

FOR PUBLIC DISTRIBUTION

THE ECONOMICS OF ACER EDUCATION TRANSFORMATION

*An Appendage to Volume 1 of the KR5 Framework
For Public Distribution, Review and Comment*

Document Type: Internal Companion Volume

Classification: Public Review — KR5 Initiative | For Public Distribution

Date: August 2026

Prepared by: KR5 Initiative Research Team — Economics Working Group

Companion to: KR5 Framework Circulation Document, Volume 1, Section 2A

"The cost of education is not what we spend. It is the gap between what we spend and what we get — and that gap, in American higher education today, has become a moral emergency." —

KR5 Economics Working Group, August 2026

CONTENTS

Table of Contents

Prefatory Note	2
Section 1 — The Economic Argument Is the Moral Argument	3
Section 2 — The Sticker Price and the Real Price	5
2.1 The Published Prices and Why They Matter	
2.2 The Room and Board Question	
2.3 Books, Supplies, Transportation, and Personal Expenses	
2.4 The Net Price: What Actually Gets Paid, and by Whom	
Section 3 — Four Categories of Payer, Four Distinct Crises	9
Payer 1 — The Federal Government: Spending More, Buying Less	
Payer 2 — State and Local Governments: The Long Retreat	
Payer 3 — Household Families: The Impossible Arithmetic	
Payer 4 — The Working Student: The Race That Is Designed to Be Lost	
Section 4 — The KR5 ACER Cost Restructuring Proposals	15
4.1 Proposal One: Compress to Two Years Plus a Six-Month Employer-Funded Practicum	
4.2 Proposal Two: Seventy Percent Home/Local-Based Enrollment	
Section 5 — The Combined ACER Savings Model	21
Section 6 — The Employer Practicum: A Closer Economic Analysis	24

Section 7 — The Student Debt Question: Structural Root, Not Symptomatic Relief .. 27

Section 8 — Transition Costs and the Net Present Value of Transformation 30

Closing Reflections — The Arithmetic of Justice 32

Appendix A — Data Sources and Methodological Notes 33

Appendix B — Glossary of ACER Economic Terms 34

Prefatory Note

This document is not a policy prescription. Volume 1 of the KR5 Framework Internal Circulation Document handles that responsibility. What this companion volume undertakes is something different and, in its own way, more demanding: a careful, transparent, and fully reasoned accounting of the economic realities of American higher education as they exist in 2024–25, followed by a rigorous and explicitly argued modeling of what the ACER transformation would cost, save, and redistribute across every category of payer in the system.

The distinction matters. A policy prescription can afford to assert. An economic accounting must demonstrate. Every figure presented in this document is accompanied by the reasoning that produced it, the source that grounds it, and the assumptions — stated plainly and defended at length — that connect raw data to the conclusions we draw. Where the data is imprecise, we say so. Where assumptions are made — particularly about the costs and savings associated with ACER's implementation — those assumptions are explicitly labeled as such and their logic is explained as though to a thoughtful, skeptical reader who deserves to understand not merely what the number is but why it was chosen and what it represents for real human beings.

This document is intended for internal review by KR5 working group members, institutional partners, and invited economists. It is not, and does not pretend to be, the final word. The figures and assumptions presented here represent the best available analysis as of August 2026, and they are offered in the spirit of honest inquiry rather than advocacy dressed as arithmetic. They invite challenge, critique, and refinement. A reader who finds a better number, a stronger assumption, or a more precise source is not a critic of this document — they are its intended collaborator.

— *KR5 Economics Working Group, August 2026, Rio Rancho, NM*

SECTION 1

The Economic Argument Is the Moral Argument

There is a temptation, in conversations about educational transformation, to treat the economics as a secondary concern — the necessary but somewhat unglamorous companion to the more inspiring work of reimagining what learning can look like, what students can become, and what a truly equitable educational system might achieve. This temptation must be resisted with particular firmness in the case of the ACER model, because the economics of educational transformation are not a constraint on the vision. They are, in the most direct sense, the vision itself. Every ambitious proposal for how students should learn, what they should learn, and over what period of time they should learn it will be measured, ultimately, by a budget committee. And the question that budget committee asks is neither philosophical nor pedagogical. It is: who pays, and how much less than they pay now?

The KR5 framework is built on a conviction that is both analytical and moral: the ACER transformation is not only pedagogically superior to the current four-year model — it is dramatically cheaper for

KR5 Rajeev Kasturi All Rights Reserved August 12, 2026 © Page 4 of 52

every meaningful category of payer in the system. This claim is not made casually, and it is not asserted as a marketing proposition. It is a claim that must be earned through rigorous analysis, verified against authoritative data, and defended against the most skeptical objections that informed readers can mount. This document is that defense.

We begin with undergraduate education specifically — not graduate or professional education, not community college, not online-only programs — because undergraduate education is the largest, most expensive, and most inequitable layer of the American postsecondary system. It affects the greatest number of people. It generates the greatest volume of debt. It produces the widest variance in outcomes between students who enter with economic advantage and those who do not. It is the layer of the system where the gap between the promise of higher education and the economic reality of accessing it is most visible, most painful, and most consequential for the individuals who must navigate it.

"Every vision of educational transformation — no matter how intellectually compelling — dies in a budget committee if it cannot answer the question: who pays, and how much less than they pay now? The ACER model answers that question. This document proves it."

The economics of education reform have historically been neglected by the very people most passionate about educational change. Reformers — and the history of American educational innovation is full of passionate, brilliant, and well-intentioned reformers — have consistently focused their energy on curriculum redesign, pedagogical innovation, and institutional culture change, while leaving the cost structures of the institutions they sought to transform essentially untouched. This is not primarily a failure of intention. It reflects the structural reality that cost structures are embedded in institutional architecture — in faculty tenure systems, in physical plant commitments, in

accreditation requirements, in financial aid bureaucracies — in ways that make them extraordinarily resistant to change even when the case for changing them is compelling.

The result has been a half-century of well-intentioned reform efforts that were absorbed by institutions without altering their fundamental economic architecture. Problem-based learning, service-learning, competency-based education, flipped classrooms, MOOCs — each of these innovations produced genuine pedagogical improvements in the contexts where they were adopted, and each was ultimately annexed by an institutional cost structure that remained unchanged. Students who benefited from innovative pedagogy still paid four years of tuition. They still paid four years of room and board. The Pell Grant still covered the same declining share of those costs. The loans were still necessary. The debt was still accumulated.

The KR5 position is, at its core, a refusal to allow this pattern to repeat itself. You cannot transform what a student learns without simultaneously transforming what that learning costs. These are not separate questions that can be addressed in sequence — curriculum first, then costs. They are a single question, and the ACER model treats them as such. The Two-in-a-Box GenAI-plus-human-teacher model is not merely a pedagogical innovation. It is a cost restructuring mechanism — one that, by compressing the time required to deliver high-quality undergraduate education, directly attacks the duration-based cost architecture that has made American higher education so expensive for so long.

A note on the limits of this analysis is appropriate before we proceed. The cost and enrollment figures referenced throughout this document are 2024–25 data drawn from the College Board's annual *Trends in College Pricing* and *Trends in Student Aid* reports, the National Center for Education Statistics, the Federal Student Aid office's quarterly portfolio summaries, the Pell Institute's equity indicators research, and the Society for Human Resource Management's benchmarking studies, among other authoritative sources. All primary sources are listed in Appendix A. Where assumptions are made about the ACER model's implementation costs and savings — particularly in Sections 4 and 5 — those assumptions are explicitly labeled, their derivation is explained in full prose, and the reader

is encouraged to interrogate and refine them. The numbers in this document are our best current estimates. They are not dogma.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

SECTION 2

The Sticker Price and the Real Price: Understanding What American Families Actually Pay

2.1 — The Published Prices and Why They Matter Even Though Few Pay Them

Every year, the College Board compiles and publishes what it calls the "published prices" for American higher education — the sticker prices that institutions set before financial aid is applied. For 2024–25, these figures are as follows. At public four-year institutions for in-state students: tuition and fees of **\$11,610**, plus room and board of **\$13,310**, plus books and supplies of **\$1,240**, plus transportation and personal expenses of **\$3,760**, producing a total annual cost of attendance of **\$29,920**. At public four-year institutions for out-of-state students, tuition rises dramatically to **\$30,780**, producing a total annual COA of **\$49,090**. At private nonprofit four-year institutions, tuition reaches **\$43,350** and room and board **\$15,250**, for a total annual COA of **\$63,600**. Extrapolated to a four-year program duration, the sticker-price totals are: public in-state, **\$119,680**; public out-of-state, **\$196,360**; private nonprofit, **\$254,400**.

The instinctive response to these figures — particularly from those familiar with financial aid systems — is to note that most students do not pay them. This is true, and the concept of net price, addressed below in Section 2.4, is essential to understanding what students actually pay. But the dismissal of

sticker prices as irrelevant because they are rarely paid in full is one of the most consequential errors in public discourse about higher education affordability, and it must be challenged directly and at length, because the ACER model's economic case rests in part on understanding why sticker prices matter even when most students pay something less than the sticker price.

First: sticker prices set the psychological ceiling of perceived cost. They determine what families believe college will cost before they have done any research into financial aid, before they have spoken to an admissions counselor, and before they have filed a FAFSA. Research has consistently demonstrated that many students from first-generation, low-income, and minority backgrounds rule themselves out of four-year colleges entirely based on sticker prices — not net prices, not post-aid prices, but the raw published numbers they encounter on institution websites and in media coverage of college costs. They never reach the financial aid application stage, which means they never discover what they would actually owe. The sticker price, in other words, is not merely a pricing signal for students who enroll. It is a gatekeeping signal that silently deters a significant population of prospective students before any enrollment conversation begins.

Second: sticker prices determine the maximum debt ceiling for students who do enroll and receive financial aid. A student receiving a grant that covers half of the sticker price may still borrow against the remaining half. The sticker price is the upper bound from which all debt calculations begin, which means that even partially subsidized students are taking on debt proportional to an inflated published price rather than to a more rational cost of instruction.

