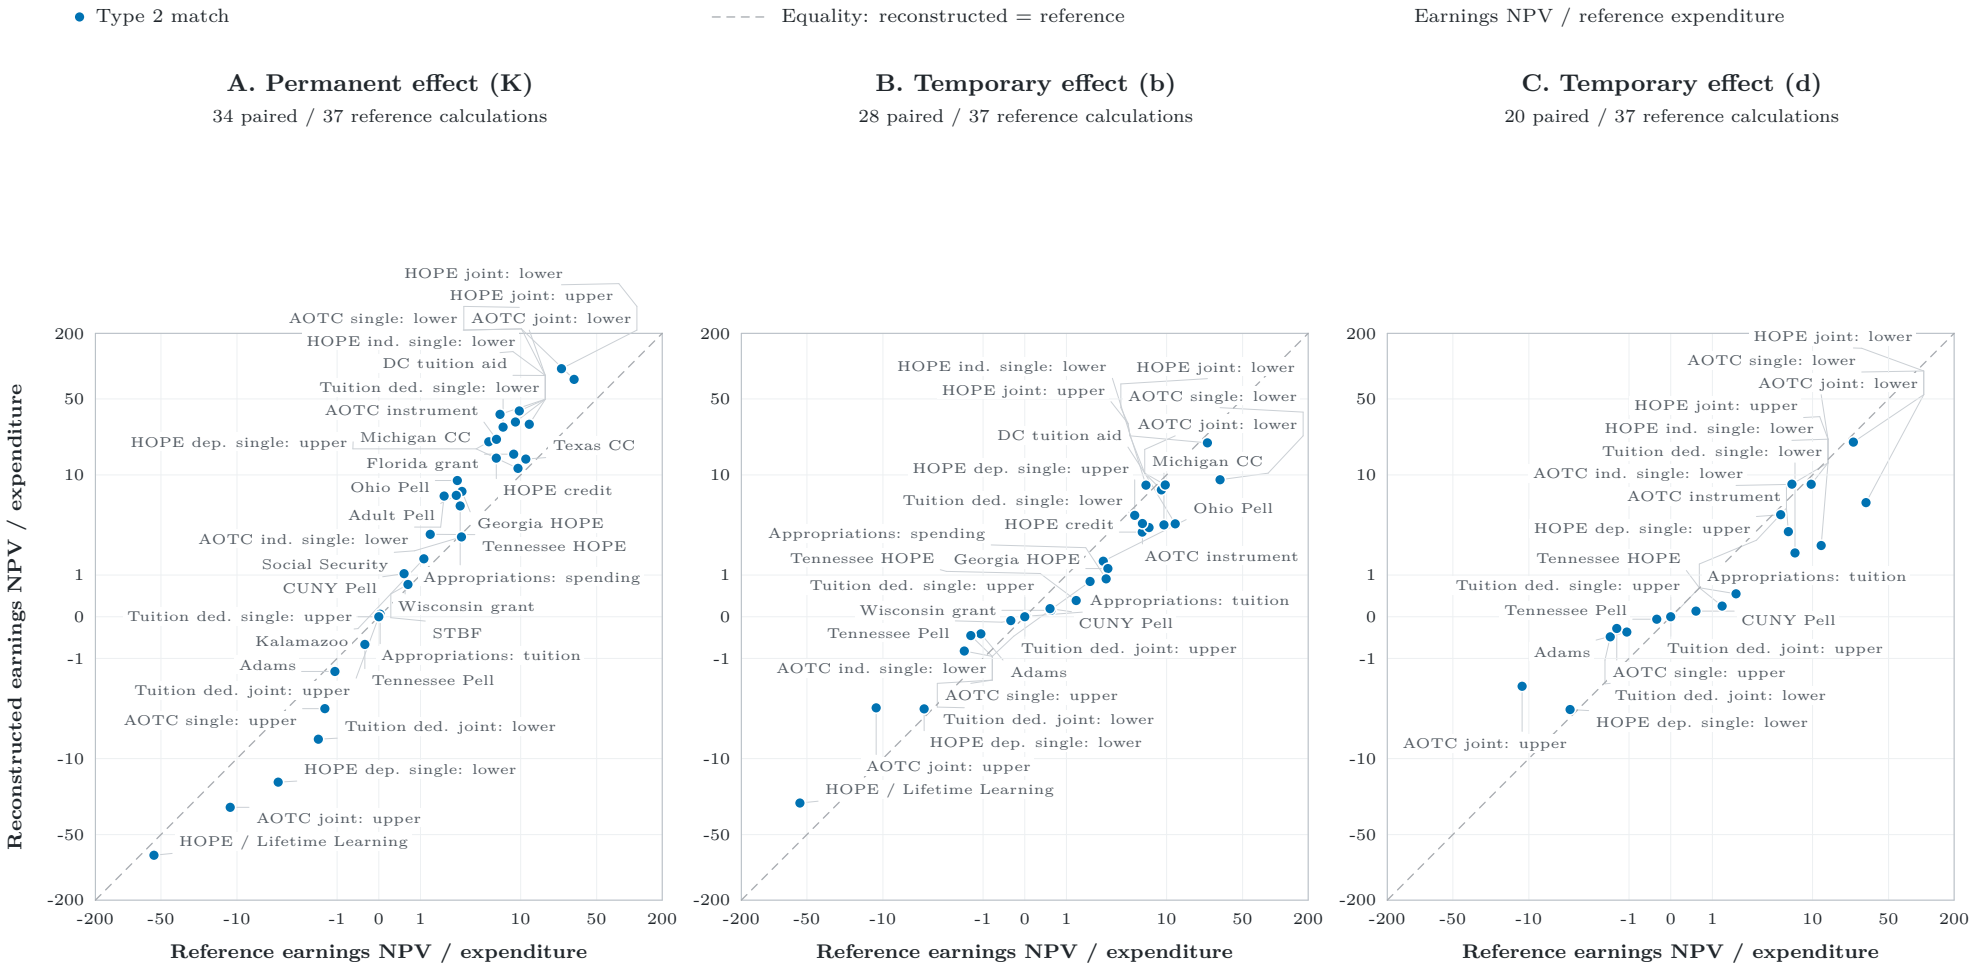
Notes: Panels A–C use hypothetical attainment shares: one parameter varies around $(K, b, d) = (0.60, 0.50, 25)$, with the common age-17 level. Black, blue, and orange denote reference, smaller, and larger values. Panel D uses Years of Schooling for the ACS 1985–1987 baseline and the plateau-only illustration in Figure 4. Blue is H_c^0 and orange is H_c^1, with $(\Delta K_c, \Delta b_c, \Delta d_c) = (0.57, 0, 0)$. Lines connect annual values; marked gaps are at 25, 35, and 65. Values above age 35 are extrapolated.