Third: for a substantial and often overlooked segment of American families — those earning in the **\$60,000–\$100,000** annual household income range — the gap between sticker price and actual financial aid received is uncomfortably close to the sticker price itself. These families earn too much to qualify for maximum Pell Grant awards and need-based institutional aid, but they do not earn enough to write checks for \$30,000 per year out of discretionary income. They occupy what might fairly be called the financial aid desert: above the poverty-based assistance threshold, below the wealth threshold, and effectively paying near-full price for an education they cannot fully afford.

2.2 — The Room and Board Question

Of all the component costs within the standard cost of attendance, room and board deserves the most sustained analytical attention in this document, because it is the ACER model's primary cost target. Understanding why requires understanding the fundamental difference between tuition and room and board as economic categories, and the ACER model's differential ability to address each.

Tuition is a negotiated price. It is set by institutions, but it is also offset — sometimes dramatically — by the financial aid system: Pell Grants, institutional grants, state grant programs, and subsidized loans all operate against the tuition line in ways that can substantially reduce the out-of-pocket tuition burden. Tuition, in other words, is the part of the college cost equation that the aid system was designed to address. It is not addressed adequately — we will demonstrate this in Section 3 — but it is at least the intended target of the subsidy architecture.

Room and board is something different entirely. It is a physical cost — the cost of shelter and food for a human being living away from home — and it is, in the most straightforward sense, non-negotiable for any student who relocates to attend a residential campus. It cannot be Pell-granted away. It cannot be substantially subsidized by institutional financial aid programs without the institution absorbing a cost it has no mechanism to offset. It cannot be reduced by working harder, studying more efficiently, or choosing a more affordable major. It must simply be paid, in full, every semester, by someone.

The 2024–25 national average of **\$13,310** per year in room and board at public four-year institutions is, moreover, a conservative floor rather than a representative experience for students in major metropolitan areas. In New York, San Francisco, Boston, Chicago, and Los Angeles — cities that are home to a substantial share of the nation's four-year public universities and where a significant proportion of the student-age population lives — on-campus room and board routinely exceeds **\$18,000 to \$22,000** per year. The College Board's weighted national average masks this regional variance in ways that can lead policymakers to underestimate the room-and-board burden for students in high-cost markets.

"Room and board is the cost that cannot be Pell-granted away. It must simply be paid, in full, every semester, by someone. Over four years at the national average, it accumulates to \$53,240 — more than seven times the maximum Pell Grant award."

Over a four-year program, the cumulative room and board burden at the national average is **\$53,240**. For context: the maximum Pell Grant award for 2024–25 is **\$7,395** per year, or approximately **\$29,580** over four years. Even if a student qualified for the maximum Pell award every year — and many do not, due to income fluctuations, satisfactory academic progress requirements, and enrollment status changes — the grant would cover only 55.6% of the room and board cost alone, leaving a gap of more than **\$23,660** in room and board costs that the grant does not touch, not counting tuition. This arithmetic is at the heart of why eliminating room and board as a cost category — the ACER home-based enrollment model — is the most powerful single lever in the entire cost restructuring framework.

2.3 — Books, Supplies, Transportation, and Personal Expenses: The Invisible Costs

The College Board's figures of **\$1,240** for books and supplies and **\$3,760** for transportation and personal expenses are weighted national averages that have the virtue of methodological consistency and the deficiency of substantially understating actual expenditure for many students. We use these figures as our baseline throughout this document — because they are the most widely accepted, most consistently sourced, and most institutionally recognized estimates available — but readers should understand that they represent the lower end of what many students actually experience.

Books and supplies, to take the more dramatic example, have become one of the fastest-growing cost components in undergraduate education, driven by a textbook publishing industry that routinely

charges **\$200–\$350** for a new proprietary textbook with semester-limited online access codes that expire at course's end and cannot be resold. A student carrying five courses in a STEM field requiring laboratory manuals, discipline-specific software licenses, and upper-division texts may spend **\$400–\$600 per semester** on course materials — nearly double the College Board's annual figure in a single term. The \$1,240 annual figure is therefore best understood as the average across all disciplines, including those — humanities, social sciences, some arts — where open-source and library-accessible texts can bring costs down substantially.

Transportation costs are equally subject to dramatic geographic variance. A student commuting from a suburban home to a state university campus in a region with limited public transit — which describes much of the American South, Midwest, and Mountain West, including large portions of New Mexico — may spend upward of **\$200 per month** on fuel alone, plus vehicle maintenance, insurance, and parking, for a total transportation cost well exceeding the \$3,760 annual average. The personal expenses component similarly varies with local cost of living, family circumstances, and the degree to which a student can rely on family support for clothing, personal care, and incidental costs. We note these variances not to undermine the College Board's figures — they remain the appropriate baseline — but to ensure that readers of this document understand that the cost of attendance estimates used throughout tend to understate real-world costs for many students in many markets.

2.4 — The Net Price: What Actually Gets Paid, and by Whom

After establishing the sticker price landscape, we must introduce the concept of net price — the cost of attendance remaining after grant aid (specifically, grants and scholarships that do not require repayment, as distinguished from loans and work-study) has been subtracted. The average undergraduate student received approximately **\$16,810** in total financial aid in 2024–25. It is critical to note, however, that this "total aid" figure includes both grants and loans — the latter of which are not aid in any economically meaningful sense, because they must be repaid with interest. When only grant aid is considered, the picture is considerably less generous.

For the public in-state student, average grant aid substantially reduces the effective annual cost of attendance. After aid, the net price for a typical in-state public university student hovers around **\$13,110** per year — or **\$52,440** over four years — depending on institutional selectivity, state grant program generosity, and household income. This is still a very large sum for most American families, representing more than 60% of a median household's pre-tax annual income committed to a single child's four-year education. More critically, the average conceals an extreme distribution: students from households earning under **\$30,000** per year may receive institutional and federal grant packages that bring their net price near zero at some institutions; students from households earning **\$75,000–\$100,000** per year — thoroughly middle-class by any definition, but above the threshold for need-based federal and institutional aid — often receive little or no grant aid and face costs approaching the sticker price.

Table 1: Annual and Four-Year Cost of Attendance by Institution Type (2024–25)

Cost Component	Public 4-Year In-State	Public 4-Year Out-of-State	Private Nonprofit 4-Year
Tuition & Fees	\$11,610	\$30,780	\$43,350
Room & Board	\$13,310	\$13,310	\$15,250
Books & Supplies	\$1,240	\$1,240	\$1,240
Transportation & Personal	\$3,760	\$3,760	\$3,760
Total Annual COA	\$29,920	\$49,090	\$63,600
Four-Year Total COA (Sticker)	\$119,680	\$196,360	\$254,400

Source: College Board, Trends in College Pricing 2024–25. All figures are weighted national averages.

Four Categories of Payer, Four Distinct Crises

The cost of American higher education is not borne by a single payer. It is distributed — unevenly, inequitably, and in ways that the published figures consistently obscure — across four distinct categories of economic actor: the federal government, state and local governments, household families, and the students themselves, many of whom work to partially offset costs that no realistic amount of labor at prevailing wages can fully cover. Each of these four payers faces a distinct and compounding crisis. Understanding each crisis individually — its origins, its current severity, its relationship to the other three — is the indispensable prerequisite for understanding why the ACER model's restructuring of costs is not merely an efficiency optimization. It is a systemic intervention in a system that has been failing all four categories of payer simultaneously, for decades.

Payer 1 — The Federal Government: Spending More, Buying Less

The federal government is the largest single payer in the American higher education system. Its total commitment to higher education aid in 2024–25 reached approximately **\$275.1 billion**, encompassing **\$38.6 billion** in Pell Grants distributed to approximately **7.1 million students**, an estimated **\$93–\$103 billion** in new federal student loan disbursements, and a range of supplemental programs including the Federal Supplemental Educational Opportunity Grant, Federal Work-Study, and campus-based aid. These are very large numbers, and they are frequently cited as evidence of federal generosity toward higher education. They are better understood as evidence of federal response to a cost structure that has grown beyond the capacity of any subsidy program to contain it.

The Pell Grant is the most important single program in this portfolio and the one whose trajectory most clearly illustrates the widening gap between federal commitment and actual affordability. The Pell Grant was created in 1972 under the Higher Education Act with a founding promise that was, in its conception, both simple and revolutionary: no American student should be barred from college by inability to pay. At its peak purchasing power in the mid-1970s, the maximum Pell Grant award

covered approximately **79 to 80 percent** of the average total cost of attendance — tuition, fees, room, and board — at a public four-year university. The equation was relatively straightforward: the federal government covered four-fifths of the cost; the student covered the remainder with modest part-time work or family contribution. The program did not promise to make college free, but it came very close to doing so for students from the lowest-income households, which was its stated intent.

"The Pell Grant once covered 80% of public university costs. By 2024–25, the maximum award covers less than one-third. The 49-percentage-point gap in purchasing power has been filled, year after year, with debt."

Over the subsequent fifty years, this equation deteriorated so completely that it is almost unrecognizable. By 2024–25, the maximum Pell Grant award of **\$7,395** covers approximately **27 to 31 percent** of the average public university cost of attendance — a figure that varies slightly depending on methodology, with some analysts placing it closer to one-quarter and others closer to one-third. The arithmetic of the loss is stark: a program that once covered four-fifths of the cost now covers less than one-third. The gap — approximately **49 to 52 percentage points** of purchasing power, eroded over half a century — has been filled, consistently and predictably, by student debt. This erosion is not primarily the result of policy neglect, though legislative inattention played a role. It is the mathematically inevitable outcome of tuition inflation that has consistently and dramatically outpaced both general inflation and Pell Grant increases across five consecutive decades. From 1980 to 2025, average published tuition at public four-year institutions increased by approximately **1,200 percent** in nominal terms; the Consumer Price Index increased by approximately **280 percent** over the same period. The Pell Grant, despite periodic legislative increases, could not keep pace with an inflation rate running three to four times the general price level.

The federal student loan commitment — the **\$93 to \$103 billion** in annual new loan disbursements — must be understood not as generosity but as structural complicity. When the Pell Grant covered

80% of costs, borrowing was optional for many students and modest for most. When it covers 27%, borrowing is near-mandatory for any student who is not from a family of significant means. The federal loan program has become the pressure valve that prevents a broken cost structure from producing outright mass exclusion from higher education — but the pressure it releases is redirected into the lives of tens of millions of individual borrowers who will carry it for years or decades. The federal government, in other words, has replaced a grant program that paid for college with a loan program that finances the debt college creates. This is not a solution to the cost problem. It is a financing mechanism for the cost problem, and it has become the primary instrument through which the federal government's higher education commitment operates.

Analytical Note — Federal Aid at Scale

The total federal student loan portfolio reached approximately

\$1.77–\$1.79 trillion

as of early 2025, held by an estimated

42.5 to 45.6 million

borrowers. This portfolio represents the accumulated consequence of fifty years of grant purchasing power erosion filled by loan growth. It is not a crisis that emerged suddenly — it is a structural outcome that was predictable and, in retrospect, predicted.

Payer 2 — State and Local Governments: The Long Retreat

State and local governments contributed approximately **\$149.2 billion** to higher education in FY2025. This investment underwrites the operational subsidy that allows public universities to charge in-state students **\$11,610** in annual tuition rather than the true cost of instruction, which typically runs **\$25,000 to \$35,000 per student-year** at a research university when faculty salaries, facilities

depreciation, administrative infrastructure, library systems, and research support costs are fully and honestly accounted for. The state subsidy is, in effect, a very large discount coupon that residents redeem at enrollment — but it is a coupon whose face value has been declining in real terms for the better part of three decades, and the mechanism by which it declines is worth understanding in detail.

State higher education appropriations are discretionary expenditures. They compete, in every annual state budget cycle, with Medicaid, K–12 public education, transportation infrastructure, corrections, and the full range of other state obligations. Of these competing claims, higher education is uniquely politically vulnerable in a recurring budget crisis, for a structural reason that legislators understand intuitively even when they cannot articulate it: the consequence of cutting a university appropriation is felt not immediately, but months later, after the legislative session has closed, when the university raises tuition to compensate for the reduced state subsidy. No legislator faces an immediate political consequence for a vote that results in higher tuition. The consequence falls on students, who are rarely organized, rarely vote in the constituencies most affected by the cut, and whose tuition bill arrives long after the appropriations vote has been forgotten.

This structural incentive for disinvestment has operated with remarkable consistency across party lines and across states for thirty years. The result has been a slow but accelerating privatization of institutions that were conceived as public goods and built on the premise that the state would bear the majority of instructional costs in exchange for a broadly educated citizenry and the economic benefits associated with it. As state subsidies fall, universities raise tuition. As tuition rises, the effective cost difference between attending a public and attending a private university narrows. As that difference narrows, the founding rationale for public higher education — that the people's government will pay most of the cost of the people's education — erodes into something that barely resembles the founding vision. The ACER model's two-year compression addresses this structural problem with a directness that no other proposed reform can match: a two-year program requires the state to subsidize two years of instruction rather than four, cutting the per-student state obligation roughly in half, without any reduction in the quality or rigor of the education the student receives.

Payer 3 — Household Families: The Impossible Arithmetic

For a median American household — earning approximately **\$80,000** per year before taxes in 2024–25, the figure the U.S. Census Bureau's most recent data supports as the national median — sending a single child to a public four-year university means committing roughly **16 to 22 percent** of gross annual income per year to education costs after average grant aid. Over four years, this represents approximately **65 to 88 percent** of a single year's pre-tax household income. For families in the \$60,000 to \$90,000 household income range — solidly working and middle class by any reasonable definition, but systematically underserved by need-based aid programs calibrated for the poor and merit-based programs calibrated for the academically exceptional — the four-year cost of a public university education may represent the largest financial commitment they will make after purchasing a home. And they make this commitment without the same collateral protections, appreciation potential, long amortization structures, and tax-deductibility features that a mortgage provides.

The arithmetic of the median-income family deserves to be worked out in full, because abstractions about "affordability" consistently obscure the actual lived reality of what these families are being asked to do. Consider a family earning **\$75,000** per year — solidly within the middle of the income distribution, neither poor nor prosperous. After federal income tax at an effective rate of approximately 12%, state income taxes averaging 4–5%, and FICA contributions of 7.65%, total withholding removes approximately 24–25% of gross income, leaving net take-home pay of roughly **\$56,000–\$57,000** per year, or approximately **\$4,700** per month. From this take-home, a mortgage or rent payment of \$1,500–\$2,000 per month represents 32–43% of net income — a standard and unavoidable commitment. Add utilities, food, transportation, insurance, and healthcare, and the realistic discretionary income remaining for savings, education, and non-essential expenditure may be **\$1,500–\$2,500** per month at most. Contributing **\$1,083 per month** (\$13,000 annually) toward a child's college education from this budget is not comfortable. It is a stretch that requires sacrifice. Contributing **\$29,920 per year** — the full public in-state sticker price — is simply not possible from current income without significant debt or asset liquidation. Many families approach this reality not through rational financial planning but through a combination of wishful thinking, Parent PLUS loans

at interest rates of 9% or more, retirement account withdrawals, and the hope — often reasonable but not always fulfilled — that the degree will generate sufficient earnings to justify the cost over a career.

"For the median American family, sending one child to a public four-year university costs the equivalent of 65 to 88 percent of a single year's pre-tax household income. For the second child, the arithmetic does not change. For the third, it becomes actively impossible."

Payer 4 — The Working Student: The Race That Is Designed to Be Lost

The working student is the most economically precarious actor in American higher education, and the one whose situation most powerfully and most viscerally illustrates the structural irrationality of the cost system we have allowed to persist. An estimated **40 to 50 percent** of American undergraduates work part-time during the academic year — in fast food, retail, food service, campus operations, childcare, and a range of other minimum-wage or near-minimum-wage employment environments. They work not to fund a richer social life or a more comfortable apartment, but because the mathematics of their college finances leave them no alternative. This section works through the arithmetic of the representative working student with deliberate and complete transparency, because the numbers are the argument.

We use **\$13.00 per hour** as our base wage assumption — not the federal minimum wage of \$7.25, which has not been increased since 2009 and which represents an outlier floor that does not reflect actual prevailing wages in most labor markets where American undergraduates actually work. The \$13.00 figure reflects the practical fast-food and service industry wage in states that have enacted minimum wage increases — lower than California's \$16.50 or New York's \$16.00 state minimums, but above the federal floor, representing a reasonable national working average for the demographic

in question. We then assume a **20-hour work week** during the 36-week academic year. This 20-hour figure is not arbitrary: it is the threshold that student affairs research consistently identifies as the point beyond which additional work hours begin producing measurable negative effects on GPA, retention, and graduation rates. We use 20 hours not because it is necessarily the actual hours worked by most working students — many work more — but because it is the economically and academically defensible upper bound. Using 20 hours is therefore both an earnings estimate and an implicit acknowledgment that working more would cost these students in academic currency what they gain in wages.

The gross academic-year earnings at these parameters are: $\$13.00 \times 20 \text{ hours} \times 36 \text{ weeks} = \$9,360$. After federal income tax at an effective rate of approximately 10–12%, state income tax averaging 3–4%, and FICA contributions of 7.65% — total withholding of approximately 21 to 24% — net take-home academic-year pay is approximately **\$7,119 to \$7,395**. Adding a twelve-week summer at 40 hours per week: $\$13.00 \times 40 \times 12 = \$6,240$ gross, net approximately **\$4,934 to \$5,050** after taxes. Total annual net earnings across academic year and summer: approximately **\$12,053 to \$12,445**. Call it \$12,200 for the purpose of analysis.

Now place those earnings against the costs this student must meet. Room and board alone at a public four-year university averages **\$13,310** per year. The working student's full annual net earnings — including summer wages, worked at the maximum academically sustainable intensity during the school year and at full-time during the summer — do not cover room and board alone. The annual shortfall on room and board, before a single dollar of tuition or books or transportation has been addressed, runs approximately **\$865 to \$1,257**. Adding books and supplies (\$1,240) and transportation and personal expenses (\$3,760) brings the annual non-tuition cost of attendance to \$18,310. Against net annual earnings of approximately \$12,200, the non-tuition shortfall is approximately **\$6,110 per year**. Adding tuition of \$11,610 not covered by Pell or other grants — which for a student earning \$12,200/year would typically be partially offset by a Pell Grant of perhaps \$6,000–\$7,000, leaving a tuition gap of \$4,610–\$5,610 — the total annual financing need for this

student approaches **\$10,000–\$12,000** above what working and the Pell Grant can cover. This gap is typically filled with loans. Every year. For four years.

The Central Finding:

A working student earning \$13/hour — a realistic, prevailing wage in most U.S. labor markets — working the maximum hours consistent with academic success, cannot cover room and board alone at a public four-year university. Before a single dollar of tuition is considered, this student is structurally guaranteed to finish the year with a financing gap. The system does not ask this student to sacrifice comfort. It asks the impossible.

We want to be clear about what this arithmetic demonstrates, because it is easy to read it as a sympathetic portrait and miss its analytical significance. The working student's situation is not a failure of individual effort or prudence. This hypothetical student has done everything they are supposed to do: they work diligently, they limit their hours to what is academically sustainable, they spend the summer working full-time, they apply for and receive Pell Grant aid. And the arithmetic still does not work. This is not a story about a student who made bad decisions. It is a story about a cost structure so badly misaligned with the earning capacity of the population it is supposed to serve that motivated, hardworking young people — people who are doing everything right — cannot mathematically succeed within it. When we model the ACER transformation's savings for the working student, we are not modeling an abstract percentage reduction in costs. We are modeling the difference between a system in which this student cannot survive without debt and one in which they can.