Figure B2: Comparison-path fit: ACS birth cohorts 1985–1987



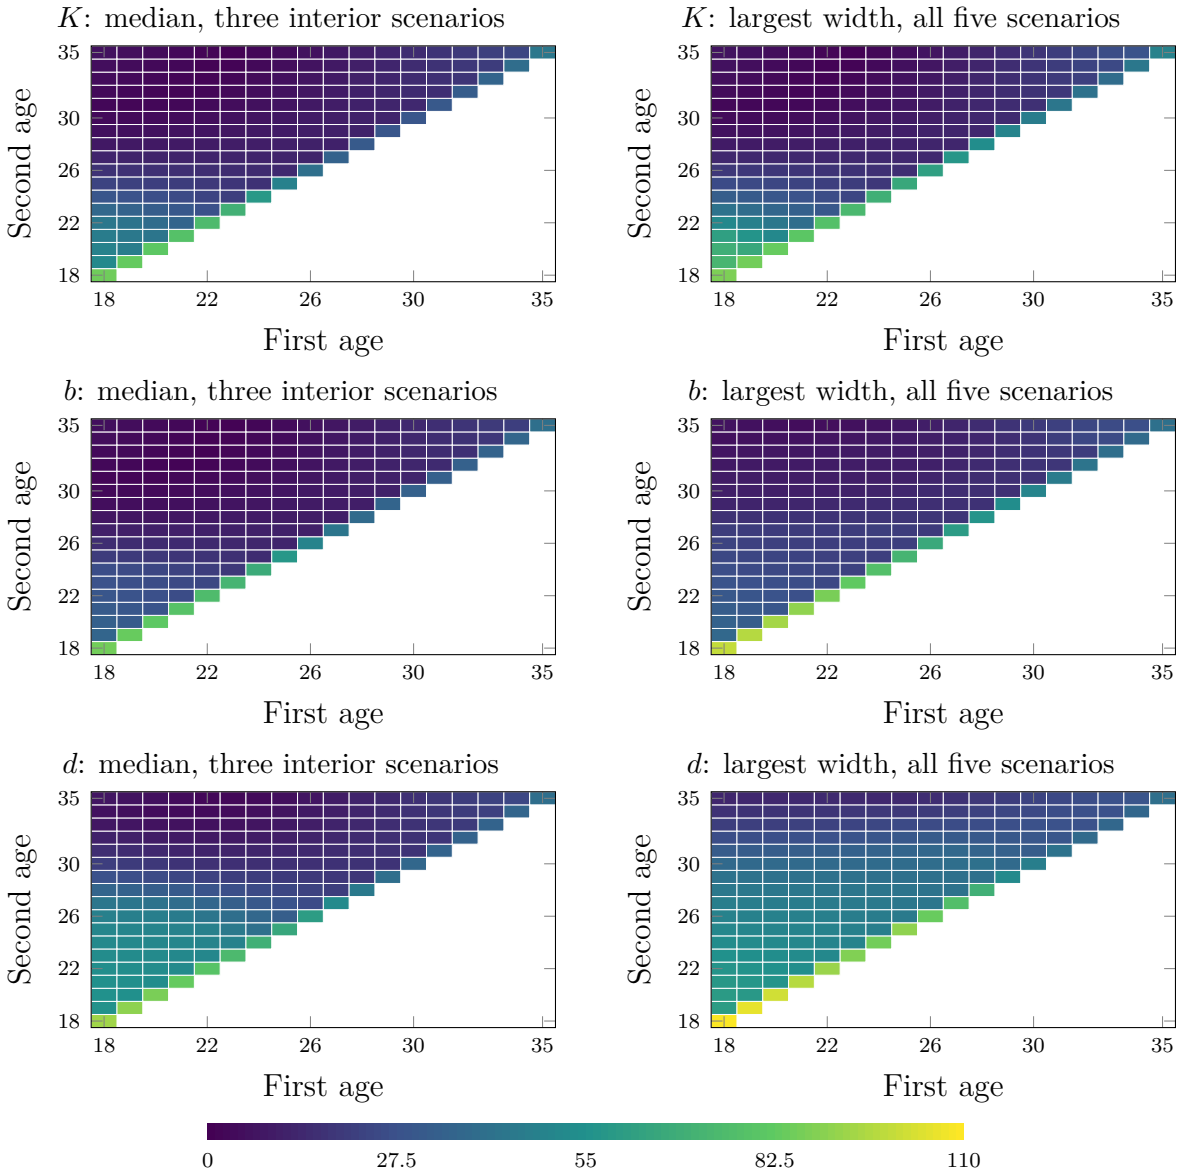
Notes: The graph covers observed ages 16–35 only. Each cohort is fitted to its person-weighted age means using normalized cell population weights. The three fitted curves and observed cell means receive fixed one-third cohort weights. This in-sample fit diagnostic does not validate extrapolation beyond age 35.

Figure B3: Reference and reconstructed earnings NPVs under permanent and temporary effects



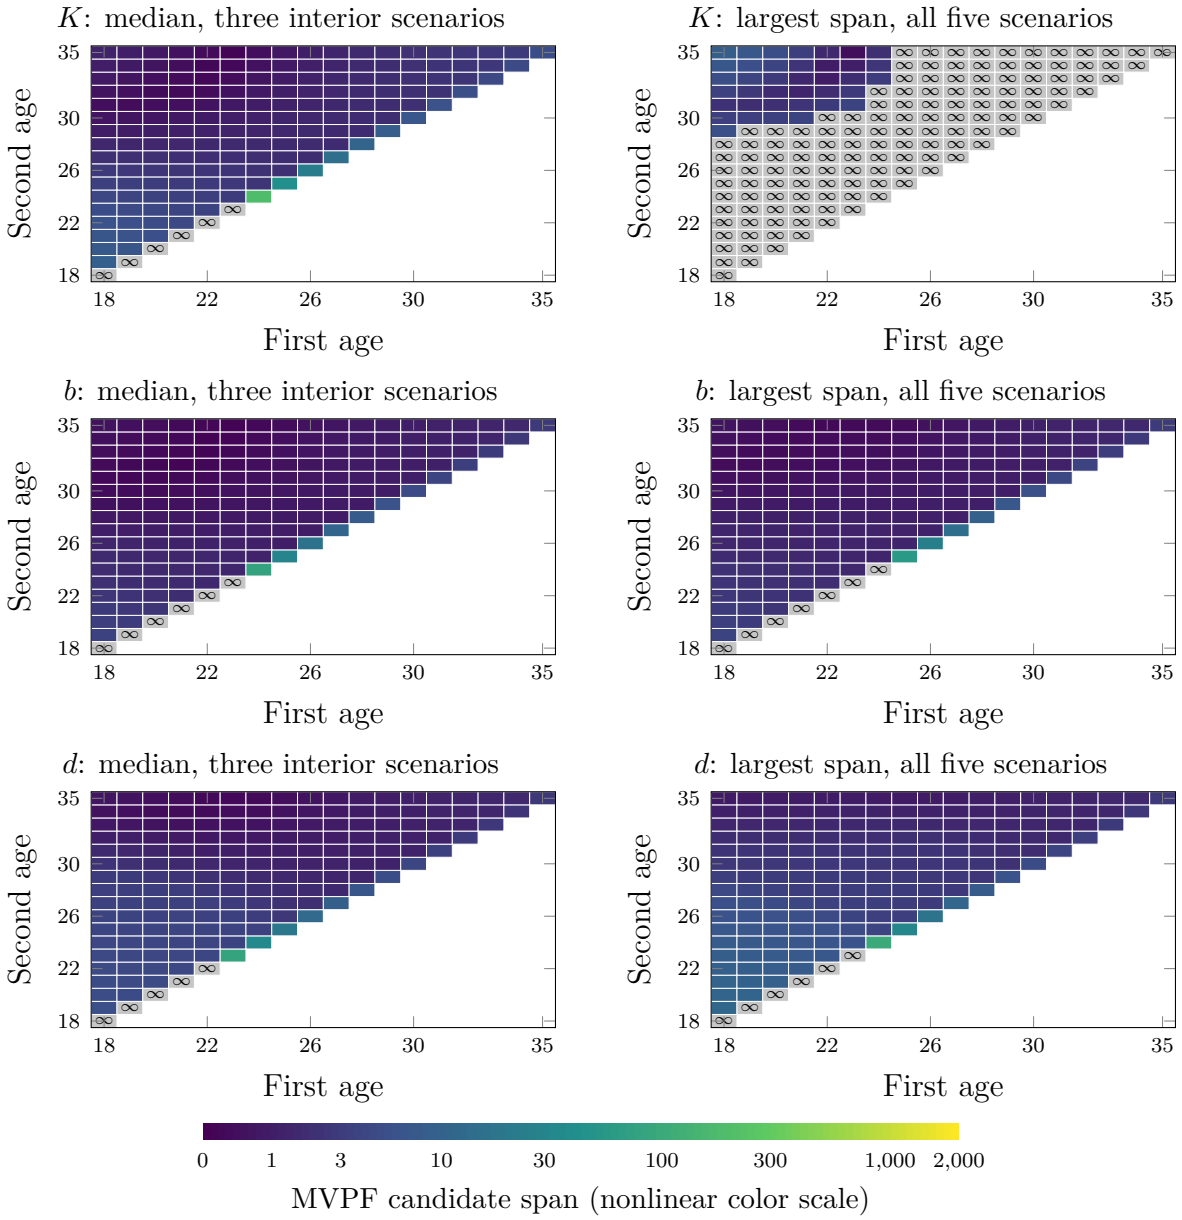
Notes: Both NPVs use matched account units and the same positive reference expenditure. Blue points match Type 2 inputs through the named parameter family; the other two families remain fixed. None also matches every additional Type 3 point input exactly. Axes have identical asinh scales, with ticks in actual normalized units; the diagonal denotes equality. Labels identify policies. Table B3 reports values and unpaired cases; Table 5 reports parameter changes. Early Head Start and Tennessee Promise lack verified common monetary years, and algorithmic placement lacks a reference earnings projection. These cases are excluded from the paired plot. Appendix B.5 explains joint inputs and the comparison.

Figure B4: Observation timing and remaining education-NPV ranges



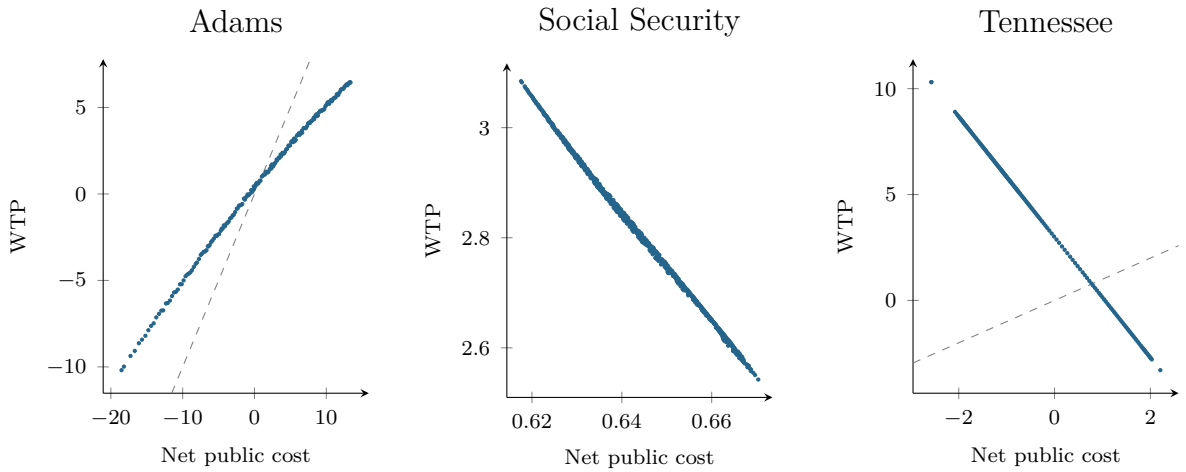
Notes: Diagonal cells represent single-age observations; upper-triangle cells represent two distinct ages. All panels use the same width scale, in thousands of 2000 dollars per person. The widths come from verified numerical candidates under the common baseline, path restrictions, and valuation rule; they are not confidence intervals or certified global bounds. Left panels report medians across the three interior scenarios; right panels report largest widths across all five scenarios.

Figure B5: Observation timing and remaining MVPF ranges under a fixed grant account