Table 2: Working Student Annual Economics — Current System (2024–25, Public 4-Year In-State)

Metric	Figure
Hourly wage assumed	\$13.00
Work hours/week (academic year)	20 hours
Academic-year weeks	36 weeks
Academic-year gross earnings	\$9,360
Net after taxes (est. 21–24% withholding)	\$7,119–\$7,395
Summer gross (12 weeks × 40 hrs × \$13.00)	\$6,240
Net summer earnings (after taxes)	\$4,934–\$5,050
Total annual net earnings (academic + summer)	~\$12,053–\$12,445
Annual Room & Board (public 4-yr avg.)	\$13,310
Annual Books + Transportation + Personal	\$5,000 (\$1,240 + \$3,760)
Annual Tuition (sticker, in-state)	\$11,610
Total Annual COA (sticker)	\$29,920
Annual shortfall (net earnings vs. total COA)	~\$17,720–\$17,867
Annual shortfall after max Pell Grant (~\$7,000)	~\$10,720–\$10,867

Sources: College Board Trends in College Pricing 2024–25; BLS wage data; IRS withholding tables; Federal Student Aid 2024–25 Pell Grant figures. All figures are estimates based on stated assumptions.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

SECTION 4

The KR5 ACER Cost Restructuring Proposals: The Reasoning Behind Every Assumption

The ACER model proposes two structural interventions in the economics of undergraduate education. These are not incremental optimizations — they are not attempts to make the existing four-year residential model somewhat cheaper at the margins. They are redesigns of the fundamental cost architecture. The first proposal compresses program duration from four years to two, adding a six-month employer-funded practicum as the bridge to professional life. The second proposes that 70% of ACER students will be home or locally based, eliminating on-campus room and board as a cost category for that population. Each proposal, taken individually, produces substantial savings. Taken together, they produce a transformation in the economics of undergraduate education that is unlike anything previously proposed within the mainstream of American higher education reform.

This section explains each proposal in detail — not merely what it proposes but why, not merely what it saves but how the savings were calculated, not merely what assumptions it embeds but why those particular assumptions were chosen over the alternatives. A reader who disagrees with any assumption should be able to identify it precisely here, articulate the alternative assumption they would substitute, and calculate the effect of that substitution on the model's outputs. That is the standard of transparency this document intends to meet.

4.1 — Proposal One: Compress the Four-Year Program to Two Years Plus a Six-Month Employer-Funded Practicum

The Compression Assumption — Why Two Years?

The choice of two years as the target duration for the ACER accelerated program is not arbitrary, and it requires more careful explanation than almost any other assumption in this model, because it is the assumption most likely to draw initial skepticism from thoughtful readers. The conventional argument against duration compression runs roughly as follows: four years of undergraduate education represents not merely a credential delivery mechanism but a developmental arc — students need time to mature intellectually and personally, to explore ideas across disciplines, to fail and recover and integrate, to form identities in relationship to a community of peers. This argument is not wrong. But it is importantly incomplete, because it conflates two distinct questions: how long

students need to develop as human beings, and how long institutions need to deliver the academic content and credential that constitute an undergraduate degree. These are related but not identical questions, and the ACER model's compression proposal engages the second question while explicitly preserving and prioritizing the first.

Consider where the time actually goes in a conventional four-year degree program. A standard four-year program involves approximately **120 credit hours** of academic work. At roughly three hours of in-class instruction per credit unit per week across a fifteen-week semester, a student in a conventional program receives approximately **1,800 hours** of formal classroom instruction over eight semesters. But what does research on learning efficiency tell us about those 1,800 hours? Consistently, studies of lecture-based instruction — the dominant delivery modality in American undergraduate education — find that effective pedagogical density runs at roughly **50 to 60 percent** of available instructional time. The remainder is occupied by administrative activities (syllabi review, attendance, logistical announcements), repetition of material covered in prior sessions to serve the median student who did not fully absorb it the first time, pacing delays inherent in delivering uniform instruction to a heterogeneous audience, and transitional dead-time between content blocks. These are not criticisms of individual faculty members. They are structural features of a delivery system designed for groups, not individuals.

The Two-in-a-Box model — pairing adaptive GenAI instruction calibrated to demonstrated individual mastery with human teacher mentorship focused on integration, depth, and non-linear intellectual development — fundamentally changes this equation. AI-assisted adaptive instruction does not pace to the median student. It paces to the individual student's demonstrated mastery at every point in the learning sequence, accelerating through content the student has effectively absorbed and slowing to provide additional pathways through content the student has not. Research on adaptive learning systems consistently finds efficiency gains of 30 to 60 percent in time-to-mastery for equivalent academic content — meaning that content requiring 15 weeks to deliver at lecture pace may require 9 to 11 weeks under a well-designed adaptive system. Compounding this efficiency across a full undergraduate program, and adding the human teacher's mentorship function — which

enriches rather than merely delivers — the ACER model's argument that two years is sufficient for equivalent academic content delivery is not wishful thinking. It is the conservative reading of a substantial and growing empirical literature.

We set two years as the target — rather than the 12 to 18 months that a purely efficiency-maximizing argument might support — precisely because we are not trying to maximize compression. We are trying to achieve compression while preserving the developmental and integrative functions that make an undergraduate education more than a credential. The two-year program includes time for mentorship, for reflection, for the cross-domain synthesis that is the defining intellectual feature of the KR5 polydomain approach, and for the kind of non-linear intellectual growth that cannot be accelerated without being degraded. Two years is the number at which we believe — and at which the evidence we have reviewed suggests — that the efficiency gains of the Two-in-a-Box model can be captured without sacrificing the developmental richness that makes a degree educationally meaningful. It is the conservative estimate of what is achievable, not the optimistic one.

The Practicum Assumption — Why Six Months, and Why Employer-Funded?

The six-month practicum represents the designed and explicit bridge between formal education and professional life in the ACER model. Its parameters were not chosen casually, and each deserves individual explanation. Six months was chosen as the practicum duration based on three converging lines of reasoning: first, that six months represents sufficient time for an ACER graduate to demonstrate genuine professional competency in an applied setting — not the simulated competency of a short internship, but the real operational competency that comes from sustained engagement with an actual professional environment; second, that six months represents sufficient time for an employer to form a well-evidenced hiring judgment, having observed the practicum student across a full project cycle in most knowledge-economy and professional service environments; and third, that six months is long enough to represent a meaningful talent pipeline investment for employers but short enough to remain operationally manageable for small and medium-sized enterprises that cannot absorb a full-year trainee without significant disruption.

The employer-funding assumption is grounded in a specific and fully defensible economic argument: the practicum is not an act of corporate philanthropy but a rational economic substitution of practicum investment for conventional onboarding cost. The Society for Human Resource Management's Human Capital Benchmarking Study estimates average direct recruitment costs at **\$4,683 per new hire** — covering job posting, recruiter time, interview logistics, and background verification. Total onboarding costs, including the productivity ramp-up period during which a new hire is present but not yet operating at full effectiveness, are estimated by SHRM to frequently exceed **\$28,000** in knowledge-economy roles when the value of lost productivity, supervision time, training material costs, and colleague distraction during the ramp-up period are honestly accounted for. An ACER practicum graduate is not a new hire who needs onboarding. They are a student who has been doing embedded, supervised, real-work for six months within a specific company. They know its systems, its culture, its workflows, and its people. The employer hiring them is not placing a bet on a credential — they are formalizing a relationship with a known, assessed, observed performer. The practicum wage investment of approximately **\$15,600 to \$20,800** (at \$13 to \$20/hour for 26 weeks at 40 hours/week) is not an additional cost layered on top of hiring costs. It directly replaces the typical onboarding investment at a substantially lower price.

"The practicum is not charity. It is a rational economic substitution: the employer invests \$15,600 in practicum wages rather than \$28,000 in onboarding costs — and ends the six months with a known, assessed, job-ready hire rather than a speculative one."

On the 150% tax deduction mechanism: we propose a deduction of 150% of documented practicum wages — rather than 100%, which would merely make the investment tax-neutral — because the policy goal is not to make practicum participation equivalent to conventional hiring but to make it actively financially preferable. A small business paying **\$15,600** in practicum wages receives a deductible amount of **\$23,400**. At an effective combined federal and state corporate tax rate of 21%,

KR5 Rajeev Kasturi All Rights Reserved August 12, 2026 © Page **25** of **52**

the tax benefit is **\$4,914**, making the net cost of the practicum **\$10,686**. At a 28% combined rate — reflecting many states' combined burden on small business — the net cost drops to **\$9,432**. This figure — roughly **\$9,000 to \$11,000** net — compares favorably with every alternative mechanism for securing a similarly qualified hire: substantially below typical agency recruitment fees (\$6,000–\$12,000 for a single placement), below typical first-90-day productivity gap costs in professional roles, and a fraction of total SHRM-estimated onboarding costs. The 150% figure was chosen by modeling the incentive threshold at which SME participation becomes financially preferable to conventional hiring across a range of business sizes and tax positions. The analysis converged on a range of approximately 140 to 155%, making 150% the defensible midpoint.

Table 3: Per-Student Cost Comparison — Four-Year Traditional vs. Two-Year ACER (Public In-State, 2024–25 Rates)

Cost Component	4-Year Traditional	2-Year ACER	Dollar Saving	% Saving
Tuition & Fees	\$46,440	\$23,220	\$23,220	50%
Room & Board (on-campus, residential)	\$53,240	\$26,620	\$26,620	50%
Books & Supplies	\$4,960	\$2,480	\$2,480	50%
Transportation & Personal	\$15,040	\$7,520	\$7,520	50%
Practicum cost to student	N/A	\$0 (employer-funded)	—	—
Total Student COA (residential)	\$119,680	\$59,840	\$59,840	50%
Federal aid commitment (per student, est.)	~\$16,810/yr × 4 = ~\$67,240	~\$16,810/yr × 2 = ~\$33,620	~\$33,620	50%
State/local subsidy (per student, est.)	~\$20,000/yr × 4 = ~\$80,000	~\$20,000/yr × 2 = ~\$40,000	~\$40,000	50%

Sources: College Board Trends in College Pricing 2024–25; Federal Student Aid data; NCES institutional cost estimates. ACER practicum figures reflect employer-funded model assumptions. All per-student state subsidy figures are estimates derived from instructional cost minus published tuition.