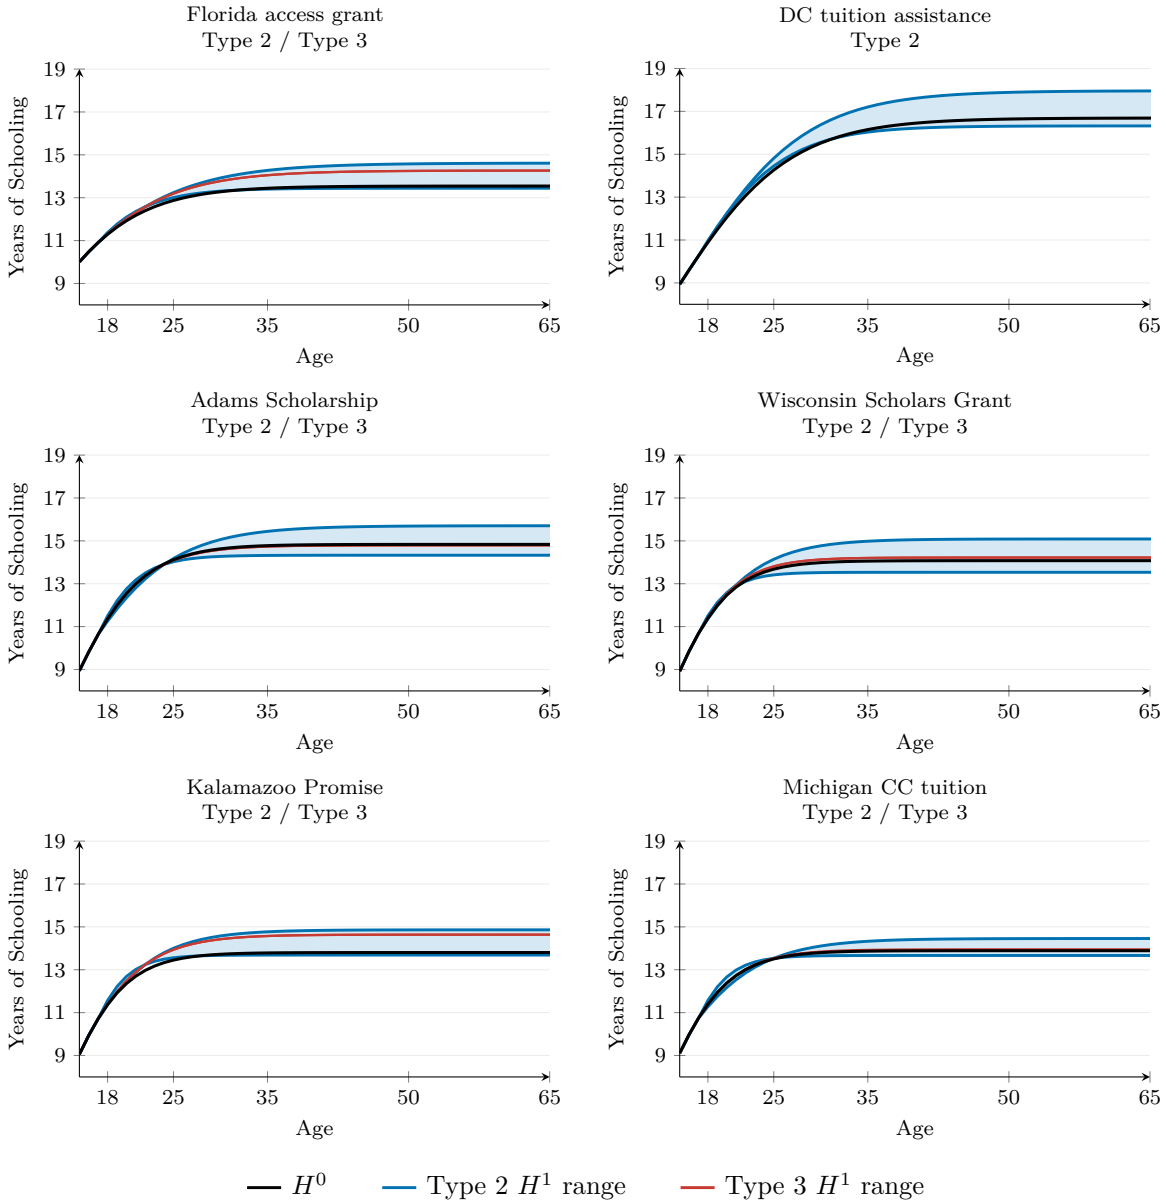
Notes: Diagonal cells use one observed age; upper-triangle cells use two distinct ages. The fixed account has a $10,000 grant, valued dollar-for-dollar by recipients, and a 20 percent earnings tax rate. Dollars and discounting follow the NPV simulation. Other non-earnings components are zero. Colors show the largest minus the smallest candidate MVPF, using a common $\log(1 + \text{span})$ scale; tick labels give actual spans. Gray cells mark infinite summary spans that contain both finite values and paths under which the policy pays for itself. Left panels report the median over three interior scenarios; right panels report the largest span over all five. Sets containing only paths under which the policy pays for itself have zero dispersion under the convention, but no displayed aggregate cell has zero span. Results are conditional on the specified paths and account, not confidence intervals or certified bounds.

Figure D1: Paired welfare components under the source-specific path mappings



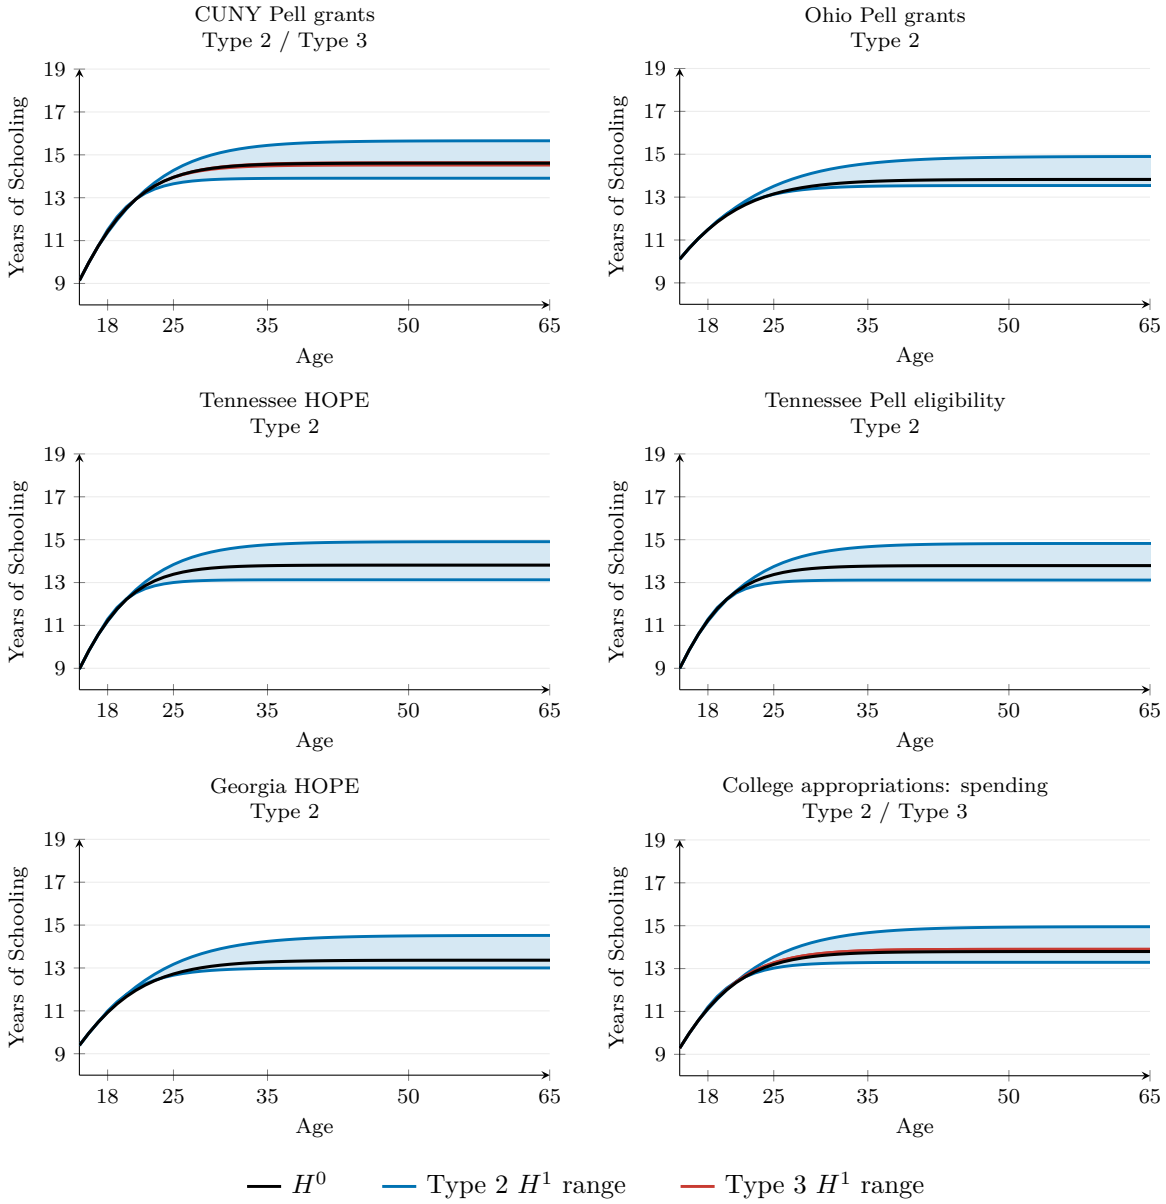
Monetary components are divided by the same fixed reference expenditure within each panel. Points are complete Type 2 paths, sampled across the earnings range. The dashed line is WTP equal to net public cost. Axis ranges vary by policy; coordinates from different monetary vintages are not dollar comparisons. No separate opportunity cost is deducted from net earnings.

Figure D2: Comparison and treated schooling paths (1 of 6)



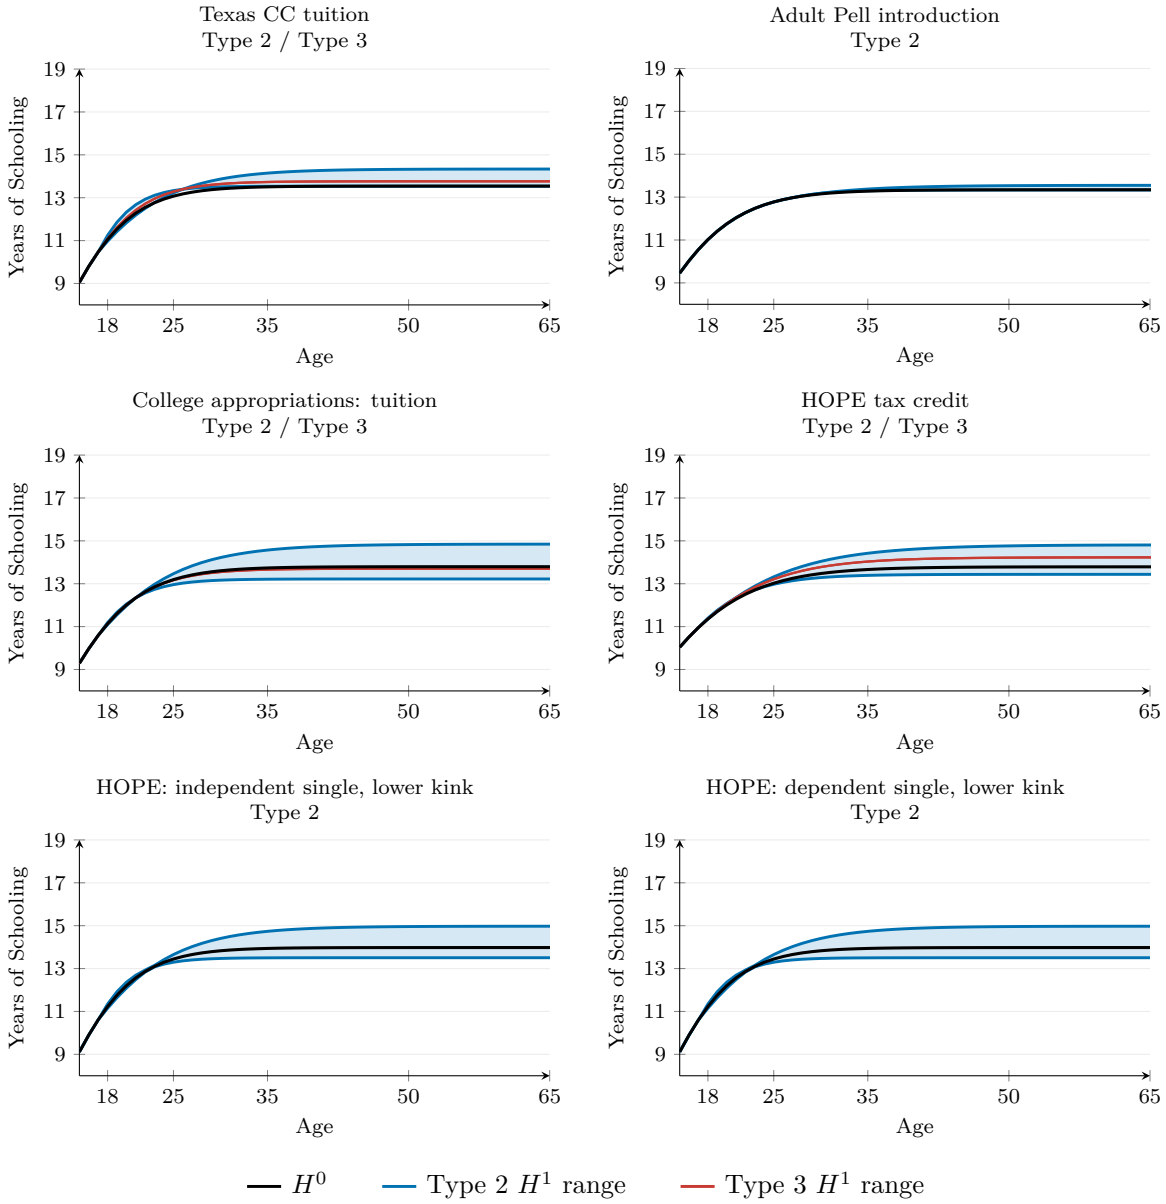
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D3: Comparison and treated schooling paths (2 of 6)