Two aspects of this table require explicit prose explanation. First, the tuition and fee savings in the ACER column reflect a simple proportional reduction — two years rather than four at the same annual rate — not an assumption that the ACER program will cost less per year to deliver. In fact, the per-year instructional cost of the Two-in-a-Box model may be modestly higher in the near term than conventional lecture delivery, because AI platform licensing, enhanced faculty-student mentorship ratios, and adaptive curriculum development represent real upfront investments. We do not project a per-year tuition reduction under ACER — we project a per-program reduction driven entirely by duration compression. This is an important methodological note: the savings are real, but they derive from time, not from cheaper instruction. Second, the practicum cost to the student is shown as zero, which reflects the employer-funded model assumption. This zero is conditional on the policy mechanism — the 150% tax deduction — being enacted and operationalized. It is an achievable zero, not a naive one.

4.2 — Proposal Two: Seventy Percent Home/Local-Based Enrollment

The Seventy Percent Assumption — How It Was Derived and Why It Is Defensible

The ACER model proposes that **70 percent** of students in the two-year accelerated program will be home-based or locally enrolled — living at home or within commuting distance of a KR5-aligned institution — fully eliminating on-campus room and board as a cost category for that population. This 70% figure is the most consequential single assumption in the entire cost model. It produces more savings, for more people, across more payer categories, than any other element of the ACER design. It therefore demands the most careful and fully disclosed explanation of its derivation.

Three distinct and independent strands of reasoning converge on the 70% figure, each sufficient individually to establish plausibility, and collectively sufficient to establish it as a reasonable central estimate rather than an aspirational projection.

The first strand is demographic: current enrollment patterns already support a commuter-majority undergraduate population. As of the most recent comprehensive data available, approximately **56**

percent of all American undergraduates already commute to campus rather than living in institutional housing. This 56% figure has been growing steadily for more than a decade, driven by the rising cost of campus housing, the increasing share of the undergraduate population enrolled at community colleges with primarily commuter populations, the growing proportion of undergraduate students who are over 25, employed, or have family caregiving responsibilities, and the gradual normalization of hybrid and online learning that reduces the practical necessity of residential proximity to campus. The ACER model's 70% home-based assumption is therefore not a radical departure from current demographic reality — it is a **14-percentage-point incremental increase** over the current commuter share, made achievable by the GenAI-enabled Two-in-a-Box model that delivers high-quality, fully interactive, personalized instruction independent of physical campus proximity.

The second strand is targeting: the ACER model is specifically designed for and targeted at the student populations that are already most likely to choose commuter enrollment when the option is genuinely available. First-generation students — for whom campus housing represents an overwhelming financial barrier — are natural home-based enrollees. Low-to-middle-income students for whom the \$13,310 annual room-and-board cost is the single largest barrier to affordable completion are natural home-based enrollees. Students from minority communities that are geographically concentrated near urban institutions are natural home-based enrollees. Students with family obligations — adult dependents, younger siblings, caregiving responsibilities — who cannot or will not relocate for education are natural home-based enrollees. These are the populations the ACER model is designed to serve, and they are the populations for whom home-based enrollment is not a second-best option but the most rational and accessible choice. The 70% assumption is therefore not a general demographic claim about all students everywhere. It is a specific claim about the demographic profile of the ACER student population, which skews strongly toward groups already predisposed to commuter enrollment.

The third strand is structural: the 30% residential assumption is deliberately retained because some students genuinely benefit from, need, or strongly prefer residential campus life in ways that are not merely preferences but reflect real educational and personal needs. Students from geographically

isolated rural communities — a significant demographic in New Mexico, Wyoming, Montana, and much of the Mountain West — may have no accessible KR5-aligned institution within commuting distance and may require relocation as a precondition of participation in any degree program. Students from difficult or unsafe home environments for whom campus housing represents access to stability and safety, not merely amenity, are not candidates for home-based enrollment — and the ACER model explicitly does not propose that they should be. The 70/30 split is designed not to mandate a single enrollment modality but to shift the economic center of gravity of the program toward home-based enrollment while preserving genuine residential access for students who need it.

What the Seventy Percent Saves — The Full Arithmetic

Room and board at the public four-year national average is **\$13,310** per year, or **\$26,620** over the two-year ACER program. This is the cost that disappears entirely for the 70% of ACER students who enroll in the home-based modality. The system-wide arithmetic is: total enrolled undergraduate students of approximately **18.4 million** (Spring 2025 figures); 70% home-based = approximately **12.88 million students**; annual room and board cost avoided per student = **\$13,310**; total annual system-wide room and board cost avoidance = **\$171.4 billion**. Over the two-year ACER program duration, this cumulates to approximately **\$342.8 billion** in total room and board cost avoidance — costs that never arise, that are never borrowed against, and that never become debt.

We deliberately use the term *cost avoidance* rather than *savings* throughout this analysis, and the distinction is analytically important. A savings implies that money was spent and then recovered — that it appeared as a line-item expenditure in someone's budget and was subsequently reduced. Cost avoidance means the expenditure never arises in the first place. It does not appear in any government budget as a recoverable surplus. It does not generate a fiscal dividend that can be reallocated to other programs. What it does is leave money in the bank accounts and wallets of approximately 12.88 million students and their families — money that is never borrowed, never earns interest for a lender, never adds to the aggregate loan portfolio, and never contributes to the \$1.79 trillion in student debt that the current system has accumulated. This is, arguably, more valuable than a government savings

of equivalent magnitude, because it prevents financial hardship at the household level where that hardship is actually experienced.

Table 4: Room and Board Cost Avoidance — System-Wide Impact of 70% Home-Based Enrollment

Metric	Figure
Total enrolled undergraduates (Spring 2025, est.)	~18.4 million
70% home-based under ACER	~12.88 million students
Annual R&B avoided per home-based student	\$13,310
Annual system-wide R&B cost avoidance	~\$171.4 billion
2-year program total R&B cost avoidance	~\$342.8 billion
Per-student saving over 2-year ACER program	\$26,620
Est. reduction in annual student loan need (system-wide)	~\$100–\$120 billion/year (est.)

Sources: NCES enrollment data; College Board R&B figures 2024–25; KR5 Economics Working Group modeling. Loan reduction estimate assumes 70–80% of R&B costs not covered by grants would otherwise be financed by federal student loans.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

SECTION 5

When Both Proposals Work Together: The Compounding Effect

The full economic power of the ACER model becomes visible only when its two principal proposals are examined not in isolation but in combination — applied simultaneously to the same student population, producing savings that reinforce and amplify each other in ways that neither proposal achieves independently. This compounding effect is not mathematical compounding in the

exponential sense; it is economic compounding in the structural sense, meaning that each proposal removes a cost category that the other proposal cannot address, so that together they eliminate a far larger share of total program cost than either removes alone.

To understand the compounding logic, consider three distinct student scenarios under the two proposals applied individually versus jointly. For an on-campus residential student in a two-year ACER program — benefiting from proposal one (duration compression) but not proposal two (home-based enrollment) — the savings against the traditional four-year on-campus baseline of \$119,680 are exactly 50%: the student pays \$59,840 over two years instead of \$119,680 over four. This is a substantial saving — \$59,840 in total COA reduction — but room and board remains a significant component of the two-year cost at \$26,620. For a home-based student in a traditional four-year program — benefiting from proposal two but not proposal one — the savings over the on-campus four-year baseline consist of four years of room and board avoidance at \$13,310 per year, or **\$53,240**. This is also a substantial saving, but four years of tuition, books, and transportation still accumulate to **\$66,440** (\$46,440 + \$4,960 + \$15,040).

But for the home-based student in the two-year ACER program — benefiting simultaneously from both proposals — the arithmetic changes entirely. This student pays only two years of tuition (\$23,220), two years of books (\$2,480), and two years of transportation and personal expenses (\$7,520) — with no room and board at all — for a total two-year program cost at sticker price of **\$33,220**. Against the four-year on-campus sticker-price baseline of **\$119,680**, this represents a reduction of **\$86,460**, or approximately **72 percent**. The traditional on-campus four-year program costs \$119,680. The ACER home-based two-year program costs \$33,220. That is not a marginal improvement. That is a structural transformation of the economic reality of undergraduate education.

"The traditional four-year on-campus public university education costs \$119,680 at sticker price. The ACER home-based two-year program costs \$33,220. The difference — \$86,460, or 72% — is not a policy adjustment. It is a redesign of the possible."

For the Working Student — The Most Dramatic Transformation

Returning to the working student from Section 3, the compounding effect of the two ACER proposals produces what is, from an analytical standpoint, the most striking transformation in the entire model. Under the traditional four-year on-campus system, this student's annual net earnings of approximately \$12,200 cannot cover room and board (\$13,310) alone — they fall short by \$865 to \$1,257 before a dollar of tuition is considered. The student is structurally guaranteed to accumulate debt. Under the ACER home-based model, room and board disappears as a cost category entirely. The student's annual cost of attendance falls to tuition (\$11,610) plus books (\$1,240) plus transportation and personal expenses (\$3,760), totaling **\$16,610 per year**, or **\$33,220** over two years at sticker price. Against this, the student's Pell Grant eligibility — for a student earning approximately \$12,200 in annual income, eligibility would typically fall in the **\$6,000 to \$7,000** annual range — reduces the annual net cost to approximately **\$9,610 to \$10,610**. Against annual net earnings of \$12,200, this is the first moment in this entire analysis where the arithmetic becomes possible. For the first time, a working student earning fast-food wages, applying for the Pell Grant, and limiting work hours to the academically sustainable maximum can approach a debt-free or near-debt-free completion. The margin is not wide — it requires discipline and consistency — but it exists. Under the current system, it does not exist at all.

For the Federal Government

If the ACER model reduces the duration of federal student aid eligibility from four years to two years per student — and if the per-year federal aid commitment remains constant across that shorter

period — the per-student federal aid outlay is halved. The **\$275.1 billion** in total annual federal higher education aid, applied to a uniformly two-year program at equivalent per-year intensity, would yield approximately **\$137.5 billion** in annual federal savings. This is the theoretical maximum, conditional on full ACER adoption across the system and on the assumption that per-year federal aid costs do not increase to compensate for reduced duration — neither of which is realistic in a transition scenario. A more conservative modeling assumption — 60% adoption over a five-year transition, with per-year aid rates unchanged — would yield approximately **\$82 to \$90 billion** in annual federal savings at scale. This range represents one of the most significant potential reductions in federal higher education expenditure ever modeled, and it is achievable not by cutting benefits per student but by reducing the number of years for which benefits must be paid, because the educational program genuinely takes only two years to complete.