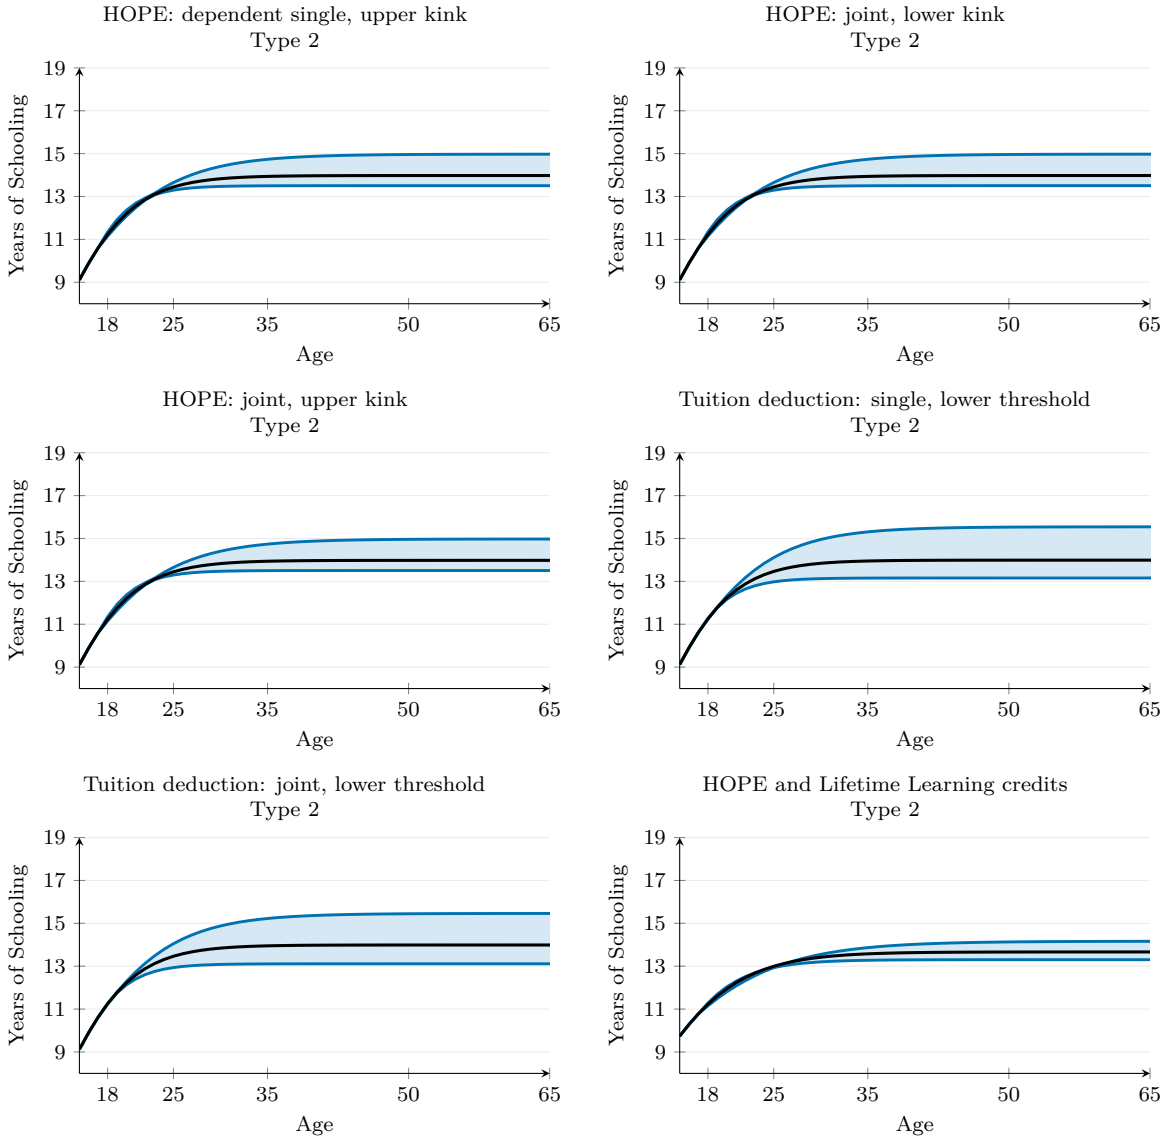
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D4: Comparison and treated schooling paths (3 of 6)



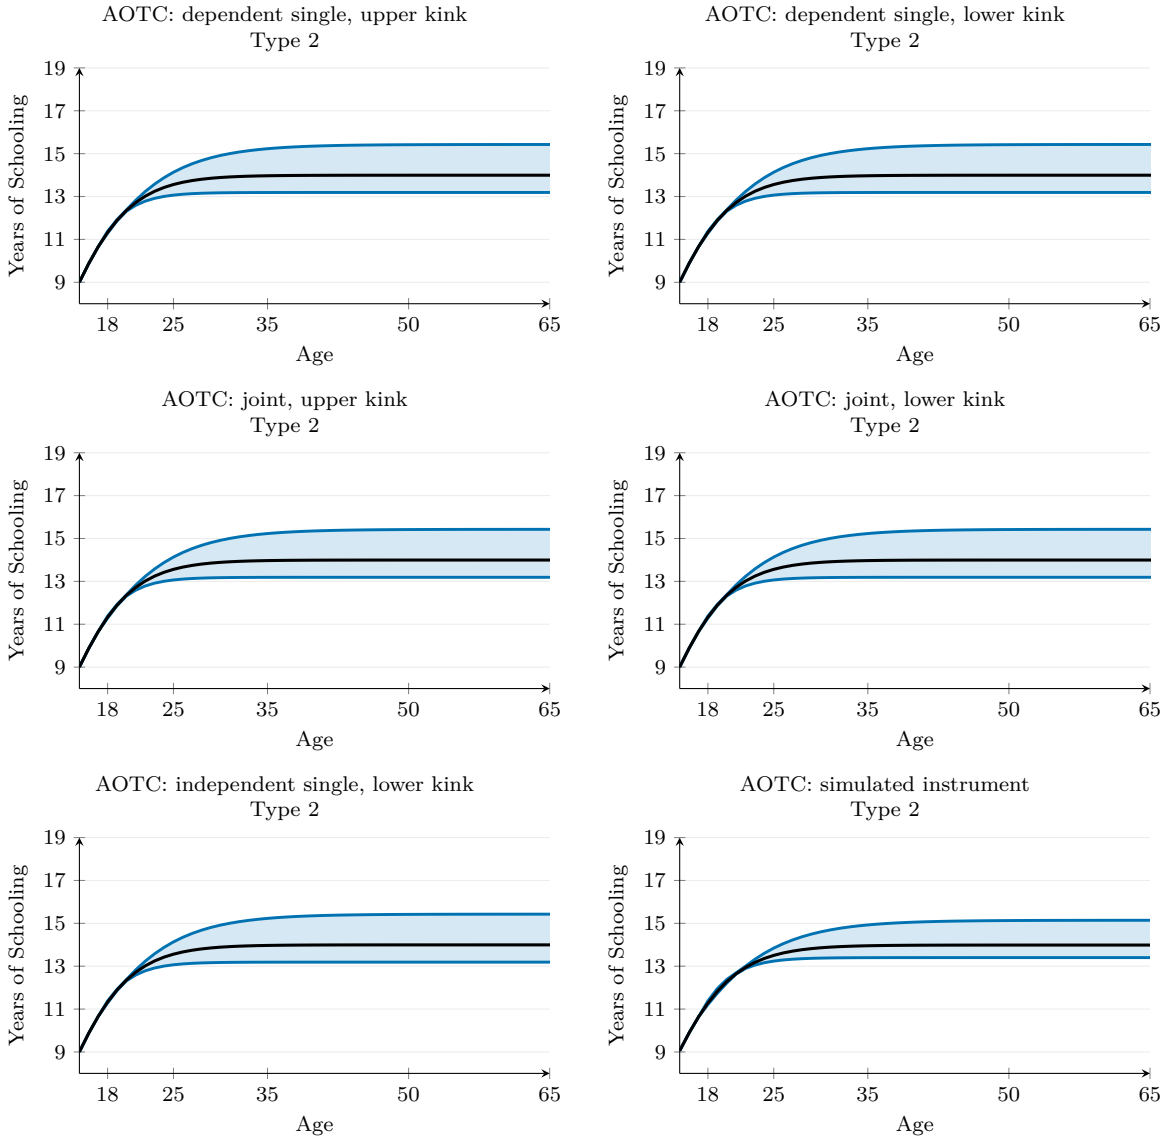
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D5: Comparison and treated schooling paths (4 of 6)