For State and Local Governments

The state subsidy per student — the gap between true per-student instructional cost and the published in-state tuition at public universities — is similarly halved when program duration is halved. Using a conservative estimate of approximately **\$20,000** per student-year in average effective state subsidy (derived from the widely cited differential between true instructional cost and published tuition at flagship public universities), a two-year program reduces per-student state commitment from approximately **\$80,000** over four years to approximately **\$40,000** over two years. Applied against the **\$149.2 billion** in current annual state and local higher education appropriations, full ACER adoption would theoretically yield approximately **\$74.6 billion** in annual state savings. As with the federal figures, transition costs, continued support for graduate and professional programs, and the capital requirements of institutional restructuring would moderate this figure in practice — but the directional magnitude is not in dispute.

Table 5: Combined ACER Savings — All Payer Categories (Full Adoption Scenario)

Payer Category	Current 4-Year Annual Burden (Est.)	ACER 2-Year + Home-Based Annual Equivalent	Annual Saving (Est.)	% Reduction
Federal Government	\$275.1 billion	~\$137.6 billion	~\$137.5 billion	~50%
State & Local Governments	\$149.2 billion	~\$74.6 billion	~\$74.6 billion	~50%
Students & Families — R&B (cost avoidance, home-based)	~\$171.4 billion/year (system-wide)	\$0 (home-based)	~\$171.4 billion	~100%
Working Student — individual annual gap	~\$10,720–\$10,867	~\$0–\$2,410	~\$8,310–\$10,867	~78–100%
Total System-Wide Annual Economic Benefit (Est.)	—	—	~\$383 billion/year	—

Sources: KR5 Economics Working Group modeling from College Board, Federal Student Aid, NCES, and U.S. Department of Education FY2025 Budget data. All figures represent full-adoption estimates and should be interpreted as ceiling values. Transition scenarios produce moderated figures as described in Section 8.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

SECTION 6

The Practicum as Economic Infrastructure: Why Employers Should Invest and Why They Will

The employer practicum is the third economic leg of the ACER model, and it is the component most vulnerable to the assumption that it depends on employer goodwill or corporate social responsibility — a fragile and inherently unreliable foundation for an economic model. This assumption is incorrect,

and it is important to establish precisely why. The ACER practicum model does not ask employers to be generous. It offers them a transaction that is, by any reasonable calculation of business economics, more advantageous than the alternatives currently available to them for securing qualified, job-ready early-career talent. This section makes that case, works through the full arithmetic, and then — with equal transparency — identifies the real risks and limitations of the practicum model that must be addressed if it is to function at scale.

6.1 — The Onboarding Cost Replacement Argument

The fundamental economic argument for employer practicum participation is not complicated, though it requires some care in its arithmetic to make fully compelling. American employers currently spend very large amounts of money to recruit and onboard conventional new hires. The Society for Human Resource Management estimates direct recruitment costs — job posting fees, recruiter time, interview logistics, assessment tools, and background verification — at approximately **\$4,683 per new hire** in direct costs alone. This figure is the floor: it covers only the explicit out-of-pocket recruitment expenditure and does not include the cost of the hiring manager's time, the opportunity cost of vacant positions during the search period, or any element of the onboarding and ramp-up investment that follows the hire.

Total onboarding cost — a more honest measure of what a new hire actually costs before they reach full productivity — is substantially larger. SHRM's comprehensive benchmarking data places total onboarding costs in knowledge-economy and professional service roles frequently in excess of **\$28,000** when the full picture is counted: initial training and orientation, the productivity gap during the first 60 to 90 days when the new hire is physically present but not yet operating at speed, the time colleagues and supervisors spend answering questions and providing direction that would not be required for an experienced team member, and the cost of errors and suboptimal decisions made during the learning curve period. None of these costs appear on a hiring budget as a clean line item; they are embedded in the operating costs of every team that absorbs a new hire, and they are real.

An ACER practicum graduate arrives at the practicum employer genuinely onboarded before their formal employment begins. What does "genuinely onboarded" mean in this context, and why does it matter economically? It means that after six months of embedded, supervised, structured real-work engagement, the practicum student knows the company's systems, has navigated its internal workflows, has built working relationships with colleagues and supervisors whose names and roles are familiar, has encountered and learned from the company's specific operational challenges, and has been assessed — not on academic performance alone but on actual professional performance in actual professional situations — by the people who will, if a hiring offer is extended, become their permanent colleagues and managers. The employer making a post-practicum hiring decision is not placing a speculative bet on a credential from an institution they may not know well. They are converting a six-month relationship with a known performer into a permanent employment agreement. The information asymmetry that makes conventional hiring risky — neither employer nor candidate can know before the hire whether the candidate will actually perform — is substantially eliminated.

6.2 — The Talent Pipeline Argument

Beyond the onboarding cost argument — which applies to any individual employer's decision about any individual hire — the practicum model offers something the conventional labor market cannot provide: structural, predictable, co-designed access to a talent pipeline. For industries facing chronic, systemic talent shortages — healthcare support services, advanced manufacturing, logistics and supply chain management, technology services, early childhood education, public administration, and environmental services, among others — the conventional labor market operates as a competition for a fixed and inadequate supply of qualified candidates. Employers in these sectors routinely leave positions unfilled for months, absorbing the costs of reduced capacity, overtime for existing staff, service degradation, and customer attrition that accompany vacancy at scale.

Employers who participate in KR5 ACER practicum partnerships gain something categorically different from what the open labor market provides: first-access to graduates whose multidomain

competencies have been shaped, at least in part, by the employer's own participation in the practicum curriculum design. An employer who engages with a KR5-aligned institution's practicum planning process — specifying the competencies it most needs, the operational challenges it most wants practicum students to engage with, the professional culture it is trying to build — ends the practicum period with a cohort of graduates who are not merely generically qualified but specifically prepared for that employer's environment. This is a level of talent development customization that is simply not available through any conventional recruitment channel, and its value is particularly high in sectors where generic credentials fail to capture the firm-specific or domain-specific competencies that actually drive performance.

6.3 — The SME Accessibility Analysis with Full Arithmetic

The concern most frequently raised about employer-funded practica at scale is that they will inevitably favor large employers — corporations with dedicated HR departments, established internship program infrastructure, and the administrative capacity to manage a structured practicum relationship — while remaining inaccessible to the small and medium-sized enterprises that employ the majority of Americans and that represent the most geographically diverse slice of the employer landscape. The 150% tax deduction mechanism is specifically designed to address this concern, and it is worth working through the arithmetic in detail for a realistic SME scenario.

Consider a small professional services firm — an accounting practice, a marketing agency, a small medical office, a local engineering consultancy — with 15 to 25 employees and annual revenues of \$2 to \$5 million. This firm pays a practicum student **\$15.00 per hour** for 26 weeks at 40 hours per week: gross practicum wage cost = **\$15,600**. Under the 150% deduction proposal, the deductible amount = **\$23,400**. At a combined federal and state effective tax rate of 28% — a reasonable estimate for a profitable SME in most states, combining the federal 21% rate with a typical state corporate rate — the tax benefit is **\$6,552**. The net cost of the practicum to this firm is therefore **\$15,600 minus \$6,552 = \$9,048**. For less than \$10,000 net, this firm has had a competent, supervised practicum student for six months, has assessed that person's professional fit for its specific environment, has

either prepared for a highly informed hiring decision or made a contribution to the broader talent pipeline, and has done so at a cost below a single agency recruiter placement fee (typically \$6,000–\$12,000 in direct fees), well below the cost of the productivity gap associated with a conventional new hire (\$8,000–\$15,000 for a professional role), and at roughly one-third of SHRM's estimated total onboarding cost.

6.4 — Risks and Limitations of the Practicum Model

This document would not meet its own standard of intellectual honesty if it presented the employer practicum model without identifying the genuine risks and limitations that must be addressed before it can function reliably at scale. We identify four primary risk categories.

Administrative capacity constraints among small businesses. Many SMEs lack the administrative infrastructure to manage a structured practicum relationship without institutional support. A small business owner running a five-person operation simply does not have the bandwidth to process practicum agreements, track competency assessments, maintain the documentation required for the tax deduction, and coordinate with institutional advisors — all while running a business. This argues strongly for a KR5-aligned placement and support office that assumes the administrative burden of the practicum on behalf of participating employers: maintaining standard agreement templates, managing the assessment documentation, serving as the intermediary between the student, the institution, and the employer, and ensuring that the tax deduction paperwork is complete and compliant. The practicum model will not scale to SME participation without this institutional infrastructure.

Industry-specific regulatory constraints. Not all industries are equally amenable to six-month structured practica. Highly regulated sectors — clinical healthcare, financial services involving licensed activities, legal services, aviation, and certain construction and engineering specialties — have licensing requirements, liability frameworks, and scope-of-practice regulations that significantly constrain what a practicum student without professional licensure can do in an operational role. The ACER model's practicum design must be modular and sector-specific, with carefully differentiated

engagement frameworks for regulated and non-regulated environments. A practicum student in a hospital system is not the same as a practicum student in a retail business, and the model cannot treat them identically.

Geographic mismatch. A student in rural New Mexico — to take an example that is geographically immediate to this document's drafting location — may not have access within reasonable commuting distance to employers with the sophistication, the staffing capacity, or the industry alignment to offer a productive ACER practicum experience. This is a real and structurally important gap that requires two parallel responses: first, the development of robust remote and hybrid practicum frameworks for industries where professional work can be meaningfully conducted at a distance; and second, a practicum mobility fund — potentially funded from a portion of the federal and state savings generated by the ACER model — that enables students in geographically underserved areas to temporarily relocate for practicum placements in regional or urban markets.