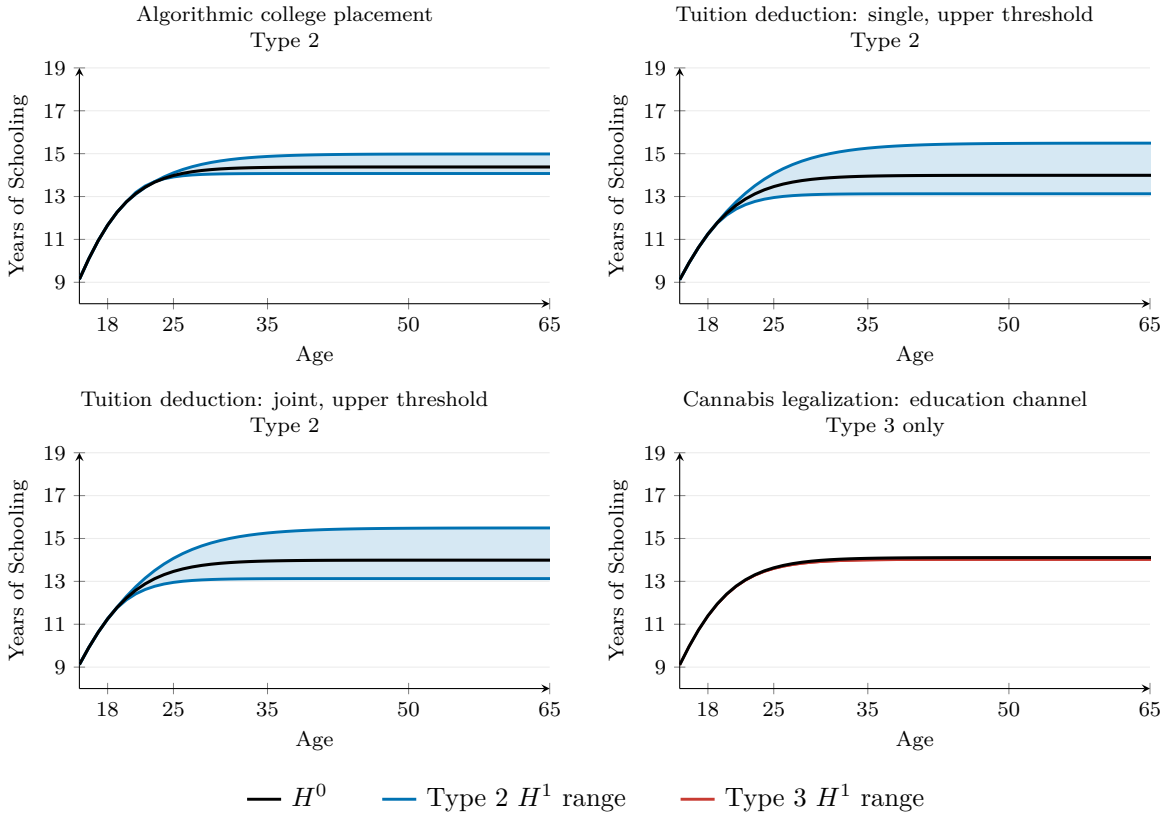
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D6: Comparison and treated schooling paths (5 of 6)