Quality variance across practicum sites. A six-month practicum at a well-organized, mentorship-oriented employer who takes their educational partnership responsibility seriously is a transformative professional development experience. A six-month practicum at an employer who views the practicum student primarily as cheap labor and provides minimal structured supervision is something considerably less valuable. Quality assurance across a large and geographically distributed practicum network requires robust institutional oversight — regular check-in protocols, student feedback mechanisms, practicum site assessments, and the institutional willingness to remove a site from the network when quality standards are not met. This is not a trivial administrative investment, but it is essential if the practicum is to deliver on its educational and economic promises.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

SECTION 7

Why the ACER Model Is the Only Debt Solution That Actually Works

The student debt crisis is, at this point, so thoroughly documented and so widely discussed that further description of its dimensions risks seeming redundant. The aggregate figures are familiar: approximately **\$1.77 to \$1.79 trillion** in outstanding federal student loans held by approximately **42.5 to 45.6 million** borrowers as of early 2025. The Class of 2023 graduated with an average of approximately **\$29,300** in cumulative federal student loan debt per borrower — a figure that understates the reality for students at private institutions, graduate and professional programs, and students who attended but did not complete their degrees, who represent some of the most financially distressed borrowers in the entire portfolio. Average monthly loan payments for bachelor's degree holders in standard repayment run approximately **\$523** per month — a figure that, for a new graduate earning \$45,000 to \$55,000 annually (the approximate median entry-level salary range for recent bachelor's graduates across fields), represents **11 to 14 percent** of gross monthly income, or roughly **15 to 18 percent** of net take-home pay. This is a significant and sustained constraint on the financial capacity of a generation of young adults — constraining housing access, delaying family formation, suppressing retirement savings, and reducing the consumer spending that the broader economy requires for growth.

But the scale of the crisis is less analytically interesting than its cause, because the cause determines what kind of response can actually resolve it rather than merely defer it. The persistent misframing of the student debt crisis — in political discourse, in media coverage, and in policy debate — is that it reflects excessive risk-taking by individual students who borrowed more than their degrees could justify. This framing implies that the solution is better individual decision-making: students should choose cheaper schools, borrow less, pick higher-earning majors, or attend community college rather than four-year universities. The framing is not supported by the data. The overwhelming majority of the **\$1.79 trillion** portfolio is held by students who made individually rational decisions: they enrolled in regionally accredited institutions that the federal government had approved for student aid

eligibility, they borrowed within federal loan limits set by the same government that approved their institutions, they chose programs in fields with reasonable expected labor market returns, and they still graduated with debt that the economy's wage structure — particularly for early-career workers without prior experience — could not comfortably service. The debt crisis is not primarily a story of individual imprudence. It is the story of a cost structure that has inflated faster than wages for fifty consecutive years, combined with a grant system whose purchasing power has eroded so severely that debt has become the default financing mechanism for higher education access.

"Loan forgiveness is relief without prevention. It addresses the symptom — existing debt — without touching the cause: a cost structure that will regenerate equivalent debt within a generation. The ACER model is the only intervention that changes the cause."

Loan forgiveness programs — whatever their merits as relief measures for current borrowers who made educational decisions under the cost conditions that prevailed at the time they enrolled — cannot solve the debt crisis because they do not change the cost structure that created it. A student who begins a four-year public university program the day after a comprehensive forgiveness program is enacted faces exactly the same published tuition, the same room and board costs, the same Pell Grant purchasing power gap, and the same working student arithmetic as the student who graduated yesterday with \$29,300 in debt. The system that produces debt at a particular rate — driven by the difference between what education costs and what Pell Grants, family income, and student labor can collectively cover — will continue producing debt at that rate regardless of what happens to past debt. Without structural cost reform, within twenty years the system will have regenerated a debt portfolio of comparable magnitude to whatever was forgiven. This is not an argument against forgiveness as a measure of equity and relief for current borrowers. It is an argument that forgiveness without structural reform is debt relief without debt prevention — treating the fever without curing the infection.

The ACER model is different because it operates at the level of the structural cause. It addresses debt not by forgiving what has been borrowed but by making the conditions under which future debt accumulates structurally less severe. It does so through four simultaneous mechanisms. First, it halves program duration, halving the total tuition liability that any student accumulates. Second, it eliminates room and board as a cost category for 70% of students, removing the largest single non-tuition cost from the equation for the majority of the student population. Third, it produces graduates two years earlier — at approximately 22 rather than 24 years of age — giving them two additional years of full professional earnings to build financial stability, savings, and asset acquisition before the median age at which traditional graduates are entering the workforce. Fourth, the employer-funded practicum ensures that the transition from education to employment is seamless and financially supported: the practicum student is paid for the bridge period, eliminating the period of unemployment or underemployment between graduation and first professional employment that often forces new graduates into additional borrowing or into financial dependence at a moment of acute vulnerability.

The ACER debt outcome for the representative home-based working student — working 20 hours per week at \$13/hour during the academic year, qualifying for near-maximum Pell Grant at approximately \$7,000 annually — can be modeled as follows. Total two-year sticker-price COA: **\$33,220** ($\$16,610 \times 2$). Less Pell Grant over two years: approximately **\$14,000** ($\$7,000 \times 2$). Remaining COA after Pell: approximately **\$19,220**. Less net earnings contribution over two academic years (conservatively $\$7,200 \times 2 = \$14,400$, net of taxes and keeping some for living expenses): approximately **\$14,400**. Residual financing need: approximately **\$4,820** over two years — or approximately **\$2,410 per year**. This is the debt exposure of the ACER graduate under the home-based model. Compare it to the average current graduate's **\$29,300** in federal loan debt, to say nothing of the **\$50,000 or more** that private university graduates frequently carry. The difference is not a marginal improvement achievable through better financial aid advising. It is a structural transformation of who can realistically complete a college education without incurring life-altering debt — and who cannot.

SECTION 8

What the Transformation Costs, and Why It Is Still Worth It

An economic analysis that presents only the savings and benefits of a proposed transformation without an honest accounting of its transition costs is advocacy, not analysis. The ACER model generates very large potential savings across every payer category in the system. It also requires very real investments — in infrastructure, in human capital development, and in institutional revenue adjustment — that must be estimated honestly, even where the estimates are inherently uncertain, and weighed against the benefits with the same rigor applied to the savings side of the ledger.

First transition cost: GenAI infrastructure investment. The Two-in-a-Box model requires significant upfront investment in GenAI platform deployment, institutional data privacy infrastructure, faculty-facing training and interaction interfaces, and integration with existing learning management systems across a large and heterogeneous institutional landscape. Current market estimates for institutional AI platform deployment at meaningful scale range from approximately **\$8,000 to \$15,000** per institution in Year 1, with marginal costs declining rapidly as shared platform infrastructure matures and as competition among AI platform vendors drives unit pricing down. Applied across the approximately **3,896** degree-granting institutions identified in NCES data, full deployment at the upper-bound \$15,000 per institution produces a one-time platform investment of approximately **\$58.4 million**. This figure assumes — and this assumption is important — a shared national ACER platform architecture in which a common system is developed once and licensed to member institutions rather than each institution building its own proprietary solution. This model is both substantially more cost-efficient and pedagogically superior: a shared platform allows adaptive

KR5 Rajeev Kasturi All Rights Reserved August 12, 2026 © Page **43** of **52**

learning models to improve continuously from cross-institutional learning data, producing better individual student outcomes over time while keeping per-institution costs low. Against the \$383 billion in annual system-wide economic benefit that full ACER adoption generates, a \$58.4 million one-time platform investment is not a meaningful cost barrier. It is a rounding error.

Second transition cost: faculty retraining and professional development. Transitioning faculty from the lecture-centered instructional model that dominates American higher education to the Two-in-a-Box mentorship and facilitation model that the ACER program requires is not a trivial professional development exercise. It requires fundamental shifts in pedagogical orientation — from knowledge transmission to learning facilitation, from content delivery to competency mentorship, from group pacing to individualized guidance. Research on comparable faculty development initiatives suggests that approximately **40 to 80 hours** of structured professional development per faculty member is required for a genuine rather than superficial transition. The approximately **1.5 million** full-time faculty at degree-granting institutions in the United States cannot be retrained simultaneously — the ACER model envisions a phased rollout that prioritizes institutions and faculty most aligned with the program's objectives. Targeting **200,000** faculty in Year 1 at an average training cost of **\$2,000** per faculty member — covering professional development programming, curriculum redesign support, and ongoing coaching — produces a Year 1 faculty development investment of **\$400 million**. This is a significant figure, but it is one-time and declining: faculty trained in Year 1 require only ongoing light-touch updates in subsequent years rather than full retraining, and the per-faculty cost declines as institutional expertise in facilitating the transition accumulates.

Third transition cost: institutional revenue adjustment. This is the most complex, most politically sensitive, and most institutionally significant transition cost in the ACER model, and it is the one where honest acknowledgment of uncertainty is most important. Residential colleges and universities — particularly those whose revenue models are substantially dependent on room and board income — currently derive meaningful net revenue from their housing and dining operations. Many institutions run at a surplus of **\$2,000 to \$4,000 per residential student per year** after actual housing and dining costs, using this surplus to cross-subsidize academic programs, financial aid, and

operational infrastructure. A 70% shift to home-based enrollment would substantially reduce this revenue stream — not immediately or uniformly, because adoption will be uneven across institution types and geographies, but meaningfully and over time. Institutions heavily dependent on residential revenue will need transition support as they restructure their facilities — converting underutilized residential buildings to academic, research, or community use — and develop new revenue models appropriate to a predominantly commuter population. A dedicated ACER transition fund, drawn from a portion of the federal and state savings the model generates, would need to support this institutional revenue adjustment for a period of approximately five to seven years, providing a bridge during which institutions can restructure without facing the acute financial distress that would otherwise accompany a rapid enrollment-model shift. The precise scale of this fund requires institutional-level financial modeling that this document does not attempt, but a reasonable initial estimate — based on the surplus revenue at risk and the five-to-seven-year bridge period — would range from **\$5 to \$15 billion** in total transition support spread across the adoption period.