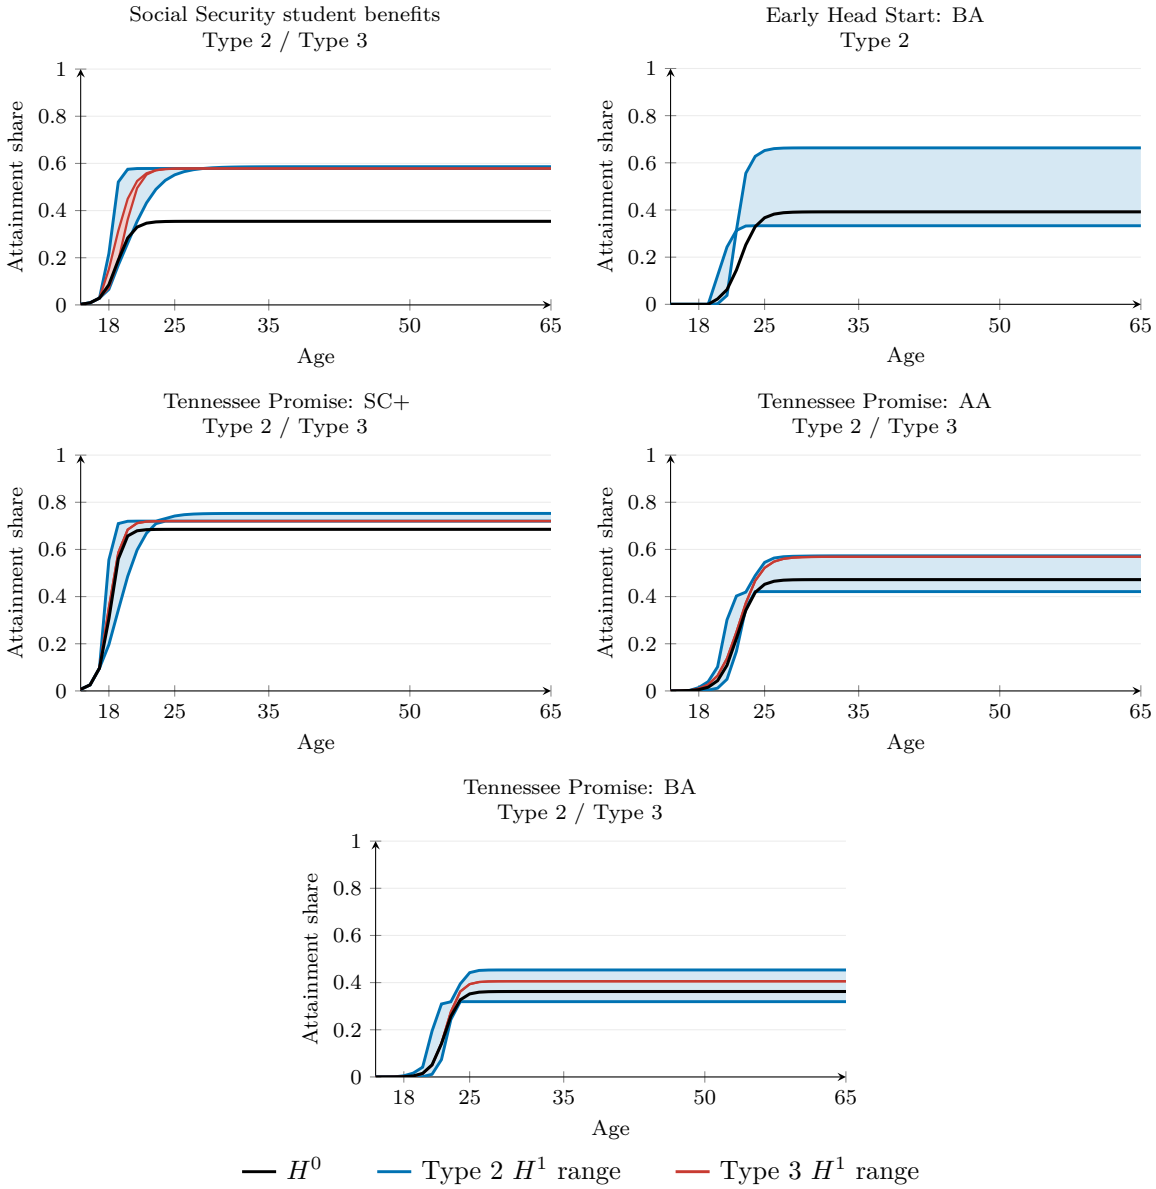
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D7: Comparison and treated schooling paths (6 of 6)



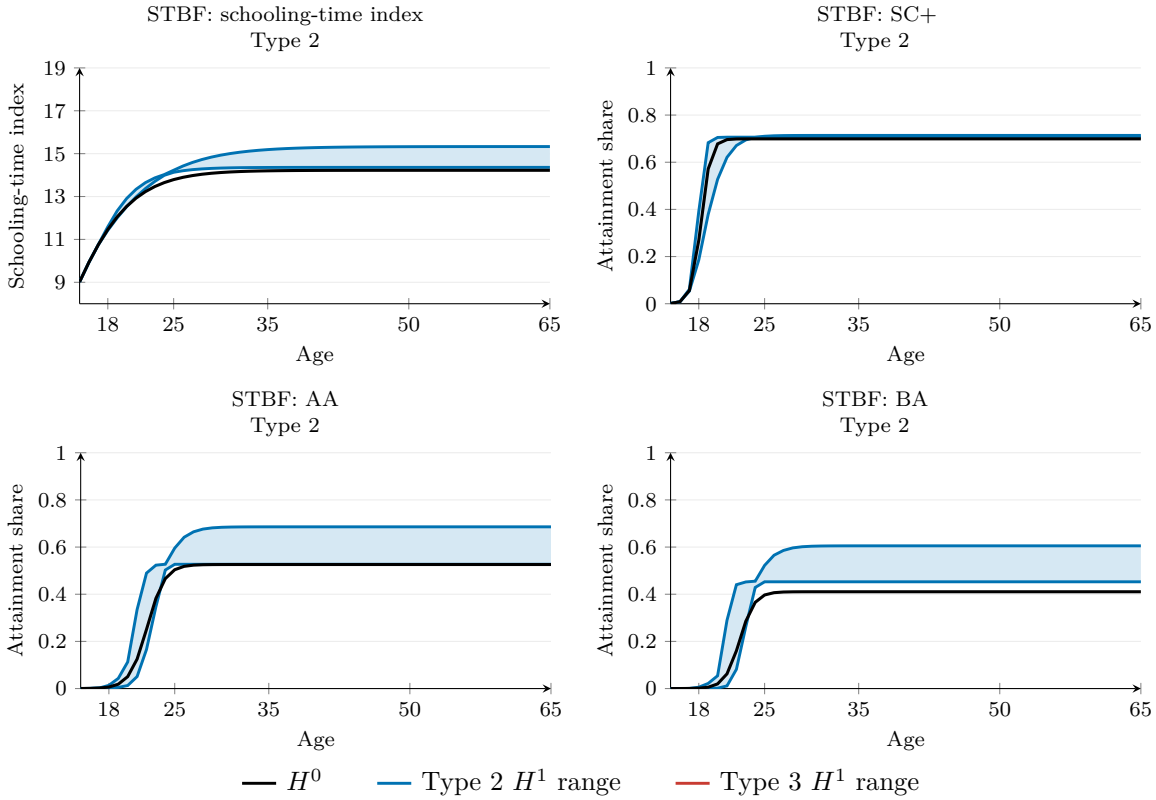
Notes: Black is the matched comparison profile H^0. Blue (Type 2) and red (Type 3) boundaries and shading show the pointwise minimum and maximum treated level H^1 among the candidates. An edge may combine different paths across ages. Cannabis shows its primary Type 3 only. All panels share scales. The ranges summarize the finite search, not confidence bands or path probabilities. Appendix D.2 records source windows and age mappings.

Figure D8: Comparison and treated credential paths



Notes: Black is H^0; blue and red boundaries and shading show the pointwise H^1 ranges for Type 2 and Type 3. All panels share age and attainment-share scales. Tennessee ranges come from joint candidates satisfying credential nesting and the restrictions in Appendix D.2. Envelope edges may combine different candidates across ages and credentials; they need not form a feasible joint path. These finite-search ranges are not confidence bands or probability distributions. Targets and calibrations differ across the three applications.

Figure D9: Comparison and treated paths: STBF



Notes: Black is the comparison profile H^0; blue boundaries and shading show the pointwise Type 2 treated range H^1. All four outcomes are matched jointly at the same six-year follow-up. The schooling-time index uses the Years of Schooling calibration scale and is not relabeled as completed schooling in the income model. Its scale differs from the three credential shares. Boundaries can combine different candidates across ages and outcomes; they are not confidence bands. No additional comparable schooling-time or credential window is admitted as Type 3.