Net Present Value Summary

The annual system-wide economic benefit of full ACER adoption approaches

\$383 billion per year

, across federal government, state governments, and student/family cost avoidance. One-time and transitional costs are unlikely to exceed

\$10–\$15 billion in any single transition year

. Even at conservative adoption rates and substantial discount rates, the net present value of the ACER transformation is strongly positive within a three-to-five-year horizon. This represents the largest potential economic efficiency gain in the history of American higher education — achievable while improving educational quality, not reducing it.

The net present value argument for the ACER transformation is, even under conservative assumptions, straightforward. Annual system-wide benefit at full adoption: approximately \$383 billion. Total transition costs across the adoption period: unlikely to exceed \$10 to \$15 billion in any single year, or \$50 to \$75 billion cumulative over a five-year transition. The payback period for the full transition investment, measured against the annual benefits, is well under one year at full adoption. Even at a 60% adoption rate — a reasonable five-year target for a transformation of this scale — annual benefits of approximately \$230 billion against cumulative transition costs of \$50 to \$75 billion produce an NPV that is strongly positive within three years at any discount rate below 30%. This is not a marginal benefit that requires heroic assumptions to justify. It is one of the most economically compelling cases for educational investment ever assembled, and it requires only the political will to act on it.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

CLOSING REFLECTIONS

Numbers Are Not Neutral — What the Economics of ACER Are Really Saying

It is tempting, at the conclusion of a document of this length and analytical density, to summarize the numbers — to present a final table of savings and costs, to restate the percentage reductions and the dollar figures, to leave the reader with the quantitative takeaways in sharp relief. We will resist that temptation, not because the numbers are unimportant — they are the entire point of this document — but because the numbers are not neutral, and ending with them without naming what they mean would be a failure of the intellectual responsibility this document has set for itself from its opening pages.

The working student at \$13.00 per hour who cannot cover room and board alone is not an economic abstraction. She is a young woman in Albuquerque whose family cannot contribute to her education and who has committed herself, against significant odds, to the project of changing her economic circumstances through a degree that the arithmetic of the current system has quietly but conclusively arranged for her to fail to complete without debt she cannot afford to carry. The family in the financial aid desert, earning \$78,000 per year, approaching retirement age with a college savings account drained by two semesters of a child's undergraduate education, is not a demographic category. They are people who did everything right — they saved, they earned, they believed in the promise — and who are now discovering that the promise did not extend to them. The forty-two and a half million Americans carrying a combined \$1.79 trillion in student loan debt are not a macroeconomic variable. They are individuals making daily decisions — about whether to rent or buy, whether to start a family, whether to save for retirement — constrained by a monthly payment that represents a contract they signed at eighteen, with information they did not have, under conditions they could not fully assess.

"The ACER model's economics are not an efficiency argument for its own sake. They are the evidence that a more just educational system is also a less expensive one — that equity and economy are, in this case, the same project pursued by the same means."

The economics of the ACER transformation, as this document has argued at length and with full transparency, are not ambiguous. The model is cheaper for every payer. It is cheaper for the federal government by tens of billions of dollars per year. It is cheaper for state governments by comparable magnitudes. It is dramatically cheaper for families, and most dramatically cheaper — potentially transformatively cheaper — for the working students and first-generation strivers who are the moral core of the case for educational reform. The arithmetic works, and this document has shown it working, in detail, with every assumption named and defended.

The question this document therefore leaves with its readers is not whether the ACER model is economically viable. The evidence is clear that it is. The question is whether the institutions, the governments, the accreditation bodies, the textbook publishers, the residential housing operations, and the full ecosystem of economic actors who currently benefit from a cost structure that is simultaneously impoverishing students and their families will permit the transformation to occur. That is not an economic question. It is a political and moral question, and it involves the exercise of power — specifically, the power to choose between a system that works for the people who control it and a system that works for the people who need it.

That question — the political economy of transformation, the institutional resistance to structural change, and the mechanisms by which the KR5 framework proposes to navigate and ultimately overcome that resistance — is the subject of the next volume.

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution

APPENDIX A

Data Sources and Methodological Notes

All primary sources used in the construction of this document are listed below. Where figures from different sources showed minor variance, this document used the most conservative figure and noted the range. No figure has been fabricated or extrapolated beyond the defensible range of the underlying data.

Table A-1: Primary Data Sources

Source	Data Used	Citation Basis
College Board, <i>Trends in College Pricing 2024–25</i>	Tuition, R&B, books, transportation, total COA by institution type	All Section 2 pricing figures
College Board, <i>Trends in Student Aid 2024–25</i>	Pell Grant totals, average student aid, aid by type	Sections 2.4, 3, 5, 7
Pell Institute, <i>Indicators of Higher Education Equity 2024</i>	Pell purchasing power historical data (79–80% in 1975; 27–31% in 2024–25)	Section 3 (Payer 1)
Federal Student Aid, <i>FY2025 Q1 Portfolio Summary</i>	Total loan portfolio (\$1.77–1.79T), borrower count (42.5–45.6M)	Sections 3, 7
U.S. Department of Education, <i>FY2025 Budget Justification</i>	Federal higher education appropriations; state/local figures	Sections 3, 5
SHRM, <i>Human Capital Benchmarking Study</i> (most recent edition)	Average recruitment cost per hire (\$4,683); total onboarding cost estimates (\$28,000+)	Sections 4.1, 6.1, 6.3
NCES, <i>Digest of Education Statistics 2022–23</i>	Number of degree-granting institutions (3,896); full-time faculty (1.5M)	Section 8
NCES enrollment data (Spring 2025)	Total undergraduate enrollment (~18.4 million)	Section 4.2, Table 4
U.S. Census Bureau, <i>American Community Survey 2023</i>	Median household income (~\$80,000); family income distribution	Section 3 (Payer 3)
Bureau of Labor Statistics, <i>May 2025 OES / CPI Data</i>	Wage benchmarks; inflation comparison (CPI 1980–2025)	Sections 3, 4, 6
LendEDU / MeasureOne / Federal Student Aid	Average debt per borrower, Class of 2023 (\$29,300); average monthly payment (\$523)	Section 7
IRS Publication 15-T (2024–25 withholding tables)	Federal income tax withholding calculations for working student model	Section 3 (Payer 4), Table 2

All figures are as of the most current available data at time of drafting (August 2026). Where figures may have been updated after the referenced publication date, the Working Group has used the most recently available comparable data.

APPENDIX B

Glossary of ACER Economic Terms

The following terms appear throughout this document and are defined here both for the convenience of new readers and to ensure that recurring terms are used with consistent meaning across all KR5 internal circulation documents.

Term	Definition
ACER	Accelerated Competency-based Educational Reimagination. The KR5 two-year polydomain accelerated undergraduate program model that combines GenAI adaptive instruction with human teacher mentorship, home-based enrollment, and an employer-funded six-month practicum.
Cost of Attendance (COA)	The total annual cost of attending a college or university as defined by federal financial aid regulations, including tuition and fees, room and board, books and supplies, transportation, and personal expenses. COA is the basis for determining financial aid eligibility.
Sticker Price	The published, pre-aid cost of attendance as set by an institution — the COA before any grant aid, scholarships, or other discounts are applied. Sticker prices set psychological and legal cost ceilings even when most students do not pay them in full.
Net Price	The cost of attendance remaining after all grant aid (grants and scholarships that do not require repayment) has been subtracted. Net price is the most accurate measure of what a specific student or family actually pays, though it varies significantly by income, institution, and state.
Pell Grant Purchasing Power	The percentage of average public four-year university total cost of attendance covered by the maximum Pell Grant award in a given year. In 1975, this was approximately 79–80%. In 2024–25, it is approximately 27–31%. The decline in Pell purchasing power is the primary driver of the growth in student borrowing.
Cost Avoidance vs. Savings	A cost avoidance is a reduction in expenditures that never arise — money that is never spent and therefore never needs to be recovered. A saving is a reduction in expenditures that were previously occurring. In this document, the elimination of room and board costs for home-based students is a cost avoidance, not a savings: the cost never arises and therefore does not appear as a recoverable surplus in any budget.

Term	Definition
Two-in-a-Box Model	The KR5 ACER instructional model pairing an adaptive GenAI learning system (which delivers personalized, mastery-paced academic content) with a human teacher mentor (who facilitates integration, depth, non-linear intellectual development, and relational educational support). The two components are designed to be complementary rather than competitive.
Practicum	The six-month employer-funded structured professional experience that constitutes the bridge between the two-year ACER academic program and formal professional employment. Practicum students are paid by the employer at prevailing wages, are subject to structured competency assessment, and are engaged in real operational work rather than simulated or observational activities.
Home-Based Enrollment	The ACER enrollment modality in which students attend a KR5-aligned institution while living at home or within commuting distance, fully eliminating on-campus room and board as a cost category. Home-based enrollment is the proposed modality for 70% of ACER students.
Net Present Value (NPV) of Transformation	The present-dollar value of the stream of future economic benefits generated by ACER adoption, discounted at an appropriate rate to account for the time value of money, minus the present-dollar value of the stream of transition costs. A positive NPV indicates that the transformation generates more economic value than it costs, even accounting for the timing of costs and benefits.
150% Practicum Tax Deduction	The proposed federal tax mechanism under which employers who document qualifying ACER practicum wages may deduct 150% of those wages from taxable income — creating an active financial incentive (rather than merely a tax-neutral arrangement) for employer practicum participation, particularly among small and medium-sized enterprises.
Financial Aid Desert	The household income band (approximately \$60,000–\$100,000 annually) in which families earn too much to qualify for substantial need-based federal and institutional grant aid, but too little to pay near-sticker-price college costs from discretionary income. Families in this band disproportionately bear near-full sticker-price costs and represent a significant share of Parent PLUS loan borrowers.

End of Document

This document contains **8 sections**, **5 data tables**, and **2 appendices**.

Drafted: August 2026 | KR5 Initiative Research Team — Economics Working Group

Location of Drafting: Albuquerque, NM, United States

Internal — KR5 Initiative | Appendage to Volume 1 | August 2026 | Not for Public Distribution